

ZEMAX[®]

Optical Design Program

User's Manual

July 8, 2011



Radiant ZEMAX LLC
support@zemax.com
www.zemax.com

Important notice

ZEMAX® is a registered trademark of Radiant ZEMAX LLC Copyright © Radiant ZEMAX LLC 1990-2011. All rights reserved.

All other product names or trademarks are property of their respective owners.

Information in this manual is subject to change without notice and does not represent a commitment on the part of the vendor. The software described in this manual is furnished under a license agreement and may be used or copied only in accordance with the terms of the agreement.

Radiant ZEMAX LLC (RZ) provides this publication "as is" without warranty of any kind, either express or implied, including but not limited to the implied warranties or conditions of merchantability or fitness for a particular purpose. In no event shall RZ be liable for any loss of profits, loss of business, loss of use or data, interruption of business, or for indirect, special, incidental, or consequential damages of any kind, even if RZ has been advised of the possibility of such damages arising from any defect or error in this publication or in the Software.

TABLE OF CONTENTS

INTRODUCTION

Chapter 1

About this document	37
What does ZEMAX do?	37
What doesn't ZEMAX do?.....	37
Learning to use ZEMAX.....	38
System requirements	38
Multiple processor computers	38
Installation procedure.....	38
Policy on the key.....	38
Definition of support.....	38
Getting technical support	39
Policy on bug fixes	39

USER INTERFACE

Chapter 2

Introduction	41
Types of windows	41
Main window operations	41
Editor windows operations	42
Graphic windows operations.....	42
Using the annotation feature	44
Using pan and zoom	45
Text windows operations	45
Dialog operations	46
Aborting long computations	46
Summary of useful shortcuts	47
Using the Windows clipboard.....	48
ZEMAX file types by extension	48

CONVENTIONS AND DEFINITIONS

Chapter 3

Introduction	51
Active configuration.....	51
Angular magnification	51
Apodization	51
Back focal length.....	51
Cardinal planes	51
Chief ray.....	51
Coordinate axes.....	52
Diffraction limited	52
Edge thickness.....	52
Effective focal length.....	53
Entrance pupil diameter	53
Entrance pupil position.....	53
Exit pupil diameter	53
Exit pupil position	53
Extra data.....	53
Field angles and heights.....	53
Float by stop size	53
Ghost reflections	54
Glasses	54

Hexapolar rings.....	54
Image space F/#.....	54
Image space numerical aperture (NA).....	54
Lens units.....	54
Marginal ray.....	54
Maximum field.....	54
Non-paraxial systems.....	55
Non-sequential ray tracing.....	55
Normalized field coordinates.....	55
Radial field normalization.....	55
Rectangular field normalization.....	56
Normalized pupil coordinates.....	57
Object space numerical aperture.....	57
Parameter data.....	57
Paraxial and paraxial rays.....	57
Paraxial image height.....	58
Paraxial magnification.....	58
Paraxial working F/#.....	58
Primary wavelength.....	58
Radii.....	58
Real propagation.....	58
Sagittal and Tangential.....	59
Semi-diameters.....	59
Sequential ray tracing.....	59
Special characters.....	60
Strehl ratio.....	60
Surface apertures.....	60
System aperture.....	60
Tangential.....	60
Thicknesses.....	60
Total internal reflection (TIR).....	60
Total track.....	60
Units.....	61
Vignetting factors.....	61
Virtual propagation.....	62
Wavelength data.....	62
Working F/#.....	62

FILE MENU

Chapter 4

New.....	65
Open.....	65
Save.....	65
Save As.....	65
Use Session Files.....	65
Backup To Archive File.....	65
Restore From Archive File.....	66
Program Mode.....	67
Insert Lens.....	67
Preferences.....	68
Address.....	68
Folders.....	68
Graphics.....	69
Miscellaneous.....	70
Editors.....	71

Printing	72
Colors 1-12, Colors 13-24	73
Buttons 1-16, Buttons 17-32, Buttons 33-48	73
Status Bar	74
Exit	74
Recently used files	74

EDITORS MENU

Chapter 5

Lens Data	75
Inserting and deleting surfaces	75
The surface number display	75
Cutting, copying, and pasting surface data	75
Entering surface comments	76
Entering radii data	76
Entering thickness data	76
Entering glass data	76
Entering semi-diameter data	76
Entering conic data	76
Entering parameter data	76
The Surface Properties dialog box	76
Surface properties type tab	77
Surface Type	77
Surface DLL	77
Surface Color	77
Surface Opacity	77
Row Color	77
Make Surface Stop	77
Make Surface Global Coordinate Reference	77
Surface Cannot Be Hyperhemispheric	77
Ignore This Surface	78
Surface properties draw tab	78
Hide Rays To This Surface	78
Skip Rays To This Surface	78
Do Not Draw This Surface	78
Do Not Draw Edges From This Surface	78
Draw Local Axis	78
Draw Edges As	78
Mirror Substrate and Thickness	79
Surface properties aperture tab	79
Aperture type and other aperture controls	79
Aperture type codes and parameters	80
Aperture projection on non-plane surfaces	80
Aperture decenters and pickups	80
User defined apertures and obscurations	80
The UDA file format	81
UDA Examples	82
Surface properties scattering tab	84
Surface scattering settings	84
Surface tilt/decenter tab	84
Using the Coordinate Return	85
Surface physical optics tab	86
Surface coating tab	86
Setting and removing solves	86
Setting and removing variables	86

Menu options.....	86
Edit	86
Solves.....	87
View.....	88
Help.....	88
Merit Function	88
Menu Options.....	88
Edit	88
Tools.....	89
View.....	89
Help.....	89
Multi-Configuration.....	89
Menu Options.....	89
Edit	89
Solves.....	90
Tools.....	90
View.....	91
Help.....	91
Tolerance Data	92
Menu Options.....	92
Edit	92
Tools.....	92
View.....	93
Help.....	93
Extra Data.....	93
Menu Options.....	93
Edit	93
Solves.....	93
Tools.....	93
View.....	94
Help.....	94
Non-Sequential Components.....	94
Menu options.....	94
Edit	94
Solves.....	95
Tools.....	95
View.....	95
Help.....	96
NSC Editor Tools	96
Replicate Object.....	96
Create Polygon Object	96
Combine Objects.....	97
Modify Reference Objects	98
Create Source Ray From Last Geometry Error Tool.....	98
Insert/Delete Z Point Tools.....	98
Export Polar Detector Data as Source File	98
Save Detector Data	99
Load Detector Data	99
Playback ZRD on Detector.....	99
Ignore NSC Trace Errors.....	99
Undo, Redo, and Recover	100
Undo: None.....	100
Undo: Memory 1 Step	100
Undo: Disk Multi Step.....	100

Update	101
Update All	101
General	101
Aperture	101
Aperture Type.....	101
Aperture Value	102
Apodization Type.....	102
Uniform apodization	102
Gaussian apodization.....	102
Cosine cubed apodization	102
User defined apodization.....	103
Apodization Factor	103
Telecentric Object Space	103
Afocal Image Space	103
Iterate Solves When Updating.....	104
Units	104
Lens Units	104
Source Units.....	105
Analysis Units.....	105
Afocal Mode Units	105
MTF Units.....	106
Title/Notes	106
Lens Title.....	106
Notes	106
Glass Catalogs.....	106
Ray Aiming.....	106
Ray Aiming	106
Use Ray Aiming Cache	107
Robust Ray Aiming (slow)	107
Pupil Shift, Pupil Compress.....	107
Environment	108
Adjust Index Data To Environment.....	108
Temperature in degrees C	108
Pressure in ATM.....	108
Polarization (ZEMAX-EE only)	108
Unpolarized	109
Jx, Jy, X-Phase, Y-Phase.....	109
Convert thin film phase to ray equivalent	109
Method	109
Files.....	109
Coating File	109
Scatter Profile.....	109
ABg Data File	109
GRADIUM Profile	110
Miscellaneous	110
Reference OPD	110
Paraxial Rays	110
Semi Diameter Margin (lens units).....	111
Semi Diameter Margin %	111
Global Coordinate Reference Surface	111
Method To Compute F/#	112
Fast Semi-Diameters.....	112

Check GRIN Apertures.....	113
Turn Off Threading.....	113
Don't Print Coordinate Break Data.....	113
OPD Modulo.....	113
Non-Sequential.....	113
Maximum Intersections Per Ray.....	113
Maximum Segments Per Ray.....	113
Maximum Nested/Touching Objects.....	114
Minimum Relative Ray Intensity.....	114
Minimum Absolute Ray Intensity.....	114
Glue Distance In Lens Units.....	114
Missed Ray Draw Distance in Lens Units.....	114
Maximum Source File Rays In Memory.....	114
Simple Ray Splitting.....	114
Retrace Source Rays Upon File Open.....	115
Fields.....	115
Vignetting factors.....	115
Saving and loading field data.....	116
Wavelengths.....	116
Saving and loading wavelength data.....	116
Next Configuration.....	116
Last Configuration.....	116

ANALYSIS MENU

Chapter 7

Introduction.....	117
Layout.....	117
2D Layout.....	117
3D Layout.....	118
The orientation indicator.....	119
Ray errors.....	120
Configuration data.....	120
Raylist file format.....	120
Shaded Model.....	120
ZEMAX Element Drawing.....	122
Special characters.....	123
ISO Element Drawing.....	123
Summary of ISO 10110 symbols and codes.....	124
Diameters and bevels.....	125
Tolerance data.....	125
NSC 3D Layout.....	125
NSC Shaded Model.....	126
NSC Object Viewer.....	126
Fans.....	127
Ray Aberration.....	127
Evaluating results at intermediate surfaces.....	128
Optical Path.....	129
Pupil Aberration.....	129
Spot Diagrams.....	129
Standard.....	129
Through Focus.....	131
Full Field.....	132
Matrix.....	132
Configuration Matrix.....	132
MTF.....	133

FFT MTF	133
FFT Through Focus MTF	134
FFT Surface MTF	135
FFT MTF vs. Field	135
Comment about vignetting factors	136
FFT MTF Map	136
Huygens MTF	137
Huygens Through Focus MTF	138
Huygens Surface MTF	138
Huygens MTF vs. Field	139
Geometric MTF	139
Geometric Through Focus MTF	140
Geometric MTF vs. Field	140
Geometric MTF Map	141
PSF	141
FFT PSF	141
Assumptions used in the FFT PSF calculation	142
Discussion of the FFT method and sampling issues	143
FFT PSF Cross Section	144
FFT Line/Edge Spread	145
Huygens PSF	146
Huygens PSF Cross Section	147
Wavefront	148
Wavefront Map	148
Interferogram	148
Foucault Analysis	150
Surface	151
Surface Sag	151
The Contour Format String	151
Surface Phase	152
Surface Curvature	152
Surface Sag Cross Section	152
Surface Phase Cross Section	153
Surface Curvature Cross Section	153
RMS	154
RMS vs. Field	154
Comments about RMS computation methods	155
Comments about RMS wavefront computations	155
RMS vs. Wavelength	155
RMS vs. Focus	156
RMS Field Map	157
Encircled Energy	158
Diffraction	158
Geometric	159
Geometric Line/Edge Spread	160
Extended Source	161
Image Simulation	162
Image Simulation	162
Suggestions for use	164
Comments on using field angles	165
Comments on using paraxial and real image height	165
Comments on JPG and BMP files	165
Geometric Image Analysis	165
The IMA format	167

The BIM format.....	168
Comments on using field angles	168
Comments on using paraxial and real image height	168
How rays are chosen for analysis	168
Calculating efficiency of multi-mode fibers	169
Text output	169
Geometric Bitmap Image Analysis	169
Comments on field units	171
Selection and tracing of rays	172
Partially Coherent Image Analysis	172
Comment about using the IMA file format	174
Comment about using the ZBF file format.....	175
The optical transfer function algorithm	175
Comment about the coherent optical transfer function.....	176
Computing partially coherent images	176
Computing MTF using Cross Section X/Y	177
Extended Diffraction Image Analysis	177
IMA/BIM File Viewer.....	179
JPG/BMP File Viewer.....	180
Biocular Analysis.....	180
Field of View.....	180
Dipvergence/Convergence.....	181
Miscellaneous	182
Field Curvature/Distortion	182
Computing distortion when using real image height as the field type	184
Grid Distortion	184
SMIA-TV Distortion.....	186
Limitations of Grid Distortion	186
Relative Illumination	187
Effective F/#	188
Vignetting Plot.....	188
Footprint Diagram	188
Longitudinal Aberration	189
Lateral Color.....	190
Y-Ybar Diagram	190
Chromatic Focal Shift.....	191
System Summary Graphic	191
Power Field Map	191
Power Pupil Map.....	193
Incident Angle vs. Image Height	194
Aberration coefficients	194
Seidel Coefficients	194
Seidel Diagram.....	196
Zernike Fringe Coefficients	196
Subaperture computations	197
Strehl ratio approximation	198
Zernike Standard Coefficients.....	199
Zernike Annular Coefficients	202
Calculations	205
Ray Trace.....	205
Fiber Coupling Efficiency	206
YNI Contributions	207
Sag Table.....	208
Cardinal Points.....	209

Glass and Gradient Index	209
Dispersion Diagram.....	209
Glass Map.....	210
Athermal Glass Map.....	210
Internal Transmittance vs. Wavelength.....	211
Grin Profile	211
Gradium™ Profile.....	212
Universal Plot.....	212
New Universal Plot 1D.....	212
New Universal Plot 2D.....	214
Polarization	216
Polarization Ray Trace.....	216
Polarization Pupil Map	216
Modeling interference between configurations.....	217
Transmission.....	218
Phase Aberration	218
Transmission Fan.....	219
Coatings.....	219
Reflection vs. Angle	219
Transmission vs. Angle.....	220
Absorption vs. Angle	220
Diattenuation vs. Angle	220
Phase vs. Angle	220
Retardance vs. Angle.....	221
Reflection vs. Wavelength.....	221
Transmission vs. Wavelength	221
Absorption vs. Wavelength	222
Diattenuation vs. Wavelength	222
Phase vs. Wavelength	222
Retardance vs. Wavelength	222
Physical Optics	222
Paraxial Gaussian Beam.....	222
Limitations of the analysis	223
Overview of Gaussian beams	223
Default beam parameters.....	223
Propagating the embedded beam	223
The quality factor.....	224
Interactive analysis.....	224
Skew Gaussian Beam.....	225
Physical Optics Propagation	225
Comments about beam projection	228
Beam File Viewer	229
Non-Sequential	229
Ray Trace Control.....	229
Detector Viewer.....	229
Ray Database Viewer	229
Ray Database Analysis	229
Flux vs. Wavelength.....	229
Obsolete.....	230
Wireframe.....	230
Solid Model	231
Illumination XY Scan.....	231
Illumination 2D Surface	233
Export IGES Line Work	233

Export 2D DXF File	234
Export 3D DXF File	235
Conjugate Surface Analysis	236

TOOLS MENU

Chapter 8

Optimization	239
Optimization	239
Global Search	239
Hammer Optimization	239
Find Best Asphere.....	239
Merit Function Listing	240
Remove All Variables.....	240
Glass Substitution Template	240
Tolerancing	242
Tolerancing	242
Tolerance Listing.....	242
Tolerance Summary.....	242
Test Plates	242
Test Plate Fitting	242
Test Plate Lists.....	243
Catalogs.....	244
Glass Catalogs.....	244
Glass Compare	244
Glass Fitting	245
Lens Catalogs	246
Make Private Lens Catalog	248
Coatings.....	249
Edit Coating File.....	249
Reload Coating File.....	249
Add Coatings to All Surfaces	249
Coating Listing	249
Export Encrypted Coating	249
Scattering.....	250
ABg Scatter Data Catalogs	250
Wavelength scaling of ABg data	251
Scatter Function Viewer	252
Apertures	253
Convert Semi-Diameters to Circular Apertures.....	253
Convert Semi-Diameters to Floating Apertures	253
Convert Semi-Diameters to Maximum Apertures.....	254
Remove All Apertures	254
Replace Vignetting With Apertures	254
Coordinates.....	254
Tilt/Decenter Elements.....	254
Add Fold Mirror	255
Limitations of reversing surfaces following the fold mirror.....	255
Delete Fold Mirror	255
Local To Global	256
Global To Local.....	256
Export Data	257
Export IGES/SAT/STEP Solid.....	257
Limitations of exported data	259
Export Source Data.....	259
Export ZEMAX Black Box Data.....	260

Miscellaneous	262
Reverse Elements	262
Scale Lens	262
Make Focal.....	263
Ghost Focus Generator.....	263
Make Double Pass	264
Performance Test.....	265
Lock/Unlock All Windows	265
Quick Focus	265
Quick Adjust.....	265
Slider	266
Convert to NSC Group	267
Convert File Formats.....	268
*.INT Zernike files.....	269
*.INT Grid Files.....	269
OptiWave *F3D files	270

REPORTS MENU

Chapter 9

Introduction	271
Surface Data	271
System Data	271
Prescription Data	271
Comments on computing element volumes	272
System Check.....	272
Report Graphics 4/6.....	273

MACROS AND EXTENSIONS MENUS

Chapter 10

Edit/Run ZPL Macros.....	275
Refresh Macro List.....	275
Macro Names.....	275
Extensions	275
Refresh Extensions List	276
Extension Names.....	276

SURFACE TYPES

Chapter 11

Introduction	277
Parameter data	277
Extra data.....	277
Summary of surface types	277
User defined surfaces	277
Surface types	278
ABCD	280
Alternate Even	281
Alternate Odd.....	281
Atmospheric Refraction.....	281
Biconic	282
Biconic Zernike	282
Binary Optic 1	283
Binary optic coefficients sign conventions.....	284
Binary Optic 2	284
Binary optic coefficients sign conventions.....	285
Binary Optic 3	286
Binary optic coefficients sign conventions.....	287

Birefringent In and Birefringent Out	287
Defining the substrate shape	289
Defining the index	289
Defining the crystal axis for homogeneous materials.....	289
Defining the crystal axis for inhomogeneous materials.....	289
Determining which ray is traced	291
Accounting for phase rotation	291
Transmission and other properties of birefringent surfaces.....	291
Sample ZEMAX files	291
Black Box Lens	292
Features not supported when using Black Box surfaces	292
Conjugate.....	293
Coordinate Break	293
The order flag.....	294
Cubic Spline.....	294
Comments about spline surfaces.....	294
Cylinder Fresnel.....	295
Diffraction Grating	296
Elliptical Grating 1	296
Elliptical Grating surface shape.....	297
Elliptical Grating 2	297
Even Asphere	298
Extended Asphere	298
Extended Cubic Spline.....	299
Extended Fresnel.....	299
Extended Odd Asphere.....	301
Extended Polynomial	301
Extended Toroidal Grating	302
Filter	303
Fresnel	303
Generalized Fresnel.....	304
Gradient 1	305
Discussion on maximum step size for GRIN surfaces	305
Restrictions on surfaces following GRIN surfaces	305
Gradient 2	305
Gradient 3	306
Gradient 4	306
Gradient 5	307
Gradient 6	308
Gradient 7	309
GRADIUM™	309
GRADIUM profile file format.....	310
Gradient 9	311
Gradient 10	312
Grid Phase	313
Using the shear distance.....	313
Phase coefficients sign conventions	313
Grid Sag.....	313
Importing grid data	314
Bicubic spline vs. linear interpolation	315
Suggestions for using the Grid Sag surface.....	315
Hologram 1	316
Hologram 2	317
Irregular.....	317

Jones Matrix.....	318
Lenslet Array.....	318
Non-Sequential Components.....	318
Odd Asphere.....	318
Odd Cosine.....	318
Optically Fabricated Hologram.....	319
Defining the substrate shape.....	320
Defining the construction optics.....	320
Optimizing the construction and playback optics.....	320
Selecting the hologram type.....	321
Paraxial.....	321
Setting the OPD Mode.....	321
Paraxial XY.....	322
Periodic.....	322
Polynomial.....	323
Q-Type Asphere.....	323
Radial Grating.....	323
Radial NURBS.....	324
Retro Reflect.....	325
Standard.....	325
Modeling an ellipse with the Standard surface.....	326
Modeling an axicon with the Standard surface.....	326
Superconic.....	326
Tilted.....	327
Toroidal.....	328
Toroidal Grating.....	329
Toroidal Hologram.....	329
Toroidal NURBS.....	330
User Defined.....	331
The UDS DLL.....	331
Refractive and reflective UDS DLLs.....	332
Gradient index UDS DLLs.....	333
Diffractive UDS DLLs.....	333
Lenslet arrays using DLLs.....	333
User defined surface apodization using DLLs.....	333
Polarization and coating data using DLLs.....	334
Error handling and UDS.....	334
Sample DLLs.....	334
Variable Line Space Grating.....	339
Zernike Fringe Phase.....	340
Zernike Fringe Phase coefficients sign conventions.....	340
Zernike Fringe Sag.....	341
Zernike Standard Phase.....	341
Zernike Standard Phase coefficients sign conventions.....	342
Zernike Standard Sag.....	343
Zone Plate.....	343

NSC - OVERVIEW

Chapter 12

Introduction.....	345
Paraxial data and ray tracing with NSC.....	345
Optimization with NSC.....	346
The two methods of using NSC ray tracing.....	346
NSC with ports.....	346
NSC without ports.....	346

Combining NSC with and without ports.....	346
Overview of NSC ray tracing with ports	347
The entry port.....	347
The exit port	347
Getting rays in.....	348
Tracing rays within the NSC.....	348
Getting rays out.....	348
Overview of NSC ray tracing without ports	348
Getting rays in.....	349
Tracing rays within the NSC.....	349
Detecting rays	349
Launching rays for analysis.....	349
Object Placement.....	349
The object coordinate system	350
Reference objects	350
Placing objects inside of, adjacent to, or overlapping other objects.....	351
Glue distance	351
Nesting object limits	351
Nesting volumes.....	351
Nesting surfaces	351
Refraction and reflection from NSC objects.....	352
No material defined.....	352
MIRROR or ABSORB	352
Material defined by a glass name	352
Material defined by a table glass.....	352
Material defined by MIL number.....	352
Material defined by a glass model.....	352
Object properties by material	352
Diffraction from NSC objects.....	353
The object properties dialog box.....	353
Type tab	353
Coat/Scatter tab	355
Scatter To tab.....	355
Face tab	355
Bulk Scatter tab.....	356
GRIN tab	356
Birefringence tab.....	357
Diffraction tab.....	357
Sources tab.....	358
Random vs. Sobol sampling.....	359
Modeling arrays of sources	360
Defining the color and spectral content of sources	360
The spectrum fitting algorithm	361
Defining a spectrum file.....	362
How rays are selected from the spectrum.....	362
Using ABg scattering with source color models	362
Draw tab.....	362
Comments about drawing resolution and bounding regions	363
Defining paths using the Consider Objects list	363
Defining an Ignore Objects list.....	364
User defined apertures	364
Object faces.....	364
Polarization and thin film coatings	364
Coatings on surfaces in contact.....	364

Scattering.....	364
Fraction to scatter and number of scatter rays.....	365
Scatter models	365
Bi-Directional Scatter Distribution Function.....	366
Available scatter models	366
No scattering.....	366
Lambertian scattering.....	366
Gaussian scattering	366
ABg model scattering.....	367
Defining ABg data	368
User defined scattering	368
Defining an object to use DLL defined surface scattering.....	368
DLL parameters.....	368
Creating a new DLL.....	368
Sample surface scattering: TwoGaussian.DLL	369
Modeling measured BSDF data with the RI_BSDF scattering DLL	369
Sample surface scattering: K-correlation.DLL.....	369
How to model scattering efficiently.....	369
The Scatter To list	369
Importance Sampling	370
Bulk scattering	370
No Bulk Scattering	371
Angle Scattering.....	371
DLL Defined Scattering.....	371
Sample bulk scattering: Poly_bulk_scat.DLL	371
Sample bulk scattering: Henyey-Greenstein-bulk.DLL.....	371
Sample bulk scattering: Rayleigh.DLL	372
Sample bulk scattering: Mie.DLL.....	372
Sample bulk scattering: Phosphor.DLL	372
Defining an object to use DLL defined bulk scattering	372
DLL parameters.....	372
Creating a new DLL.....	372
Coherence length modeling.....	373
Defining GRIN media for non-sequential ray tracing	373
Defining an object to be of gradient index material.....	373
Discussion on maximum step size for GRIN objects	373
GRIN DLL parameters	373
Creating a new GRIN DLL	374
Sample GRIN DLLs.....	374
Defining DLLs for ray splitting at diffractive surfaces.....	375
Defining an object to use the diffraction DLL	375
Diffraction DLL parameters	375
Creating a new Diffraction DLL	376
Ray splitting	376
Ray splitting and polarization	376
Putting it all together	377
The Ray Trace Control.....	377
Saving ray data to a file.....	379
Lost energy	379
Ray database (ZRD) files.....	380
ZRD Format And Version Numbers	380
The ZRD Uncompressed Full Data Format (UFD).....	380
The Ray Database Viewer	382
The Detector Viewer	384

Comments regarding Detector Surface objects	386
The filter string	386
Filter string examples	389
Saving and Loading Detector Data	390
Detector Data File Formats	390
Header Block Format	390
Data Block Format.....	391
Pixel Ordering.....	391
Special considerations for faceted objects.....	391
Comments about DLLs	392

NSC Objects

Chapter 13

NSC Objects	393
Annular Aspheric Lens.....	396
Annular Axial Lens	397
Annular Volume	398
Annulus	398
Array	399
Array Ring	401
Aspheric Surface.....	402
Aspheric Surface 2.....	402
Axicon Surface.....	403
Biconic Lens.....	403
Biconic Surface	404
Making a hyperhemispheric surface	405
Biconic Zernike Lens.....	405
Biconic Zernike Surface	406
Binary 1	407
Binary 2.....	408
Binary 2A	408
Boolean.....	409
Compound Parabolic Concentrator (CPC).....	411
CPC Rectangular	412
Cone	412
Cylinder Pipe.....	413
Cylinder Volume.....	413
Cylinder 2 Pipe.....	413
Cylinder 2 Volume.....	414
Diffraction Grating.....	414
Dual BEF Surface	414
Ellipse	415
Elliptical Volume.....	415
Even Asphere Lens.....	415
Extended Polynomial Lens	416
Extended Polynomial Surface.....	417
Extruded.....	417
Freeform Z	417
Fresnel 1	418
Fresnel 2	419
Grid Sag Lens	420
Grid Sag Surface	421
Hexagonal Lenslet Array.....	421
Hologram Lens.....	421
Hologram Surface	422

Imported.....	423
Imported objects and ZOF files.....	424
Comments about imported objects.....	424
Ray tracing accuracy for imported objects.....	425
Ray tracing speed for imported objects.....	425
Limitations on imported objects.....	425
Jones Matrix.....	425
Lenslet Array 1.....	426
Lenslet Array 2.....	427
Micro Electro Mechanical System (MEMS).....	427
Null Object.....	428
Odd Asphere Lens.....	428
Paraxial Lens.....	429
Polygon Object.....	429
Defining Polygon Objects.....	429
The comment symbol: !.....	430
The face name symbol: C.....	430
The invisible symbol: I.....	430
The vertex symbol: V.....	430
The triangle symbol: T.....	430
The rectangle symbol: R.....	431
Maximum triangles in the polygon object.....	431
Example POB file.....	431
Ray Rotator.....	432
Rectangular Corner.....	432
Rectangle.....	433
Rectangular Pipe.....	433
Rectangular Pipe Grating.....	433
Rectangular Roof.....	434
Rectangular Torus Surface.....	434
Rectangular Torus Volume.....	434
Rectangular Volume.....	435
Comment about tilted faces.....	435
Rectangular Volume Grating.....	435
Slide.....	435
Sphere.....	436
Standard Lens.....	436
Standard Surface.....	437
STL Object.....	437
Defining STL Objects.....	438
Maximum triangles in STL objects.....	438
Example STL files.....	438
Swept Object.....	438
Tabulated Faceted Radial.....	439
Tabulated Faceted Toroid.....	440
Tabulated Fresnel Radial.....	441
Toroidal Hologram.....	441
Toroidal Lens.....	442
Toroidal Surface.....	442
Torus Surface.....	443
Torus Volume.....	443
Triangular Corner.....	444
Triangle.....	444
User Defined Object.....	444

Sample Object DLLs	445
Wolter Surface	446
Zernike Surface.....	446

NSC Sources

Chapter 14

Sources.....	449
Placing sources inside objects	449
Adding new source types	449
Parameters common to all source objects	450
Source Diode	450
Source DLL	452
Source DLL parameters	452
Creating a new Source DLL	452
Comments on Source DLL relative intensity	452
Sample Source DLLs	453
Source Ellipse	453
Source EULUMDAT File	454
Source Filament.....	454
Source File.....	454
Text vs. binary format files	455
Memory requirements for source file objects	455
Restrictions on the number of rays selected	455
Binary Source File format.....	456
Text Source File format.....	457
Intensity normalization in source files	457
Source Gaussian	457
Source IESNA File	457
Source Imported	457
Source Object	458
Source Point	459
Source Radial	459
Source Ray	459
Source Rectangle	459
Source Tube	459
Source Two Angle.....	459
Source Volume Cylinder	460
Source Volume Ellipse.....	460
Source Volume Rectangle	460

NSC Detectors

Chapter 15

Detectors.....	463
Detector Color object	463
Comments on Detector Color pixel numbering	464
Comments on Pixel Interpolation	464
Detector Polar object:	464
Exporting Detector Polar Data	465
Comments on Detector Polar Pixel Numbering	465
Comments on Pixel Interpolation	465
Detector Rectangle object.....	466
Comments on coherent data computations	467
Comments on Detector Rectangle pixel numbering.....	468
Comments on Pixel Interpolation	468
Detector Surface object	468

Comments about detector triangles	469
Comments on Detector Surface pixel numbering	469
Detector Volume object.....	470
Comments on Detector Volume pixel numbering.....	470
Objects as detectors	470
Comments on general objects as detectors pixel numbering	471

SOLVES

Chapter 16

Introduction	473
Curvature: Marginal ray angle.....	474
Curvature: Chief ray angle	475
Curvature: Pickup	475
Curvature: Marginal ray normal	475
Curvature: Chief ray normal.....	475
Curvature: Aplanatic	475
Curvature: Element power	475
Curvature: Concentric with surface.....	475
Curvature: Concentric with radius.....	475
Curvature: F number.....	475
Curvature: ZPL Macro.....	476
Thickness: Marginal ray height.....	476
Thickness: Chief ray height.....	476
Thickness: Edge thickness	476
Thickness: Pickup.....	476
Thickness: Optical path difference.....	476
Thickness: Position	477
Thickness: Compensator	477
Thickness: Center of curvature	477
Thickness: Pupil Position	477
Thickness: ZPL Macro	477
Glass: Model	477
Glass: Pickup	477
Glass: Substitute.....	477
Glass: Offset	478
Semi-Diameter: Pickup	478
Semi-Diameter: Maximum	478
Semi-Diameter: ZPL Macro	478
Conic: Pickup	478
Conic: ZPL Macro	478
Parameter: Pickup	478
Parameter: Chief ray.....	479
Parameter: ZPL Macro.....	479
TCE: Pickup.....	479
Extra Data Value: Pickup.....	479
Extra Data Value: ZPL Macro	479
Solve restrictions.....	479
Using ZPL Macro solves.....	479
Important considerations for ZPL Macro solves.....	480
Integer codes for column numbers	480
Suggestions for use	481

OPTIMIZATION

Chapter 17

Introduction	483
--------------------	-----

Selecting variables.....	483
Defining the default merit function	483
Selecting the type of optimization	483
Physically significant RMS merit functions.....	485
Selecting the pupil integration method.....	485
Rings.....	486
Arms.....	486
Sampling considerations for Rings and Arms	486
Grid	486
Delete Vignetted.....	486
Setting thickness boundary values.....	487
Start At	487
Assume Axial Symmetry	487
Ignore Lateral Color	487
Relative X Weight.....	488
Overall Weight.....	488
Pitfalls with the default merit function.....	488
Optimization with apodized beams	488
Modifying the merit function	488
Hx, Hy, Px, and Py	488
Target, Value, and Weight	488
Percent Contribution	489
Operand weights less than zero.....	489
Operand weights equal to zero	489
Operand weights greater than zero.....	489
Merit function definition	489
Optimization operands.....	489
Understanding boundary operands.....	524
Using MTF operands	525
Performing an optimization	526
Defining complex operands	527
Optimizing glass selection	528
Using model glasses	528
Using glass substitution	529
Optimizing zoom and multi-configuration lenses	529
Optimizing extra data	530
Optimizing tolerance sensitivity.....	530
Optimizing objects in a non-sequential group with sequential rays	531
Optimizing with sources and detectors in non-sequential mode.....	531
Comments about random numbers and NSTR.....	532
Pixel numbering for detectors	532
Optimizing with the IMAE operand.....	532
Using gradient index operands	532
DLTN.....	532
LPTD	532
User defined operands.....	533
Optimizing with ZPL macros	533
Changes made to the lens from within the ZPLM macro	534
ZPLM and the default merit function	534
Optimizing with externally compiled programs.....	534
Suggestions for use	536
The global optimum	537

Introduction	539
Capabilities of ZEMAX	539
The Global Search algorithm	540
The Hammer algorithm	541
Optimizing glass selection	541
Using glass substitution	542
Restricting selected glasses.....	542
Suggestions for use	543
Summary.....	543

TOLERANCING

Chapter 19

Introduction	545
The basic procedure	545
Tolerance operands	546
TRAD: Tolerance on radius.....	548
TCUR: Tolerance on curvature	548
TFRN: Tolerance on fringes.....	548
TTHI: Tolerance on thickness	548
TCON: Tolerance on conic.....	549
TSDI: Tolerance on semi-diameter	549
TSDX, TSDY: Tolerance on surface decenters	549
TSTX, TSTY: Tolerance on surface tilts.....	549
TIRX, TIRY: Tolerance on surface TIR	549
TIRR: Tolerance on surface irregularity	550
TEXI: Tolerance on surface irregularity using the Fringe Zernike model	550
TEZI: Tolerance on surface irregularity using the Standard Zernike model.....	551
TPAI: Tolerance on the inverse of parameter data	552
TPAR: Tolerance on parameter data	552
TEDV: Tolerance on extra data values	552
TIND: Tolerance on index	552
TABB: Tolerance on Abbe	553
TCMU: Tolerance on coating multiplier	553
TCIO: Tolerance on coating index offset.....	553
TCEO: Tolerance on coating extinction offset.....	553
TEDX, TEDY: Tolerance on element decenters	553
TETX, TETY, TETZ: Tolerance on element tilts.....	554
TOFF: Tolerance off (can be used for comments).....	554
TUDX, TUDY, TUTX, TUTY, TUTZ: Tolerance on user defined tilts and decenters.....	554
TNPS, TNPA, TNMA: Tolerances on non-sequential data	554
TMCO: Tolerance on multi-configuration data	555
Tolerancing with the Irregular surface type.....	555
Tolerance control operands	555
CEDV: Define extra data value compensator	556
General comments about min and max values on compensators	556
CMCO: Define multi-configuration operand compensator	556
CNPA: Define non-sequential parameter compensator	556
CNPS: Define non-sequential position compensator	556
COMP: Define compensator	556
CPAR: Define parameter compensator.....	556
SAVE: Save sensitivity analysis lenses	557
SEED: Seed the random number generator	557
STAT: Define statistics.....	557
TWAV: Test wavelength.....	557
Defining default tolerances	557

Surface tolerances	557
Element tolerances	558
Index tolerances.....	558
Other controls.....	558
Defining compensators	559
Performing the tolerance analysis.....	559
Set-Up tab	559
Criterion tab.....	561
Monte Carlo tab.....	563
Display tab	563
Other buttons	564
How ZEMAX computes the tolerance analysis	564
Evaluating compensators.....	564
Sensitivity analysis.....	565
The RSS estimated change	565
Inverse sensitivity analysis.....	566
Monte Carlo analysis	566
Normal statistical distribution	567
Uniform statistical distribution	567
Parabolic statistical distribution	567
User defined statistical distribution.....	568
Discussion of Monte Carlo analysis method	568
Nesting rules for Monte Carlo analysis	569
Using Tolerance Scripts.....	569
Tolerance Script overview.....	570
The Tolerance Script commands	570
!	570
CEDV	570
CLEARCOMP.....	571
CLOSEFILE.....	571
CMCO	571
CNPA	571
CNPS	571
COMP.....	571
CPAR	572
FORMAT	572
GETMERIT	572
HEADER	572
LOADMERIT	573
NOMINAL	573
OPENFILE.....	573
OPTIMIZE	574
OPTIMIZE-OD.....	574
PERTURB	574
QUIET	575
REPORT	575
SAVE.....	576
UPDATE.....	576
Tolerance Script example	576
Tolerancing multi-configuration (zoom) lenses	577
Tolerancing with solves.....	577
Trouble shooting the tolerance results.....	577
Optimizing for tolerance sensitivity	577
Pitfalls when tolerancing	577

Summary.....	578
MULTI-CONFIGURATIONS	Chapter 20

Introduction	579
The first step	579
Summary of multi-configuration operands	579
Comment about operands that define character strings	583
Defining the number of configurations	584
Defining each configuration	584
Adding and removing elements	584
Changing configurations	584
Optimization with multi-configurations	584
Suggestions for organizing multiple-configuration merit functions	584
Solves for multi-configuration data.....	585
Fixed	585
Variable (numerical value operands only).....	585
Substitute (glass operands only).....	585
Pickup	585
Thermal pickup.....	585

USING GLASS CATALOGS	Chapter 21
-----------------------------	-------------------

Introduction	587
Specifying which glass catalogs to use.....	587
Editing and reviewing glass catalogs	587
Description of catalog data	588
Creating a new catalog	589
Copying or moving glass catalog files.....	589
The glass dispersion formulas	590
The Schott formula.....	590
The Sellmeier 1 formula	590
The Sellmeier 2 formula	590
The Sellmeier 3 formula	590
The Sellmeier 4 formula	591
The Sellmeier 5 formula	591
The Herzberger formula	591
The Conrady formula	591
The Handbook of Optics 1 formula	591
The Handbook of Optics 2 formula	591
The Extended formula.....	592
The Extended 2 formula.....	592
The Extended 3 formula.....	592
General comments on using dispersion formulas	592
Fitting index data.....	592
Fitting melt data	593
Discussion of melt fitting method:	594
Defining Transmission Data.....	594
Modeling gases and liquids.....	595
Finding a glass quickly.....	595
Glass catalog sources.....	595
Obsolete catalog data.....	599
The AGF and BGF file formats	599
Alternate methods of defining dispersion data.....	600
Using MIL number glasses.....	600

Using table glasses	601
Using model glasses	601

THERMAL ANALYSIS

Chapter 22

Introduction	603
Defining temperature and pressure	603
Defining wavelengths	603
Index of refraction computation	603
Environmental effects on index for gradient index, MIL number, and model glasses	604
Defining multiple temperature and pressure values	605
Defining which parameters consider thermal effects	605
Radius of curvature (CRVT) values	605
Thickness (THIC) values	606
Parameter values	606
Extra data values	606
All other values	606
Defining multiple environments within a single configuration	606
Automatic thermal setup	607
Adding TCE data	607
Modeling gases and liquids	607
Adding thermal index variation data	607
Optimizing athermal lenses	608
Limitations of thermal analysis	608

POLARIZATION ANALYSIS

Chapter 23

Introduction	609
Review of polarization concepts	609
The electric field vector	609
Field vs. ray phase conventions	610
The polarization ellipse	610
Definition of terms	610
Intensity, I	610
Phase, P	610
Path length through medium,	611
Internal absorption per lens unit,	611
Internal Transmittance, IT	611
Propagation Phase Factors, pc, ps	611
Amplitude Reflectance,	611
Amplitude Transmittance,	611
Intensity Reflectance, Rs, Rp	611
Intensity Transmittance, Ts, Tp	611
Intensity Absorption, As, Ap	612
Diattenuation, D	612
Coating Phase, Ps, Pp	612
Retardance, S	612
Phase Difference Between X and Y, Pxy	612
Major and Minor semi axis of the Polarization Ellipse, EM, Em	612
Angle of XY Polarization Ellipse, Ap	612
Properties of uncoated surfaces	612
Defining coatings in ZEMAX	613
Coating file data syntax	613
The MATE data section	615
The TAPR data section	615

The COAT data section	617
Defining replicated groups of coating layers	618
Defining simple ideal coatings	618
The ENCRYPTED data section	618
The IDEAL data section	619
The IDEAL2 data section	619
The TABLE data section	619
Adding comments to the coating file	620
Limits on the amount of coating data	620
Editing the coating file	621
Optimizing coatings with ZEMAX	621
Default materials and coatings supplied with ZEMAX	621
Specifying coatings on surfaces	622
What ZEMAX does if no coating is specified	622
Defining the initial polarization	622
Defining polarizing components	623
What ZEMAX can compute using polarization analysis	624
Bulk absorption and transmission	625
Modeling birefringent materials	625
Modeling frustrated total internal reflection	625
Limitations of polarization analysis	626

PHYSICAL OPTICS PROPAGATION

Chapter 24

Introduction	627
Support for multiple processors	628
Diffraction propagation	628
Representation of the electric field	628
The Fresnel number	628
Near and far field	629
Angular spectrum propagation	629
Fresnel diffraction	631
Selecting the correct propagator	632
Fraunhofer diffraction	632
The pilot beam	632
Sign conventions for phase data	633
Propagating in and out of the Rayleigh range	634
Separation of X and Y propagation	635
Comments about point spacing and sampling	636
Propagation through arbitrary optical surfaces	636
Propagating through non-sequential surfaces	637
Accounting for polarization	637
Memory requirements	637
Defining the initial beam	638
Gaussian Waist	638
Gaussian Angle	639
Gaussian Size+Angle	639
Top Hat	640
File	640
ZEMAX Beam File (ZBF) binary format	640
ZEMAX Beam File (ZBF) text format	641
DLL	641
Beam DLL parameters	641
Creating a new Beam DLL	641
Multimode	642

!, Comment.....	642
COHERENT	642
DLL.....	642
GW	642
INCOHERENT.....	643
PHASE	643
TH.....	643
Using random values.....	643
Using the scale factor.....	643
Surface specific settings	643
Considerations when using rays to propagate	645
Computing Fiber Coupling	645
Where the integral is computed	645
Defining the fiber mode	645
Decenters and tilts	646
Choosing the location for the receiving fiber	646
Quantitative beam analysis.....	646
Beam coordinates and pilot beam properties.....	646
Peak Irradiance and total power	646
Centroid locations	646
Beam width and M-squared	646
Wavefront error and RMS beam deviations	647
Encircled energy	647
Suggestions for use	647
Algorithm assumptions.....	648
Samples	649
Free space propagation	649
A pinhole aperture.....	650
A lens array	650
Talbot imaging.....	650
Fresnel lens.....	650

ZEMAX PROGRAMMING LANGUAGE

Chapter 25

Introduction	651
Creating ZPL macros	651
Running ZPL macros	651
Assigning ZPL macros to buttons	652
An overview of ZPL.....	652
Assignments.....	652
Keywords	652
Comments.....	653
Creating graphics	653
Numeric variables	653
Array variables	653
Numeric operations.....	654
Numeric logical operators	654
String variables	655
String operations.....	655
String logical operators	656
Numeric functions	656
Using the FICL() function	668
String functions	668
Keywords	670
APMN, APMX, APTP, APXD, APYD.....	670

ATYP, AVAL.....	670
BEEP.....	671
CALLMACRO.....	671
CALLSETDBL.....	671
CALLSETSTR.....	671
COAT.....	671
CLOSE.....	671
CLOSEWINDOW.....	672
COLOR.....	672
COMMAND.....	672
COMMENT.....	672
CONI.....	672
CONVERTFILEFORMAT.....	672
COPYFILE.....	673
CURV.....	673
DECLARE.....	673
DEFAULTMERIT.....	673
DELETE.....	674
DELETECONFIG.....	674
DELETEFILE.....	674
DELETEMCO.....	674
DELETEMFO.....	675
DELETEOBJECT.....	675
DELETETOL.....	675
EDVA.....	675
END.....	675
EXPORTBMP.....	675
EXPORTCAD.....	676
EXPORTJPG.....	677
EXPORTWMF.....	677
FINDFILE.....	677
FLDX, FLDY, FWGT, FVDX, FVDY, FVCX, FVCY, FVAN.....	678
FOR, NEXT.....	678
FORMAT.....	678
FTYP.....	679
GCRS.....	679
GDATE.....	679
GETEXTRADATA.....	679
GETGLASSDATA.....	679
GETLSF.....	680
GETMTF.....	681
GETPSF.....	682
GETSYSTEMDATA.....	683
GETTEXTFILE.....	683
GETVARDATA.....	684
GETZERNIKE.....	685
GLAS.....	686
GLASSTEMPLATE.....	686
GLENSNAME.....	687
GLOBALTOLOCAL.....	687
GOSUB, SUB, RETURN, and END.....	687
GOTO.....	688
GRAPHICS.....	688
GTEXT.....	689

GTEXTCENT	689
GTITLE.....	689
HAMMER	689
IF-THEN-ELSE-ENDIF.....	690
IMA.....	690
IMAGECOMBINE.....	691
IMAGEEXTRACT	691
IMASHOW.....	692
IMASUM.....	692
IMPORTEXTRADATA.....	692
INPUT	692
INSERT	693
INSERTCONFIG.....	693
INSERTMCO.....	693
INSERTMFO.....	693
INSERTOBJECT.....	694
INSERTTOL.....	694
LABEL.....	694
LINE	694
LOADARCHIVE	695
LOADCATALOG	695
LOADDETECTOR.....	695
LOADLENS.....	696
LOADMERIT	696
LOADTOLERANCE	696
LOCALTOGLOBAL.....	696
LOCKWINDOW.....	697
MAKEFACETLIST.....	697
MAKEFOLDER	697
MODIFYSETTINGS	697
NEXT.....	702
NSTR	702
NUMFIELD.....	702
NUMWAVE	702
OPEN.....	702
OPENANALYSISWINDOW.....	703
OPTIMIZE	703
OPTRETURN.....	704
OUTPUT	704
PARM.....	705
PARAXIAL.....	705
PAUSE.....	705
PIXEL.....	705
PLOT.....	706
PLOT2D	707
POLDEFINE.....	708
POLTRACE.....	709
POP.....	710
PRINT	710
PRINTFILE.....	711
PRINTWINDOW.....	711
PWAV.....	711
QUICKFOCUS	711
RADI.....	712

RANDOMIZE.....	712
RAYTRACE.....	712
RAYTRACEX.....	713
READ.....	713
READNEXT.....	714
READSKIP.....	714
READSTRING.....	715
RELEASE.....	715
RELOADOBJECTS.....	715
REM, !, #.....	716
REMOVEVARIABLES.....	716
RENAMEFILE.....	716
RETURN.....	716
REWIND.....	716
SAVEARCHIVE.....	717
SAVEDETECTOR.....	717
SAVELENS.....	717
SAVEMERIT.....	717
SAVETOLERANCE.....	718
SAVEWINDOW.....	718
SCATTER.....	718
SDIA.....	718
SETAIM.....	718
SETAIMDATA.....	719
SETAPODIZATION.....	719
SETCONFIG.....	719
SETDETECTOR.....	720
SETMCOPERAND.....	720
SETNSCPARAMETER.....	720
SETNSCPOSITION.....	721
SETNSCPROPERTY.....	721
SETOPERAND.....	725
SETSTDD.....	725
SETSURFACEPROPERTY, SURP.....	725
SETSYSTEMPROPERTY, SYSP.....	728
SETTEXTSIZE.....	731
SETTITLE.....	732
SETTOL.....	732
SETUNITS.....	732
SETVAR.....	732
SETVECSIZE.....	733
SETVIG.....	733
SHOWBITMAP.....	733
SHOWFILE.....	733
SOLVERRETURN.....	734
SOLVETYPE.....	734
STOPSURF.....	737
SUB.....	737
SURFTYPE.....	737
TELECENTRIC.....	737
TESTPLATEFIT.....	737
THIC.....	737
TIMER.....	737
TOLERANCE.....	738

UNLOCKWINDOW	738
UPDATE.....	738
VEC1, VEC2, VEC3, VEC4.....	739
WAVL, WWGT	739
XDIFFIA	739
ZBFCLR	740
General comments about all ZBF related keywords	740
ZBFMULT.....	740
ZBFPROPERTIES	741
ZBFREAD	741
ZBFRESAMPLE.....	742
ZBFSHOW	742
ZBFSUM	742
ZBFTILT	743
ZBFWRITE.....	743
ZRDAPPEND.....	743
ZRDFILTER	744
ZRDPLAYBACK.....	744
ZRDSAVERAYS	744
ZRDSUM.....	745
Example Macro 1	745
Example Macro 2	746
Calling a Macro from within a Macro.....	747

ZEMAX EXTENSIONS

Chapter 26

Introduction	749
Requirements for writing Extensions	749
Establishing the link	750
Comments on the DDE timeout interval.....	750
Terminating the link.....	750
Extracting data from ZEMAX	751
The data items	751
CloseUDOData	751
DeleteConfig	751
DeleteMCO	751
DeleteMFO.....	751
DeleteObject	751
DeleteSurface	751
ExportCAD	751
ExportCheck.....	752
FindLabel	752
GetAddress	752
GetAperture.....	752
GetApodization.....	753
GetAspect	753
GetBuffer.....	753
GetComment.....	753
GetConfig	753
GetDate.....	753
GetExtra	753
GetField.....	754
GetFile.....	754
GetFirst	754
GetGlass	754

GetGlobalMatrix	754
GetIndex.....	754
GetLabel.....	755
GetMetaFile.....	755
GetMulticon	755
GetName.....	756
GetNSCData	756
GetNSCMatrix.....	756
GetNSCObjectData.....	756
GetNSCObjectFaceData.....	757
GetNSCParameter.....	757
GetNSCPosition	757
GetNSCProperty	758
GetNSCSettings.....	758
GetOperand	758
GetPath	758
GetPolState.....	758
GetPolTrace	758
GetPolTraceDirect.....	759
GetPupil	759
GetRefresh.....	759
GetSag	760
GetSequence	760
GetSerial	760
GetSettingsData.....	760
GetSolve	760
GetSurfaceData	760
GetSurfaceDLL	762
GetSurfaceParameter	762
GetSystem	762
GetSystemAper.....	762
GetSystemProperty.....	762
GetTextFile.....	762
GetTol	763
GetTrace	763
GetTraceDirect.....	764
GetUDOSystem	764
GetUpdate.....	765
GetVersion	765
GetWave	765
Hammer	765
ImportExtraData	765
InsertConfig.....	765
InsertMCO.....	765
InsertMFO	766
InsertObject.....	766
InsertSurface.....	766
LoadDetector.....	766
LoadFile	766
LoadMerit	766
LoadTolerance	767
MakeGraphicWindow.....	767
MakeTextWindow.....	767
ModifySettings.....	768

NewLens	768
NSC coherentData	768
NSC detectorData	768
NSC Trace	768
OpenWindow	769
OperandValue	769
Optimize	769
PushLens	769
PushLensPermission	769
QuickFocus	770
ReleaseWindow	770
SaveDetector	770
SaveFile	770
SaveMerit	770
SaveTolerance	770
SetAperture	771
SetBuffer	771
SetConfig	771
SetExtra	771
SetField	771
SetFloat	772
SetLabel	772
SetMulticon	772
SetNSCObjectData	772
SetNSCObjectFaceData	772
SetNSCPosition	772
SetNSCParameter	773
SetNSCProperty	773
SetNSCSettings	773
SetOperand	773
SetPolState	773
SetSettingsData	773
SetSolve	773
SetSurfaceData	774
SetSurfaceParameter	774
SetSystem	774
SetSystemAper	774
SetSystemProperty	774
SetTol	774
SetUDOData	774
SetUDOItem	775
SetVig	775
SetWave	775
WindowMaximize	775
WindowMinimize	775
WindowRestore	775
Tracing large numbers of rays	775
Step 1: Placing the ray data in the array	776
Mode 0: Similar to GetTrace	776
Mode 1: Similar to GetTraceDirect	777
Mode 2: Similar to GetPolTrace	778
Mode 3: Similar to GetPolTraceDirect	779
Mode 5: For tracing non-sequential rays	780
Step 2: Pass the array to ZEMAX	781

How ZEMAX calls the client.....	781
Generating a text window	782
Generating a graphic window	782
ADDRESS.....	782
BOX.....	783
DATA.....	783
DATE.....	783
FRAME.....	783
GRID	783
LENSNAME	783
LINE	783
NOFRAME	783
PENSTYLE	783
TEXT	784
TITLE	784
A sample Extension program.....	784
The DDE_DEMO code.....	784
A simplified technique for implementing Extensions.....	785
The ZCLIENT program.....	785
A sample Extension program using ZCLIENT	786
Writing Extensions from the PhasePlot template	786
Getting analysis data using GetTextFile and GetMetaFile.....	786

About this document

ZEMAX is available in two different editions: ZEMAX-SE (Standard Edition) and ZEMAX-EE (Engineering Edition). This manual covers both editions, however certain features are unique to ZEMAX-EE, as noted in the text. If a feature is available in ZEMAX-EE, but not ZEMAX-SE, the following message will usually be shown in the manual where the feature is described:



This feature is only available in the EE edition of ZEMAX.

Note that ZEMAX-EE does everything ZEMAX-SE does, plus additional features. This manual covers both editions of ZEMAX running on Microsoft operating systems. ZEMAX® is a registered trademark of Radiant ZEMAX LLC.

What does ZEMAX do?

ZEMAX is a program which can model, analyze, and assist in the design of optical systems. The interface to ZEMAX has been designed to be easy to use, and with a little practice it can allow very rapid interactive design. Most ZEMAX features are accessed by selecting options from either dialog boxes or pull-down menus. Keyboard shortcuts are provided for quickly navigating or bypassing the menu structure. This manual provides explanations of the conventions used in ZEMAX, descriptions of procedures, and a description of available features.

What doesn't ZEMAX do?

Neither the ZEMAX program nor the ZEMAX documentation will teach you how to design lenses or optical systems. Although the program will do many things to assist you in designing and analyzing optical systems, you are still the designer. The ZEMAX documentation is not a tutorial on optical design, terminology, or methodology. Technical support available to ZEMAX users includes assistance in using the program, but does not include tutoring on fundamental optical design principles. If you have little or no experience in optical design, you may want to read up on any of the many good books available on the subject. The following table lists some (but by no means all) of the books which will aid in your education.

REFERENCES ON LENS DESIGN

Author	Title	Publisher
Bass	Handbook of Optics	McGraw-Hill
Born & Wolf	Principles of Optics	Pergamon Press
Fischer & Tadic-Galeb	Optical System Design	McGraw-Hill
Geary, Joseph M.	Introduction to Lens Design: With Practical ZEMAX Examples	Willmann-Bell
Hecht	Optics	Addison Wesley
Kingslake, Rudolph	Lens Design Fundamentals	Academic Press
Laikin, Milton	Lens Design, Third Edition	Marcel Dekker
Mahajan, Virendra	Aberration Theory Made Simple	SPIE Optical Engineering Press
O' Shea, Donald	Elements of Modern Optical Design	John Wiley and Sons
Rutten and van Venrooij	Telescope Optics	Willmann-Bell
Shannon, Robert	The Art and Science of Optical Design	Cambridge University Press

Author	Title	Publisher
Smith, Gregory Hallock	Practical Computer-Aided Lens Design	Willmann-Bell, Inc.
Smith, Warren	Modern Optical Engineering	McGraw-Hill
Smith, Warren	Modern Lens Design	McGraw-Hill
Welford	Aberrations of Optical Systems	Adam Hilger Ltd.
Welford	Useful Optics	University of Chicago Press

Most importantly, ZEMAX is not a substitute for good engineering practices. No design should ever be considered finished until a qualified engineer has checked the calculations performed by the software to see if the results are reasonable. This is particularly important when a design is to be fabricated and significant costs are involved. It is the engineers responsibility to check the results of ZEMAX, not the other way around.

Learning to use ZEMAX

There are many more tutorials and application articles available from the ZEMAX Knowledge Base. The Knowledge Base is available at www.zemax.com/kb. Short courses which cover beginner to advanced use of ZEMAX are offered throughout the year, please check our web site, www.zemax.com, for details.

System requirements

ZEMAX requires the current version of the Windows operating system, 200 Mb (or more) of available storage space, a minimum display resolution of 1024 x 768 pixels, and Internet and e-mail access for program installation, program updates, and technical support. A USB port is required for the key (see below).

The amount of RAM ZEMAX requires varies substantially depending upon the nature of the optical system being modeled and the type of analysis being done. For conventional imaging systems, 256 Mb of RAM is adequate and should be considered a minimum requirement. For modeling of very complex objects, extensive physical optics, or scattering and illumination analysis, substantially more memory may be required for optimal performance. For these types of analysis, 512 Mb of RAM should be considered the minimum, and better performance will be achieved with 1 Gb or more of RAM. It is possible to perform physical optics analysis with ZEMAX using very large arrays that require substantially more RAM than this, see "Memory requirements" on page 637 for details.

Multiple processor computers

Many ZEMAX features are designed to run on multiple processors simultaneously. If more than one CPU is installed in the computer system, ZEMAX will automatically detect and use the additional processors.

Installation procedure

For complete installation instructions please visit our web site, www.zemax.com.

Policy on the key

The key supplied with the ZEMAX software is worth the full purchase price of the software. If the key is lost or stolen, it will not be replaced without payment of the full purchase price. Insure the key as you would any other business or personal asset of comparable value. Defective keys, or keys that become inoperable, will be replaced at no charge if the key has a current support contract. Defective keys must be returned before a replacement key will be provided.



If the key is lost or stolen, it will not be replaced without payment of the full purchase price.

Definition of support

The purchase of a new license to use ZEMAX includes a period of one year of support. Additional years of support may be purchased at the time the new license is purchased or at a later date. The support period is

associated with a specific ZEMAX key. To "have support" for a given ZEMAX key means the support period for that key has not expired. Support consists of two benefits: technical support and program updates.

Technical support is defined as Radiant ZEMAX LLC responding to questions that require an engineer to resolve. This includes answering telephone calls, e-mail, faxes, or letters regarding how to use ZEMAX, understanding ZEMAX features, problems with the key, program installation, or bug reporting.

Program updates include new features as they are developed, along with any bug fixes recently incorporated in the program. Updates also include new electronic documentation. Program updates are available for download from www.zemax.com, for all users who have support.

Getting technical support

If you have any questions about installing or using ZEMAX, please try the following suggestions to find the information you require:

- 1) Look in the Table of Contents to see if there is a chapter or section on the subject.
- 2) Check the index (at the back of this document) to see if it is referenced.
- 3) Look for a sample file of the type of lens you are trying to set up, if appropriate.
- 4) Visit the ZEMAX Knowledge Base at www.zemax.com/kb to see if articles have been written on the subject.

If you still cannot find the information you require, you may need to call, FAX, e-mail, or write for technical support. The phone numbers and address are on the title page of this manual. If you telephone, try to be in front of your computer when you call. Note that technical support is available only if you have support for the ZEMAX key you are using, as described in the previous section.

Policy on bug fixes

All non-trivial computer programs have bugs. Every effort is made to find and fix all bugs before a new version is released. However, the program is sufficiently complex such that even a talented team of pre-release program testers will not find every bug. Radiant ZEMAX LLC therefore may provide periodic bug fix releases.

If you find an apparent bug, please report it to us. Try to isolate the exact actions which cause or lead up to the bug. Try to see if the error is unique to the lens file you are working with. Bugs are usually easier to find if they also appear on the example files provided with the program.

If you find a bug which causes significant degradation of program performance while you have support, the bug fix will be provided free of charge. Radiant ZEMAX LLC reserves the right to determine what bugs constitute significant degradation. If you do not have support, you may be required to renew your support before receiving the bug fix; in this case you will receive the latest program version.

Bug fixes may not be provided if the problem is traced to incompatible, non-standard, or insufficient hardware, or if you are using an outdated version of the operating system.

Introduction

This chapter describes the conventions used by the ZEMAX user interface, and describes some keyboard shortcuts for common window operations.

Types of windows

ZEMAX has different types of windows, each of which serves a different purpose. The types of windows are:

The main window: This window has a large blank area, with a title bar, menu bar, and toolbar at the top. The commands available on this menu bar apply in general to the current optical system as a whole.

Editor windows: There are six different editors: the Lens Data Editor, the Merit Function Editor, the Multi-Configuration Editor, the Tolerance Data editor, and in ZEMAX-EE only, the Extra Data Editor and the Non-Sequential Components Editor.

Graphic windows: These windows are used to display graphic data such as layouts, ray fans, and MTF plots.

Text windows: Text windows are used to display text data such as prescription data, aberration coefficients, and numerical data.

Dialogs. Dialogs are pop-up windows which cannot be resized. Dialogs are used to change options or data such as field angles, wavelengths, apertures, and surface types. They are also used extensively to change options on graphic and text windows, such as for changing the number of rays on a layout plot.

All windows can be moved or resized (except dialogs) using the standard mouse or keyboard commands. If you are not familiar with these procedures, see any good book on using Windows or the Windows documentation.

Main window operations

The main window bar has several menu headings. Most menu headings correspond to a chapter of the same name later in this manual. See the specific chapter for detailed instructions on using the features available from each menu option. These are the menu bar headings:

File: Used to open, close, save, and rename (save as) lens files.

Editors: Used to invoke (make appear) any of the editor windows.

System: Used to define properties of the optical system as a whole.

Analysis: Analysis groups together features which do not change the lens data, but compute numerical or graphical results from the lens data. These features include layouts, ray fans, spot diagrams, diffraction calculations, and more.

Tools: Tools are features which can change the lens data or perform advanced computations on the system as a whole. These include optimization, tolerancing, test plate fitting, and more.

Reports: Reports are used to document a lens design. These features include system data summaries, surface data summaries, and report graphics.

Macros: Used to edit and run ZPL macros.

Extensions: Provides access to ZEMAX Extensions, which are compiled features added to ZEMAX.

Window: Selects from the list of currently open windows which window to move to the front of the display.

Help: Provides access to the program documentation.

Most of the commonly used menu options have keyboard equivalents which may be faster to use. For example, typing Ctrl-Q will quit ZEMAX. The keyboard shortcuts are listed along side the menu options.

A handy shortcut key for switching between windows within the main window is Control-Tab. This will advance control to the next window in the list of windows maintained by ZEMAX.

The main window also displays a row of buttons just below the menu bar. This row of buttons is used to quickly select common operations. All of the buttons represent functions available on the menus. The button definitions may be changed on the File, Preferences dialog box, on the button bar tabs. The three letter mnemonics used

by the buttons are defined on that dialog box tab. A display resolution of 1024 x 768 or higher is recommended for displaying buttons.

Editor windows operations

The editor windows are used primarily for entering lens and merit function data. Each editor is similar to a spreadsheet, with rows and columns. The intersection of a row and a column forms a cell. If the editor is the active window (the one with a highlighted title bar) then one cell will be shown highlighted or in reverse colors. This cell is called the active cell, and it has what is called the input "focus". The inverted color of the cell is called the cursor, although it is not a cursor in the usual sense.

Having the input focus means that any data typed on the keyboard will be sent to the active cell. The exceptions are control commands such as cursor keys or control key combinations, which are sent directly to the main window. To modify the data in the active cell, type in the new data and press the Enter key when finished.

To add an incremental value to a cell, type a plus sign and then the increment, then Enter. For example, to change a 12 to 17, type "+5" and Enter. The "*" multiply and "/" divide symbols also work. To subtract a value, type a minus sign and a space followed by the value to subtract. The space is required to distinguish between subtraction and entering a negative number.

To modify a portion of the contents of a cell without retyping the entire value, first highlight the cell, then press the back space or the F2 function key. The left cursor, right cursor, home, and end keys can then be used to navigate within the cell for editing. The mouse may also be used to select and replace portions of text. Once changes have been made to the data in the cell, pressing Enter will complete editing and leave the cursor on that cell. Pressing up or down cursors will also complete editing, and move the cursor accordingly. Pressing Tab or Shift-Tab will complete editing and move the cursor right or left.

To abort editing of any cell, press the escape key.

The left, right, up and down cursor keys will move the cursor accordingly. Pressing the control key and the left, right, up and down keys simultaneously will move the editor display one page at a time in the appropriate direction. The Tab and Shift-Tab keys also move the cursor right and left.

The page up and page down keys move the cursor one page at a time. Control-page up and control-page down move the cursor to the top and bottom of the current column. The home and end keys will move the cursor to the first column first row and first column last row, respectively. Control-home and control-end will move the cursor to first and last row of the last column, respectively.

Clicking once in any cell will move the cursor to that cell. Double clicking on a cell will invoke the solve dialog box for that cell if one exists. Right mouse clicking will also invoke the solve dialog box for that cell if one exists.

All editors support adjustable column widths. To change the size of the column, place the mouse pointer in the top row of the spreadsheet in the title area directly on one of the vertical column separator lines. When the mouse pointer changes to the resize icon, press the left mouse button and then drag the column width left or right to the desired size. Menu options for sizing, hiding, and unhiding columns are available in the "View" menu of all editors.

Graphic windows operations

Graphic windows have the following menu items:

Update: This recomputes the data displayed in the window with the current settings.

Settings: Invokes a dialog box which controls the options for that window.

Print: Prints the contents of the window.

Window: The submenus under Window are:

Annotate: See "Using the annotation feature" on page 44 for details. The items under Annotate are:

Line: Draws a single line on the graphic window.

Arrow: Draws a line segment with a directional arrow at the end point.

Text: Prompts for and then draws text on the graphic window.

Box: Draws a box on the graphic window.

Edit: Allows extended editing of annotations.

Copy Clipboard: Copies the window contents to the Windows clipboard. See "Using the Windows clipboard" on page 48 for details.

Export: Exports the displayed graphic as a Windows Metafile, BMP, or JPG file.

Lock: If selected, the window will be converted into a "static" window whose data cannot change. The locked window contents may be printed, copied to the clipboard, or saved in a file. The application for this feature is for comparing the results of different lens files. Once a window is locked, it cannot be updated, and so any new lens files which are subsequently loaded may be analyzed and compared to the locked window results. Once the window is locked, it cannot be updated until unlocked. See also "Lock/Unlock All Windows" on page 265.

Unlock: Unlocks a window previously locked. See also "Lock/Unlock All Windows" on page 265.

Clone: This selection will open a new window whose settings and displayed data are initially identical to the current window. This feature is useful for creating a new window based upon the settings in the original window. The cloned window acts like any other window after it is created, so it may be updated or have its settings changed independent of the first window.

Aspect Ratio: The aspect ratio may be selected to be 3 x 4 (height x width) which is the default, or 3 x 5, 4 x 3, or 5 x 3. The latter two are taller than they are wide. The default aspect ratio may be set on the Graphics tab of the File, Preferences dialog box.

Active Cursor: The active cursor displays the value of the coordinates the cursor is currently pointing to in the title bar of the window when the cursor covers an "active" region of the displayed graphic. On most X vs. Y type graphs, the meaning of the displayed values is obvious. On some graphics, such as the 3D Layout, the displayed image is a projection of a 3D object on a 2D plane. The projection of the image renders the coordinate data displayed by the active cursor less meaningful if the image has been rotated. Not all graphics support the active cursor. The active cursor is by default "off" but can be turned on and off by choosing this menu option. The active cursor can be set to automatically be on or off when a new graphic window is created on the Graphics tab of the File, Preferences dialog box.

Configuration: Selects either the current or any specific configuration for the data to be displayed. The default is "current" which means the data in the window is displayed for the active configuration. Some analysis windows, such as the 3D layout, Report Graphics, and Spot Diagrams, independently allow selection of one or more configurations to be displayed on the "settings" dialog box. That option, if present on the settings box, will override this menu option, and this option will be grayed out.

Overlay: Provides a list of all open graphical windows; any of these may be selected for overlaying with the currently displayed data. The overlay feature is useful for comparing two similar graphs or layouts to detect small changes.

Text: Displays a text listing of the data in a new window. Not all graphics windows support this option. Some ZEMAX graphic windows compute both a text and graphic at the same time, and store the text version for later use. These features will display the text version of the window nearly instantly, without the need to recompute the data. Other windows do not yet support this feature, and the data will be recomputed and then displayed in a text window.

Zoom: Controls the "zooming" in on smaller areas of the graphic. See "Using pan and zoom" below. When in Zoom Mode, to zoom in on an area of a graph, click with the left mouse button and drag to select the region to zoom in on. Zoom Mode is active when Spin Mode (see below) is Off. The submenus under the Zoom menu are:

In: Zooms in by 2X centered on the current center position.

Out: Zooms out by 2X.

Last: Restores the previous zoom setting.

Unzoom: Restores the complete view of the graphic.

Spin: Controls the rotation of the view when using the Shaded Model and Object Viewer windows.

Off: Turns Spin Mode off. Dragging the mouse with the left mouse button down will zoom.

XY, XZ, YZ: Dragging the mouse with the left button down will rotate the view around the specified axes.

Reset: Restores the view to the unrotated perspective.

There are two mouse shortcuts available when using graphic windows:

Double clicking anywhere in the body of graphic window will update the contents. This is the same as selecting Update.

Clicking with the right mouse button anywhere in the body of a graphic window will invoke the Settings dialog box.

Using the annotation feature

There are several ways to annotate graphics windows with custom lines, arrows, boxes, and text notes. One way is to choose Window, Annotate from any graphic window menu bar, then select Line, Arrow, Text, or Box. To draw a line, choose the Line menu item, then click at the starting position of the line, holding down the left mouse button, then drag the crosshairs to the ending position of the line, then release the mouse button. A similar procedure will draw an arrow or box.

To add text to the window, select Window, Annotate, Text. A text entry dialog box will appear. Type in the desired text, then click "OK", then click on the desired location for the text on the window display area.

For more precise control over the exact locations of lines and text, as well as control over the text font, and the ability to add more complex annotations, choose Annotate, Edit from the graphic menu. This will invoke the annotation editor, which consists of a text editor and several buttons. There is also a single checkbox for enabling or disabling annotations for the graph.

The text edit field is used to define the annotations to be applied to the graph. To insert a new line, use the keyboard command Ctrl-Enter.

There are several supported commands, each with a specific syntax:

```
TEXT "string" x y angle fontx fonty
```

The TEXT command will write any text within the double quotes at the location specified by x and y, at an angle in degrees given by angle, using a fixed font whose width and height are given by fontx and fonty. The coordinates are in normalized units, where the left edge of the graph has a coordinate of x = 0.0, the right edge x = 100.0, the bottom edge is y = 0.0, and the top edge is y = 100.0. The origin is the lower left hand corner. The "angle" value is in degrees. The fontx and fonty units are in arbitrary units. The angle, fontx, and fonty values may all be left undefined, and then default values will be used.

```
LINE x1 y1 x2 y2
```

The LINE command draws a straight line from x1, y1 to x2, y2. The units and coordinate system are as described for x and y in the description of the TEXT command.

```
ARROW x1 y1 x2 y2 SIZE
```

The ARROW command draws a one-headed arrow pointing from x1, y1 to x2, y2. The units and coordinate system are as described for x and y in the description of the TEXT command. If Size is 1.0 or is omitted, the arrow head is drawn at the default size. To scale the default arrow head size, use any other floating point value for size. For example, a size value of 2.0 will make the head twice as large as the default, a value of 0.5 will make it half the default size, etc.

```
BOX x1 y1 x2 y2
```

The BOX command draws a box with opposite corners from x1, y1 to x2, y2. The units and coordinate system are as described for x and y in the description of the TEXT command.

```
ELLIPSE x y rx ry
```

ELLIPSE draws an ellipse centered at x and y with x half width of rx and half height of ry. If ry is the same as rx or if ry is omitted, a circle of radius rx results.

There are also several buttons on the annotation dialog box:

OK: Accepts the annotations as displayed and exits.

Cancel: Reverts back to the last annotations and exits.

Save: Opens a "Save As" type box where the file name may be specified. The annotations are saved in the user named file.

Load: Opens a "Load: type box where the file name to load may be selected. The loaded file contains the annotations to use.

Reset: Clears the edit buffer.

Help: Invokes the help function.

Using pan and zoom

Any graphic window may be panned (scrolled left, right, up and down) or zoomed (magnified). To activate the pan and zoom, choose any ZEMAX graphic window, then click the left mouse button and hold the button down for 1/2 second anywhere in the window. The cursor will change from an arrow to a cross. Now drag the mouse down and right to define a rectangle of the desired size which covers the area to zoom in on. Now, let go of the left mouse button. The selected area will be magnified to fill the window, while maintaining the aspect ratio of the plot.

To pan, drag the scroll bars on the sides of the window. Panning is only allowed after a graphic has been zoomed.

To restore the graphic to the original size, select Zoom, Unzoom from the Graphic window menu bar.

Keyboard shortcuts are supported when using pan and zoom, see "Summary of useful shortcuts" on page 47 for details.

Text windows operations

Text windows have the following menu items:

Update: This recomputes the data displayed in the window with the current settings.

Settings: Invokes a dialog box which controls the options for that window.

Print: Prints the contents of the window.

Window: There are five submenus under Window:

Copy Clipboard: Copies the window contents to the Windows clipboard. See "Using the Windows clipboard" on page 48 for details.

Save Text: Saves the displayed text data in a file.

Lock: If selected, the window will be converted into a "static" window whose data cannot change. The locked window contents may be printed, copied to the clipboard, or saved in a file. The application for this feature is for comparing the results of different lens files. Once a window is locked, it cannot be updated, and so any new lens files which are subsequently loaded may be analyzed and compared to the locked window results. Once the window is locked, it cannot be updated until unlocked. See also "Lock/Unlock All Windows" on page 265.

Unlock: Unlocks a window previously locked. See also "Lock/Unlock All Windows" on page 265.

Clone: This selection will open a new window whose settings and displayed data are initially identical to the current window. This feature is useful for creating a new window based upon the settings in the original window. The cloned window acts like any other window after it is created, so it may be updated or have its settings changed independent of the first window.

Configuration: Selects either the current or any specific configuration for the data to be displayed. The default is "current" which means the data in the window is displayed for the active configuration.

There are two mouse shortcuts available when using text windows:

Double clicking anywhere in the body of graphic window will update the contents. This is the same as selecting Update.

Clicking with the right mouse button anywhere in the body of a graphic window will invoke the Settings dialog box.

Dialog operations

Most dialog boxes are self explanatory. Usually there are buttons for OK and Cancel which are common in Windows dialog boxes.

Analysis features, such as ray fan plots, have dialog boxes which allow selection of various options. All of these dialog boxes have six buttons:

OK: Causes the window to recompute and redisplay the data with the currently selected options.

Cancel: Reverts to those options selected before the dialog was invoked, and does not recompute the data.

Save: Saves the currently selected options for use as the default settings in the future. See below.

Load: Loads the default values that were previously saved. See below.

Reset: Resets the default values to the "factory default" settings.

Help: Calls the ZEMAX help system. The help page displayed will contain information about the options on the active dialog box.

The Save and Load buttons have a dual functionality. When Save is pressed, the settings are saved for the current lens file, as well as all lenses without specific settings of their own. For example, if lens "A" is loaded, and the number of rays on the layout plot for "A" is set to 15, and then Save is pressed, the new default number of rays for "A" will be 15. Also, any new lenses created, or old lenses without their own settings, will also use 15 as the new default. Now suppose that a subsequent lens "B" is loaded, and the number of rays is changed to 9, and then Save is pressed again. For "B", and all lenses that have no specific settings of their own, 9 will be the new default. However, the original "A" lens will retain the setting of 15 because it has its own private settings.

The Load button works in a similar way. When Load is pressed, ZEMAX checks to see if that specific lens has any previous settings saved. If it does, then those settings are loaded. If not, then ZEMAX loads the last saved settings for all lenses. From the previous example, a new lens "C" will load up the settings of 9 rays, since that was the last setting saved; while lenses "A" and "B" would load up settings of 15 and 9, respectively, because they have private settings.

The Save and Load private settings are stored in a file with the same name as the lens file, except the extension is CFG rather than ZMX. No lens data is stored in the CFG file, just the user defined settings for each analysis feature.

The other options that appear on dialog boxes can be selected using either the keyboard or the mouse. For keyboard control, use the Tab and Shift-Tab key combinations to move from option to option. The space bar can be used to toggle the currently selected checkbox. The cursor keys can be used select items within a drop down list. Pressing the first letter of an entry in a drop down box will also select that entry.

Aborting long computations

Certain ZEMAX tools may require relatively long computation times. For example, the optimization, global optimization, and tolerancing tools may run from several seconds to many days. To terminate the execution of these tools, there is a "Terminate" button displayed which can be clicked. After the terminate button is clicked, ZEMAX gracefully exits the computation and returns control to the main program. Usually, the results of the computation are not available, and are not displayed.

Some analysis features, such as MTF and the image analysis feature, will run for long periods of time in some circumstances. For example, large MTF grids or large image analysis ray densities require longer computation times. However, analysis features do not display a status box or terminate button, because analysis features display their output directly in a window. For this reason, the keyboard command "Escape" is used to terminate lengthy analysis computations. There is no mouse equivalent for this feature; only the Escape key is used.

The Escape key will terminate MTF, PSF, encircled energy, and other diffraction computations. If the Escape key is pressed, control will return to the main program (it may take 1 or 2 seconds) and the data displayed in the

window will be invalid. For the image analysis feature, the escape key will terminate the tracing of new rays, however, rays that have already been traced will be displayed, and data from those rays is accurate, if incomplete.

Summary of useful shortcuts

The following tables summarize useful shortcuts, using both the keyboard and the mouse.

ZEMAX SHORTCUTS

Action	Result
Ctrl + Tab	Moves the input focus from window to window.
Shift + Ctrl + Tab	Moves the input focus from window to window (backwards).
Ctrl + letter	Shortcut for many ZEMAX tools and functions. For example, Ctrl-L invokes the 2D layout plot. All the keyboard shortcuts are listed alongside the menu items.
F1..F10	Function keys are also used as shortcuts for many functions. All the keyboard shortcuts are listed alongside the menu items.
Backspace	When an editor window has the input focus, the highlighted cell can be edited using the backspace key. Once the backspace key is pressed, the mouse and left-right cursors can be used for editing.
Left double click	If the mouse is positioned over any graphic or text window, double clicking will cause the window contents to be recalculated and redrawn. This is the same as selecting Update. For editor windows, invokes the solves dialog box.
Right click	If the mouse is positioned over any graphic or text window, right clicking will invoke the settings box for that window. This is the same as selecting Settings. For editor windows, invokes the solves dialog box.
Tab	Moves to the next cell in the editor windows, or the next field in the dialog boxes.
Shift + Tab	Moves to the previous cell in the editor windows, or the previous field in the dialog boxes.
Home/End	On spreadsheet editors, moves to the top left/bottom left of the current editor. On text windows, moves to the top/bottom of the window. On graphic windows, Zoom in/out.
Ctrl + Home/End	On spreadsheet editors, moves to the top right/bottom right of the current editor. On graphic windows, Last Zoom/Unzoom.
Cursor Key (Left, Right, Up, Down)	On spreadsheet editors, moves one cell at a time. On 3D graphic windows, rotates the view around X and Y.
Ctrl + Cursor Key (Left, Right, Up, Down)	On spreadsheet editors, moves one page at a time. On graphic windows, Pan left, right, up, down.
Page Up/Down	On spreadsheet editors, moves up or down one page at a time. On 3D graphic windows, rotates the view around Z.
Ctrl + Page Up/Down	On spreadsheet editors, moves to top or bottom of column.

WINDOWS SHORTCUTS

Action	Result
Alt + Tab	Toggles between all currently running applications. Extremely useful for quickly moving between ZEMAX and other running applications.
Ctrl + Esc	Invokes the Windows task list. The task list allows selection of another running application.
Alt	Selects the top menu bar of the current application.
Alt + letter	Selects the menu option with that letter. For example, Alt + F selects the File menu.
Tab	Moves to the next option or field.
Shift + Tab	Moves to the previous option or field.
Space bar	Toggles selected checkboxes on or off.
Enter	Equivalent to pressing the highlighted or default button in a dialog box.
letter	Pressing the underlined letter of a word in a drop down box will select that entry in the drop down box.

Using the Windows clipboard

One of the most useful features in Windows is the clipboard. The clipboard is "holding area" for graphics and text. The advantage to using the clipboard is that virtually all Windows programs can either import or export to the clipboard.

Since ZEMAX is primarily used to generate graphical and text data, ZEMAX only supports exporting to the clipboard. Once the desired data has been copied to the clipboard, it is easy for another application, such as a word processor, graphics editor, or desktop publishing system to retrieve the data. For example, the graphics within this manual were generated within ZEMAX, copied to the clipboard, and then pasted into a desktop publishing program from the clipboard.

To get ZEMAX graphics and text into the clipboard, choose the graphic or text window desired, then select Window, Copy Clipboard. Nothing will appear to happen (the data transfer is extremely fast), however the data will now be available to other applications.

To now get the clipboard data into a document processing application, run that application, and choose Paste, usually from the Edit menu of that application. See the documentation for that application for details.

Some Windows applications cannot import ZEMAX graphics, even if they appear correctly in the Windows Clipboard Viewer. In this case, the work around is to use the "Export Metafile" option described in "Graphic windows operations" on page 42. Once the metafile is created, most Windows applications can import the graphic.

Another way of getting ZEMAX graphics into other applications is to perform a screen capture which creates a bitmap image of either the entire display or any single window. To capture the entire screen as a bitmap image, press Ctrl-Print Screen on the keyboard. To capture a single window, select that window and press Alt-Print Screen. Once the image has been captured, the image can usually be pasted into another application using Ctrl-V or Edit, Paste, depending upon the application.

ZEMAX file types by extension

The following table lists some of the extensions used by ZEMAX to indicate different types of data files.

ZEMAX FILE TYPES BY EXTENSION

Extension	Description
AGF	ANSI Glass Format. These are text files which contain the data for the glass catalogs. A text format is used so that additional data may be incorporated in future releases of ZEMAX.
ANN	Annotation files. These binary files store user defined annotations for ZEMAX graphics.
BGF	Binary Glass Format. To accelerate the loading of AGF files, ZEMAX converts AGF files to BGF, which are the version specific binary images of the glass catalogs.
C	C language source code files.
CFG	Configuration file. ZEMAX.CFG is the main configuration file which describes the user-selected options on the environment dialog box. There are numerous other files that may be present which end in CFG, such as RAY.CFG, which contains the user-defined defaults for the ray fan plot.
DLL	Dynamic Link Library files. These files are externally compiled programs linked into ZEMAX at run time. Typically, DLLs are used to implement user-defined surfaces, objects, or other features.
EXE	Executable file. This includes the main program, ZEMAX.EXE and other executables and utilities used by ZEMAX.
IGS, IGES	Initial Graphics Exchange Specification (IGES) file.
IMA	Image file, used by the image analysis feature.
NOT	Element drawing note files.
POB	Polygon object definition file for NSC object.
SES	These files are the "session" files, which define the display configuration and open windows and settings for each saved ZMX lens file.
STP, STEP	STEP files are CAD format files for describing solid model data.
TPD	Test Plate Data files. These contain test plate lists supplied by various optical fabricators, and are used in the test plate fit feature.
ZBF	ZEMAX Beam Files. These are used by the Physical Optics Propagation feature.
ZMF	ZMF are compilations of ZMX format files which are used to define lenses in the stock lens catalogs.
ZMX	ZMX files are the files used to store lens data. The ZMX file is a text file which contains the complete description of the lens, including apertures, wavelengths, prescription data, and the merit function.
ZPL	ZEMAX Programming Language macros.

Introduction

This chapter describes conventions and defines terminology used throughout this manual. Most of the conventions and terms ZEMAX uses are common in the optics industry, however there may be some important differences.

Active configuration

The active configuration is the configuration currently being displayed in the lens data editor. For details see the Chapter "Multi-Configurations".

Angular magnification

The ratio of the paraxial image space chief ray angle to the paraxial object space chief ray angle. The angles are measured with respect to the paraxial entrance and exit pupil locations.

Apodization

Apodization refers to the uniformity of illumination in the entrance pupil of the system. By default, the pupil is always illuminated uniformly. However, there are times when the pupil should have a non-uniform illumination. For this purpose, ZEMAX supports pupil apodization, which is a variation of amplitude over the pupil.

Three types of pupil apodization are supported: uniform, Gaussian, and tangential. For each type (except uniform), an apodization factor determines the rate of variation of amplitude in the pupil. See the discussion on apodization types and factors in the chapter "System Menu".

ZEMAX also supports user defined apodizations, which may be placed on any surface. Surface apodizations behave differently than pupil apodizations, because surfaces need not be located at a pupil. For more information on surface apodizations, see "User Defined" on page 331.

Back focal length

ZEMAX defines the back focal length as the distance along the Z axis from the last surface made of glass to the paraxial image surface for the object at infinite conjugates. If no surfaces are made of glass, the back focal length is the distance from surface 1 to the paraxial image surface for the object at infinite conjugates.

Cardinal planes

The term cardinal planes (sometimes called cardinal points) refers to those special conjugate positions where the object and image surfaces have a specific magnification. The cardinal planes include the principal planes, where the lateral magnification is +1, the anti-principal planes, where the lateral magnification is -1, the nodal planes, where the angular magnification is +1, the anti-nodal planes, where the angular magnification is -1, and the focal planes, where the magnification is 0 for the image space focal plane and infinite for the object space focal plane.

Except for the focal planes, the cardinal planes are conjugates with each other, that is, the image space principal plane is conjugate with the object space principal plane, etc. If the lens has the same index in both object space and image space, the nodal planes are identical to the principal planes.

ZEMAX lists the distance from the image surface to the various image space planes, and lists the distance from the first surface to the various object space planes.

Chief ray

If there is no vignetting, and there are no aberrations, the chief ray is defined to be the ray that travels from a specific field point, through the center of the entrance pupil, and on to the image surface. Note that without vignetting or aberrations, any ray passing through the center of the entrance pupil will also pass through the center of the stop and the exit pupil.

When vignetting factors are used, the chief ray is then considered to be the ray that passes through the center of the vignetted pupil, which means the chief ray may not necessarily pass through the center of the stop.

If there are pupil aberrations, and there virtually always are, then the chief ray may pass through the center of the paraxial entrance pupil (if ray aiming is off) or the center of the stop (if ray aiming is on), but generally, not both.

If there are vignetting factors which decenter the pupil, then the chief ray will pass through the center of the vignetted entrance pupil (if ray aiming is off) or the vignetted stop surface (if ray aiming is on).

The common convention used is that the chief ray passes through the center of the vignetted pupil, while the principal ray passes through the center of the unvignetted stop. ZEMAX never uses the principal ray. Most calculations are referenced to the chief ray or the centroid. Note the centroid reference is generally superior because it is based upon the aggregate effect of all the rays that actually illuminate the image surface, and not on the arbitrary selection of one ray which is "special".

Coordinate axes

The optical axis is the Z axis, with the initial direction of propagation from the object being the positive Z direction. Mirrors can subsequently reverse the direction of propagation. The coordinate system is right handed, with the sagittal X axis being oriented "into" the monitor on a standard layout diagram. The tangential Y axis is vertical.

The direction of propagation is initially left-to-right, down the positive Z axis. After an odd number of mirrors the beam physically propagates in a negative Z direction. Therefore, all thicknesses after an odd number of mirrors should be negative.

Diffraction limited

The term diffraction limited implies that the performance of an optical system is limited by the physical effects of diffraction rather than imperfections in either the design or fabrication. A common means of determining if a system is diffraction limited is to compute or measure the optical path difference. If the peak to valley OPD is less than one quarter wave, then the system is said to be diffraction limited.

There are many other ways of determining if a system is diffraction limited, such as Strehl ratio, RMS OPD, standard deviation, maximum slope error, and others. It is possible for a system to be considered diffraction limited by one method and not diffraction limited by another method.

On some ZEMAX plots, such as the MTF or Diffraction Encircled Energy, the diffraction limited response is optionally shown. This data is usually computed by tracing rays from a reference point in the field of view. Pupil apodization, vignetting, F/#'s, surface apertures, and transmission may be accounted for, but the optical path difference is set to zero regardless of the actual (aberrated) optical path.

For systems which include a field point at 0.0 in both x and y field specifications (such as 0.0 x angle and 0.0 y angle), the reference field position is this axial field point. If no (0, 0) field point is defined, then the field coordinates of field position 1 are used as the reference coordinates instead.

Edge thickness

ZEMAX defines the edge thickness of a surface as:

$$E_i = Z_{i+1} - Z_i + T_i$$

where Z_i is the sag of the surface, Z_{i+1} is the sag of the next surface, and T_i is the axial thickness of the surface. The sag values are computed at the +y semi-diameter of their respective surfaces; Note that edge thickness is computed for the +y radial aperture, which may be inadequate if the surface is not rotationally symmetric, or if surface apertures have been placed upon either of the surfaces.

Edge thickness solves use a slightly different definition of edge thickness. For edge thickness solves only, the sag of the i+1 surface is computed at the semi-diameter of surface i. This method avoids a possible infinite loop in the calculation, since changing the thickness of surface i may alter the semi-diameter of surface i+1, if the latter semi-diameter is in "automatic" mode and the value is based upon ray tracing. See "Thickness: Edge thickness" on page 476 and "Entering semi-diameter data" on page 76 for more information.

Effective focal length

The distance from the rear principal plane to the paraxial image surface. This is calculated for infinite conjugates. Principal plane calculations are always based upon paraxial ray data. The effective focal length is always referenced to an index of refraction of 1.0, even if the image space index is not unity.

Entrance pupil diameter

The diameter in lens units of the paraxial image of the stop in object space.

Entrance pupil position

The paraxial position of the entrance pupil with respect to the first surface in the system. The first surface is always surface 1, not the object surface, which is surface 0.

Exit pupil diameter

The diameter in lens units of the paraxial image of the stop in image space.

Exit pupil position

The paraxial position of the exit pupil with respect to the image surface.

Extra data

Extra data values are used to define certain non-standard surface types. For example, extra data values are used to define the phase of diffractive optic surfaces such as the Binary 1 surface type. For a complete discussion of the extra data values see “Extra data” on page 277.

Field angles and heights

Field points may be specified as angles, object heights (for systems with finite conjugates), paraxial image heights, or real image heights. Field angles are always in degrees. The angles are measured with respect to the object space z axis and the paraxial entrance pupil position on the object space z axis. Positive field angles imply positive slope for the ray in that direction, and thus refer to negative coordinates on distant objects. ZEMAX converts x field angles (α_x) and y field angles (α_y) to ray direction cosines using the following formulas:

$$\tan \alpha_x = \frac{l}{n}$$

$$\tan \alpha_y = \frac{m}{n}$$

$$l^2 + m^2 + n^2 = 1$$

where l, m, and n are the x, y, and z direction cosines.

If object or image heights are used to define the field points, the heights are measured in lens units. When paraxial image heights are used as the field definition, the heights are the paraxial image coordinates of the primary wavelength chief ray on the paraxial image surface, and if the optical system has distortion, then the real chief rays will be at different locations. When real image heights are used as the field definition, the heights are the real ray coordinates of the primary wavelength chief ray on the image surface.

ZEMAX uses normalized field coordinates for many features. For information on how field coordinates are normalized, see “Normalized field coordinates” on page 55. To set the field type and values, see “Fields” on page 115.

Float by stop size

Float by stop size is one of the system aperture types supported by ZEMAX. This phrase refers to the fact that the entrance pupil position, object space numerical aperture, image space F/#, and stop surface radius all are specified if just one of them is specified. Therefore, setting the stop radius, and then allowing the other values to

be whatever they are, is a perfectly valid way of defining the system aperture. It is particularly handy when the stop surface is a real, unchangeable aperture buried in the system, such as when designing null corrector optics.

Ghost reflections

Ghost reflections are spurious, unwanted images formed by the small amount of light which reflects off of, rather than refracts through a lens face. For example, the multiple images of the aperture stop visible in photographs taken with the sun in the field of view are caused by ghost reflections. Ghost images can be problematic in imaging systems and in high power laser systems.

Glasses

Glasses are entered by name in the glass column. Available glasses may be reviewed, and new ones entered using the glass catalog tool. See the Chapter "Using Glass Catalogs" for details.

Blanks are treated as air, with unity index. Mirrors can be specified by entering "MIRROR" for the glass type, although this name will not appear in the glass catalog. The index of refraction of the mirror space is always equal to the index of refraction of the media before the mirror.

For information on the affect of temperature and pressure on index of refraction data, see "Defining temperature and pressure" on page 603.

Hexapolar rings

ZEMAX usually selects a ray pattern for you when performing common calculations such as spot diagrams. The ray pattern refers to how a set of rays is arranged on the entrance pupil. The hexapolar pattern is a rotationally symmetric means of distributing a set of rays. The hexapolar pattern is described by the number of rings of rays around the central ray. The first ring contains 6 rays, oriented every 60 degrees around the entrance pupil with the first ray starting at 0 degrees (on the x-axis of the pupil). The second ring has 12 rays (for a total of 19, including the center ray in ring "0"). The third ring has 18 rays. Each subsequent ring has 6 more rays than the previous ring.

Many features which require a sampling parameter to be specified (such as the spot diagram) use the number of hexapolar rings as a convenient means of specifying the number of rays. If the hexapolar sampling density is 5, it does not mean that 5 rays will be used. A sampling of 5 means $1 + 6 + 12 + 18 + 24 + 30 = 91$ rays will be used.

Image space F/#

Image space F/# is the ratio of the paraxial effective focal length calculated at infinite conjugates over the paraxial entrance pupil diameter. Note that infinite conjugates are used to define this quantity even if the lens is not used at infinite conjugates.

Image space numerical aperture (NA)

Image space NA is the index of image space times the sine of the angle between the paraxial on-axis chief ray and the paraxial on-axis +y marginal ray calculated at the defined conjugates for the primary wavelength.

Lens units

Lens units are the primary unit of measure for the lens system. Lens units apply to radii, thicknesses, apertures, and other quantities, and may be millimeters, centimeters, inches, or meters.

Marginal ray

The marginal ray is the ray that travels from the center of the object, to the edge of the entrance pupil, and on to the image surface.

If there is vignetting, ZEMAX extends this definition by defining the marginal ray to be at the edge of the vignetted entrance pupil. If ray aiming is on, then the marginal ray is at the edge of the vignetted stop.

See also the definition of the chief ray.

Maximum field

The maximum field is the minimum radial coordinate that would enclose all the defined field points if the x and y values of each field point were plotted on an Cartesian XY plot. The maximum field is measured in degrees if the field type is angles, or in lens units for object height, paraxial image height, or real image height. The field type is described in "Field angles and heights" on page 53. To set the field type and values, see "Fields" on page 115.

Non-paraxial systems

The term non-paraxial system refers to any optical system which cannot be adequately represented by paraxial ray data. This generally includes any system with tilts or decenters, strong aspheres, axicons, holograms, gratings, cubic splines, ABCD matrices, gradient index, diffractive components, or non-sequential surfaces.

A great deal of optical aberration theory has been developed for systems with conventional refractive and reflective components in rotationally symmetric configurations. This includes Seidel aberrations, distortion, Gaussian beam data, and virtually all first order properties such as focal length, F/#, and pupil sizes and locations. All of these values are calculated from paraxial ray data.

If the system being analyzed contains any of the non-paraxial components described, then any data computed based upon paraxial ray tracing cannot be trusted. ZEMAX will generally use exact real rays rather than paraxial rays for ray tracing through these surfaces and components.

A system which is well described by paraxial optics will have the general property that the real and paraxial marginal ray data converge as the radial entrance pupil coordinate of the rays being traced tends toward zero.

Non-sequential ray tracing

Non-sequential ray tracing means rays are traced only along a physically realizable path until they intercept an object. The ray then refracts, reflects, or is absorbed, depending upon the properties of the object struck. The ray then continues on a new path. In non-sequential ray tracing, rays may strike any group of objects in any order, or may strike the same object repeatedly; depending upon the geometry and properties of the objects.

See also “Sequential ray tracing” on page 59.

Normalized field coordinates

Normalized field coordinates are used in both the ZEMAX program and documentation. There are two normalized field coordinates: H_x and H_y . Normalized field coordinates are convenient because useful field locations may be defined in a manner that does not change with the individual field definitions or field of view of the optical system. For example, the normalized field coordinate (0, 1) is always at the top of the field of view, whether the field points are defined as angles or heights, and regardless of the magnitude of the field coordinates.

There are two methods used to normalize fields: radial and rectangular. The choice of field normalization method is made on the Field Data dialog; for a description see “Field angles and heights” on page 53.

Radial field normalization

If the field normalization is radial, then the normalized field coordinates represent points on a unit circle. The radius of this unit circle, called the maximum radial field, is given by the radius of the field point farthest from the origin in field coordinates. The maximum radial field is then used to scale all fields to normalized field coordinates. Real field coordinates can be determined by multiplying the normalized coordinates by the maximum radial field value:

$$f_x = H_x F_r, \text{ and } f_y = H_y F_r,$$

where F_r is the maximum radial field and f_x and f_y are the field coordinates in field units.

For example, suppose 3 field points are defined in the (x, y) directions using object height in lens units at (0.0, 0.0), (10.0, 0.0), and (0.0, 3.0). The field point with the maximum radial coordinate is the second field point, and the maximum radial field is therefore 10.0. The normalized coordinate (Hx = 0, Hy = 1) would refer to the field coordinates (0.0, 10.0). The normalized coordinate (Hx = 1, Hy = 0) would refer to the field coordinates (10.0, 0.0). Note the normalized field coordinates can define field coordinates that do not correspond to any defined field point. The maximum radial field is always a positive value.

If a fourth field point at (-10.0, -3.0) were added in the above example, the maximum radial field would become

$$F_r = \sqrt{3^2 + 10^2},$$

or approximately 10.44031. Normalized coordinates should always be between -1 and 1, and should also meet the condition:

$$H_x^2 + H_y^2 \leq 1.$$

Otherwise, the field point lies outside of the maximum radial field.

Rectangular field normalization


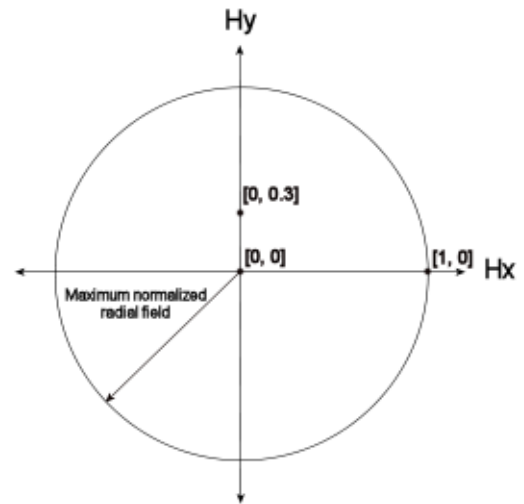
If the field normalization is rectangular, then the normalized field coordinates represent points on a unit rectangle. The x and y direction widths of this unit rectangle, called the maximum x field and maximum y field, are defined by the largest absolute magnitude of all the x and y field coordinates. The maximum x and y field is then used to scale all fields to normalized field coordinates. Real field coordinates can be determined by multiplying the normalized coordinates by the maximum radial field value:

$$f_x = H_x F_x, \text{ and } f_y = H_y F_y,$$

where F_x and F_y are the maximum x and y fields and f_x and f_y are the field coordinates in field units.

For example, suppose 3 field points are defined in the (x, y) directions using object height in lens units at (0.0, 0.0), (10.0, 0.0), and (0.0, 3.0). The field point with the maximum x coordinate is the second field point, and the maximum x field is therefore 10.0. The field point with the maximum y coordinate is the third field point, and the maximum y field is therefore 3.0. The normalized coordinate (Hx = 0, Hy = 1) would refer to the field coordinates (0.0, 3.0). The normalized coordinate (Hx = 1, Hy = 0) would refer to the field point (10.0, 0.0). The normalized coordinate (Hx = 1, Hy = 1) would refer to the field coordinates (10.0, 3.0). Note the normalized field coordinates can define field coordinates that do not correspond to any defined field point. The maximum x and y field are always positive values.

If a fourth field point at (-10.0, -3.0) were added in the above example, the maximum x and y field values would not change. Normalized rectangular coordinates should always be between -1 and 1.



Normalized pupil coordinates

Normalized pupil coordinates are often used in both the ZEMAX program and documentation. There are two normalized pupil coordinates: P_x and P_y . Normalized pupil coordinates are convenient because useful pupil locations may be defined in a manner that does not change with the aperture size or position. For example, the normalized pupil coordinate (0.0, 1.0) is always at the top of the pupil, and therefore defines a marginal ray. The normalized pupil coordinate (0.0, 0.0) always goes through the center of the pupil, and therefore defines a chief ray.

The normalized pupil coordinates represent points on a unit circle. The radial size of the pupil is defined by the radius of the paraxial entrance pupil, unless ray aiming is turned on, in which case the radial size of the pupil is given by the radial size of the stop. For more information on ray aiming see "Ray Aiming" on page 106.

For example, if the entrance pupil radius (not diameter) is 8 mm, then ($P_x = 0.0$, $P_y = 1.0$) refers to a ray which is aimed to the top of the entrance pupil. On the entrance pupil surface, the ray will have a coordinate of ($x = 0.0$ mm, $y = 8.0$ mm).

Note that the normalized pupil coordinates should always be between -1 and 1, and that

$$P_x^2 + P_y^2 \leq 1.$$

A significant advantage of using normalized pupil coordinates is that rays defined in normalized coordinates remain meaningful as the pupil size and position changes. Suppose prior to optimizing a lens design, a ray set is defined to compute the system merit function. By using normalized coordinates, the same ray set will work unaltered if the entrance pupil size or position or object size or position is changed later, or perhaps even during the optimization procedure.

Object space numerical aperture

Object space numerical aperture is a measure of the rate of divergence of rays emanating from the object surface. The numerical aperture is defined as the index of refraction times the sine of the paraxial marginal ray angle, measured in object space. The marginal ray defines the boundary of the cone of light diverging from an object point.

Parameter data

Parameter data values are used to define certain non-standard surface types. For example, parameter data may include aspheric coefficients, grating spacings, or tilt and decenter data. For a complete discussion of the parameter data values see "Parameter data" on page 277.

Paraxial and parabal rays

The term paraxial means "near the axis". Paraxial optics are optics that are well described by the linear form of Snell's law. Snell's law is:

$$n \sin \theta = n' \sin \theta'.$$

For small angles this becomes

$$n \theta = n' \theta'.$$

Many definitions in optics are based upon this assumption of linearity. Aberrations are deviations from this linearity, and so the paraxial properties of optical systems are often considered the properties the system has in the absence of aberrations. Paraxial rays are traced using formulas which assume the optical surface power is based only upon the vertex radius of curvature, ignoring local linear tilts and higher order curvature of the surface.

Paraxial data is computed on a plane tangent to the surface vertex, assuming the vertex radius of curvature is an acceptable approximation to the surface power over the entire aperture of the surface. Certain unusual surface types do not have a paraxial analog, so the real ray tracing calculations are made at these surfaces, even for a paraxial ray.

ZEMAX computes many paraxial entities, such as focal length, F/#, focal position, entrance pupil diameter, and others. These values should be used with caution when the optical system has components which violate the assumption that the vertex curvature is an acceptable approximation to the surface power over the entire aperture of the surface.

For many analysis features, paraxial data is required, typically as a reference against which real rays are measured. To ensure these features work properly, even for optical systems that do not meet the paraxial assumption above, ZEMAX traces "parabasal" rays which are real (real means using Snell's law explicitly) that make small angles with respect to a reference ray, which is usually an axis or chief ray. The parabasal rays are used to compute the limiting properties of the system as the stop size is decreased, which provides a good estimate of the paraxial properties.

The reason ZEMAX uses parabasal rays rather than paraxial formulas is because many optical systems include non-paraxial components. Non-paraxial means these components are not well described by conventional axial first-order theory. This includes tilted or decentered systems, systems using holograms, diffractive optics, general aspheres, and gradient index lenses.

In summary, paraxial ray data is computed using first order approximations to the surface power for tracing rays, while parabasal rays are real, exact ray traces close to a chief or reference ray. Most paraxial data, such as EFL, F/#, and magnification, use paraxial rays and the data is invalid if the optical system is not well described by the vertex power of every surface. Most analysis features in ZEMAX use parabasal rays, to allow these features to work with a greater range of optical systems, including those with optical surfaces not well described solely by their vertex surface power.

Paraxial image height

The paraxial radial size of the image in lens units of the full field image at the paraxial image surface.

Paraxial magnification

The radial magnification, being the ratio of paraxial image height to object height. The paraxial magnification is measured at the paraxial image surface. The value is always zero for infinite conjugate systems.

Paraxial working F/#

The paraxial working F/# is defined as

$$W = \frac{1}{2n \tan \theta},$$

where θ is the paraxial marginal ray angle in image space and n is the index of refraction of image space. The paraxial marginal ray is traced at the specified conjugates. For non-axial systems, this parameter is referenced to the axis ray and is averaged over the pupil. The paraxial working F/# is the effective F/# ignoring aberrations. See also the definition for working F/#.

Primary wavelength

The primary wavelength in micrometers is displayed. This value is used for calculating most other paraxial or system values, such as pupil positions.

Radii

The radius of curvature of each surface is measured in lens units. The convention is that a radius is positive if the center of curvature is to the right (a positive distance along the local z axis) from the surface vertex, and negative if the center of curvature is to the left (a negative distance along the local z axis) from the surface vertex. This is true independent of the number of mirrors in the system.

Real propagation

A real propagation means the rays are propagated in the direction energy would actually flow. See also "Virtual propagation" on page 62 and "Thicknesses" on page 60.

Sagittal and Tangential

The term "tangential" refers to data computed in the tangential plane, which is the plane defined by a line and one point: the line is the axis of symmetry, and the point is the field point in object space. The sagittal plane is the plane orthogonal to the tangential plane, which also intersects the axis of symmetry at the entrance pupil position.

For typical rotationally symmetric systems with field points lying along the Y axis, the tangential plane is the YZ plane and the sagittal plane is the plane orthogonal to the YZ plane which intersects the center of the entrance pupil.

The problem with this definition is that it is not readily extended to non-rotationally symmetric systems. For this reason, ZEMAX instead defines the tangential plane to be the YZ plane regardless of where the field point is, and tangential data is always computed along the local y axis in object space. The sagittal plane is the orthogonal to the YZ plane, and intersects the center of the entrance pupil in the usual way, and sagittal data is always computed along the x axis in object space.

The philosophy behind this convention is as follows. If the system is rotationally symmetric, then field points along the Y axis alone define the system imaging properties, and these points should be used. In this case, the two different definitions of the reference planes are redundant and identical. If the system is not rotationally symmetric, then there is no axis of symmetry, and the choice of reference plane is arbitrary.

One feature, the computation of Fast Semi-Diameters (see "Fast Semi-Diameters" on page 112), does use the "true" tangential plane, which ZEMAX defines as the plane that contains the actual field point and the z axis in object space.

Semi-diameters

The size of each surface is described by the semi-diameter setting. The default setting is the radial distance to the aperture required to pass all real rays without clipping any of the rays. Typing a value for the semi-diameter column results in the character "U" being displayed next to the value. The "U" indicates that the semi-diameter is user defined. When a user-defined semi-diameter is placed on a surface with refractive power (which is done by typing in a value in the appropriate column), and no surface aperture has been defined, ZEMAX automatically applies a "floating" aperture to the surface. A floating aperture is a circular aperture whose radial maximum coordinate is always equal to the semi-diameter of the surface. For more information on surface aperture types, see "Surface properties aperture tab" on page 79.

Semi-diameters on any surface for axial symmetric systems are computed very accurately, as long as the surface does not lie within the caustic of the ray bundle (note this usually occurs at or near the image surface). ZEMAX estimates semi-diameters for axial systems by tracing a few marginal pupil rays. For non-axial systems, ZEMAX estimates the required semi-diameters using either a fixed number of rays or by an iterative technique, which is slower but more accurate. See "Fast Semi-Diameters" on page 112 for details. It is important to note that the "automatic" semi-diameter computed by ZEMAX is an estimate, although it is generally a very good one.

Some surfaces may become so large in aperture that the surface z coordinate becomes multiple valued; for example, a very deep ellipse may have more than one z coordinate for the same x and y coordinates on the surface. For the case of spherical surfaces, this condition is called "hyperhemispheric" and ZEMAX uses this term even if the surface is not a sphere. Hyperhemispheric surfaces are denoted by an asterisk "*" in the semi-diameter column. The indicated semi-diameter is of the outer edge of the surface, which will have a smaller radial aperture than the maximum radial aperture.

Sequential ray tracing

Sequential ray tracing means rays are traced from surface to surface in a predefined sequence. ZEMAX numbers surfaces sequentially, starting with zero for the object surface. The first surface after the object surface is 1, then 2, then 3, and so on, until the image surface is reached. Tracing rays sequentially means a ray will start at surface 0, then be traced to surface 1, then to surface 2, etc. No ray will trace from surface 5 to 3; even if the physical locations of these surfaces would make this the correct path.

See also "Non-sequential ray tracing" on page 55.

Special characters

There are many places in ZEMAX where user defined file, material, glass, or other names may be provided. Generally, ZEMAX allows any characters to be used in these names, except for a few reserved "special characters" The special characters are space, semi-colon, single quote, and tab.

Strehl ratio

The Strehl ratio is one commonly used measure of optical image quality for very high quality imaging systems. The Strehl ratio is defined as the peak intensity of the diffraction point spread function (PSF) divided by the peak intensity of the diffraction point spread function (PSF) in the absence of aberrations. ZEMAX computes the Strehl ratio by computing the PSF with and without considering aberrations, and taking the ratio of the peak intensity. The Strehl ratio is not useful when the aberrations are large enough to make the peak of the PSF ambiguous, or for Strehl ratios smaller than about 0.1.

Surface apertures

Surface apertures include circular, rectangular, elliptical, and spider shaped apertures which can vignette rays. There are also user defined shapes for surface apertures and obscurations; and a "floating" aperture that is based upon the current semi-diameter value. Surface apertures do not effect ray launching or tracing, except for the termination of a ray if it does not pass the surface aperture. Surface apertures have no effect on the system aperture. For more information, see page 79.

System aperture

The system aperture is the overall system F/#, Entrance Pupil Diameter, Numerical Aperture, or Stop Size. Any of these 4 quantities is sufficient to define the other 3 for a particular optical system. The system aperture is used to define the object space entrance pupil diameter, which in turn is used to launch all rays. The system aperture is always circular. Rays may be vignetted after being launched by various surface apertures. There is only one system aperture, although there may be many surface apertures.

Tangential

See "Sagittal and Tangential" on page 59.

Thicknesses

Thicknesses are the relative distance to the next surface vertex in lens units. Thicknesses are not cumulative, each one is only the offset from the previous vertex along the local z axis. The orientation of the local z axis can change using coordinate breaks (see "Coordinate Break" on page 293) or surface tilts and decenters (see "Surface tilt/decenter tab" on page 84).

Thicknesses corresponding to real propagation (see "Real propagation" on page 58) always change sign after a mirror. After an even number of mirrors (including zero mirrors), thickness are positive for real propagations and negative for virtual propagations (see "Virtual propagation" on page 62). After an odd number of mirrors, thicknesses are negative for real propagations and positive for virtual propagations. This sign convention is independent of the number of mirrors, or the presence of coordinate breaks. This fundamental convention cannot be circumvented through the use of coordinate rotations of 180 degrees.

Total internal reflection (TIR)

TIR refers to the condition where a ray makes too large an angle with respect to the normal of a surface to meet the refraction condition as specified by Snell's Law. This usually occurs when a ray with a large angle of incidence is refracting from a high index media to a lower index media, such as from glass to air. When doing sequential ray tracing, rays which TIR are considered errors, and are terminated. Physically, the ray would reflect rather than refract from the boundary, but ZEMAX does not consider this effect when doing sequential ray tracing. For non-sequential ray tracing, rays which TIR are properly reflected.

Total track

Total track is the length of the optical system as measured by the vertex separations between the "left most" and "right most" surfaces. The computation begins at surface 1. The thickness of each surface between surface 1 and the image surface is considered, ignoring any coordinate rotations. The surface which lies at the greatest

z coordinate defines the "right most" surface, while the surface with the minimum z coordinate defines the "left most" surface. Total track has little or no value in non-axial systems.

Units

See "Units" on page 104

Vignetting factors

Vignetting factors are coefficients which describe the apparent entrance pupil size and location for different field positions. ZEMAX uses five vignetting factors: VDX, VDY, VCX, VCY, and VAN. These factors represent decenter x, decenter y, compression x, compression y, and angle, respectively. The default values of all five factors are zero, which indicates no vignetting.

Both the field of view and the entrance pupil of an optical system can be thought of as unit circles. The normalized field and pupil coordinates, defined in "Normalized field coordinates" on page 55, are the coordinates on these two unit circles. For example, the pupil coordinates ($p_x = 0$, $p_y = 1$) refer to the ray which is traced from some point in the field to the top of the entrance pupil. If there is no vignetting in the system, ZEMAX will trace rays to fill the entire entrance pupil during most computations.

Many optical systems employ deliberate vignetting. This means a portion of the rays are intentionally "clipped" by apertures other than the stop surface. There are two common reasons for introducing vignetting in an optical system. First, vignetting decreases the size of the lenses, particularly in wide angle lenses. Second, vignetting may remove a portion of the beam which would be excessively aberrated. Vignetting usually increases the F/# as a function field angle (which darkens the image), but the image quality may improve if the most severely aberrated rays are clipped.

Vignetting factors redefine the entrance pupil for a specific field position. The normalized pupil coordinates are modified using two successive transformations. First, the coordinates are scaled and shifted:

$$P'_x = VDX + P_x(1 - VCX), \text{ and}$$

$$P'_y = VDY + P_y(1 - VCY).$$

The scaled and shifted coordinates are then rotated by the vignetting angle:

$$P''_x = P'_x \cos \theta - P'_y \sin \theta, \text{ and}$$

$$P''_y = P'_x \sin \theta + P'_y \cos \theta,$$

where θ is the vignetting angle VAN. The VDX term can shift the apparent pupil left and right, while VCX makes the pupil larger or smaller in the x direction. Similar results hold for the VDY and VCY values. Note that if the vignetting factors are all zero, the pupil coordinates are left unmodified. Vignetting factors provide a convenient way of designing optics which employ vignetting. However, there are restrictions to using vignetting factors that must be understood.

Some ZEMAX features are capable of tracing rays from arbitrary field positions where no vignetting factors have been assigned. These features may not provide completely accurate results for data computed at field positions other than the defined fields. Some features will remove the vignetting factors for these computations by placing a clear aperture on each surface that vignettes the rays an equivalent amount. Features that automatically remove the vignetting factors are described in the Analysis chapter.

Some features in ZEMAX do not automatically remove vignetting factors for intermediate field positions, such as ray operands in the merit function (operands like REAX that can launch a single ray, for example) or ZPL macros. If the vignetting factors are not removed, ZEMAX will attempt to interpolate the vignetting factors. For rotationally symmetric systems, or systems with field points entirely along the y axis, ZEMAX interpolates between adjacent field points to estimate the vignetting factors to use at intermediate field points. For more general optical

systems with X field values, ZEMAX uses the closest defined field point for determining the vignetting factors for an arbitrary field point.

Once a vignetting factor is defined, it is up to the designer to ensure that rays beyond the apparent pupil are in fact vignetted! If the vignetting factor is used to shrink the size of the lenses, then the lenses should be made no larger than is required to pass the rays which are at the edge of the apparent pupil. If rays from beyond the vignetted aperture are allowed to pass in the real optical system, then the lens performance will not correlate with the computer model.

Identical or nearly identical field coordinates may not be defined with different vignetting factors. Field coordinates must be different by roughly 1E-06 times the maximum field coordinate if two neighboring fields use different vignetting factors. This is required because ZEMAX must determine vignetting factors for any field coordinates, not just those at defined field positions; and identical field coordinates with different vignetting factors have no physical interpretation. The proper way to set up this sort of system is to use multiple configurations, and change the vignetting factors via the multi-configuration editor.

The vignetting factors work with and without ray aiming turned on. If ray aiming is off, then the paraxial entrance pupil is remapped according to the equations given earlier. If ray aiming is turned on, then the remapping is done at the stop surface.

One possible application of vignetting factors is to account for pupil aberration without using the ray aiming feature. This is an advanced trick which can be used to speed up ray tracing in wide angle systems.

Vignetting factors may be defined on the "Field Data" dialog box. See the Chapter "System Menu" for more information. Vignetting factors may also be zoomable parameters; see the Chapter "Multi-Configurations". For more information on the use of vignetting as a design tool, see any of the good books referenced in the first chapter.

Virtual propagation

A virtual propagation means the rays are propagated in a direction opposite to the direction energy would actually flow. Virtual propagations are often useful for placing virtual sources or pupils. See also "Real propagation" on page 58 and "Thicknesses" on page 60.

Wavelength data

Wavelength data are always measured in micrometers referenced to "air" at the current system temperature and pressure. The default system temperature is 20 degrees Celsius, and the default air pressure is 1.0 atmospheres. If the system temperature and/or pressure is modified, or under the control of multi-configuration operands, care must be taken to adjust the wavelengths to the new air temperature and pressure.

Wavelength data is entered on the "Wavelength Data" dialog box; see the "Wavelengths" on page 116 for details.



Wavelength data are always measured in micrometers referenced to "air" at the system temperature and pressure.

Working F/#

Working F/# is defined as

$$W = \frac{1}{2n \sin \theta},$$

where θ is the real marginal ray angle in image space and n is the index of refraction of image space. The marginal ray is traced at the specified conjugates. Working F/# ignores surface apertures but considers vignetting factors.

For off axis field points or non-axial systems, working F/# is determined by the average of the square of the numerical aperture between the axis ray and four marginal rays, at the top, bottom, left, and right side of the vignetted pupil. The average of the square of the numerical aperture of the four rays is converted back to equivalent F/#.

Working $F/\#$ is generally much more useful than image space $F/\#$ because it is based upon real ray data at the actual conjugates of the lens. See also "Paraxial working $F/\#$ " on page 58.

If the marginal rays cannot be traced (due to ray errors) then a smaller pupil is temporarily used to estimate the working $F/\#$. In this case, ZEMAX scales the data to estimate the working $F/\#$ at the full pupil size, even though rays are not traceable at the full aperture.

If the marginal rays are nearly parallel to the chief ray, the resulting $F/\#$ may become so large as to be inaccurate. ZEMAX will automatically "cap" the $F/\#$ at 10,000 if the calculated $F/\#$ becomes larger than 10,000. This result simply means the $F/\#$ cannot be calculated accurately using rays. Such a large $F/\#$ means the output beam is nearly collimated, and many assumptions ZEMAX makes are not valid in this case. The two solutions are to bring the nearly collimated beam to a focus using a paraxial lens, or to use the exit pupil size and position to estimate the $F/\#$. For more information, see "Method To Compute $F/\#$ " on page 112.

For a generalized definition of $F/\#$ useful for evaluating image brightness, see "Effective $F/\#$ " on page 188.

New*Purpose:*

Clears the current lens.

Discussion:

This option restores ZEMAX back to the start-up condition. If the current lens has not been saved, ZEMAX will prompt to save the lens before it is destroyed.

Open*Purpose:*

Opens an existing ZEMAX lens file.

Discussion:

This option opens a previously saved lens file.

Save*Purpose:*

Saves lens file.

Discussion:

This option saves the current lens file. To save the lens under a different name or folder, use Save As.

Save As*Purpose:*

Saves lens file, optionally with another name.

Discussion:

This option saves the current lens file with the option to change the current file name or folder.

Use Session Files*Purpose:*

Toggles the environment setting between using and not using session files. If the menu option is checked, then session files are used.

Discussion:

For a complete discussion of session files, see "Editors" on page 71.

Backup To Archive File*Purpose:*

Saves the *.ZMX lens file, and all associated files to a single *.ZAR archive file.

Discussion:

The main "lens" file in ZEMAX ends in the extension ZMX. However, many other files are used to define critical data for the optical design. These files include glass and gradient index data catalogs, coating data files, user defined aperture files, CAD format files, bulk and surface scatter data, externally defined DLL files, ZPL macros, and many others. The possible large number of these data files, and the fact that these files may be stored in a wide range of folder locations, makes backing up and archiving all the individual files cumbersome and error prone. The same problem occurs if the lens design needs to be sent to a colleague via email. The files must all be collected, backed up, and then stored or sent. To restore the complete design and all the referenced files requires all the files be restored to the correct folder so ZEMAX can find the files.

This feature solves this data backup and restoration problem. The backup process gathers all of the data and files used by the design, including all data and files used by any analysis windows open at the time the backup

archive is created, and stores the data in a single archival file. The ZEMAX Archive format (ZAR) is a one-file solution to making a complete backup of the optical design.

The archive works by storing a complete copy of each file the design uses in the ZAR file. The original folder information is not stored, just the file type, name, and data. For example, the ZAR will contain any AGF glass catalogs the design was using. When the design is restored from the ZAR file, the AGF file will be created and placed in the *current* folder for glass catalogs. This procedure allows the ZAR to be restored to another user's machine that may have different drive and folder names for the glass catalogs than the machine used to create the ZAR was using. For a complete list of the files that will be stored in the ZAR before creating the ZAR file, see "Files Used" under "Prescription Data" on page 271.

When the "Backup To Archive" feature is selected, ZEMAX first prompts for the file name to create. This may be an existing file or a new name may be entered. The ZAR file defaults to the current lens file name, but may be given any valid name. The file name does not need to match the lens file name, however, if the ZAR file is renamed, it is important to understand that when restored the *original* lens file name is restored.

When the file name is selected, a dialog box appears with the options "OK" and "Cancel". Select OK to start the archival process or Cancel to abort the archive. Note that the amount of time it takes to create the ZAR depends upon how many files need to be backed up, how large those files are, and whether data compression is applied or not. Most files used by ZEMAX are modest in size, but some files, such as ZRD, source ray DAT, and CAD files are potentially quite large. If any individual file exceeds 50.0 MB in size, ZEMAX will prompt the user to decide if the file should still be placed in the archive file. Note if "No" is selected the file WILL NOT be archived and will not be able to be restored from the ZAR file. In this case a separate backup must be made as the file will not be stored in the archive.

The Backup To Archive feature also supports optional data compression. This data compression is applied to the archive to yield a smaller file size. Older versions of ZEMAX may not have the ability to read these compressed files, and so the use of data compression may be disabled to provide compatibility with older versions of ZEMAX if required. The data compression algorithm used by ZEMAX is designed to be compatible with future versions of ZEMAX and provides reasonably good compression with little overhead. However, the algorithm used does not generally provide the highest possible levels of compression. Compression is not applied to ZRD files, nor any other files larger than 50.0 MB in size, as the compression time becomes unreasonably long on these large files.

At the end of the archive process, the status box will report how many files were saved. Press "Close" to close the archive feature. Once the ZAR is generated, the file may be renamed, backed up, or sent to a colleague, like any other data file. To restore the lens design files from the ZAR, see "Restore From Archive File" below.

Restore From Archive File

Purpose:

Opens an existing ZEMAX archive file and restores all the individual files in the archive.

Discussion:

This feature opens a previously created archive file in ZAR format and (optionally) restores all the files contained in the archive. For more information on ZAR files and how to create a ZAR archive file, see "Backup To Archive File" on page 65.

When this feature is selected, ZEMAX first prompts for the ZAR file name to restore. Then a dialog box with the following options appears:

To Folder: This button allows the user to select the master folder in which the archived files will be restored. The initial default folder is the "Lens Folder" (see "Folders" on page 68). However, after the first archive file is restored, the most recently used folder will be the default, but any valid folder for which the user has read and write permission can be selected. This is the folder where the main ZEMAX file (*.ZMX), config file (*.CFG), session file (*.SES), and any ray database (*.ZRD) files are stored. Most other files, such as glass catalogs (*.AGF) and Dynamic Link Libraries (*.DLL) are stored in the folders the current running ZEMAX is using, as defined by "Folders" on page 68.

Existing Files: This feature tells ZEMAX what to do if the archive contains a file of the same name as an existing file. The options are Overwrite All (the file in the ZAR will overwrite the existing file without warning or saving the

old file), Skip All (the existing files will be retained and the files in the ZAR will be ignored), and Prompt Every File (ZEMAX will ask you whether to overwrite or skip every single file when a duplicate exists).

OK: Begins the restoration of files from the ZAR.

Cancel: Exits the dialog without restoring the files from the ZAR.

List Files: Creates a window that lists all of the files in the ZAR archive, along with the file size and date. If the file is compressed, the compressed file size and not the original file size is listed.

When the restoration is complete, press Close to close the dialog box. ZEMAX will then immediately load the ZMX file that was just restored.

Program Mode

Purpose:

Selects the user interface and program mode.

Discussion:

There are two distinct modes to the ZEMAX User Interface:

Sequential or Mixed Sequential/Non-Sequential Design ("Sequential mode")

Non-Sequential Design ("Non-Sequential mode" or "NSC mode")

In Sequential mode, all program features are available, including the ability to place a non-sequential group upon any surface. In Sequential mode, ZEMAX will trace rays which leave the object surface through all defined surfaces and non-sequential groups. This is the preferred mode for designing imaging systems, or any system requiring optimization, tolerancing, and detailed image analysis.

In Non-Sequential mode, multiple changes are made to the user interface and to the lens data to simplify the analysis of these systems:

Only surface 1 is used, and all NSC objects are placed in a single non-sequential group placed on surface 1. The object and image surfaces are not used.

The Lens Data Editor and Extra Data Editors are not used. The Non-Sequential Components Editor (NSCE) is used instead as the primary editor.

The field dialog box is not used. The wavelengths dialog box is still used to define the wavelengths to be traced.

Only rays which originate from sources placed within the NSCE are traced.

Features which are not used in Non-Sequential mode are removed from the menus and button bars to simplify the user interface. These features include ray fans, MTF plots, spot diagrams, and many other features unique to sequential ray tracing. The primary means of analysis in Non-Sequential mode is tracing rays to one or more detector objects.

Using Non-Sequential mode makes the operation of the program easier for systems which are non-imaging.

It is always possible to switch from Non-Sequential mode to Sequential mode without loss of data. Switching from Sequential mode to Non-Sequential mode will cause all the sequential data to be deleted.

Insert Lens

Purpose:

Inserts surface data from a previously stored lens data to the current lens file.

Discussion:

This option is similar to the "Open" command, but the current lens data is not disturbed. After the file to insert is selected, ZEMAX will prompt for the surface number where the new lens data should be inserted. New surfaces will be inserted into the lens prescription to make room for the new data. The dialog box also has an "Ignore Object" checkbox which by default will ignore the object thickness of the new lens. Therefore, the new lens data to be inserted will start at surface 1 rather than surface 0.

Although this feature can save a lot of typing, the resulting lens may have some extra surfaces, or may need a little hand editing to achieve the desired result.

Preferences

ZEMAX allows users a number of options that can be set and saved so that these options are selected automatically when ZEMAX is run. The main configuration file is "ZEMAX.CFG"; this file can be deleted to return the start-up configuration to the default state.

To save the current settings in a file for future use, choose "Save" and select the file name to store the settings. A copy of these settings will also be saved to the default "ZEMAX.CFG" file and these will be the new preference settings. To load preference settings from a previously saved file, choose "Load" and select the desired file. A copy of the newly loaded settings will be placed in the default file "ZEMAX.CFG" and these will be the new preference settings. Note that both Save and Load will override the current ZEMAX.CFG settings and the displayed settings will be the new preference settings until further changes are made.

The preference options are divided into groups as follows.

Address

Purpose:

These settings determine how the "address" box will appear. The address box can be used to display user defined text such as a company name or drawing number. The address box appears on the bottom right corner of most graphics.

Settings:

Item	Description
Address Line 1	The first line of text to appear in the address box.
Address Line 2	The second line of text to appear in the address box.
Address Line 3	The third line of text to appear in the address box.
Address Line 4	The fourth line of text to appear in the address box, unless file name or zoom position is chosen instead.
Address Line 5	The fifth line of text to appear in the address box, unless file name or zoom position is chosen instead.
Show Line 4 As	Choose either the entered text, the lens file name, or the zoom position.
Show Line 5 As	Choose either the entered text, the lens file name, or the zoom position.
Hide Address	If checked the address box will not appear.

Folders

Purpose:

These settings determine where ZEMAX will place or search for certain types of files.

ZEMAX uses two main folders for storing programs and data. The ZEMAX program, Dynamic Link Library (DLL) files both supplied with ZEMAX and added for user-defined capabilities, and executable files for ZEMAX extensions, are stored under the Program Files folder, typically in C:\Program Files\ZEMAX. This folder is determined upon installation and should generally not need to be changed by the user. This folder is called the "Program" folder and is referred to as <program> elsewhere in this manual.

All other user data is by default stored in the My Documents\ZEMAX folder, where "My Documents" is generally operating system dependent. This folder will usually be C:\Documents and Settings\Username\Documents\ZEMAX on Windows XP and earlier versions of Windows, or C:\Users\Username\Documents\ZEMAX on Vista and Windows 7. ZEMAX will display this root folder as the "Data" folder, and is referred to as <data> elsewhere in this manual. This folder name cannot be modified.

The folders for lens, object, macro, glass catalog, and coating files may optionally be modified from the default using these preference settings. The "?" button may be used to browse for and to select the desired folder. Note that if the default folder names are modified and alternate folders used, then any updated files distributed with

ZEMAX will NOT automatically be updated in the new folder names. For this reason, it is recommended that if the folder names are modified, that none of the files included with the ZEMAX distribution (such as SCHOTT.AGF) are used in the modified folders.

Settings:

Folder Name	Description
Lens	The folder for lens files.
Objects	The folder for Non-Sequential Component object, source, aperture, and other files.
ZPL	The folder for ZPL Macros. Macros may also be placed in subfolders within the specified folder.
Glass	The folder for glass catalog files.
Coating	The folder for coating definition files.
POP	The folder for POP ZBF beam files.
Images	The folder for Image Simulation IMA, BMP, and JPG files.

Graphics

Purpose:

These settings determine the size, color, and behavior of most ZEMAX graphics windows.

Settings:

Item	Description
B/W Display	By default, ZEMAX draws color graphics. If this checkbox is selected, then all graphics will be black and white.
B/W Plots	By default, ZEMAX prints color graphics to printers. If this checkbox is selected, then all printed graphics will be black and white. Only printers which support color will print in color, otherwise, gray scales will be used to simulate color.
Show Options First	If this checkbox is selected, then the settings box will appear before any analysis graphic or text is computed and displayed.
Win X,Y Size	The default x and y size for graphics and text windows in pixels. This can be adjusted to suit the monitor size and resolution.
Background	The background color on graphics windows can be selected from this drop down list.
Metafiles	ZEMAX generates several different types of Windows Metafile formats. Metafiles are used to copy graphics to the clipboard, or to copy graphics to files, where the graphics may then be pasted or imported into other Windows programs. Most 16 bit Windows 3.1 applications used the "16 Bit Standard" format. However, some Windows 3.1 applications use a variation called the "16 Bit Placeable" format. Newer, 32 bit applications use the "32 Bit Enhanced" format. When a 32-bit format is used, the extension used is EMF, for enhanced metafile format.
Metafile Pen Width	Determines the pen width in device dependent pixels for graphics which are exported as Metafiles, via either the clipboard or as files. Larger values yield thicker lines.
Aspect Ratio	The default aspect ratio for ZEMAX graphic windows is 3 x 4, which fits well on a standard 8.5 x 11 inch paper printer. For 11 x 17 inch printing, a 3 x 5 aspect ratio fits better. The 4 x 3 and 5 x 3 aspect ratios are taller than they are wide. This option selects the default aspect ratio for both printing and display graphics. Each graphic window may have it's own aspect ratio modified by using the Window, Aspect Ratio setting.

Item	Description
Use Active Cursor	This setting determines whether or not the active cursor is initially on when creating a new window. Each window may be set to display or not to display the active cursor individually once the window is opened using the Window menu pull down.
Frame Zoomed Graphics	If selected, zoomed graphics which are pasted to the clipboard will be framed; otherwise, the frame will not be shown.
Layouts Rotate Z, Y, X	If selected, all 3D Layout type plots will rotate the displayed lens first in Z, then Y, then X about the window's coordinates. This selection results in rotations that are similar to ZEMAX Coordinate Break conventions if the order flag is zero. If this option is not selected, then the displayed lens will be rotated first in X, then Y, then Z about the window's coordinates. This selection results in rotations that are similar to ZEMAX Coordinate Break conventions if the order flag is not zero.
Spin Mode On By Default	If checked, Shaded Model and Object Viewer windows have Spin Mode enabled when the windows are created. If unchecked, Zoom Mode is enabled.
Highlight Layouts	If selected, then the current surface (in the LDE or EDE) or object (in the NSCE) will be highlighted in all open layout type plots. Moving the cursor from row to row will cause the layout plots to redraw with the selected surface or object highlighted. The highlight color is defined on the Colors tab. Highlighting does slow down navigation of the editors somewhat, because the graphics are continuously redrawn.
Highlight While Scrolling	If selected, then highlighting (see Highlight Layouts above) will occur while scrolling with the vertical scroll bars on the right side of the editors. On complex lens drawings this may interfere with proper updating of the display because the lens drawing updates are not fast enough to keep up with the scroll bar controls.
Fletch Size	Determines the relative size of fletch marks drawn on the various layout plots. The default value of 1.0 draws fletches at the default size; a value of 2.0 would draw the fletches twice as large as normal, and 0.5 would draw fletches at half size. Any value between 0.0 and 100.0 may be used.
Orientation Indicator Size	Determines the relative size of the orientation indicator drawn on the various layout plots. The default value of 1.0 draws the indicator at the default size; a value of 2.0 would draw the indicator twice as large as normal, and 0.5 would draw the indicator at half size. Any value between 0.0 and 10.0 may be used. A value of zero will remove the indicator.
False Color	Selects the "spectrum" or the original ZEMAX "classic" look for false color plots.
Show Grid Lines On Plots	If checked, most ZEMAX graphics will draw grey grid lines on data plots. The intensity of the grey lines is controlled by the Grid Line Intensity setting.
Grid Lines Intensity	An integer value between 0 and 255 that indicates the intensity of the grey grid lines, if shown. A value of 0 will make lines that are pure white, while a value of 255 will make the grid lines pure black.

Miscellaneous

Purpose:

Miscellaneous settings.

Settings:

Item	Description
Text Window Font	The font size in points to be used for text displayed in an analysis window, such as a ray fan text listing. To change the size of text displayed in the editors see "Editors" below.
Date/Time	Select either None for no date or time, Date for the date only, and Date/Time for both the date and the time to be printed on plots and graphics.

Item	Description
Check For Updates	Select how often ZEMAX should check to see if a new version of ZEMAX is available. Updates may be downloaded from www.zemax.com . An active Internet connection is required for ZEMAX to be able to check for updates. The "Check For Updates" option listed under the Help menu can be used to check for updates anytime.
Interface Language	Choose English, Japanese, Chinese, Russian, or Korean. This setting changes the user interface language of ZEMAX. Not all features fully support languages other than English.
ZMX File Encoding	To support user interface languages other than English, ZEMAX by default uses Unicode file encoding for most files, including ZMX files. This control allows ZMX files to be written using ANSI (also called ASCII) encoding instead. The use of ANSI may allow older versions of ZEMAX to read ZMX files created by newer versions of ZEMAX. However, it is never advisable to use older versions of ZEMAX to read newer ZEMAX files as the newer version data format may not be fully compatible with older versions. Session and CFG files are usually not compatible with older versions of ZEMAX. ANSI encoding may only be used if the interface language is English.
TXT File Encoding	This setting controls whether text data saved using the "Save Text" menu option (see "Text windows operations" on page 45) uses ANSI or Unicode file encoding. See "ZMX File Encoding" above for a discussion about Unicode. ANSI encoding may only be used if the interface language is English.

Editors

Purpose:

These settings determine characteristics of the spreadsheet editors. If the editor cell sizes are too narrow to display all the data in each cell, an asterisk "*" symbol will be printed instead of the truncated data.

Settings:

Item	Description
Decimals	Decimals will alter the number of decimal digits displayed on the lens data editor. Selecting "Compact" will vary the number of decimals displayed to minimize the space required to display numbers.
Exponential Above/ Below	ZEMAX displays numeric data in either fixed decimal or exponential (scientific) notation. If the absolute value of the number is either a) greater than or equal to the "Exponential Above" value, or b) not zero but less than or equal to the "Exponential Below" value, then the number is displayed in exponential notation. Otherwise fixed notation is used.
Font Size	The font size in points to be used for text.
Choose Font	Pressing this button opens a font selection dialog that displays the mono spaced fonts currently available. The current font name is displayed on the button for reference.
Auto Update	Controls how and when ZEMAX performs updates on the editor data. The "None" setting means pupil positions, solves, and other lens data editor data are not updated until "Update" is selected from the "System" menu. The "Update" setting will cause an update to be performed whenever new data is typed in to the lens, extra, or multi-configuration editors. "Update All" causes all windows to be updated whenever new data is typed in the editors. See the "System menu" chapter for details on Update and Update All.
Undo	The Undo feature has three selectable states: None, Memory 1 Step, and Disk Multi Step. See the Chapter "Editors Menu" for details on the Undo feature.
LDE Cell Size	The width in characters of a single edit cell in the Lens Data Editor. Larger cells will mean fewer columns can be displayed, but the data is generally more readable.
MFE Cell Size	The width in characters of a single edit cell in the Merit Function Editor.

Item	Description
MCE Cell Size	The width in characters of a single edit cell in the Multi-Configuration Editor.
EDE Cell Size	The width in characters of a single edit cell in the Extra Data Editor.
TDE Cell Size	The width in characters of a single edit cell in the Tolerance Data Editor.
NSC Cell Size	The width in characters of a single edit cell in the Non-Sequential Components Editor.
Default Catalog	The name of the default glass catalog to use. See "Glass Catalogs" on page 106.
Show Comments	If checked, then the surface comments column is displayed in the Lens Data Editor. Otherwise, the column is hidden from view.
Use Session Files	If checked, all open windows will be closed, and new windows will be opened and restored to their original locations when new lens files are loaded. If unchecked, files will be loaded without changing the arrangement of open windows.
Allow Extensions To Push Lenses	If checked, then any ZEMAX extension can replace the lens data in the Lens Data Editor with new data, without warning or prompt. This option should not be checked unless it is intended that an extension control ZEMAX. See "Extensions" on page 275 for more information on extensions.
Color Editor Rows	If checked, then rows in the spreadsheet are color coded. See "Row Color" on page 77.

Printing

Purpose:

These settings determine characteristics of the printed output.

Settings:

Item	Description
Skip Print Dialog	If this box is checked, then ZEMAX will not present the "Print" dialog box which allows selection of printer type and other options when "Print" is selected from any window. If the box is unchecked, then the default printer settings will be used as specified by the Windows Print Manager.
Rotate Plots	If selected, this will cause all printed graphics to be rotated by 90 degrees. This allows a landscape type orientation while printing in portrait mode. See the discussion below.
Plot Width	See the discussion below.
Pen Width	The width of the pen in pixels. Values of 0 result in fine lines, larger values yield thicker lines.
Print Font Size	The font size in points to be used for text when a text window is printed.
Left Graphic Margin%	The left margin of the plot as a percentage of the maximum possible plot width. Only affects graphics printing.
Right Graphic Margin %	The right margin of the plot as a percentage of the maximum possible plot width. Only affects graphics printing.
Top Graphic Margin %	The top margin of the plot as a percentage of the maximum possible plot height. Only affects graphics printing.
Bottom Graphic Margin %	The bottom margin of the plot as a percentage of the maximum possible plot height. Only affects graphics printing.

Item	Description
Left Text Margin	The left margin to be used when printing text files, measured in average character widths. Note that when printing text files in portrait mode, this is the left margin, allowing room for hole punching and binding of the printed page.
Top Text Margin	The top margin to be used when printing text files, measured in lines. Note that when printing text files in landscape mode, this is the left margin, allowing room for hole punching and binding of the printed page.

Discussion:

The "Plot Width" control is different from most other settings in ZEMAX, because it is actually used to tell ZEMAX how wide the plots are rather than how wide they should be. Every printer will print ZEMAX graphics at a different size. To get accurate scaling on layout plots and scale bars, ZEMAX must be informed as to how big the plots are when printed. With this information, ZEMAX can print exactly at 1:1, or 2:1, and so on.

Pressing the "Plot Width" button will invoke the Print dialog box. This is the same box that is displayed when printing graphics from ZEMAX. This box allows selection of the printer driver, and usually allows selection of printer specific items such as resolution, orientation, and perhaps other options which vary from printer to printer. Using the Print dialog box, select the driver and mode you normally like to print in. Note ZEMAX prints graphics in "landscape mode" by rotating the plot 90 degrees and then using the portrait setting. This is done because virtually all printers default to portrait mode. It is therefore better to stay in portrait mode all the time, and use the rotate plots feature. ZEMAX will use these settings to determine how wide the actual plot will be when printed, and place this width, in inches, in the "plot width" edit box. Note that the plot width is automatically recalculated for the default printer settings whenever the print orientation or margins are adjusted, or the "Reset" button is pressed.

Once the actual plot width has been computed, the scaling on layout plots will be very accurate. However, it will only be accurate if the same printer driver and mode settings are actually used for printing. If a different printer driver or mode is selected at print time, then the plot width scaling will not automatically be recalculated. To get correct scaling when printing to a new printer or using a different printer mode, the plot width environment variable must be reset using the procedure described above.

Lastly, there are times when it may be required to override the default settings for the plot width. For example, if the final printed output is to be reduced to a certain size for inclusion in another document, then the known final size can be used to get exact scaling at the final image size. To accomplish this, enter the known final image width in inches in the plot width edit box, and press Save. All subsequent printing will be done assuming the final image size specified.

Note that the precise scaling controls specified only affect the Layout and Element Drawing features, since all other plots are scale independent.

Colors 1-12, Colors 13-24

These settings define the colors used for drawing graphics and for the background color of rows in the various spreadsheet editors. Each color is used by the corresponding wavelength or field position when graphing ray fan data, spot diagrams, and other data. For example, typically wavelength or field number 1 uses color 1, field or wavelength 2 uses color 2, and so on. The Red, Green, and Blue values define the color. Each entry must be between 0 and 255. Colors are defined using the RGB values for 24 bit color (16 million colors), but will only display up to the resolution provided by the graphics hardware present. The resulting colors are shown to the right of each RGB value. When using the defined colors in the spreadsheet editors, ZEMAX automatically blends the color with white to make the foreground text more readable.

The Highlight Color is the color ZEMAX uses to draw lens surfaces and NSC objects currently selected in the LDE and NSCE, respectively.

The Gradient Top and Gradient Bottom colors are used to define the colors for the gradient fill background in the Shaded Model displays.

Buttons 1-16, Buttons 17-32, Buttons 33-48

Purpose:

These settings determine which functions are available from the ZEMAX button bar. The buttons are only user-definable in sequential mode. In NSC mode, ZEMAX uses a fixed set of button definitions because many of the sequential button features are unavailable in NSC mode. See “Program Mode” on page 67.

Discussion:

There are 48 buttons which can be set to invoke any of the ZEMAX main menu features. Each button has an identical list of menu options which can be associated with the button. Selecting "Off" will cause the button to disappear. To make the interface more manageable, the available buttons are arranged in groups of 16 on separate tabs.

ZPL macros may be assigned to buttons. The name on the button will be the first 3 letters of the macro name. For more information on ZPL macros, see “ZEMAX PROGRAMMING LANGUAGE” on page 651. ZPL macros must be placed in the appropriate folder, see “Folders” on page 68.

Status Bar

Purpose:

These settings determine which values are displayed in the ZEMAX status bar.

Discussion:

There are 4 regions which can be set to display various data, such as EFL, EPD, F/#, etc.

Exit

Purpose:

Exits ZEMAX.

Discussion:

If the lens has been modified, ZEMAX will prompt to save the lens. If not, the program terminates.

Recently used files

The most recently used lens files are listed at the bottom of the file menu. Selecting one of these files from the list will cause the file to load. This is a shortcut for using the File, Open command.

Lens Data

The Lens Data Editor is the primary spreadsheet where the majority of the lens data will be entered. This data includes the radius of curvature, thickness, and glass for each surface in the system. Single lenses are defined by two surfaces (a front and a back), and the object and image also each require one surface. This fundamental data can be entered directly on the spreadsheet. When the Lens Data Editor is displayed, data can be entered on the spreadsheet by typing in the required values at the highlighted bar. Each column is labeled by the data type, and each row represents a single optical (or perhaps a dummy) surface. The cursor keys move the highlighted bar to whichever column is desired. Moving the cursor continuously to the right or left causes the display to scroll, which provides access to the other columns of data such as semi-diameters, conic constants, and "parameter" terms whose values depend upon the surface type. The display will scroll left to right or right to left. The page up and page down keys will move the bar to the top and bottom of the display. The display will also scroll up and down if required, when the number of lens surfaces is sufficiently large.

Inserting and deleting surfaces

Note that initially (unless a lens has been loaded) three surfaces are shown; the object, stop, and image surfaces. The object and image surfaces are permanent, and cannot be deleted. However, other surfaces can be inserted and deleted by using the Insert and Delete keys. No surface can be inserted in front of the object, or behind the image. In this context "in front of" means a smaller surface number, and "behind" means a larger surface number, in the sense that the light reaches the various surfaces sequentially. ZEMAX numbers surfaces from the object, being surface zero, upwards to the last surface, being the image surface.

To enter values in the spreadsheet, move the cursor to the correct cell, and start typing. To edit values currently displayed, press the backspace key. Once you are editing the cell contents, you can use the left/right cursor keys and the home and end keys to navigate around in the text. When the data has been changed, press any cursor key, or click on any other cell, or press enter.

There are also a few shortcuts available. To add a value to the current value, type a plus sign before the number. For example, if the displayed number is 10, typing "+5" and pressing enter will change it to 15. The symbols "*" and "/" also work. To subtract a value, use the minus sign followed by a space. For example entering "- 5" would change a 17 to a 12. Note the space between the "-" and the "5". If no space was entered, the program assumes you are entering a new value which happens to be negative. Entering "*-1" will change the sign of the number.

The surface number display

The left most column of the Lens Data Editor displays the surface number and type for each surface. The surface number is zero for the object surface, 1 for the first surface, and so on to the image surface. Three "special" surfaces, the object, stop, and image, display "OBJ", "STO", and "IMA", respectively, instead of the surface number. The object surface is always surface zero, and the image surface is always the last surface, however the stop surface may be any surface number in between.

Some additional information is displayed adjacent to the surface number. When an aperture is defined on a surface, ZEMAX will display an asterisk "*" symbol next to the surface number. If surface tilt/decenter data is defined on a surface, the "+" symbol will be displayed. If both aperture and tilt/decenter data are defined the "#" symbol is displayed.

To define the surface type and other properties, see "The Surface Properties dialog box" on page 76. To define the stop surface, see "Make Surface Stop" on page 77. For information on defining surface aperture data, see "Aperture type and other aperture controls" on page 79. For information on defining tilt/decenter data, see "Surface tilt/decenter tab" on page 84.

Cutting, copying, and pasting surface data

See the discussion on the "Edit" menu below.

Entering surface comments

Each surface has a comment field which can be used to enter text. The comment column is used for improving the readability of the lens prescription, and has no effect on ray tracing. Some analysis features also display the surface comments. The entire comment column can be hidden, see the "Options" menu description below.

Entering radii data

To enter or change the radius of curvature of a surface, move the cursor to the desired cell, and type in the new value. Radii data is always entered and displayed in lens units, which have dimensions of length.

Entering thickness data

To enter or change the thickness of a surface, move the cursor to the desired cell, and type in the new value. Thickness data is always entered and displayed in lens units, which have dimensions of length. The thickness of a surface is the distance to the next surface. The only thickness that is not used is the thickness of the image surface.

Thicknesses always change sign after a mirror. After an odd number of mirrors, all thicknesses should be negative. This sign convention is independent of the number of mirrors, or the presence of coordinate breaks. This fundamental convention cannot be circumvented through the use of coordinate rotations of 180 degrees.

Entering glass data

The material used for each surface is usually specified by entering the name of the glass in the "Glass" column on the Lens Data Editor. The glass name entered must be in one of the currently loaded glass catalogs. The default catalog is the "Schott" catalog; others are available. To use multiple glass catalogs, or to review, edit, or append the glass catalogs, see "USING GLASS CATALOGS" on page 587. To specify that a particular surface be a mirror, use the glass name "mirror".

There is an optional "/P" command that may be appended to the glass name when entering a new glass. This option will cause ZEMAX to alter the curvatures before and/or after the surface to attempt to maintain a constant power for the lens element. For example, if the glass type is already BK7, entering a new glass type of "SF1/P" will change the glass type to SF1 while adjusting the radii of the surface and the surface after, to maintain constant power. ZEMAX considers the front and back vertex powers, as well as the element power correction due to the lens thickness. The algorithm adjusts both front and back curvatures if the lens is in air. If either the front or back surfaces are adjacent to another glass element, then only the curvature adjacent to the air is adjusted.

Entering semi-diameter data

The default semi-diameter is automatically calculated to be the radial clear aperture required to pass all rays from all field points. If any value is entered for the semi-diameter, then the new number will be retained and a "U" will be placed next to the value. This indicates the semi-diameter is user defined. For information on the use of semi-diameters, see "Semi-diameters" on page 59. For information on vignetting rays with apertures, see "Surface properties aperture tab" on page 79.

Entering conic data

Conic data is permitted on many different surface types. To enter or change the conic constant of a surface, move the cursor to the desired cell, and type in the new value. Conic constants are always dimensionless. See "Standard" on page 325 for a discussion of how conic surfaces are defined.

Entering parameter data

Parameter data consists of additional numeric values which define the properties of certain surface types. For more information on parameter data, See "Parameter data" on page 277.

The Surface Properties dialog box

Double clicking on the left most column (the one with the surface numbers listed) on the surface will invoke the surface properties dialog box. The previous/next surface buttons allow rapid navigation through all surfaces. The following surface properties may be defined on this dialog box.

Surface properties type tab

Surface Type

ZEMAX models planes, spheres, and conics; all of these surface types are grouped under the category of standard surface. On the surface properties dialog box is a list of surface types. Select the appropriate surface type from the drop down list. ZEMAX supports many different types of surfaces in addition to the standard surface. The types are discussed in detail in the Chapter "Surface Types". Many optical designs only use the standard surface type.

Surface DLL

If the surface type is "User Defined" then the surface shape and ray trace properties are defined in an external program linked into ZEMAX called a Dynamic Link Library, or DLL. This control selects which DLL the surface uses. See "User Defined" on page 331 for more information.

Surface Color

By default, on the Shaded Model Layout plot, ZEMAX draws mirrors in green, and refractive and dummy surfaces in blue. The color of the surface as drawn on shaded models may alternatively be selected to be any of the colors defined on the Colors tab of the File, Preferences dialog box.

Surface Opacity

If the opacity is set at 100%, then the surface will be rendered on the shaded model plot as a solid color, and the surface may fully obscure other surfaces from view. If the opacity is less than 100%, then the surface is partially transparent, which allows other surfaces to be visible through the partially obscure surface.

Row Color

This control chooses the color of the row in the Lens Data Editor for the surface. By default, glass surfaces, coordinate breaks, mirrors, and paraxial lenses are color coded. Any surface may use either no color, the default color, or a user defined color. The user defined colors are described in "Colors 1-12, Colors 13-24" on page 73. The coloring of rows may be disabled, see "Editors" on page 71. The entire spreadsheet may be set to show default or no colors, see "Menu options" on page 86.

Make Surface Stop

The stop surface may be any surface in the system, except for the object and image surfaces. To change the stop surface, select the "Type" tab and then click on the checkbox labeled "Make surface stop" and press OK. The dialog box will disappear and the surface now will display the "STO" label instead of the surface number. This control is grayed out if the surface is the object, image, or is already the stop surface.

It is important to define the stop surface such that the entrance pupil is on the same axis as the object surface. You can ensure this condition by placing the stop surface of the optical system before any coordinate breaks, obscuration decenters, holograms, gratings, or other components which can alter the optical axis. If your system is symmetric for rotations about the optical axis, then this limitation does not apply. Only systems that use surfaces which can tilt or decenter the optical axis should have the stop placed before any such surfaces. If coordinate breaks are used, but only to implement fold mirrors in an otherwise axial system, then the pupil locations will be correctly computed even if the stop is placed after the fold mirrors.

In certain systems it is not possible to place the aperture stop before coordinate breaks. In this case, ray aiming must be used. Ray aiming is discussed in "Ray Aiming" on page 106.

Make Surface Global Coordinate Reference

The global coordinate reference surface may be any surface in the system. To make the currently selected surface the global coordinate reference surface, choose the "Type" tab and check this box on. See "Global Coordinate Reference Surface" on page 111.

Surface Cannot Be Hyperhemispheric

ZEMAX normally detects if a surface must be hyperhemispheric (the surface continues past the maximum possible radial aperture and curves back toward the vertex to fill more than a hemisphere) to pass all rays. If this option is checked, the surface is not allowed to become hyperhemispheric. This switch should be used in conjunction with a floating or circular aperture to vignette any rays which intercept beyond the desired aperture.

Ignore This Surface

If checked, the surface will be ignored for most purposes. Ray tracing, analysis, layout plots, optimization, tolerancing and most other features will produce results as if the surface had been deleted. The purpose of this feature is to completely ignore surfaces, or groups of surfaces, particularly in multi-configuration lenses where optics are present in some configurations but not others. The multi-configuration operand IGNR (see page 581) can be used to change the ignore surface setting for each configuration.

When a surface is ignored, all properties of the surface are ignored; including radius, thickness, glass, apertures, surface tilts and decenters, coordinate break data, and refraction. However, the surface is not actually deleted, and therefore subsequent surfaces are not renumbered. The ignored surfaces retain their specific surface numbers. Some features which are or may be surface number specific, such as spot diagrams, can still be computed on the ignored surfaces. The resulting data will be for the first surface that is not ignored following the requested surface. Some features that list text data on a surface by surface basis may still list data for ignored surfaces, but the data on subsequent surfaces that are not ignored will be as if the ignored surfaces had been deleted.

Surface properties draw tab

Hide Rays To This Surface

If checked, no rays will be drawn to the surface in the various types of layout plots. See Skip Rays below for an alternate means of hiding rays to dummy surfaces.

Skip Rays To This Surface

If checked, rays will skip the surface. This means ray segments will be drawn from the last surface that was not a coordinate break and did not have the Skip Rays option checked, all the way to the next surface that is not a coordinate break and does not have the Skip Rays option checked. For example, if Skip Rays is selected on surfaces 3, 4, and 5, rays will be drawn directly from the coordinates on surface 2 to their coordinates on surface 6, as long as surfaces 2 and 6 are not coordinate breaks.

The purpose of the Skip Rays option is to suppress the drawing of rays to and from dummy surfaces. Although this can be useful in some cases, this option can cause strange, misleading layout plots if the skipped surfaces can refract or diffract the rays. Therefore, the option should only be used on dummy (optically ineffective) surfaces. ZEMAX does not check or verify that surfaces with the Skip Rays option are in fact dummy surfaces, so care is required when using this feature.

Do Not Draw This Surface

If checked, the surface will not be drawn on layout plots.

Do Not Draw Edges From This Surface

If this option is checked, then no edges will be drawn from this surface to the next surface. This feature makes for cleaner drawings in some systems, especially those using liquids or other non-air fillers between optical surfaces.

Draw Local Axis

If checked, an arrow will be drawn to indicate the orientation of the local Z axis at the surface on 3D Layout plots. The size of the arrow is determined by the fletch size, see "Fletch Size" on page 70.

Draw Edges As

This setting controls how edges that connect a surface to the next surface are drawn on the various layout and element drawings. Usually this setting is only used to connect edges on surfaces with circular apertures defined by the radial semi-diameter. Other cases may use another, usually simpler algorithm than the options described here, or may not draw the edge at all. The available options are described below:

Squared To Next Surface: A flat face is drawn radially outward from the smaller of the current and next surfaces to match the semi-diameter of the larger surface. The outer edge of the surface is then formed by a cylinder extending to connect the surface edges.

Tapered To Next Surface: The current and next surface are connected by a single tapered cylinder.

Flat To Next Surface: The current surface edge is drawn as a cylinder at the current semi-diameter value

extending to the contact point on the next surface.

The 3D Layout, Wireframe, Shaded Model, and Solid Model plots currently do not support the "Flat To Next Surface" option. Surfaces with this setting selected will be drawn using the "Squared To Next Surface" method.

Mirror Substrate and Thickness

Sequential surfaces with a glass type of "MIRROR" may be drawn as either a thin surface, or as a solid object with a substrate. The back of the substrate may be a flat surface that is perpendicular to the local z axis, or as a curved surface that follows the contour of the actual mirror surface. The thickness of the substrate is measured along the local z axis at the vertex of the surface, even if the surface aperture is decentered or user defined. ZEMAX will draw the edges of the substrate by connecting the mirror surface and the substrate back. All surface and aperture types are supported on all layout plots.

By default, ZEMAX will draw the back of the mirror substrate as a curved surface with a thickness of 4% of the surface semi-diameter, or 0.5 lens units if the semi-diameter is zero. The default thickness will be used if the Mirror Substrate Thickness parameter is set to zero.

Mirror substrates will not be drawn on surfaces which are mirrors in some, but not in all defined configurations. Mirror substrates are also not drawn on surfaces which are Mangin mirrors (mirrors on the back of glass substrates) in any configuration. To draw mirror substrates on surfaces which are present in some configurations but not in others, use the multi-configuration operand IGNR to ignore the mirror surface rather than GLSS to change the surface from a mirror to glass or air. See "Adding and removing elements" on page 584 for details.

Surface properties aperture tab

Aperture type and other aperture controls

Surface apertures are used to account for the effects of vignetting. The supported types of surface apertures in ZEMAX are: no aperture, circular aperture, circular obscuration, rectangular aperture, rectangular obscuration, elliptical aperture, elliptical obscuration, spider obscuration, user defined aperture, user defined obscuration, and floating aperture. The aperture and obscuration types define regions which pass or stop rays, respectively. More than one aperture may be described at a given optical element by inserting a dummy surface with zero thickness at the desired location, and then setting additional apertures on that surface. This is useful for constructing complex apertures. Multiple simultaneous apertures and obscurations may also be defined on a single surface using the user defined apertures and obscurations feature, see "User defined apertures and obscurations" on page 80.

Setting aperture data for each surface is done on the surface data dialog box. The surface data dialog box can be reached by double clicking on the left most column in the Lens Data Editor. When "none" is selected for the aperture type (the default), all rays which can be refracted or reflected by this surface are allowed to proceed. To clear any aperture to this default state, or to change the current aperture type, select another aperture type on the surface data dialog box.

The individual aperture types are described below.

Circular Aperture/Obscuration: A Circular Aperture defines an annular region which vignettes all rays which strike the surface inside of the minimum radius, and outside of the maximum radius. If the ray is between the minimum and maximum radii, then the ray will be allowed to proceed. The Circular Obscuration is the complement of the Circular Aperture.

Rectangular Aperture/Obscuration: Rays are vignettted which intercept the surface outside the rectangular region defined by the half widths in x and y. The Rectangular Obscuration is the complement to the Rectangular Aperture.

Elliptical Aperture/Obscuration: Rays are vignettted which intercept the surface outside the elliptical region defined by the half widths in x and y. The Elliptical Obscuration is the complement to the Elliptical Aperture.

Spider: The spider is defined by the width of each arm, and by the number of arms. ZEMAX assumes the arms are all the same width, and that they are spaced in equal radial angles. The first arm starts at a radial position of zero degrees, which is along the local positive x-axis. More complex spiders, which contain arms of different widths at unequal angles can be constructed using several spiders on adjacent dummy surfaces. The coordinate break surfaces can be used to rotate the spider(s) to any desired angle.

User Defined Aperture/Obscuration: See "User defined apertures and obscurations" on page 80.

Floating Aperture: A floating aperture is very similar to the circular aperture, except the minimum radius is always zero, and the maximum radius is always equal to the semi-diameter of the surface. Since the semi-diameter value may be adjusted by ZEMAX (when in automatic mode) the aperture value "floats" as the semi-diameter value. The floating aperture is useful when macros or external programs use ZEMAX to trace rays that may lie outside of the default semi-diameters, and these rays are to be vignetted.

Aperture type codes and parameters

Different types of apertures use the two available aperture parameters for different purposes as described in the following table. The codes are used by the ZPL macro language, see "SETSURFACEPROPERTY, SURP" on page 725 for details.

APERTURE TYPES, CODES, AND PARAMETER USAGE

Aperture type and code	Param 1	Param 2
None, 0	-	-
Circular aperture, 1	Minimum aperture radius	Maximum aperture radius
Circular obscuration, 2	Minimum aperture radius	Maximum aperture radius
Spider, 3	Width of each spider arm	Number of spider arms
Rectangular aperture, 4	X-half width	Y-half width
Rectangular obscuration, 5	X-half width	Y-half width
Elliptical aperture, 6	X-half width	Y-half width
Elliptical obscuration, 7	X-half width	Y-half width
User aperture, 8	-	-
User obscuration, 9	-	-
Floating aperture, 10	-	-

Aperture projection on non-plane surfaces

All apertures are modeled as a surface projected from the vertex tangent plane onto the optical surface. The actual ray-surface x and y intercept coordinates are used to determine vignetting, the z coordinate is ignored. Different results may be calculated for steep optical surfaces if the aperture is placed on a dummy surface in front of the optic, instead of being directly entered on the curved surface. This will only occur if the ray incidence angles are steep. Usually it is best to place the apertures directly on the optical surface, unless of course the dummy surface more accurately represents the situation.

Aperture decenters and pickups

All types of apertures may be decentered from the current optical axis by specifying either X- or Y-decenters, or both. The decenters are given in lens units. If the identical aperture is used on more than one surface, the aperture "Pickup From" is a useful feature. Any aperture may be picked up from any previous surface.

User defined apertures and obscurations

The circular, rectangular, and elliptical apertures and obscurations are easy to use and cover common cases. However, there are times when a more general aperture model is required.

ZEMAX supports a general user specified aperture defined by a series of line segments, arcs, circles, polygons, and rectangles. The aperture may be closed in a simple or complex way, and multiple aperture regions may be defined which are either nested or not nested. To define a user defined aperture or obscuration, select the desired type (aperture or obscuration) from the list of aperture types. The "Aperture File" control will list all available user defined aperture (UDA) data files. The UDA files are text files that may be created and edited using any text editor. The UDA files are stored in the <objects>\Apertures folder (see "Folders" on page 68). The UDA file format is described in the next section.

The button "Edit Aperture File" will invoke a text editor to allow user editing of the selected UDA file. The UDA file needs to be saved and the lens file updated to make the changes effective. The UDA Scale is a dimensionless multiplier that scales the aperture defined in the UDA file. This control allows scaling of the UDA without need to modify the UDA file.

The UDA file format

The UDA file consists of a series of simple text commands that define pieces of the boundary shape of the aperture. The individual pieces are called entities. All UDA entities are defined as they lie on the XY plane. The Z coordinate is determined by the sag equation of the surface to which the UDA is applied. The commands which define the entities consist of three letter mnemonics followed by numerical values stored in a text file whose extension ends in UDA. The file name may not include any spaces.

In the entity definitions, all spatial coordinates are in lens units and all angles are in degrees. The available entity types are described in the following table.

USER DEFINED APERTURE (UDA) ENTITIES

Entity Name/Syntax	Description
ARC cx cy angle n	Arc. Appends an arc to the current end point. The starting point of the arc is the current end point, defined by a previous ARC or LIN command. The cx and cy values are the center of the circle on which the arc lies. Note that the current end point, along with cx and cy, define the radius of the arc. The angle (in degrees) defines the length of the arc. Positive angles are clockwise on the XY plane. ZEMAX internally represents the arc as a series of line segments. The parameter n determines how many line segments are used to approximate the arc. A recommended value is roughly 1 segment for each 6 degrees of arc. The maximum allowed value is 64 segments.
BRK	Break. Ends the current aperture definition, closing any previously defined aperture. Entering the coordinates: 0.0 0.0 or LIN 0.0 0.0 will also be interpreted as a break.
CIR cx cy radius n	Circle. Defines a circle centered on cx and cy, with the specified radius. ZEMAX draws the circle as a series of n line segments. The minimum value of n is 8. The maximum value of n is 64. The default value is 32 if n is omitted. The actual aperture applied to the rays is an exact circle, independent of the value for n.
ELI cx cy rx ry angle n	Ellipse. Defines an ellipse centered on cx and cy, with the specified radial (half) widths in the x and y directions. ZEMAX draws the ellipse as a series of n line segments. The minimum value of n is 8. The maximum value of n is 64. The default value is 32 if n is omitted. The actual aperture applied to the rays is an exact ellipse, independent of the value for n. The parameter angle (in degrees) may be used to rotate the ellipse clockwise to any desired position.
LIN x y n or x y	Line. Appends a line segment from the current end point to the specified x, y coordinates. Use BRK to close a polygon formed by LIN commands. ZEMAX optionally represents the line as a series of shorter line segments; this is useful when the line lies on a steeply contoured part. The parameter n determines how many line segments are used to construct the total line. If n is omitted then 1 line segment is used. The mnemonic LIN is not required; if two coordinates are listed with no mnemonic, LIN is assumed. This is for backward compatibility with older versions of ZEMAX.

Entity Name/Syntax	Description
POL cx cy radius n angle	Polygon. Defines an "n" sided regular polygon centered on cx and cy, with each vertex the specified radial distance from the center point. The POL command may be used to define triangles, squares, pentagons, hexagons, and other polygons with equal length sides. The first point on the polygon starts at the angle 0.0 (along the local +x axis) and proceeds clockwise from there. The parameter angle (in degrees) may be used to rotate the polygon to any desired position. The maximum allowed number of segments is 64, the minimum is 3.
REC cx cy xhw yhw angle nx ny	Rectangle. Defines a rectangle centered on cx and cy with the specified half widths in the x and y directions. The parameter angle (in degrees) may be used to rotate the rectangle clockwise to any desired position. The optional arguments nx and ny define how many segments to use in the x and y directions. Using a larger number of segments will yield better rendering of NSC object surfaces that are steeply curved over the rectangle aperture.
! comment	Comment. Any line starting with the ! symbol is ignored, and may be used for placing comments in the UDA.

The end of a polygon is indicated by either a BRK entity or a LIN entity with both X and Y being set to zero. For this reason, a polygon may not be defined which has a vertex at (0,0). If a vertex must be defined at (0,0), a work around is to instead use some very small value for one of the points, such as (1e-6, 0). As long as at least one of the coordinates is not zero, the point will be considered a vertex rather than indicating the end of a polygon. The last listed vertex in the polygon is assumed to then connect to the first point. Note that because the last vertex is automatically connected to the first vertex, the last vertex should not be defined to be identical to, and redundant with, the first point. Defining multiple points that are within round-off error of the same location may result in incorrect ray tracing or aperture rendering.

Note that the UDA entities REC and CIR define "stand alone" apertures not connected to any other aperture entities, while LIN and ARC define pieces of a polygon whose definition ends with a BRK. When using LIN and ARC entities, the first entity defined should be an LIN, which defines the first point of the polygon. Multiple apertures may be defined separated by BRK entities.

The maximum number of line segments defined within the UDA file is limited only by available memory. However, ray tracing and rendering of surfaces and objects is slower when using very complex UDA data files. See the following sections for UDA examples.

UDA Examples


To define a square 20 lens units on the side, the UDA file is:

```
LIN -10, -10
LIN -10, 10
LIN 10, 10
LIN 10, -10
BRK
```

Note the BRK indicates the end of the polygon definition, and the last point is assumed to join back to the first point, thus defining the last side of the square. This same aperture may be defined with a single REC command:


```
REC 0 0 10 10 0
```

Note the REC command does not require a BRK command to define the end of the aperture.



Multiple polygons may be defined by separating them with a single BRK. For example, to define an aperture consisting of two slits, each 5 lens units wide, with an inner separation of 10 lens units, the UDA file is:

```
LIN -10, -10
LIN -10, 10
LIN -5, 10
LIN -5, -10
BRK
LIN 10, -10
LIN 10, 10
LIN 5, 10
LIN 5, -10
BRK
```




Or the equivalent

```
REC -7.5 0 2.5 10 0
REC +7.5 0 2.5 10 0
```

To create a square with rounded edges, use the ARC and LIN commands together. Here is a UDA that creates a square, 4 lens units on a side, with rounded corners. Each rounded corner is formed by one ARC with a radius of 1 lens unit and 12 segments along the arc:

```
LIN -1 2
LIN 1 2
ARC 1 1 90 12
LIN 2 -1
ARC 1 -1 90 12
LIN -1 -2
ARC -1 -1 90 12
LIN -2 1
ARC -1 1 90 12
```




The first LIN command defines the starting point near the upper left corner. The subsequent commands define each of the 8 parts of the aperture. Note a BRK command is not needed if only a single aperture region is defined.

Multiple polygons may also be nested. If a ray intercepts a point within a polygon that is within another polygon, then the point is considered to be outside the aperture. This convention allows "islands" to be defined within apertures that become obscurations, and vice-a-versa. Any number of nesting levels is allowed, and each level toggles the inside/outside status of the point.

Note that when multiple polygons are defined, the polygons may be separated or nested. However, they may not intersect or share adjacent boundaries.

Here is a UDA with the rounded corner square from the previous example, with 5 small sub apertures nested within; a pentagon, a hexagon, an ellipse, a rectangle, and a circle. Note the BRK command is not needed to separate the REC, ELI, POL, and CIR entities from the ARC and LIN entities because REC, ELI, POL, and CIR all define stand alone aperture regions:



```
LIN -1 2
LIN 1 2
ARC 1 1 90 12
LIN 2 -1
ARC 1 -1 90 12
LIN -1 -2
ARC -1 -1 90 12
LIN -2 1
ARC -1 1 90 12
REC 1 1 .3 .5 -30
```

```
ELI -1 -1 .6 .2 15
POL 1 -1 .5 5 0
POL -1 1 .5 6 0
CIR 0 0 .3
```

Surface properties scattering tab



This feature is only available in the EE edition of ZEMAX.

Highly polished optical surfaces typically have a small amount of scattered light over a small, roughly angular cone around the specular refracted or reflected ray direction. Surface scattering properties are useful for modeling this small amount of scattered light, to analyze the effects of the scattering on MTF or other quality measures. Scattering is modeled in ZEMAX by randomly deviating the refraction or reflection angle of some or all of the rays leaving a surface.

There are two fundamentally different ways of modeling scattering in ZEMAX: sequential and non-sequential.

For the sequential model, the scattering is assumed to deviate the ray angle only very slightly. The primary effect is a "blurring" of the spot diagram or ray pattern at the image surface. Sequential scattering is not intended for modeling the effects of back scatter, wide angle scatter (such as Lambertian) or other scatter where the scattered rays may not follow a well behaved sequential path to the image surface.

For the non-sequential model, scattering may occur over any range of angles, and rays will be traced correctly no matter in what direction they propagate. The non-sequential model is the better choice if it is necessary to model scatter from surfaces that are not part of the sequential ray path, such as from lens mounts, baffles, or other objects. See the Chapter "Non-Sequential Components" for details on this method of modeling scatter.

It is not necessary to use the non-sequential components feature of ZEMAX to model scattering, if the system is otherwise sequential and the only scattering of interest is small angle scatter from the small scale roughness of the optical surfaces, and the only desired analysis is the degradation in image quality due to scattering.

Surface scattering settings

Only ZEMAX-EE supports this capability.

ZEMAX uses the same terminology and scattering models for both sequential surfaces and non-sequential components. See the Chapter "Non-Sequential Components" on page 568 for a detailed technical description of the scattering models. For information on defining ABg data, see "ABg Scatter Data Catalogs" on page 250.

The non-sequential ABg scattering model supports both reflection and refraction coefficients, while the sequential ABg scattering model uses a single set of coefficients since sequential surfaces either reflect or refract, but not both. Whether a sequential surface reflects or refracts depends only upon the whether the surface is a mirror or not.

Although Lambertian is a valid option on the sequential surface dialog, rays which scatter using this model may go in any forward direction, which may cause rays to scatter at large enough angles so that they do not correctly propagate through the rest of the optical system.

Many sequential analysis features have a "Use Polarization" option to consider ray transmission. When using polarization, all sequential surfaces that scatter rays should be defined to have small-angle scatter. Small-angle scatter means the rays are only perturbed in direction by an amount that would not significantly change the polarization properties of the ray. The reason for this requirement is that the polarization transmission computation is performed only on the *unscattered* ray. The unscattered polarization transmission is then applied to the actual scattered ray. This limitation only applies to sequential systems. For systems that require large angle scatter and polarization analysis, use the non-sequential component capability as described in "NON-SEQUENTIAL COMPONENTS" on page 477.

Only some ZEMAX analysis features use the scattering data. Most features, such as optimization and layout diagrams, ignore surface scattering. Features that use the scattering data, such as the spot diagrams have a "Scatter Rays" checkbox.

Surface tilt/decenter tab

Surface tilts and decenters allow a change in the coordinate system to be implemented both before and after ray tracing to the surface. Applications include decentering a surface and returning to the original coordinate

system, tilting a fold mirror and tilting again to follow the beam, tilting a surface to model a wedge, and many others. Surface tilts and decenters are redundant with and very similar to coordinate breaks. For detailed information on coordinate breaks, see "Coordinate Break" on page 293. A surface tilt/decenter can be thought of as a coordinate break, followed by the surface, followed by another coordinate break.

Surface tilts and decenters may not be placed on a coordinate break surface. However, coordinate break surfaces do support a coordinate return solve, see "Using the Coordinate Return" on page 85. Surface tilts and decenters are also not allowed before the first surface if the object is at infinite conjugates.

The advantage to using surface tilts and decenters is the elimination of "dummy" coordinate break surfaces in the Lens Data Editor. This allows for a somewhat less cluttered display some users prefer. The disadvantage of using the surface tilts and decenters is the current implementation does not support optimization of surface tilt and decenter data.

The operation of surface tilts and decenters is done in a sequence:

Before the surface, the coordinate system is decentered in X, decentered in Y, and then tilted around X, tilted around Y, tilted around Z. Decenters are measured in lens units, tilts in degrees about the respective axis in a right hand direction. Alternatively, the order may be chosen to tilt then decenter. In this case, the coordinate system is tilted around Z, tilted around Y, tilted around X, decentered in Y then finally decentered in X.



The order of the decenters and tilts matters!

After the surface, the same set of operations may be performed in either order. The values for the before and after decenters and tilts, and the order in which they are done, may be independent. However, it is frequently useful to have the after tilts and decenter values be related to previous before values. The options available are:

- a) Explicitly define the tilts and decenters.
- b) Pick-up the values from the before data of the current surface (typically used for fold mirrors)
- c) Reverse the values from the before data of the current surface (sometimes called a decenter/ tilt and return)
- d) Pickup the values from the before data of a prior surface
- e) Reverse the values from the before data of a prior surface (typically used for decentering a range of surfaces)

All of these options are supported by the "After Surface" settings. Reversing values from a prior surface involves changing the order and picking up the tilt and decenter values from the target surface and reversing the sign.

The coordinate system resulting from the before and after tilt and decenters will define the coordinate system for the next surface. The thickness of a surface is the thickness in the new coordinate system after all the tilt and decenters are applied, measured along the resulting Z axis.

Using the Coordinate Return

It is frequently convenient to return to the coordinate system of a prior surface. The Coordinate Return feature can be used on any coordinate break surface to achieve this. Note that the surface must be a coordinate break surface. Use of the Coordinate Return will overrule any solve, variable, or multi-configuration status on any data controlled by the Coordinate Return. To control the Coordinate Return solve from the multi-configuration editor, see "CROR" on page 580.

The following options are supported for Coordinate Return:

None: The feature is disabled and the coordinate break tilts and decenters are not automatically determined.

Orientation Only: The tilt about x, y, and z axes are determined to return the orientation of the coordinate system to the previous surface. No adjustment is made to the position offset of the surface vertex.

Orientation XY: The tilt about x, y, and z axes, and the decenters in x and y are determined to return the orientation of the coordinate system and the x and y components of the vertex offset to the previous surface. No adjustment is made to the z position offset of the surface vertex.

Orientation XYZ: The tilt about x, y, and z axes, and the decenters in x, y, and z are determined to return the orientation of the coordinate system and vertex to the previous surface. The z decenter is placed on the

thickness of the coordinate break, and so it is the *following* surface whose orientation and position are identical to the prior surface.

To Surf: This setting controls the prior surface to return the coordinate system to.

For more information on coordinate transformations, see “Coordinate Break” on page 293.

Surface physical optics tab

See the chapter “PHYSICAL OPTICS PROPAGATION” on page 627 and “Surface specific settings” on page 643 for a description of these features.

Surface coating tab

The coating tab allows selection of the optical coating to be applied to the surface. For information on defining coatings, see “Defining coatings in ZEMAX” on page 613 and “Optimizing coatings with ZEMAX” on page 621.

This tab supports an additional capability to individually modify the thickness, index, or extinction of any layer of a defined coating applied to a particular surface without changing the underlying definition of the coating. Individual layer thicknesses may be scaled by a “coating multiplier”, a dimensionless parameter. The index and extinction coefficients may be offset by a dimensionless value. Note the index offsets are applied uniformly at all wavelengths; no dispersion in the offset is supported. These parameters may be unique for every layer and every surface, even if the same coating is applied to more than one surface. Encrypted coatings do not support this feature.

The available controls are:

Use Layer Multipliers: If unchecked, all coating layer multipliers are disabled for this surface. The coating thickness will be as defined by the coating data file. If checked, the coating multipliers are considered.

Layer: Chooses the particular layer to edit or review the coating multiplier for.

Multiplier: The dimensionless multiplier for the selected layer.

Status: The multiplier may be a fixed value, a variable for optimization, or a value picked up from a prior layer.

Index Offset: The dimensionless offset of the real part of the coating index for the selected layer.

Index Status: The index offset may be a fixed value, a variable for optimization, or a value picked up from a prior layer.

Extinction Offset: The dimensionless offset of the imaginary part of the coating index for the selected layer.

Extinction Status: The extinction may be a fixed value, a variable for optimization, or a value picked up from a prior layer.

The buttons provide a fast way of setting all multipliers and offsets to fixed or variable status or resetting the multipliers and offsets to the default unperturbed values.

Setting and removing solves

Most data columns (such as radii and thickness) support one or more solves. To set a solve on a cell, double click with the left mouse button on that cell. Alternatively, single click with the right mouse button, or select the menu option from the Lens Data Editor menu. Solves are described in “SOLVES” on page 473.

Setting and removing variables

To make a parameter variable, click on the cell containing the parameter, and press Ctrl-Z on the keyboard. Ctrl-Z also removes the variable status, and acts like a toggle.

Menu options

The Lens Data Editor menu options are used for inserting and deleting surfaces, selecting surface types, and setting solves and variables.

Edit

The edit menu offers these options:

Surface Type: This option allows the type of a surface to be changed.

Insert Surface: This inserts a new surface row in the spreadsheet at the current row. The Insert key is a shortcut for this menu option.

Insert After: This inserts a new surface row in the spreadsheet after the current row. The Ctrl-Insert key combination is a shortcut for this menu option.

Delete Surface: This deletes the current row in the spreadsheet. The Delete key is a shortcut for this menu option.

Cut Surfaces: Copies all the data for a single surface or a range of surfaces to the Windows Clipboard, then deletes the surfaces. The surface or range of surfaces must first be selected using either of the following techniques:

Using the mouse: Click on the first surface to be selected. Hold down the left mouse button, and drag the mouse cursor to cover the range of surfaces desired. The selected surfaces will be highlighted in inverse colors. To select only one surface, drag the mouse up or down from the surface desired until two surfaces are selected, then move the mouse back to the desired surface.

Using the keyboard: Move the cursor to any cell on the desired surface. Then hold down the shift key, while moving the cursor up or down until the desired range of surfaces is selected. The selected surfaces will be highlighted in inverse colors. To select only one surface, move the cursor up or down from the surface desired until two surfaces are selected, then move the cursor back to the desired surface.

Copy Surfaces: Copies all the data for a single surface or a range of surfaces to the Windows Clipboard. To select a single surface or a range of surfaces, see the discussion under "Cut Surfaces" above. Non-Sequential Component surfaces may not be copied.

Paste Surfaces: Copies all the data for a single surface or a range of surfaces from the Windows Clipboard to the current cursor location in the Lens Data Editor. The surface data must first have been copied to the Windows Clipboard using either "Cut Surfaces" or "Copy Surfaces" described above.

Copy Cell: Copies a single cell's data to the Windows Clipboard.

Paste Cell: Pastes the data for a single cell from the Windows Clipboard to the current cell. The data must first have been copied to the Windows Clipboard using "Copy Cell" described above.

Edit Cell: Puts the cell in edit mode.

Copy Spreadsheet: Copies either the highlighted range of surfaces, or the entire spreadsheet (if no range of surfaces is selected) to the Windows Clipboard in a text format suitable for being pasted into another Windows application, such as a spreadsheet or word processor. The format is tab delimited text. When using this feature ZEMAX uses curvature rather than radius to retain maximum numerical precision.

Set Row Colors: Selects either default or none for the color for each row in the spreadsheet. Individual row colors may be set using the row color setting, see "Row Color" on page 77.

Solves

Solves and variables may be placed on many Lens Data Editor values:

Radius: Sets solves on the radius of curvature.

Thickness: Sets solves on the thickness.

Glass: Sets solves on the glass.

Semi-Diameter: Sets solves on the semi-diameter.

Conic: Sets solves on the conic constant.

Parameter: Sets solves on the parameter columns.

TCE: Sets solves on the TCE value.

Variable Toggle: This will toggle the variable status of the currently selected cell. Ctrl-Z is a shortcut for this operation.

For more information on solves, see "SOLVES" on page 473.

View

Hide Current Column: Minimizes the width of the column displaying the cursor.

Hide Unused Columns: Minimizes the width of columns that are unused or only display default data for any surface in the current configuration. The Radius, Thickness, Glass, and Semi-Diameter columns are never hidden by this feature. If the surface type or configuration is changed, it may be required to use "Unhide All Columns" to make the data visible.

Unhide All Columns: Sets all hidden columns to the default width (see "Editors" on page 71).

Reset Column Widths: Sets all columns to the default width (see "Editors" on page 71).

Equal Column Widths: Sets all columns to the same width as the column displaying the cursor.

Hide Surfaces: Hides the selected surface or range of surfaces from view in the editor. When surfaces are hidden, a heavy line is drawn in place of the hidden surface range as a reminder that not all surfaces are visible. Hiding surfaces only prevents the surfaces from being displayed in the editor; it does not remove the surface from the optical system. To ignore a surface so that rays are not affected by the surface see "Ignore This Surface" on page 78. To prevent the surface from being drawn, see "Do Not Draw This Surface" on page 78.

Hide Coordinate Breaks: Hides all coordinate breaks in the selected range, or if no range is selected hides all coordinate breaks.

Show Surfaces: Shows previously hidden surfaces within the selected range, or if no range is selected shows all hidden surfaces.

Go To Surface: Opens a dialog box that allows entry of either the surface number or comment. The editor then displays the desired surface.

Help

Opens the help system for using the Lens Data Editor.

Merit Function

The Merit Function Editor is used to define, modify, and review the system merit function. The system merit function is used for optimization, and is described in the Chapter "Optimization".

Menu Options

Edit

Operand Type: This option allows the type of operand and other data to be changed. For a complete description of merit function operands, see "OPTIMIZATION" on page 483.

Insert Operand: Inserts a new row in the spreadsheet at the current row. The Insert key will perform this same function.

Insert After: Inserts a new row in the spreadsheet after the current row. The Ctrl-Insert key combination will perform this same function.

Delete Operand: Deletes the row at the current cursor location. The Delete key will perform this same function.

Delete All: Deletes all operands in the merit function.

Cut Operands: Copies all the data for a single operand or a range of operands to the Windows Clipboard, then deletes the operands. The operand or range of operands must first be selected using either of the following techniques:

Using the mouse: Click on the first operand to be selected. Hold down the left mouse button, and drag the mouse cursor to cover the range of operands desired. The selected operands will be highlighted in inverse colors. To select only one operand, drag the mouse up or down from the operand desired until two operands are selected, then move the mouse back to the desired operand.

Using the keyboard: Move the cursor to any cell on the desired operand row. Then hold down the shift key, while moving the cursor up or down until the desired range of operands is selected. The selected operands will be highlighted in inverse colors. To select only one operand, move the cursor up or down from the operand desired until two operands are selected, then move the cursor back to the desired operand.

Copy Operands: Copies all the data for a single operand or a range of operands to the Windows Clipboard. To select a single operand or a range of operands, see the discussion under "Cut Operands" above.

Paste Operands: Copies all the data for a single operand or a range of operands from the Windows Clipboard to the current cursor location in the Merit Function Editor. The operand data must first have been copied to the Windows Clipboard using either "Cut Operands" or "Copy Operands" described above.

Copy Cell: Copies a single cell's data to the Windows Clipboard.

Paste Cell: Pastes the data for a single cell from the Windows Clipboard to the current cell. The data must first have been copied to the Windows Clipboard using "Copy Cell" described above.

Edit Cell: Puts the cell in edit mode.

Copy Spreadsheet: Copies either the highlighted range of operands, or the entire spreadsheet (if no range of operands is selected) to the Windows Clipboard in a text format suitable for being pasted into another Windows application, such as a spreadsheet or word processor. The format is tab delimited text.

Set Row Colors: Selects either default or none for the color for each row in the spreadsheet. Individual row colors may be set using the row color setting on the operand type dialog box.

Tools

Update: This option recomputes the merit function. All operands are evaluated, and the new values are displayed.

Default Merit Function: Invokes the dialog box which used to define one of the default merit functions. See the Chapter "Optimization".

Save: Saves the current merit function in a *.MF file. This step is only required if the merit function is to be subsequently loaded into another lens. ZEMAX automatically stores the merit function with the lens when the entire lens is saved.

Load: Loads a merit function previously stored in a *.MF file or in a *.ZMX file. Either file type may be selected; only the merit function portion of the file will be loaded into the spreadsheet. The current merit function is destroyed.

View

Hide Current Column: Minimizes the width of the column displaying the cursor.

Unhide All Columns: Sets all hidden columns to the default width (see "Editors" on page 71).

Reset Column Widths: Sets all columns to the default width (see "Editors" on page 71).

Equal Column Widths: Sets all columns to the same width as the column displaying the cursor.

Go To Operand: Opens a dialog box that allows entry of either the operand number or comment. The editor then displays the desired operand.

Help

Opens the help system.

Multi-Configuration

The Multi-Configuration editor is very similar to the Lens Data Editor. To edit the contents of a cell, move the cursor to that cell and type in the new data. To set a solve on a cell, double click with the left mouse button, or select the solve type menu option.

Menu Options

Edit

Operand Type: This option allows the type of multi-configuration operand to be changed. For a complete description of multi-configuration operands, see "MULTI-CONFIGURATIONS" on page 579.

Insert Operand: This inserts a new row in the spreadsheet at the current row. The new operand type is "OFF" which means the operand is ignored. The Insert key is a shortcut for this menu option.

Insert After: This inserts a new row in the spreadsheet after the current row. The new operand type is "OFF" which means the operand is ignored. The Ctrl-Insert key combination is a shortcut for this menu option.

Delete Operand: This deletes the current row in the spreadsheet. The Delete key is a shortcut for this menu option.

Cut Operands: Copies all the data for a single operand or a range of operands to the Windows Clipboard, then deletes the operands. The operand or range of operands must first be selected using either of the following techniques:

Using the mouse: Click on the first operand to be selected. Hold down the left mouse button, and drag the mouse cursor to cover the range of operands desired. The selected operands will be highlighted in inverse colors. To select only one operand, drag the mouse up or down from the operand desired until two operands are selected, then move the mouse back to the desired operand.

Using the keyboard: Move the cursor to any cell on the desired operand row. Then hold down the shift key, while moving the cursor up or down until the desired range of operands is selected. The selected operands will be highlighted in inverse colors. To select only one operand, move the cursor up or down from the operand desired until two operands are selected, then move the cursor back to the desired operand.

Copy Operands: Copies all the data for a single operand or a range of operands to the Windows Clipboard. To select a single operand or a range of operands, see the discussion under "Cut Operands" above.

Paste Operands: Copies all the data for a single operand or a range of operands from the Windows Clipboard to the current cursor location in the Multi-Configuration Editor. The operand data must first have been copied to the Windows Clipboard using either "Cut Operands" or "Copy Operands" described above.

Copy Cell: Copies a single cell's data to the Windows Clipboard.

Paste Cell: Pastes the data for a single cell from the Windows Clipboard to the current cell. The data must first have been copied to the Windows Clipboard using "Copy Cell" described above.

Edit Cell: Puts the cell in edit mode.

Copy Spreadsheet: Copies either the highlighted range of operands, or the entire spreadsheet (if no range of operands is selected) to the Windows Clipboard in a text format suitable for being pasted into another Windows application, such as a spreadsheet or word processor. The format is tab delimited text.

Insert Configuration: Selecting this item will insert a new configuration at the configuration indicated by the current position of the cursor.

Insert Configuration With Pickups: Selecting this item will insert a new configuration at the configuration indicated by the current position of the cursor, with pickup solves inserted to pick up values from the original configuration.

Delete Configuration: Deletes the configuration indicated by the current position of the cursor. This eliminates the column and all the data it contained.

Set Row Colors: Selects either default or none for the color for each row in the spreadsheet. Individual row colors may be set using the row color setting on the operand type dialog box.

Solves

Solve Type: Selecting this option will invoke the solve dialog box for the cell which is currently indicated by the cursor. For more information on solves, see "SOLVES" on page 473.

Variable Toggle: This will toggle the variable status of the currently selected cell.

Tools

Make Thermal: This tool is used to perform the bulk of the tedious work in setting up a multi-configuration thermal analysis. A dialog box will appear which allows setting of the number of configurations, as well as the minimum and maximum temperatures. There are also options to delete or retain the existing multi-configuration data, and to sort the data by surfaces rather than the default sort by operand type. Be sure to read and understand how the thermal analysis works in ZEMAX. For more information, see "THERMAL ANALYSIS" on page 603 and "Limitations of thermal analysis" on page 608.

If the "Delete existing configuration data" option is selected, then 1 nominal configuration at the current temperature and pressure will be created. Additional configurations are then defined to cover the temperature range specified. If 3 configurations are requested, then there will be 1 nominal configuration (configuration 1), and 3 configurations that span the specified temperature range in equal increments, for a total of 4

configurations. The air pressure will be assumed to be the same as the nominal pressure.

If the "Use existing n configurations as nominal" option is selected, then the existing MCE data will be used to create the nominal configurations. The tool will automatically add new operands for other data typically required for thermal modeling, such as glass curvature, semi-diameters, parameter, and extra data editor values. In this mode, the total number of configurations will be determined by the number of nominal configurations times the number of new configurations, plus the original nominal configurations.

For each radius, thickness, glass, semi-diameter, parameter, and extra data value affected by temperature effects, appropriate operands will be inserted with TCE pickup solves. It is always good engineering practice to check the results of this automatic thermal set up tool carefully to make sure no important parameters have been omitted.



Check the results of the thermal set-up tool carefully!

Make Single Config: This tool will delete all the multi-configuration data, leaving the lens as a single configuration system in whatever configuration was active at the time this tool was selected.

Add All Data: This tool will add to the MCE all data currently in the system. This eliminates the need to add many data items manually. For some systems, it is easier to delete unwanted data than to insert desired data.

Make Conjugate: This tool will create a subset of the optical system which models imaging between any two conjugate surfaces as a new configuration. This feature is useful for evaluating or optimizing the performance of the optical system in between two surfaces which are not always the object and image surfaces. The subset system is defined by the object, stop, and image surface numbers as defined in a reference configuration. If no multi-configuration data exists, then two configurations, the nominal and subset configurations, will be created. If multi-configuration data already exists, then any configuration, or all configurations, may be used as the reference configuration(s). If a single reference configuration is selected, then one new configuration will be added. If all configurations are selected, then the total number of configurations will double, as one subset configuration will be created for each existing configuration.

The subset configurations are created by adding as required multi-configuration operands to alter the system aperture type, size, stop location, field data, thicknesses, curvatures, aspheric data, and other data as appropriate. The ignore surface feature (See "Ignore This Surface" on page 78) is used to eliminate surfaces not relevant to the subset configuration. This feature makes the following assumptions about how to construct the subset system:

The system aperture is set to "float by stop size" and the semi-diameter of the new stop is set to the current semi-diameter of the surface in the reference configuration.

The field type is set to object height, and the field values are set to the primary wavelength chief ray coordinates on the specified object surface in the reference configuration.

The original object and image surfaces in the reference configuration are placed under the control of multi-configuration operands to allow those surfaces to take on the correct shape of the subset object and image surfaces. For this reason, currently the original and subset object and image surfaces must all be Standard surface types.

These assumptions may not be correct in all cases. The resulting subset conjugate system should be carefully checked for validity.

View

Hide Current Column: Minimizes the width of the column displaying the cursor.

Unhide All Columns: Sets all hidden columns to the default width (see "Editors" on page 71).

Reset Column Widths: Sets all columns to the default width (see "Editors" on page 71).

Equal Column Widths: Sets all columns to the same width as the column displaying the cursor.

Go To Operand: Opens a dialog box that allows entry of either the operand number. The editor then displays the desired operand.

Help

Opens the help system.

Tolerance Data

The Tolerance Data editor is used to define, modify, and review the system tolerance values. See "TOLERANCING" on page 545 for details.

Menu Options

Edit

Operand Type: This option allows the type of operand and other data to be changed. For a complete description of tolerance operands, see "TOLERANCING" on page 545.

Insert Operand: Inserts a new row in the spreadsheet at the current row. The Insert key will perform this same function.

Insert After: Inserts a new row in the spreadsheet after the current row. The Ctrl-Insert key combination will perform this same function.

Delete Operand: Deletes the row at the current cursor location. The Delete key will perform this same function.

Cut Operands: Copies all the data for a single operand or a range of operands to the Windows Clipboard, then deletes the operands. The operand or range of operands must first be selected using either of the following techniques:

Using the mouse: Click on the first operand to be selected. Hold down the left mouse button, and drag the mouse cursor to cover the range of operands desired. The selected operands will be highlighted in inverse colors. To select only one operand, drag the mouse up or down from the operand desired until two operands are selected, then move the mouse back to the desired operand.

Using the keyboard: Move the cursor to any cell on the desired operand row. Then hold down the shift key, while moving the cursor up or down until the desired range of operands is selected. The selected operands will be highlighted in inverse colors. To select only one operand, move the cursor up or down from the operand desired until two operands are selected, then move the cursor back to the desired operand.

Copy Operands: Copies all the data for a single operand or a range of operands to the Windows Clipboard. To select a single operand or a range of operands, see the discussion under "Cut Operands" above.

Paste Operands: Copies all the data for a single operand or a range of operands from the Windows Clipboard to the current cursor location in the Tolerance Data Editor. The operand data must first have been copied to the Windows Clipboard using either "Cut Operands" or "Copy Operands" described above.

Copy Cell: Copies a single cell's data to the Windows Clipboard.

Paste Cell: Pastes the data for a single cell from the Windows Clipboard to the current cell. The data must first have been copied to the Windows Clipboard using "Copy Cell" described above.

Edit Cell: Puts the cell in edit mode.

Copy Spreadsheet: Copies either the highlighted range of operands, or the entire spreadsheet (if no range of operands is selected) to the Windows Clipboard in a text format suitable for being pasted into another Windows application, such as a spreadsheet or word processor. The format is tab delimited text.

Set Row Colors: Selects either default or none for the color for each row in the spreadsheet. Individual row colors may be set using the row color setting on the operand type dialog box.

Tools

Default Tolerances: Invokes the default tolerances dialog box. See "Defining default tolerances" on page 557 for details.

Loosen 2X: Increases all tolerance ranges by a factor of 2. This is a quick way of loosening the tolerances if they are all too tight.

Tighten 2X: Decreases all tolerance ranges by a factor of 2. This is a quick way of tightening the tolerances if they are all too loose.

Sort by Surface: Sorts all operands in ascending order first by surface number, then by type. The operands COMP and CPAR will always be placed at the top of the list. The SAVE operand will automatically be moved to stay just beneath the same operand it was beneath before sorting, because the SAVE operand is associated

with the previous operand in the list. The STAT operand, if present, will be placed at the top of the list, and this operand must manually be moved or reinserted. Since STAT affects all operands which follow in the list, sorting the list will invalidate the STAT operand. Whenever STAT is used in the body of the tolerance list (to change the statistics on the fly) then editing of the list to correctly replace the STAT operand will be required whenever sorting. Note multiple STAT operands may be required if the operands that originally followed the STAT operand are dispersed through the list by the sort operation.

Sort by Type: Sorts all operands in ascending order by type, then by surface number. See Sort by Surface.

Save: Saves the current tolerance data in a *.TOL file. This step is only required if the data is to be subsequently loaded into another lens or stored for archival purposes. ZEMAX automatically stores the tolerance data with the lens when the entire lens is saved.

Load: Loads tolerance data previously stored in a *.TOL file or in a *.ZMX file. Either file type may be selected; only the tolerance data portion of the file will be loaded into the spreadsheet. The current tolerance data is destroyed.

Validate: Checks all tolerance operands for proper surface ranges and other common errors.

View

Hide Current Column: Minimizes the width of the column displaying the cursor.

Unhide All Columns: Sets all hidden columns to the default width (see "Editors" on page 71).

Reset Column Widths: Sets all columns to the default width (see "Editors" on page 71).

Equal Column Widths: Sets all columns to the same width as the column displaying the cursor.

Go To Operand: Opens a dialog box that allows entry of either the operand number or comment. The editor then displays the desired operand.

Help

Opens the help system.

Extra Data

The extra data editor is very similar to the lens data spreadsheet editor, except only the extra data values are displayed and edited. See "Extra data" on page 277 for a discussion of the surface types which use extra data values. Surfaces cannot be inserted or deleted on the extra data editor.

Menu Options

Edit

Copy Cell: Copies a single cell's data to the Windows Clipboard.

Paste Cell: Pastes the data for a single cell from the Windows Clipboard to the current cell. The data must first have been copied to the Windows Clipboard using "Copy Cell" described above.

Edit Cell: Puts the cell in edit mode.

Copy Spreadsheet: Copies the entire spreadsheet to the Windows Clipboard in a text format suitable for being pasted into another Windows application, such as a spreadsheet or word processor. The format is tab delimited text.

Solves

Solve Type: Selecting this option will invoke the solve dialog box for the cell which is currently indicated by the cursor. For more information on solves, see "SOLVES" on page 473.

Variable Toggle: This will toggle the variable status of the currently selected cell.

Tools

Import: The import tool is used to load extra data values for extra data surfaces from a file rather than by typing the numbers in directly. This menu option will invoke a dialog box which presents a list of data files which end in the extension .DAT. The dialog box also permits specification of which surface number should receive the data. Numerical data must be in the file exactly as it appears in the extra data spreadsheet. The format of the file is a single column of free-format numbers, and the file must end in the DAT extension. Data for grid sag

and phase surfaces may also be imported using this tool. For a description of the proper file format, see “Importing grid data” on page 314. Initially, the folder ZEMAX will look for DAT files in the <data> folder (see “Folders” on page 68). However, after any import or data output operation, the most recently used folder will be used.

View

Hide Current Column: Minimizes the width of the column displaying the cursor.

Unhide All Columns: Sets all hidden columns to the default width (see “Editors” on page 71).

Reset Column Widths: Sets all columns to the default width (see “Editors” on page 71).

Equal Column Widths: Sets all columns to the same width as the column displaying the cursor.

Go To Surface: Opens a dialog box that allows entry of either the surface number or comment. The editor then displays the desired surface. Note the Extra Data Editor does not display surface comments, the comments from the Lens Data Editor are used for searching.

Help

Opens the help system for using the Extra Data Editor.

Non-Sequential Components



This feature is only available in the EE edition of ZEMAX.

The NSC Editor is only used by the Non-Sequential Components surface type supported by ZEMAX-EE. The editor is very similar to the lens data spreadsheet editor. See “NON-SEQUENTIAL COMPONENTS” on page 477 for a detailed discussion of non-sequential components and ray tracing. The editor can only be activated if a Non-Sequential Components surface is present in the Lens Data Editor, or if the program mode is set to Non-Sequential Design. For information on changing the program mode, see “Program Mode” on page 67.

Menu options

The Non-Sequential Components Editor menu options are used for inserting and deleting objects, selecting object types and properties, and setting solves and variables.

Edit

Object Properties: Activates a dialog box from which the object type, surface and bulk properties may be selected. For complete documentation on the extensive settings available, see “The object properties dialog box” on page 556.

Next Group: If there is more than one Non-Sequential Components surfaces defined in the Lens Data Editor, selecting this option will advance to the group of objects for the next non-sequential surface.

Edit Object: If the object is defined by a file, such as a Polygon Object (POB), then this menu option will be available. If selected, a text editor will be invoked so that the defining data file may be edited. After editing and saving the file; choose Reload Object to update the new information in ZEMAX.

Reload Object: This reloads and recreates an object that was defined by an external data file that was changed since the lens file was last loaded. Typically this option is used to refresh a POB or STL object if another application has modified the defining data file.

Reload All Objects: This reloads and recreates all objects listed in the editor, see “Reload Object” above.

Insert Object: This inserts a new object row in the spreadsheet at the current row. The new object type is “Null Object”. The Insert key is a shortcut for this menu option.

Insert After: This inserts a new object row in the spreadsheet after the current row. The new object type is “Null Object”. The Ctrl-Insert key combination is a shortcut for this menu option.

Delete Object: This deletes the current object row in the spreadsheet. The Delete key is a shortcut for this menu option.

Cut Objects: Copies all the data for a single object or a range of objects to the Windows Clipboard, then deletes the objects. The object or range of objects must first be selected using either of the following techniques:

Using the mouse: Click on the first object to be selected. Hold down the left mouse button, and drag the mouse cursor to cover the range of objects desired. The selected objects will be highlighted in inverse colors. To select only one object, drag the mouse up or down from the object desired until two objects are selected, then move the mouse back to the desired object.

Using the keyboard: Move the cursor to any cell on the desired object. Then hold down the shift key, while moving the cursor up or down until the desired range of objects is selected. The selected objects will be highlighted in inverse colors. To select only one object, move the cursor up or down from the object desired until two objects are selected, then move the cursor back to the desired object.

Copy Objects: Copies all the data for a single object or a range of objects to the Windows Clipboard. To select a single object or a range of objects, see the discussion under "Cut Objects" above.

Paste Objects: Copies all the data for a single object or a range of objects from the Windows Clipboard to the current cursor location in the NSC Editor. The object data must first have been copied to the Windows Clipboard using either "Cut Object" or "Copy Objects" described above.

Copy Cell: Copies a single cell's data to the Windows Clipboard.

Paste Cell: Pastes the data for a single cell from the Windows Clipboard to the current cell. The data must first have been copied to the Windows Clipboard using "Copy Cell" described above.

Edit Cell: Puts the cell in edit mode.

Copy Spreadsheet: Copies either the highlighted range of objects, or the entire spreadsheet (if no range of objects is selected) to the Windows Clipboard in a text format suitable for being pasted into another Windows application, such as a spreadsheet or word processor. The format is tab delimited text.

Set Row Colors: Selects either default or none for the color for each row in the spreadsheet. Individual row colors may be set using the row color setting on the object properties dialog box.

Solves

Solves and variables may be placed on X, Y, and Z positions, tilts about the X, Y, and Z axis, and any of the parameter data. The available solves include pickup solves and ZPL Macro solves. For more information on ZPL Macro solves see "Using ZPL Macro solves" on page 479.

X Position: Sets solves on the X position of an object.

Y Position: Sets solves on the Y position of an object.

Z Position: Sets solves on the Z position of an object.

Tilt about X: Sets solves on the tilt about the X axis of an object.

Tilt about Y: Sets solves on the tilt about the Y axis of an object.

Tilt about Z: Sets solves on the tilt about the Z axis of an object.

Material: Sets solves on the material of an object.

Parameter: Sets fixed, variable, or pickup solves on any parameter of an object. Pickup solves on parameters are only allowed if the source object is a preceding object, or if the source object is the same object and the source parameter column is to the left of the destination column.

Variable Toggle: This will toggle the variable status of the currently selected cell. Ctrl-Z is a shortcut for this operation.

Tools

See "NSC Editor Tools" below.

View

Hide Current Column: Minimizes the width of the column displaying the cursor.

Unhide All Columns: Sets all hidden columns to the default width (see "Editors" on page 71).

Reset Column Widths: Sets all columns to the default width (see "Editors" on page 71).

Equal Column Widths: Sets all columns to the same width as the column displaying the cursor.

Go To Object: Opens a dialog box that allows entry of either the object number or comment. The editor then displays the desired object.

Help

Opens the help system for the NSC Editor.

NSC Editor Tools

The following tools are available from the NSC Editor Tools menu.

Replicate Object

Purpose:

Replicates an existing object into an array of identical objects. A generally superior alternative to this tool is the Array Object, see “Array” on page 399.

Settings:

Item	Description
Object	The object to be replicated.
Number X, Y, Z	The total number of objects in the array in the X, Y, and X directions.
Delta X, Y, Z	The spacing interval along each direction.
Add Pickup Solves	If checked, all replicated objects will pickup parameter data from the original object.
Relative References	If checked, all replicated objects will use relative references to the original object, otherwise, absolute references are used See “Reference objects” on page 350.

Discussion:

This feature makes multiple copies of an existing object. Each copy is offset by a different amount to create an array of identical objects. The original object is placed at the center of the array, and all replicated objects are referenced to the original object’s position and orientation.

Create Polygon Object

Purpose:

Creates POB files of basic geometrical shapes which may be used as NSC detector objects.

Settings:

Item	Description
Object Type	The type object to be created.
Parameters	The data specific to each object type.
Output Name	The name of the POB file to create.

Discussion:

Creating simple shapes as POB files is generally straightforward using a text editor. However it is often useful to make a detector object with a large number of facets, such as a cube with hundreds of facets on each side. When the object has many facets, it is easier to create the POB using this tool.

To create the POB file, choose the object type from the list, set the appropriate parameters, enter the name of the file to create, and then click OK. Note the file name should not include a path, but should include the extension POB. The tool will automatically place the POB file in the <objects>\Polygon Objects folder (see “Folders” on page 68), ready to use by the NSC editor.

For information on POB objects, see “Polygon Object” on page 429. For information on using objects as detectors, see “Objects as detectors” on page 470. For information on the POB format, see “Defining Polygon

Objects” on page 429. To create objects using Boolean operations between existing ZEMAX objects, see “Combine Objects” on page 97.

Combine Objects

This tool combines objects using Boolean operations such as AND, OR, and XOR. See also “Boolean” on page 409.

Settings:

Item	Description
Object A, B	The first and second objects to be combined.
Operation	Defines the Boolean operation to perform on the selected objects. Intersection of objects: The region contained by both objects. Union of objects: The region contained by either object. Either, But not both objects: The region contained by either but not both objects. Subtract Object B from Object A: The part of A that is not part of B. Subtract Object A from Object B: The part of B that is not part of A.
Format	The file format to use to save the resulting object.
Spline	This setting defines the number of spline points to use for each aspheric face. A larger number of spline points generally produces more accurate results at the expense of larger files, slower ray tracing, and slower object creation. This setting has no affect on objects which are not aspheric.
Ref. Coordinate	This setting defines the coordinate system used by the combined object. If Global is selected, the object will be position relative to the origin of the NSC system. If Object A or B is selected, then the object will be positioned relative to the local coordinate frame of Object A or B respectively.
Tolerance	The dimensional tolerance in lens units. Some formats use a tolerance parameter to define the accuracy of the representation. Smaller tolerance values generally produce more accurate results at the expense of larger files, slower ray tracing, and slower object creation.
Replace Objects With New Combined Object	If checked, once the object is created, the selected objects A nd B will be replaced with the newly created object. If not checked, the object will be created but will not automatically be imported to the NSC Editor.
Preview Object	If pressed, this button will present a Shaded Model display of the object that would be created with the specified settings. This Shaded Model display may be rotated and zoomed to visually verify the correct shape will be created.

Discussion:

The combine objects tool works by applying Boolean operations on the two selected objects to "trim" the overlapping regions to form a new single object. The new object is then exported in the selected CAD format, such as IGES or STEP. The new object may then optionally be imported for ray tracing. For more information on imported objects, see “Imported” on page 423. When exporting in STL format, ZEMAX imports the new object using the STL object type. For a description of the STL object see “STL Object” on page 437.

The Boolean operations are performed by converting each component object into a NURBS based representation, and then a series of Boolean trimming and combining operations are performed to yield the resulting object. For this reason, Boolean objects have the same considerations that Imported objects do; see “Comments about imported objects” on page 424 for a discussion. Some loss of precision is possible when converting from ZEMAX native, high precision representations to NURBS representations; this is not a limitation of ZEMAX but is inherent in the NURBS representation of arbitrary surfaces. Generally, precision can be increased as desired using the Spline and Tolerance parameters described above. When extreme optical precision is required, the user should verify that the Boolean object traces to a suitable accuracy.

When the maximum possible ray trace accuracy is required, an alternative to combining objects is to create a User Defined Object. For more information, see “User Defined Object” on page 444.

Modify Reference Objects

This tool alters the reference object of any single object or range of objects. The position and rotation angle for the objects are modified to retain their current position and orientation. For more information on reference objects, see “Reference objects” on page 350.

Settings:

Item	Description
First/Last Object	The first and last object defining the range of objects to be modified. All objects within the specified range will be modified.
Refer To	Defines the new reference object. All the objects within the range defined by the first and last object settings will be modified to use the new reference object. The selected reference object must precede the first object selected.

Discussion:

This tool works by computing the x, y, and z position and the tilt about the x, y, and z axes relative to the desired reference object to maintain each objects global position. Because of the finite precision of the computer, some small amount of error may be induced from multiple changes in the reference position.

Create Source Ray From Last Geometry Error Tool

This tool is used for isolating and studying geometry errors. For more information on geometry errors, see “Lost energy” on page 379. Whenever a geometry error is displayed (not merely detected) the starting ray coordinates and cosines are stored. These starting values may be used to create a single Source Ray (see “Source Ray” on page 459) that duplicates the ray with the error. If this tool is used, all sources will be set to trace zero rays for both layout and analysis. A new Source Ray will be added to the object list with the properties of the ray that generated the geometry error.

Insert/Delete Z Point Tools

These two tools either insert or delete a single point in the list of points that define the Freeform Z object. For details on this object, see “Freeform Z” on page 417.

The insert tool defines the location of the point to be inserted by the Z coordinate entered on the tool dialog box. The Z coordinate must be greater than 0, less than the current maximum Z coordinate, and cannot be equal (within 1.0E-05 of the overall object length) to any existing Z coordinate. The Y coordinate is computed automatically to fit the current object shape. This tool will not insert a point if the total number of points would be more than the maximum allowed.

The delete tool defines the point to be deleted by the Z coordinate entered on the tool dialog box. The existing point closest to the entered Z coordinate is determined and deleted. This tool will not delete a point if the total number of remaining points would be less than 5.

Export Polar Detector Data as Source File

Saves intensity data recorded on a polar detector into an IESNA or EULUMDAT file.

Settings:

Item	Description
Format	Data may be saved in either the IESNA or EULUMDAT format. For more information on the IESNA format supported by ZEMAX, see “Source IESNA File” on page 457. For more information on the EULUMDAT format supported by ZEMAX, see “Source EULUMDAT File” on page 454.
Detector	The object number of the Detector Polar from which the intensity data will be saved.

Item	Description
File Name	The name of the file to be saved. IESNA files will be saved in the <objects>\Sources\IESNA folder, while EULUMDAT files will be saved in the <objects>\Sources\EULUMDAT folder. For more information on the <objects> folder, see "Folders" on page 68.
Smoothing	The level of smoothing to use in processing the data. The smoothing algorithm is described in "The Detector Viewer" on page 384.

Description:

Intensity values recorded on a Detector Polar object may be saved into a file format commonly used for describing photometric data from a real lamp. These data represent the angular distribution of the emission only; for more information, see "Detector Polar object:" on page 464.

Save Detector Data

Saves the data for the detector currently selected in the editor. The only supported detectors are the Detector Rectangle, Detector Color, Detector Polar, and Detector Volume objects. The file format and data stored is specific to each detector type, and the distinct extensions DDR, DDC, DDP, and DDV are used to denote the different file formats. For information on the file formats, see "Saving and Loading Detector Data" on page 390.

Load Detector Data

Loads data previously saved in a DDR, DDC, DDP, or DDV file onto the detector currently selected in the editor. If the detector size, angular range, pixel count, or other property does not match the target detector, then the detector properties are automatically modified to match that of the incoming data. If the units of the stored data do not match the current units, an error is issued and the detector is not modified. For information on the file formats, see "Saving and Loading Detector Data" on page 390.

Playback ZRD on Detector

Purpose:

Reload saved ray database data onto detectors.

Settings:

Item	Description
Detector	Data may be loaded onto an individual detector or all detectors simultaneously.
ZRD File	The name of the file to be saved. IESNA files will be saved in the <objects>\Sources\IESNA folder, while EULUMDAT files will be saved in the <objects>\Sources\EULUMDAT folder. For more information on the <objects> folder, see "Folders" on page 68.
Filter	An (optional) filter string used to isolate specific rays in the file. For more information, see "The filter string" on page 386.

Discussion:

Ray trace results are often saved to a ray database (ZRD) file for further analysis. Using this tool, the results from a saved ZRD file may be reloaded onto a particular detector or onto all detectors in the system. A corresponding ZPL keyword (ZRDPLAYBACK) is available to reload the ZRD file within a macro. Using this keyword, data saved to a ZRD file may then be filtered (using filter strings) as a part of post-processing. For more information see "ZRDPLAYBACK" on page 744.

Ignore NSC Trace Errors

If checked, ray trace errors that occur inside of a non-sequential group will not be reported if detected. These errors involve possibly improper placement of objects in a non-sequential group. However, in some systems it is possible to have properly defined object placement and still occasionally an error message will be generated. Checking this option will suppress reporting of these error messages.

Undo, Redo, and Recover

There are three different states to the Undo capability in ZEMAX: None, Memory 1 Step, and Disk Multi Step. The Undo state is set on the Editors tab of the Preferences dialog box described in the Chapter "File Menu".

Undo: None

If the Undo feature state is set to None, then no Undo capability is supported. Use this option on computers which do not have sufficient system memory or disk space to support the Undo feature.

Undo: Memory 1 Step

ZEMAX stores a copy of the current lens in memory before and after every edit or optimization. If Undo is selected, then the current lens is swapped with the previous lens. If Redo is then selected, the lenses are swapped again, which results in the edit being restored.

Memory 1 Step Undo is useful for restoring a lens when an accidental edit is made, or for restoring a lens to the prior state after an optimization. However, only a single Undo step is supported. The advantage of this option is speed; the saving of the prior lens in memory is so fast as to be unnoticeable.

Undo: Disk Multi Step

ZEMAX stores a copy of the current lens in a ZMX file on disk after every edit or optimization. These stored lens files are used to implement an infinite multi-step Undo feature, which allows the reversal of any change or series of changes made to the lens. Undo is useful for restoring a lens when an accidental edit is made, or for restoring a lens to the prior state after an optimization, or even several changes.

To reverse the changes made to a lens, select Undo from the Editors Menu. Any number of Undo's may be executed, back to the first edit made after loading the lens file. The Redo function reverses the last Undo.

ZEMAX maintains a folder of Undo files within the <data>Undo folder (see "Folders" on page 68). The Undo files are automatically deleted whenever the file is saved, a new file is opened, or ZEMAX is normally terminated. If ZEMAX abnormally terminates, the operating system fails, the computer power is disrupted, or for any other reason the lens data is lost, ZEMAX may be able to recover the lost data by restoring the last Undo file. When ZEMAX starts, a check is made to see if any of the Undo files exist. Since these files are deleted during a normal termination, the presence of Undo files indicates a previous abnormal termination. ZEMAX will issue a warning message with the option of restoring the last Undo file. If restored, this file should immediately be saved in a new file name, since the old file name is not stored with the lens.

The Undo feature does slow down the operation of the Editors slightly, since every edit is followed by a save operation. The save does not slow down ray tracing or optimization speed, only lens data editing.

If more than one ZEMAX session is running simultaneously, each session will have its own undo files. However, to recover from a crash or abnormal program termination, the same number of ZEMAX sessions will need to be run to recover all the files. For example, if 2 sessions of ZEMAX are running, and the power fails, the first new session of ZEMAX will be able to recover the former first session file. A second session of ZEMAX will need to be launched to recover the former second session file.

Update

This option only updates data in the Lens Data and Extra Data editors. Update is used to recompute first-order properties, pupil positions, semi-diameters, index data, and solves. Only data presented in the Lens Data and Extra Data editors are affected. See "Update All" in this chapter for more information.

Update All

This option updates all windows to reflect the most recent lens data. ZEMAX does not automatically change all data presented in graphic and text windows to reflect the latest lens data. The reason is that if ZEMAX constantly recalculated MTF, ray fan, spot diagram, and other data while new data is being typed in the lens data editor, then program response would be very slow. Instead, make all the required changes to the lens, and then select Update All to refresh and recompute all the data windows.

Individual graphic and text windows (not editors) can also be updated by double clicking anywhere within the window.

General

This option invokes the General System Data dialog box, which is used to define commonly used data that pertains to the lens as a whole system, rather than data associated with a single surface. The general data is divided into groups as follows.

Aperture

The aperture group consists of the following settings.

Aperture Type

The system aperture defines the size of the beam through the system on axis. To set the system aperture, you need to define the system aperture type and the system aperture value. Use the cursor keys to select from the drop down list the desired system aperture type. The available system aperture types and the associated type codes (used by ZPL macros) are defined in the following table.

SYSTEM APERTURE TYPES AND CODES

Aperture type and code	Description
Entrance Pupil Diameter, 0	The diameter of the pupil in lens units as seen from object space
Image Space F/#, 1	The infinite conjugate paraxial F/# in image space
Object Space Numerical Aperture, 2	The numerical aperture ($n \sin \theta_m$) of the marginal ray in object space
Float By Stop Size, 3	Defined by the semi-diameter of the stop surface
Paraxial Working F/#, 4	The paraxial F/# in image space for the defined conjugates.
Object Cone Angle, 5	<p>The half angle in degrees of the marginal ray in object space, which may exceed 90 degrees. Object cone angle may not be used if the entrance pupil is virtual, that is, if the distance from the object to the entrance pupil is negative. When using object cone angle, the default "Uniform" apodization of rays in the pupil is uniform in angle space, rather than uniform on a plane. If the apodization type is set to "Cosine Cubed" then the rays are uniformly distributed in solid angle, as would be correct for a point source radiating uniformly in all directions. See "Apodization Type" on page 102.</p> <p>This ray distribution may be significantly different than other aperture type settings for large cone angles.</p>

These terms are further defined in "CONVENTIONS AND DEFINITIONS" on page 51. If you select "Object Space NA" or "Object Cone Angle" for a system aperture type, the object thickness must be less than infinity. Only one type of system aperture may be defined; for example, once the Entrance Pupil Diameter is specified, all the other aperture definitions are determined by the lens data prescription.

Aperture Value

The system aperture value meaning depends upon the system aperture type selected. For example, if "Entrance Pupil Diameter" is selected as the system aperture type, the system aperture value is the entrance pupil diameter in lens units. ZEMAX uses the system aperture type and system aperture value together to determine such fundamental quantities as the entrance pupil size, and clear apertures for all components.

Selecting "Float by Stop Size" for a system aperture type is the only exception to this rule. The semi-diameter of the stop surface (set on the lens data editor) is used to define the system aperture if "Float by Stop Size" is selected for a system aperture type.

Apodization Type

By default, the pupil is always illuminated uniformly. However, there are times when the pupil should have a non-uniform illumination. For this purpose, ZEMAX supports pupil apodization, which is a variation of amplitude over the pupil. Three types of pupil apodization are supported: uniform, Gaussian, and cosine cubed. User defined apodizations are also supported.

Uniform apodization

Uniform means rays are distributed uniformly over the entrance pupil, simulating uniform illumination. This is usually the case for distant objects.

Gaussian apodization

Gaussian apodization imparts an amplitude variation over the pupil that is Gaussian in form. The apodization factor refers to the rate of decrease of the beam amplitude as a function of radial pupil coordinate. The beam amplitude is normalized to unity at the center of the pupil. The amplitude at other points in the entrance pupil is given by

$$A(\rho) = e^{-G\rho^2},$$

where G is the apodization factor and ρ is the normalized pupil coordinate. If the apodization factor is zero, then the pupil illumination is uniform. If the apodization factor is 1.0, then the beam amplitude has fallen to the 1 over e point at the edge of the entrance pupil (which means the intensity has fallen to the 1 over e squared point, about 13% of the peak). The apodization factor can be any number greater than or equal to 0.0. Values larger than about 4.0 are not recommended. This is because most computations will sample too few rays to produce meaningful results if the beam amplitude falls too quickly at the edges of the pupil.

Cosine cubed apodization

Cosine cubed apodization simulates the intensity fall off characteristic of a point source illuminating a flat plane. Note cosine cubed apodization is only useful and should only be used for point sources or field points close to the axis when compared to the entrance pupil diameter. For a point source, the intensity of a ray illuminating a differential area on a plane is given by

$$I(\theta) = (\cos\theta)^3,$$

where θ is the angle between the z axis and the ray intersecting the entrance pupil, and the relative intensity at the center of the pupil is 1.0. Converting to normalized pupil coordinates and taking the square root yields the pupil coordinate amplitude apodization

$$A(\rho) = \frac{1}{(1 + (\rho \tan\alpha)^2)^{3/4}},$$

where $\tan\alpha$ is the tangent of the angle between the z axis and the marginal ray. ZEMAX uses the entrance pupil position and size to compute $\tan\alpha$. The apodization factor is not used by the cosine cubed apodization.

If the aperture type is "Object Cone Angle" then a slightly different apodization technique is used. The apodization is achieved by careful selection of the ray cosines in object space. The resulting distribution is uniform in solid angle rather than angle space. The solid angle Ω of a sphere illuminated by a cone of angle θ is given by

$$\Omega = 2\pi(1 - \cos\theta).$$

The normalized radial pupil coordinate p is related to the fractional solid angle by

$$p^2 = \frac{1 - \cos\theta}{1 - \cos\alpha},$$

where α is the maximum cone angle defined by the system aperture value and θ is the angle of the ray with the normalized pupil coordinate p . This apodization will yield uniform ray density and power per solid angle everywhere over the defined cone.

User defined apodization

ZEMAX also supports user defined apodizations on any surface, rather than just the entrance pupil. User defined surface apodizations are implemented using the user defined surface type described in "User defined surface apodization using DLLs" on page 333.

Apodization Factor

The apodization factor determines how fast the amplitude decays in the pupil. See the previous section on the Apodization Type for details.

Telecentric Object Space

If this checkbox is selected, ZEMAX will assume the entrance pupil is located at infinity, regardless of the location of the stop surface. All chief rays leaving the object surface will be parallel to the local Z axis. Ray aiming and field points defined using angles may not be used when this option is selected. For best performance, set the stop surface to surface 1 when using telecentric mode.

Afocal Image Space

If this box is checked, ZEMAX will perform most analysis features in a manner appropriate for optical systems with output beams in image space that are nominally collimated. Strictly speaking, the term afocal means the optical beam is collimated in both object and image space. However, ZEMAX uses this term to describe collimated output in image space, whether the beam is collimated in object space or not. For brevity, the documentation uses the term "afocal mode" to indicate Afocal Image Space is checked on, and "focal mode" to indicate Afocal Image Space is checked off.

When using afocal mode, transverse, longitudinal, and MTF aberrations are computed in units appropriate to afocal systems. Transverse aberrations are computed as a function of angle relative to a reference ray rather than units of length. Longitudinal aberrations are computed as a function of defocus in diopters (inverse meters) rather than as defocus in units of length. MTF is measured in cycles per angle rather than cycles per length. The change of units is not always mentioned in the documentation for each individual analysis feature. When using afocal mode, the user should be aware that some analysis window settings described in focal units are actually in afocal units if afocal mode is on. Afocal aberrations may be measured in various units, with milliradians being the default. To change the afocal unit type, see "Afocal Mode Units" on page 105. The following table compares focal and afocal system analysis units.

COMPARISON OF FOCAL AND AFOCAL UNITS

Analysis type	Focal Unit	Afocal Unit
Transverse Aberrations	micrometers	milliradians (or other user selectable angular unit)
MTF	cycles per millimeter or cycles per milliradian (see "MTF Units" on page 106)	cycles per angular unit
Field Curvature, Longitudinal Aberrations, Defocus	units of length	diopters (inverse Meters)
Diffraction Limited Airy Radius	$1.22 \lambda F$ (length) where F is the F/#	$1.22 \lambda / D$ (angle) where D is the exit pupil diameter

Wavefront error (or Optical Path Difference) is measured somewhat differently in afocal mode than in focal mode. For afocal systems, the OPD is measured relative to a reference plane that is normal to the chief ray of the reference wavelength. If the OPD is referenced to the exit pupil (the recommended setting) then the afocal OPD is measured at the plane located at the exit pupil position relative to the image surface. If the OPD reference is set to absolute then the OPD is referenced to a plane located at the image surface. Note that a shift in the location of the plane at which the afocal OPD is referenced will in general change the OPD values, unless the beam is perfectly collimated. For more information on the OPD reference, see "Reference OPD" on page 110.

Analysis features which compute data as a function of defocus use units of length in focal mode and diopters in afocal mode. Note that a positive defocus in diopters means the image shifts toward the lens, which is a negative shift in length. For this reason through focus plots will appear to flip left-to-right when switching between a focal mode system and an afocal mode system with the same fundamental aberrations.

Other than the change of units, most ZEMAX features work exactly the same in focal and afocal mode. Some features do not have meaning for afocal systems and an error message or possibly meaningless data will result if the analysis is attempted in afocal mode. To optimize afocal systems, use one of the angle based default merit functions described in "Selecting the type of optimization" on page 483.

Iterate Solves When Updating

Solves placed on parameters in the Lens Data Editor sometimes require iteration to compute accurately. Typically this includes curvature and thickness solves that depend upon ray tracing. The tracing of specific rays, such as the marginal or chief ray, depends upon pupil positions. These pupil positions may change when the solve value is updated, creating the need for iteration. The primary symptom of the need for iteration is when the merit function increases or jumps around erratically during optimization.

ZEMAX normally detects the need for iteration. However, if ZEMAX does not iterate automatically, checking this option will force iteration on all solves when updating the lens data. Iteration does slow down calculations, particularly optimization, and so this option should only be checked on when required.

Units

Lens Units

Lens units defines the units of measure for dimensions in most of the spreadsheet editors. These dimensions apply to data such as radii, thickness, entrance pupil diameters, non-sequential position coordinates, and most other parameters in ZEMAX.

There are four choices for lens units: millimeters, centimeters, inches, or meters.

For most imaging analysis features, such as ray fans and spot diagrams, the displayed units are micrometers. The symbol for micrometers is μm . The term microns is also commonly used in optics. The unit micron is identical to micrometer. Wavelengths are always defined in micrometers. The choice of lens units has no effect on the wavelength units.

Source Units

Source units define the unit of measure for the flux (power) or energy emitted by non-sequential sources. This setting is used for sources defined in the non-sequential component editor, and for defining the power and irradiance in physical optics analysis. Source units may be either watts, lumens, or joules, with the additional choice of the prefixes femto, pico, nano, micro, milli, kilo, mega, giga, or tera. Watts are used for radiometric analysis, lumens for photometric analysis, and joules for energy analysis. The key difference between radiometric and photometric units is that photometric units are wavelength weighted to the response of the human eye. See the table in the next section for more information on units.

Analysis Units

Analysis units define the unit of measure for irradiance (radiometric) or illuminance (photometric). This setting only affects data as displayed on detectors collecting light from sources defined in the non-sequential components editor. Irradiance units are watts/area, where area is in square meters, centimeters, millimeters, feet, or inches. Illuminance units are lumens/area. Energy density units are joules/area. The unit prefixes femto, pico, nano, micro, milli, kilo, mega, giga, or tera are all supported.

Here is a summary of radiometric, photometric, and energy units by ZEMAX.

RADIOMETRIC, PHOTOMETRIC, AND ENERGY UNITS

Radiometric Unit	Photometric Unit	Energy Unit
<u>Flux</u> watt	<u>Flux</u> lumen	<u>joule</u>
<u>Radiant Intensity</u> watt/steradian	<u>Luminous Intensity</u> lumen/steradian = candela	<u>Intensity</u> joule/steradian
<u>Irradiance</u> watt/meter ² watt/centimeter ² watt/millimeter ² watt/foot ² watt/inch ²	<u>Illuminance</u> lumen/meter ² = lux = meter-candle lumen/centimeter ² = phot lumen/foot ² = footcandle	<u>Irradiance</u> joule/meter ² joule/centimeter ² joule/millimeter ² joule/foot ² joule/inch ²
<u>Radiance</u> watt/steradian-meter ²	<u>Luminance</u> lumen/steradian-meter ²	<u>Radiance</u> joule/steradian-meter ²

One unfortunate convention in the optics industry (and in ZEMAX) is to use the term "Intensity" for two distinctly different things. Intensity is commonly used in ray tracing applications and the optical design community to define the flux represented by a single ray, measured in watts or lumens. Intensity is also used in the illumination and radiometry fields, but the definition used there is flux per solid angle, measured in watts per steradian or lumens per steradian (candela).

ZEMAX keeps track of energy by associating with each ray a flux value, which is equal to the square of the electric field vector also associated with the ray. When this ray strikes a detector object, ZEMAX computes the area the pixel represents to compute irradiance (flux/area) and the solid angle the pixel represents to compute the radiometric intensity (flux/solid angle). The difference in these two uses of the word intensity can always be resolved by considering the nature of the analysis ZEMAX is performing.

For more information on radiometric and photometric units, see the Handbook of Optics referenced in "What doesn't ZEMAX do?" on page 37.

Afocal Mode Units

Afocal mode units may be microradians, milliradians, radians, arc-seconds (1/3600 of a degree), arc-minutes (1/60 of a degree), or degrees. For more information on afocal mode, see "Afocal Image Space" on page 103.

MTF Units

MTF units for focal systems may be cycles per millimeter or cycles per milliradian. When using cycles per millimeter, the MTF is computed for spatial frequencies in image space on the image surface. When using cycles per milliradian, the MTF is computed for angular frequencies in object space. The MTF units selection affects the units for all MTF computations, including analysis, optimization, and tolerancing. The MTF units selection is not used in afocal mode, as afocal mode MTF always uses cycles per Afocal mode unit.

Title/Notes

Lens Title

The lens title will appear on graphic and text output. The title is specified by typing the title in the required space. Additional text data may be placed on most graphic output, see the Chapter "File Menu" under "Preferences" for details.

Notes

The notes section allows entry of a few lines of text which is stored with the lens file.

Glass Catalogs

This control group has an edit field which lists the names of the currently used glass catalogs (without file extensions). Multiple catalogs may be specified by separating the names with a space. To add a catalog to the list, either type the name of the catalog in the list, or choose the catalog name from the drop down list, then click on "Use This Catalog". To remove a catalog from the list, either delete the name of the catalog from the list, or choose the catalog name from the drop down list and click on "Don't Use Catalog".

Glass catalogs must be placed in the glass catalog folder (see "Folders" on page 68). For more information on changing glass catalogs, see "USING GLASS CATALOGS" on page 587. To set a default catalog, see "Editors" on page 71.

Ray Aiming

The ray aiming group has the following controls.

Ray Aiming

Ray aiming is an iterative ray tracing algorithm in ZEMAX that finds rays at the object that correctly fill the stop surface for a given stop size. It is generally only required when the pupil, which is the image of the stop as seen from object space, is considerably aberrated, shifted, or tilted.

If ray aiming is off, ZEMAX will use the paraxial entrance pupil size and location determined by the aperture settings and calculated at the primary wavelength on axis for launching rays from the object surface. This means ZEMAX will ignore entrance pupil aberration. For slow systems with modest field angles, this is perfectly acceptable. However, certain systems, such as those with low F/#'s or large field angles, may have significant entrance pupil aberration. The two primary effects of pupil aberration are the shift in location of the pupil with field angle, and the anamorphic scaling of the edges of the pupil.

ZEMAX can be instructed to account for aberrations of the pupil using ray aiming. With ray aiming, every ray trace is performed iteratively, with the program adjusting the ray coordinates or cosines in object space so that the ray crosses the correct location on the stop surface. To determine the correct location on the stop surface, the stop surface radius must be computed.

The stop surface radius is computed by tracing a marginal ray from the center of the object to the stop surface at the primary wavelength. Either paraxial or real rays may be used in this trace to determine the stop radius. Paraxial rays are well behaved and paraxial definitions are commonly used for most first-order system properties such as focal length, F/#, and magnification, and so paraxial rays may also be used to determine the stop size. However, for systems that have significantly aberrated pupils, there will be a difference between the paraxial and real ray stop radius. These systems will exhibit a difference between the real and paraxial ray system aperture. For example, the paraxial object space numerical aperture may be defined as 0.4, but the actual numerical aperture of real rays may be a different value. For real rays to have the object space properties defined by the system aperture, use real rays instead of paraxial rays to determine the stop radius. Note that real ray based ray aiming will not work in systems where the stop lies in a caustic or where the real rays cannot be traced at the full entrance pupil diameter or numerical aperture. If real ray aiming causes any of these problems, set the ray aiming

to paraxial rather than real. Note that once the stop radius is determined all rays are aimed to the correct location on the stop, regardless of whether paraxial or real rays were used to determine the stop radius.

To eliminate any ambiguity in the calculation of the actual stop size, set ray aiming to paraxial or real, and then set the system aperture type to "float by stop size". This eliminates the need for any ray tracing at all to determine the stop size, and both real and paraxial rays will be aimed to the real stop exactly.

For systems with virtual stops, such as some eyepieces, the effective stop location and size may be a function of wavelength. For these systems, use the multi-configuration capability to treat each wavelength and system aperture definition separately.

Although ray aiming is more accurate than paraxial entrance pupil aiming, most ray traces will take from two to eight times as long to perform. Therefore, ray aiming should only be used when required. To determine the amount of entrance pupil aberration in your system, select ray aiming off, and then look at the pupil aberration plot (see "Pupil Aberration" on page 129). Pupil aberration of less than a few percent is generally insignificant. If your system has significant pupil aberration, select ray aiming on and repeat the calculation. The aberration will decrease to zero, or very nearly so. Ray aiming does not, of course, actually eliminate pupil aberration, it merely accounts for it.

Use Ray Aiming Cache

If checked, ZEMAX caches ray aiming coordinates so that new ray traces take advantage of previous iterations of the ray tracing algorithm. Using the cache can speed up ray tracing dramatically. However, use of the cache does require that the chief ray can be traced accurately. For some systems, the chief ray cannot be traced, and for these systems, the cache should be turned off.

Robust Ray Aiming (slow)

If checked, ZEMAX uses a more reliable, but slower algorithm for aiming rays. This switch should only be set if the ray aiming algorithm is failing even with the cache turned on. This switch has no affect unless the ray aiming cache is checked on. Robust mode goes through an additional check to make sure that if multiple ray paths to the same stop surface location exist, only the correct one is chosen. This is typically a problem in very fast, very wide angle systems where off axis fields may find a virtual path to the stop that confuses the ray aiming iteration.

Pupil Shift, Pupil Compress

For some very wide angle or highly tilted or decentered systems, the ray aiming feature will fail if unassisted. The problem is that the paraxial entrance pupil is used as a first guess to trace the ray. If the pupil aberration is severe, it is possible that even this first guess cannot be traced, which prevents the algorithm from taking a second, more refined guess.

The solution is to provide a rough guess as to how much the pupil has been shifted and compressed with respect to the paraxial pupil. There are three shift components; x, y, and z; all measured in lens units. There are two compress components; x and y, and these are dimensionless scaling factors. The default value of zero for all five may be modified to assist the algorithm in finding a successful first guess for the ray aiming.

The shifts move the center of the aim point on the paraxial entrance pupil. Positive values for the z shift indicate that the aim point is to the right of the paraxial pupil, negative values indicate the pupil is shifted to the left. Most wide angle systems have left-shifted pupils. The z pupil shift value provided is scaled linearly with the field angle of the ray being traced, so the pupil shift refers to the offset of the pupil at full field. If "Scale pupil shift factors by field" is selected, the x and y pupil shift values are also scaled with field, otherwise, the x and y shift values are used for all fields without any scaling. All shifts are in lens units.

The x and y compress values are used to change the relative coordinates on the paraxial entrance pupil to start the iteration. The pupil coordinates used to *start the ray aiming iteration* are given by:

$$P'_x = P_x(1 - C_x), \text{ and}$$

$$P'_y = P_y(1 - C_y),$$

where C_x and C_y are the ray aiming compress values and the P values are the normalized pupil coordinates. The P' values are the modified pupil coordinates used to launch the first ray, thereafter the rays are aimed to the actual pupil coordinates defined by P. By defining the compress values this way, a value of zero means no compress, while a value of 0.1 indicates the pupil is compressed 10%. The compress values are particularly useful when the real pupil is smaller than the paraxial pupil, and rays traced at the full paraxial pupil size are difficult or impossible to trace.

It is important to understand that the exact values of the pupil shift and compress values is unimportant. Once the first guess ray can be traced, the algorithm will robustly find the exact pupil location. The pupil shift and compress values are just to get the ray aiming started. Neither the shift nor the compress values actually change the size of the entrance pupil. Generally, guessing at the pupil shift and compress values is an acceptable way of determining a suitable value.

Environment

This tab allows definition of the system temperature and pressure. Note that wavelengths are always measured in micrometers referenced to "air" at the system temperature and pressure. If the system temperature and air pressure change; care should be taken to adjust the wavelength definitions to match the new environment. See "Wavelength data" on page 62 and "THERMAL ANALYSIS" on page 603 for more information.

For more information on changing the temperature and pressure of individual surfaces see "Defining multiple environments within a single configuration" on page 606.

Adjust Index Data To Environment

If checked, all index data used for ray tracing will be adjusted from the "raw" glass catalog values to be relative to the system temperature and pressure. For details see "Index of refraction computation" on page 603. Individual surfaces may have a temperature and pressure that is different from the system temperature and pressure; for these surfaces the index data will be adjusted to reflect the environment of the surface.

If unchecked, all index data will be used directly from the glass catalogs, without any adjustment made for temperature or pressure changes or discrepancies between glass data defined for different temperatures. Warnings will be listed in the index data portion of the Prescription Data report (see "Prescription Data" on page 271) if this box is unchecked and the reference temperature of any glass differs from the system temperature by more than 6 degrees Celsius. Note that when the adjust index box is unchecked the system temperature is set to 20 degrees C and the pressure to 1.0 atmospheres, and therefore all index data must be relative to that environment.

It is highly recommended that "Adjust Index Data To Environment" always be checked "on". There is no disadvantage in leaving the switch on at all times, as long as all index data has been properly defined.

Temperature in degrees C

The system temperature in degrees Celsius.

Pressure in ATM

The system air pressure in atmospheres. A value of 0.0 implies vacuum, 1.0 implies sea level.

Polarization (ZEMAX-EE only)

For a detailed discussion of polarization and thin films, see "POLARIZATION ANALYSIS" on page 609.

The polarization tab is used to set the default input polarization state for many sequential analysis computations which use polarization ray tracing. For sequential analysis features such as Spot Diagrams and RMS vs. Field that support the "Use Polarization" switch to enable polarization ray tracing and apodization, this tab is the only means for setting the initial polarization state.

For most (but not all) sequential analysis features, the polarization ray trace is only used to determine the transmitted intensity of the ray while accounting for Fresnel, thin film, and bulk absorption effects. The rays are attenuated in intensity and a weighted computation is performed. Polarization phase aberrations and the vectorial nature of polarization are ignored.

Some features consider not only transmission, but also the separate orthogonal vector components of the polarized light, and polarization phase aberrations (these polarization aberrations can be calculated using the

polarization phase aberration plot, see “Phase Aberration” on page 218). The Huygens PSF and PSF Cross Section, Huygens MTF, and Encircled Energy using the Huygens PSF all consider the full polarization vector and polarization phase aberrations. These computations work by computing data for the Ex, Ey, and Ez components of the polarized electric field separately, then incoherently summing the results. The polarization phase aberrations induced in each orthogonal component of the electric field are considered as any other phase aberration. For a discussion of how Physical Optics Propagation considers polarization, see “Accounting for polarization” on page 637.

To define the polarization state of non-sequential sources, see “Sources tab” on page 561.



Polarization ray tracing is only available in the EE edition of ZEMAX.

Unpolarized

If checked, then the polarization values Jx, Jy, X-Phase, and Y-Phase are ignored, and an unpolarized computation will be performed. An unpolarized computation is performed by tracing two rays with orthogonal polarization and averaging the resulting transmission. Note that unpolarized computations take longer than polarized computations, which take longer than computations that ignore polarization entirely.

Some ZEMAX features, such as the optimization operand CODA, require a single specific polarization state to be defined. These features will use the defined Jx, Jy, X-Phase, and Y-Phase values defined below even if the "Unpolarized" option is checked.

Jx, Jy, X-Phase, Y-Phase

The polarization is defined by four numbers: Jx, and Jy, which are the magnitudes of the electric field in X and Y, and the X-Phase and Y-Phase, which are the phase angles in degrees. Internally, ZEMAX normalizes the electric field vectors to have unity intensity. For an important discussion of the conversion from Jones vectors to 3D electric field, see “Defining the initial polarization” on page 622.

To edit these values, uncheck the "Unpolarized" box first.

Convert thin film phase to ray equivalent

If selected, ZEMAX converts the polarization phase computed using thin film conventions to phase along the ray. If unselected, the ray coefficients will not be converted from the field coefficients. For more details on the difference between these conventions, see “Field vs. ray phase conventions” on page 610. The recommended and default setting is to convert the field thin film phase to ray phase.

Method

For a discussion of this feature, see “Defining the initial polarization” on page 622.

Files

Coating File

The name of the file, in the coatings folder (see “Folders” on page 68), that contains the coating material and layer definitions used by this lens. The default name is COATING.DAT. Each lens file may use a separate coating file if desired.

Scatter Profile

The name of the file, in the <data>\Profiles folder (see “Folders” on page 68), that contains the scattering profiles used by this lens. New scatter profiles may be added or old ones deleted on the Scattering tab of the NSC Objects property dialog box. See “Coat/Scatter tab” on page 558 for details. The default name is SCATTER_PROFILE.DAT. Each lens file may use a separate scatter profile file if desired.

ABg Data File

The name of the file, in the <data>\ABg_Data folder (see “Folders” on page 68), that contains the ABg data definitions used by this lens. The default name is ABG_DATA.DAT. Each lens file may use a separate ABg data file if desired.

GRADIUM Profile

The name of the file, in the glass catalog folder (see "Folders" on page 68), that contains the GRADIUM surface profile data. The default name is PROFILE.GRD. GRADIUM profile files must end in the extension GRD. See "GRADIUM™" on page 309.

Miscellaneous

Reference OPD

The Optical Path Difference, or OPD, is of value in optical design computations because the OPD represents the phase error of the wavefront forming an image. Any deviations from zero OPD contribute to a degradation of the diffraction image formed by the optical system.

Because the exit pupil is the image of the stop in image space, the exit pupil represents the only location in image space where the beam has a clearly defined edge. The illumination at the exit pupil is generally smoothly varying in amplitude and phase, and there is a clearly defined boundary between regions of zero amplitude and non zero amplitude. A reasonable assumption is that there are no diffraction effects apparent in the wavefront when viewed in the exit pupil. This is asymptotically true if all apertures in the optical system are large compared to the stop limited beam size incident upon each aperture. Even if the exit pupil is virtual, which is often the case, the exit pupil still defines the only location in image space where the beam is diffraction free. For more information on diffraction image formation and the importance of the exit pupil, see any of the good books referenced in "What doesn't ZEMAX do?" on page 37.

As the wavefront propagates from the exit pupil towards the image surface, the beam profile becomes complex in amplitude and phase, and the wavefront extends over all space due to the effects of diffraction. For this reason, the phase error as measured in the exit pupil is uniquely and critically important to the description of the wavefront and image quality.

ZEMAX by default uses the exit pupil as a reference for OPD computations. Therefore, when the OPD is computed for a given ray, the ray is traced through the optical system, all the way to the image surface, and then is traced backward to the "reference sphere" which lies in the exit pupil. The OPD as measured back on this surface is the physically significant phase error important to diffraction computations, such as MTF, PSF, and encircled energy. The additional path length due to the tracing of the ray backwards to the exit pupil, subtracted from the radius of the reference sphere, yields a slight adjustment of the OPD called the "correction term". This computation is correct and is the desired method for all cases of practical interest.

However, ZEMAX does permit selection of other reference methods.

The reference to "Infinity" makes the assumption that the exit pupil is very far away and that the OPD correction term is given strictly by the angular error in the ray. This setting should only be used if ZEMAX is unable to correctly compute the effective exit pupil location and size. This can occur with some very unusual optics which do not form images (real or virtual) of the stop surface, or for some non-axial systems where the computed exit pupil lies too close to the image surface for the exit pupil correction to be accurate. Do not use the Infinity reference unless there is compelling reason to believe that the exit pupil computation is not correct.

The reference to "Absolute" and "Absolute 2" means that ZEMAX does not add any correction term to the OPD computation. These methods are not usually physically significant or required for focal systems.

For afocal systems, choosing "Absolute" will refer the OPD to a plane normal to the chief ray at the location of the image surface, regardless of the exit pupil location. This is a useful option if the exit pupil position cannot be determined accurately, as may be the case for non-axial afocal systems. The "Absolute 2" method is similar, except the OPD is not referenced to a plane normal to the chief ray.

In summary, always use the "exit pupil" reference, unless a compelling reason exists to use one of the other settings. Erroneous data may easily result if "exit pupil" is not selected.

Paraxial Rays

Paraxial ray properties are generally not defined for non-rotationally symmetric systems. For this reason, ZEMAX by default ignores all tilts and decenters due to coordinate breaks when tracing paraxial rays. By ignoring tilts and decenters, ZEMAX can compute the paraxial properties of an equivalent centered system, which is generally the correct approach even for systems without symmetry. The default setting "Ignore Coordinate

Breaks" is therefore highly recommended for systems where ignoring tilts and decenters results in a reasonable axial approximation to the actual system.

There are however cases where "Consider Coordinate Breaks" may be required. For ray tracing through gratings, coordinate breaks may be required even for paraxial rays, otherwise, the rays may not be able to satisfy the grating equation. Ray tracing through non-sequential objects may also require that paraxial rays consider coordinate breaks.

Semi Diameter Margin (lens units)

The semi-diameter of every surface in "automatic" mode is computed to be the radial aperture required to pass all rays without clipping. For systems with closely spaced elements in or near edge contact, this yields surface apertures that provide no clearance for finishing or mounting. Often, optical surfaces are only well finished within some fraction of the full radial aperture, typically between 90% and 98%, depending upon the part size.

The semi diameter margin control allows specification of an additional amount of radial aperture as a fixed number in lens units. The default value of zero leaves no margin, while a margin of 2.0 lens units would add 2.0 lens units to the semi-diameters of all surfaces under "automatic" control.

See "Semi Diameter Margin %" below. If both "percent" and "lens units" margin values are non-zero, the percent is added first, then the lens unit margin.

Semi Diameter Margin %

The semi-diameter of every surface in "automatic" mode is computed to be the radial aperture required to pass all rays without clipping. For systems with closely spaced elements in or near edge contact, this yields surface apertures that provide no clearance for finishing or mounting. Often, optical surfaces are only well finished within some fraction of the full radial aperture, typically between 90% and 98%, depending upon the part size.

The semi diameter margin control allows specification of an additional amount of radial aperture as a percentage. The default value of zero leaves no margin, while a margin of 5% would add 5% to the semi-diameters of all surfaces under "automatic" control. The maximum allowed margin is 50%. The semi-diameter margin is not added to the stop surface.

See "Semi Diameter Margin (lens units)" above. If both "percent" and "lens units" margin values are non-zero, the percent is added first, then the lens unit margin.

Global Coordinate Reference Surface

Global coordinates are defined by a rotation and translation from the local coordinates on each surface. The conversion can be written as

$$\begin{bmatrix} x_g \\ y_g \\ z_g \end{bmatrix} = \begin{bmatrix} x_o \\ y_o \\ z_o \end{bmatrix} + \begin{bmatrix} R_{11} & R_{12} & R_{13} \\ R_{21} & R_{22} & R_{23} \\ R_{31} & R_{32} & R_{33} \end{bmatrix} \begin{bmatrix} x_l \\ y_l \\ z_l \end{bmatrix},$$

where the "g" subscript indicates global coordinate, "o" indicates an offset coordinate (the translation) and "l" indicates local coordinate. The rotation matrix R and the offset vector can be computed for any surface using any other surface as a global reference.

The rotation matrix yields valuable insight into the orientation of the surface coordinate system with respect to the global reference surface. On the local surface, the unit vector which is oriented along the local x direction is (1, 0, 0). This vector may be rotated using the R matrix to yield the orientation of the local x axis in the global coordinate system. A matrix multiplication for each of the three local axis unit vectors yields

$$\hat{x}_l = \begin{bmatrix} R_{11} \\ R_{21} \\ R_{31} \end{bmatrix}, \hat{y}_l = \begin{bmatrix} R_{12} \\ R_{22} \\ R_{32} \end{bmatrix}, \hat{z}_l = \begin{bmatrix} R_{13} \\ R_{23} \\ R_{33} \end{bmatrix}.$$

Note that the global coordinate system orientation of the unit vectors oriented along the local coordinate axes are the columns of the R matrix. The global vertex listing on the prescription report lists the R matrix components and offset vector for each surface, all referenced to the global reference surface. If the surface is a coordinate break, then the R matrix includes the effects of the coordinate translation and rotation. If the reference surface is a coordinate break, the reference coordinate system is defined after the decentration and rotation of the local coordinate system. If the surface for which R is computed precedes the reference surface, and it is a coordinate break, the R matrix is for the system prior to the coordinate transformation. When in doubt, insert a dummy surface that is not a coordinate break at the location of interest to check the global orientation.

The reference surface will be incremented or decremented automatically as surfaces above the reference surface are inserted or deleted. Any surface may be chosen as the reference surface, unless the object is at infinity, then 0 may not be the reference surface. The reference surface is also used to define the point of overlap for multiple zoom positions on the 3D layout plot.

Method To Compute F/#

By default, ZEMAX computes the paraxial and working F/# of a system using ray tracing. For details, see "Paraxial working F/#" on page 58 and "Working F/#" on page 62. For systems with very large F/#s, the ray tracing method may be inaccurate, because the very small angles between the marginal and chief rays leads to round-off error, and even small amounts of aberrations such as spherical aberration can significantly affect the calculation of F/#. For these reasons, ZEMAX will "cap" the maximum F/# to a value of 10,000. It is recommended that rays be used to compute the F/# for all systems whose F/# is less than 10,000. The preferred method of modeling systems with very large F/#s is to use afocal mode. See "Afocal Image Space" on page 103

Optionally, the F/# can be computed using the exit pupil distance divided by the exit pupil diameter. There are some important points to consider when using this option. First, ZEMAX makes no attempt to scale the F/# with either pupil orientation or field angle. Second, the exit pupil diameter and position used are based upon the paraxial values for these parameters (see "Exit pupil diameter" on page 53). In the presence of significant pupil aberration, the paraxial values may not agree with the real, aberrated pupil sizes. The best way to avoid this possibility is to place the stop at the end of the optical system, prior to the image surface, and use ray aiming (see "Ray Aiming" on page 106) and float by stop size for the aperture type (see "Aperture Type" on page 101). This option is also useful when the on-axis field cannot be traced and the axial F/# is different from the actual F/# at the fields being traced.

Fast Semi-Diameters

ZEMAX computes "automatic" semi-diameters to estimate the clear aperture required on each surface to pass all rays at all field points and wavelengths. For axial systems, this computation can be done accurately by tracing just two rays, the top and bottom marginal rays, for each field and wavelength. By default, for axial systems ZEMAX will trace just these two rays in the "true" tangential plane (for a definition of this term see "Sagittal and Tangential" on page 59) of the vignetted pupil for each field and wavelength, then use the radial coordinates of each ray at each surface to determine the semi-diameter required.

For some non-axial systems, the resulting estimate is not accurate enough. This typically includes systems which have tight edge and clear aperture constraints. For these systems, there is no general way of accurately computing the semi-diameters other than by tracing a large number of rays in the vignetted pupil. ZEMAX will iteratively trace as many *marginal* rays around the perimeter of the pupil as required to determine every semi-diameter to an accuracy of about 0.01% (5 significant figures).

Iteration can be slow, since ZEMAX needs to update the semi-diameters frequently, especially during optimization. There is a trade-off between speed and accuracy. For axial systems, ZEMAX will use the slower iterative algorithm only if the "Fast Semi-Diameters" option is checked "off"

For non-axial systems, ZEMAX will automatically use the slower iterative algorithm. If the "Fast Semi-Diameters" option is checked "on", ZEMAX will trace only as many marginal rays as required to estimate the automatic semi-diameters to about 0.01%. The algorithm starts by tracing 2, then 4, then 8 rays, and so on, until the semi-diameter values converge to within 0.01% of the previous estimate. If the "Fast Semi-Diameters" option is checked "off", ZEMAX will trace at least 32 rays around the vignetted pupil at each field and wavelength, and more if required to estimate the automatic semi-diameters to about 0.01%. Tracing at least 32 rays is safer, but slower, and should only be used when very high precision is required in the automatic semi-diameter calculations.

Both methods used by ZEMAX to determine semi-diameters are accurate only for surfaces whose maximum size is set by the radial coordinate of a marginal ray. The methods do not work for surfaces that lie in a caustic. For surfaces in a caustic, the exact ray-based semi-diameter could be determined using a spot diagram or footprint analysis, however this level of precision is likely never required since diffraction effects, which the rays do not model, would be significant.

Check GRIN Apertures

If checked, this setting instructs ZEMAX to check all gradient index ray traces for surface aperture vignetting. Each GRIN trace within the media is checked to see if the ray ever passes outside the front surface aperture boundary, and if it does, then the ray is vignetted. If this setting is not checked, then the ray may travel outside the boundary defined on the front surface, as long as the ray passes the aperture at the surface.

Turn Off Threading

If checked, ZEMAX will not split calculations into multiple threads of execution. Threading allows computers with multiple CPU's or multi-core CPU's to calculate results more quickly. The only reason to turn off threading is if insufficient memory exists to break calculations into separate threads. This is primarily a debugging and diagnostic tool and should normally be left unchecked.

Don't Print Coordinate Break Data

If checked, selected data will not be printed for coordinate break surfaces. Use of this option can shorten and clarify some text listings, particularly in systems with many coordinate break surfaces. Supported features include ray trace listings, index data, and global coordinate data. Support for this option in other ZEMAX features may be added upon request.

OPD Modulo 2π

If checked, all OPD data will be computed as the fractional part of the total OPD. All OPD computations will return results that are between $-\pi$ and $+\pi$, or -0.5 and $+0.5$ waves. This method of computing OPD may have applications to some specific problems. However, it's use is discouraged unless the user fully understands the meaning and ramifications of computing OPD in this manner.

Non-Sequential

The controls on this tab define how rays trace within a NSC group.

Maximum Intersections Per Ray

This control defines how many times a single ray may intersect objects along any single path from the original source parent ray to the final object intersection. When ray splitting is used, this parameter controls the maximum number of generations of child rays split off from the parent ray.

In some systems, such as a source placed inside a reflective sphere with perfect reflectivity and no bulk absorption, the rays will bounce around until this limit is reached, and the energy of the ray finally will be discarded. These systems are of course not physically possible, since no reflection or propagation through anything but a vacuum is perfectly lossless.

Maximum Segments Per Ray

This is the maximum number of segments per ray launched, not the total number of segments that ZEMAX can keep track of. A segment is that portion of a ray path from one intersection to the next. When a ray is launched from the source, it travels to the first object. That is 1 segment. If the ray then splits into 2 rays, each of those are another segment (for a total of 3). If each of those rays split again, there will be 7 segments. Generally, if ray splitting is being used, the number of segments grows far faster, and needs to be set much larger, than the number of ray-object intersections does.

It is critical to understand that each ray launched from a source can have up to the maximum number of segments. If the maximum number of segments is 1,000, and 5,000 individual rays are launched, ZEMAX can store up to 5,000,000 segments. The total RAM requirement is about 140 bytes times the maximum number of segments. Setting the limit at 100,000 segments would require 14 Mb of RAM per ray. For this reason, do not arbitrarily (and needlessly) set the maximum to very large numbers; there is a limit of 2,000,000 ZEMAX will allow. Note that with even a modest number of segments, ZEMAX can still trace millions of rays... the maximum number of segments is only a limit on how many segments are allowed in one single ray. Most ZEMAX features only use

one ray at a time, so the total RAM requirement is driven by the maximum number of segments rather than the number of rays traced. The exception is the Shaded Model display, which must be able to allocate enough RAM to hold all segments for all source rays drawn.

Maximum Nested/Touching Objects

This defines an upper limit on how many objects can be inside, straddled, or in direct contact with each other. For example, if object 3 is inside of 2 which is inside of object 1, the maximum number of nested objects is 3. There may be any number of groups of objects each nested 3 deep in this case. The limit applies to the total nesting in any one collection of objects, however, there may be any number of such collections within the NSC group. If several objects share a common boundary, such as multiple volumes touching at one face, then the maximum nested objects must be set at least as high as the total number of objects that share a common point in space.

Minimum Relative Ray Intensity

As each ray splits, the energy decreases. The relative ray intensity is a lower limit on how much energy the ray can carry and still be traced. This parameter is a fraction, such as 0.001, relative to the starting ray intensity from the source. Once a child ray falls below this relative energy, the ray is terminated.

Minimum Absolute Ray Intensity

This parameter is very similar to the minimum relative ray intensity, except it is absolute in source units rather than relative to the starting intensity. If this is zero, the absolute ray intensity threshold is ignored. The initial intensity of each ray is always given by the source intensity divided by the total number of analysis rays for that source. The number of layout rays is not used to determine the initial intensity of rays, even for those rays drawn on layout plots.

Glue Distance In Lens Units

When two NSC objects are placed in contact, such as a lens touching one face of a prism, numerical roundoff will cause the ray tracing algorithm to sometimes detect a very tiny distance between the two objects. This can also occur when objects are rotated in 3D space and placed close, but not exactly, next to one another because of the finite number of digits entered in the spreadsheet editor.

The glue distance is the distance below which the objects are considered in contact. It is important that objects not be separated by distances very close to the glue distance. If it is intended that two objects be in contact, then the maximum spacing between the objects should be several times smaller than the glue distance. If it is intended that the objects be separated, then the distance between the object should be several times larger than the glue distance. Object spacings very close to the glue distance will yield inconsistent ray tracing or geometry errors. This should be avoided by adjusting either the object spacing or the glue distance.

The glue distance also determines the minimum propagation length for ray tracing. If a ray-object intersection is less than the glue distance away from the previous intercept, the intercept is ignored.

In the majority of cases, no adjustment should be made to the glue distance parameter. The glue distance must be no smaller than 1E-10 and no larger than 1.0.

Missed Ray Draw Distance in Lens Units

This parameter is the distance to draw the ray segments that miss all objects. ZEMAX will draw a short ray segment to indicate the direction the ray was traveling in. This parameter also controls how large a source indicator arrow to draw. If zero, ZEMAX will select a default value for this parameter when drawing miss rays and some sources.

Maximum Source File Rays In Memory

This parameter sets the maximum number of rays for each Source File object that will be held in memory. The recommended value is 1,000,000 rays. For a complete description of how this setting affects memory usage and ray randomization see "Source File" on page 546.

Simple Ray Splitting

When a ray strikes a refractive boundary, generally some energy will reflect and some will refract. If Simple Ray Splitting is off, then both the reflected and refracted rays are traced. Each ray gets the fraction of energy corresponding to the reflection and transmission coefficients of the interface. If Simple Ray Splitting is on, then

either the reflected or the refracted ray is traced, but not both. The decision to trace the reflected or the refracted ray is random; with the reflection and transmission coefficients being interpreted as a relative probability of taking that path. Whichever path is selected, the reflected or refracted ray gets all of the energy that would have propagated down both paths. The advantage to using Simple Ray Splitting is fewer rays are traced, so computations are faster. The disadvantage is the rays traced contain less detailed information. For more information see "Ray splitting" on page 579.

Retrace Source Rays Upon File Open

If checked, then NSC source rays will be retraced when the file is opened. This allows refreshing of detector windows automatically.

Fields

The field dialog box allows specification of field points. The available field types and the associated type codes (used by ZPL macros) are defined in the following table.

FIELD TYPES AND CODES

Field type and code	Description
Angle, 0	See "Field angles and heights" on page 53.
Object Height, 1	
Paraxial Image Height, 2	
Real Image Height, 3	

The buttons can be used to activate or deactivate field positions, as well as to sort, save and load the data. For information about field conventions, see "Field angles and heights" on page 53.

The Field Normalization control determines if fields are normalized by radial or rectangular coordinates. For more information, see "Normalized field coordinates" on page 55.

Vignetting factors

ZEMAX also provides data fields for defining vignetting factors. The vignetting factors are VDX, VDY, VCX, VCY, and VAN. These vignetting factors should be left at zero if there is no vignetting in the system. The factors are described in the Chapter "Conventions and Definitions" under "Vignetting factors" on page 61.

There is also a button labeled "Set Vig" on the fields dialog box. Clicking on this button will recompute the vignetting factors for each field based upon the current lens data. The set vignetting algorithm estimates the vignetting decenter and compression factors so that the four marginal rays in the top, bottom, left, and right edges of the pupil pass within the apertures of each surface. Only the primary wavelength is used. To clear the vignetting factors back to the default values of zero, click on "Clr Vig". The set vignetting algorithm may not work properly if there are central or non-circular obscurations or apertures in the system. In this case, the vignetting factors must be user defined.

The algorithm starts by launching a grid of rays through the pupil. At each surface with a surface aperture, the ray is tested to see if it passes within the specified aperture. All rays which pass all surfaces are then used to compute the unvignetted pupil centroid. The edge of the unvignetted pupil is then computed using an iterative method accurate to about 0.001%.

The algorithm may not work in all cases. For systems where the set vignetting algorithm fails, the vignetting factors will need to be adjusted manually. The accuracy of the set vignetting algorithm can be tested by tracing a few marginal rays.



Identical or nearly identical field coordinates may not be defined with different vignetting factors; see the Chapter Conventions and Definitions for details.

Saving and loading field data

The Save and Load buttons on the field dialog box are used to store and retrieve the field data independently from the lens data. The file format is text that can be edited or created outside of ZEMAX.

Wavelengths

The wavelengths dialog box is used to set wavelengths, weights, and the primary wavelength. The buttons can be used to activate and deactivate wavelengths, and to sort, save, and load the data. A list of commonly used wavelengths is also included. To use the entries on the list, select the desired wavelength set, and click on the "Select" button.

For more information, see the "Wavelength data" on page 62.

Wavelength data are always measured in micrometers referenced to "air" at the current system temperature and pressure. The default system temperature is 20 degrees Celsius, and the default air pressure is 1.0 atmospheres. If the system temperature and/or pressure is modified, or under the control of multi-configuration operands, care must be taken to adjust the wavelengths to the new air temperature and pressure.

Saving and loading wavelength data

The Save and Load buttons on the wavelength data dialog box are used to store and retrieve the wavelength data independently from the lens data. The file format is text that can be edited or created outside of ZEMAX.

Next Configuration

This menu option provides a quick way to change all graphics to reflect the next configuration (or zoom position). All spreadsheet, text, and graphics data will be updated.

Last Configuration

This menu option provides a quick way to change all graphics to reflect the last configuration (or zoom position). All spreadsheet, text, and graphics data will be updated.

Introduction

This chapter provides detailed descriptions of each of the analysis features ZEMAX supports. Analysis in this context means any graphical or text data computed from data defining the lens. This includes aberrations, MTF, spot diagrams, and many other computations. Program features which modify the lens data or which manipulate other data (such as glass catalog data) are described in "TOOLS MENU" on page 239.

Selecting a menu option will immediately perform the requested calculation. Once the graph or text window is displayed, you may select the Settings menu option to modify the default settings for that window. Once you have made the appropriate changes, click on "OK" and the program will recalculate and redisplay the data presented in the window. If you prefer to change the settings before the graphic or text data is displayed, use the "Show Options First" checkbox on the Graphics tab of the File, Preferences dialog box.

For a description of the OK, Cancel, Save, Load, Reset, and Help buttons present on most of the "Settings" dialog windows, see "USER INTERFACE" on page 41.

Each analysis window has an "Update" menu item. The update function forces ZEMAX to recompute and redisplay the data presented in the window. This is useful if the lens data has changed and the graph now displays obsolete data. Double clicking within the window has the same effect as selecting Update. Clicking with the right mouse button is equivalent to clicking on "Settings". For more information, see "USER INTERFACE" on page 41.



All analysis computations are done on a copy of the current lens data, and any changes made to the lens data for purposes of the computation are made only on the analysis copy, and not on the original lens data. The changes are discarded once the computation is completed.

Layout

2D Layout

Purpose:

Layout diagram. This is a YZ cross section through the lens.

Settings:

Item	Description
First Surface	The first surface to be drawn.
Last Surface	The last surface to be drawn.
Wavelength	Either any one or all wavelengths may be shown.
Field	Either any one or all field positions may be shown.
Number of Rays	The number of rays specifies the number of tangential rays to be drawn for each defined field. The rays will be evenly distributed along the fan of the pupil, unless apodization has been specified. This parameter may be set to zero.
Scale Factor	If the scale factor is set to zero, then "Fill Frame" will be selected, which will scale the range of surfaces drawn to fill the graphic page. If a numeric value is entered, then the plot will be drawn in "real" scale, times the scale factor. For example, a scale factor of 1.0 will plot the lens actual size on the printer (not the display). A factor of 0.5 will plot the lens at half scale.
Y Stretch	The relative exaggeration of the Y direction. This control is useful for drawing systems with very large aspect ratios. Ignored if negative, zero, or unity.
Upper Pupil Limit	The maximum pupil coordinate to draw rays to.

Item	Description
Lower Pupil Limit	The minimum pupil coordinate to draw rays to.
Marginal and Chief Only	Draws only the marginal and chief rays, overriding the other ray settings.
Color Rays By	Select "Field #" to use color to distinguish between each field position, or "Wave #" to distinguish between each wavelength, "Config #" to distinguish between configurations, and "Wavelength" to approximate the color of wavelengths in the visible spectrum.
Suppress Frame	Suppresses drawing of the frame on the bottom of the window, which leaves more room for the layout plot itself. No scale bar, address block, or other data will be displayed.
Delete Vignetted	If checked, rays are not drawn if they will be vignetted by any surface.
Fletch Rays	If checked, small arrows are drawn on each ray to indicate the direction of propagation.

Discussion:

This feature is not available if you use coordinate breaks, spider obscurations, obscuration decenters, X-angles, holograms, or other attributes which spoil the rotational symmetry of the lens. Use the 3D layout instead.

If rays miss a surface, then the rays will not be drawn to the surface where the error occurred. If the ray is total internal reflected, then the ray will be drawn up to but not past the surface where the error occurred. Ray failures can be evaluated in detail by using "Ray Trace" on page 205.

3D Layout

Purpose:

Draws 3D layout plots of the lens system. The algorithm draws a wireframe style representation of the lens.

Settings:

Item	Description
First Surface	The first surface to be drawn.
Last Surface	The last surface to be drawn.
Wavelength	Either any one or all wavelengths may be shown.
Field	Either any one or all field positions may be shown.
Number of Rays	The number of rays specifies the number of rays to be drawn for each selected field and wavelength. The rays will be evenly distributed along the fan of the pupil, or around the perimeter if Ring is the selected Ray Pattern, or unless apodization has been specified. This parameter may be set to zero. It is ignored if the Ray Pattern is set to List.
Ray Pattern	Choose XY Fan, X Fan, Y Fan, Ring, List, Random, or Grid to indicate what the pattern of rays to be traced should be. The List option indicates that the rays to be traced are user defined and listed in a file, see the discussion below for information on the ray list format. If List is selected the Number of Rays setting is ignored.
Scale Factor	If the scale factor is set to zero, then "Fill Frame" will be selected, which will scale the range of surfaces drawn to fill the graphic page. If a numeric value is entered, then the plot will be drawn in "real" scale, times the scale factor. For example, a scale factor of 1.0 will plot the lens actual size on the printer (not the display). A factor of 0.5 will plot the lens at half scale.
Hide Lens Faces	If checked, this option will suppress drawing of the lens faces, and only the lens edges will be drawn. This is useful because some complicated systems look cluttered with the faces drawn.

Item	Description
Hide Lens Edges	If checked, this option will suppress drawing of the outer aperture of the lens. This is useful for giving the 3D layout a 2D "cross section" appearance.
Hide X Bars	If checked, this option will suppress drawing of the X portions of the lens faces. This option is useful when "Hide Lens Edges" is checked and "Hide Lens Faces" is not checked.
Rotation About X	The angle in degrees by which the lens appears to be rotated about the X axis.
Rotation About Y	The angle in degrees by which the lens appears to be rotated about the Y axis.
Rotation About Z	The angle in degrees by which the lens appears to be rotated about the Z axis.
Color Rays By	For sequential rays, "Field #" will use color to distinguish between each field position, "Wave #" to distinguish between each wavelength, "Config #" to distinguish between configurations, and "Wavelength" to approximate the color of wavelengths in the visible spectrum. For non-sequential rays "Field #" will use color to distinguish between each source number, unless a source has a user defined color assigned to that source.
Suppress Frame	Suppresses drawing of the frame on the bottom of the window, which leaves more room for the layout plot itself. No scale bar, address block, or other data will be displayed.
Delete Vignetted	If checked, rays are not drawn if they will be vignetted by any surface.
Configuration	Select "All" to show all configurations, "Current" to show just the active configuration, or select any combination of other configurations.
Offset X, Y, Z	The X, Y, and Z direction offset between configurations in lens units. Only has an affect on the drawing if "All" configurations are being drawn.
Fletch Rays	If checked, small arrows are drawn on each ray to indicate the direction of propagation.
Split NSC Rays	If checked, rays from NSC sources will be statistically split at ray-surface intercepts. Rays entering from the entry port are not affected by this setting.
Scatter NSC Rays	If checked, rays from NSC sources will be statistically scattered at ray-surface intercepts. Rays entering from the entry port are not affected by this setting.

Discussion:

Pressing the left, right, up, down, Page Up, or Page Down keys will rotate the displayed image for a different perspective.

The orientation indicator

To create the layout, ZEMAX projects the 3D coordinates of the rotated optical model onto the 2D plane of the plot. Conceptually, the optical model is first rendered in an unrotated 3D space, with +z oriented to the right, +y up, and +x away from the viewer into the page. The coordinates and orientations of all surfaces and objects are defined by the Global Coordinate Reference Surface (see "Global Coordinate Reference Surface" on page 111). This 3D model is then rotated by the Rotation About X, Y, and Z values specified in the settings (see also "Graphics" on page 69 for information about the rotation order). The resulting 3D model is then projected onto a 2D plane by plotting the z coordinate along the horizontal direction and the y coordinate along the vertical direction; the x coordinate is ignored.

To help visually define the orientation of the model, an orientation indicator is shown in the bottom left hand corner of the drawing. This indicator consists of 3 lines extending along the directions of the +x, +y, and +z axes of the global coordinate reference surface. These 3 lines rotate as the drawing is rotated. If the coordinates of the 3D orientation indicator lines lie in the plane of the drawing, or would extend out (toward the viewer), they are drawn in black. If the orientation indicator lines would extend in (away from the viewer) then the lines are drawn in grey.

Ray errors

If rays miss a surface, then the rays will not be drawn to the surface where the error occurred. If the ray is total internal reflected, then the ray will be drawn up to but not past the surface where the error occurred. Ray failures can be evaluated in detail by using "Ray Trace" on page 205. These comments only apply to rays from the sequential object surface, and not to rays from NSC sources.

Configuration data

When drawing all configurations, an offset may be added to each configuration in the x, y, and z directions independently. The offsets may all be zero if desired. If the offsets are zero, then all the configurations are superimposed; otherwise, the configurations are all displaced from one another by the specified amount. Note that all offsets are defined from the global coordinate reference surface position. The global coordinate reference surface is defined on the Miscellaneous tab of the System, General dialog box. If all offsets are zero, the multiple configurations are all overlapped at the global coordinate reference surface.

Raylist file format

If List is chosen for the ray pattern, the rays to be traced are defined in a file. The file must be called RAYLIST.TXT and be placed in the <data>\Miscellaneous folder (see "Folders" on page 68). The file format is text, with two distinct methods for defining the rays supported, implicit and explicit. The implicit format file consists of two numbers on each line, one for the px and one for the py normalized pupil coordinates. The specified rays are traced at each defined field and wavelength selected.

Example: Four marginal rays are defined by:

```
0.0 -1.0
0.0 1.0
-1.0 0.0
1.0 0.0
```

The explicit format file consists of the word EXPLICIT followed by the values x, y, z, l, m, n, and wavenumber; where x, y, and z are the ray starting coordinates, l, m, and n are the direction cosines, and wavenumber is an integer indicating the wavelength to use. All coordinates are in object space. If the object thickness is infinity, then the spatial coordinates are relative to surface 1. If the object is not at infinity, then the coordinates are relative to surface 0. In both cases the ray itself is in the object space medium, prior to refraction into surface 1. If explicit format is used, then the field and wavelength settings are ignored, and only those rays listed in the file are traced.

Example: Three rays at wavelengths 1, 2, and 3 along the Y axis parallel to the Z axis are defined as follows:

```
EXPLICIT
0.0 -5.0 0.0 0.0 0.0 1.0 1
0.0 +0.0 0.0 0.0 0.0 1.0 2
0.0 +5.0 0.0 0.0 0.0 1.0 3
```

Shaded Model

Purpose:

Draws a shaded solid representation of the lens using OpenGL graphics.

Settings:

Item	Description
First Surface	The first surface to be drawn.
Last Surface	The last surface to be drawn.
Wavelength	Either any one or all wavelengths may be shown.
Field	Either any one or all field positions may be shown.

Item	Description
Number of Rays	The number of rays specifies the number of rays to be drawn for each selected field and wavelength. The rays will be evenly distributed along the fan of the pupil, or around the perimeter if Ring is the selected Ray Pattern, or unless apodization has been specified. This parameter may be set to zero. It is ignored if the Ray Pattern is set to List.
Draw Section	Select "Full" to draw each lens element completely. The 3/4, 1/2, and 1/4 options draw just that much of the element, yielding a cut away perspective of the lens interior. This only affects sequential surface data, and not NSC objects.
Ray Pattern	Choose XY Fan, X Fan, Y Fan, Ring, List, Random, or Grid to indicate what the pattern of rays to be traced should be. The List option indicates that the rays to be traced are user defined and listed in a file, see the discussion below for information on the ray list format. If List is selected the Number of Rays setting is ignored.
Angular Segments	The number of angular segments used to approximate the lens shapes. Larger numbers require more processing time.
Radial Segments	The number of radial segments used to approximate the lens shapes. Larger numbers require more processing time.
Rotation About X	The angle in degrees by which the lens appears to be rotated about the X axis.
Rotation About Y	The angle in degrees by which the lens appears to be rotated about the Y axis.
Rotation About Z	The angle in degrees by which the lens appears to be rotated about the Z axis.
Color Rays By	For sequential rays, "Field #" will use color to distinguish between each field position, "Wave #" to distinguish between each wavelength, "Config #" to distinguish between configurations, and "Wavelength" to approximate the color of wavelengths in the visible spectrum. For non-sequential rays "Field #" will use color to distinguish between each source number, unless a source has a user defined color assigned to that source.
Background	Selects the background color.
Opacity	The options are ignore, which ignores the opacity values and draws all surfaces and objects as opaque; consider, which draws partially transparent surfaces and objects; or All 50%, which overrides the defined opacity and uses 50% for all objects. For information on setting the opacity of surfaces and objects see "Surface Opacity" on page 77 and "The object properties dialog box" on page 556.
Brightness	This control varies the overall brightness of the image.
Delete Vignetted	If checked, rays are not drawn if they will be vignetted by any surface.
Configuration	Select "All" to show all configurations, "Current" to show just the active configuration, or select any combination of other configurations.
Offset X, Y, Z	The X, Y, and Z direction offset between configurations in lens units. Only has an affect on the drawing if "All" configurations are being drawn.
Fletch Rays	If checked, small arrows are drawn on each ray to indicate the direction of propagation.
Split NSC Rays	If checked, rays from NSC sources will be statistically split at ray-surface intercepts. Rays entering from the entry port are not affected by this setting.
Scatter NSC Rays	If checked, rays from NSC sources will be statistically scattered at ray-surface intercepts. Rays entering from the entry port are not affected by this setting.

Discussion:

When rendering NSC detector objects with more than 500,000 pixels will not show individual data for each pixel to keep the total amount of required RAM reasonable.

ZEMAX Element Drawing

Purpose:

This feature creates a mechanical drawing of surface, singlet, cemented doublet, or cemented triplet elements suitable for use in optical shop fabrication.

Settings:

Item	Description
Surface	The first surface of the element to be drawn.
Show As	Select either "Surface", "Singlet", "Doublet", or "Triplet".
Note File Name	The name of the file which contains the notes to be appended to the notes section of the element drawing. Notes should always start at number 2, since number 1 is reserved for the units specification.
Edit Note File	Clicking on this button will invoke the Windows NOTEPAD.EXE editor, which can then be used to modify the selected note file.
Radius	The radius tolerance value for the specified surface.
Irregularity	The irregularity tolerance value for the specified surface.
Clear Aper	The clear aperture diameter of the specified surface.
Thickness	The center thickness tolerance of the specified surface.
Scale Factor	If the scale factor is set to zero, then "Fill Frame" will be selected, which will scale the element to fill the right half of the element drawing. If a numeric value is entered, then the plot will be drawn in "real" scale, times the scale factor. For example, a scale factor of 1.0 will plot the element at actual size on the printer (not the display). A factor of 0.5 will plot the element at half scale.
Decimals	The number of decimals to use in numeric values. Press the "Reset all but surface, show as, and titles" button to reformat all displayed numbers.
Drawing Title	This field is for any user defined text. The default is the lens title.
Date	If left blank, the default date format is used. If any text is supplied, the exact text is used.
Drawing Name	All of these fields are for user defined text. Any text may be entered. No default value.
Approved	
Revision	
Drawn By	
Project	
Note Font Size	Choose Standard, Medium, Small, or Fine. These are in order of decreasing font size. The Note Font Size setting only affects the size of the note file that is annotated on the drawing. Smaller fonts permit larger note files to be displayed.
Save As/Load From	These two buttons allow saving and reloading of all settings displayed on the element drawing settings box. The data is stored in a binary format with the extension ELE.
Reset all but...	If selected, this button will reset all the default tolerances and apertures for the specified surfaces, but the current surface, show as, and text titles will remain as they are.

Discussion:

The element drawing settings may be stored for the specific lens file by pressing the Save button. Unlike most analysis features, the element drawing feature saves all the settings for each surface separately. For example,

the notes and tolerances for surface 1 may be saved, and then new notes and tolerances for surface 3 may be entered and then saved. To recall the settings for any specific surface, change the surface number to the desired surface, and then press the Load button. If a match is found with a previously saved surface, the settings for that surface will be displayed. This feature makes it easy to regenerate complex drawings for multiple element systems.

An important feature of the element drawing capability is the ability to load different note files and place them on the drawing. The default note file "DEFAULT.NOT" is a generic set of notes which will rarely be useful as is. However, the user can modify the note files (they are text files which any word processor or text editor can modify) and store them under different names. For example, you may want to have a .NOT file for each type of optic you design, and then load the most appropriate note file when the element drawing is generated. The .NOT files must be placed in the <data>\Miscellaneous folder (see "Folders" on page 68).

If the first line of the note files starts with the number 1, then ZEMAX will print the note file as provided, with no default notes. If the note file starts with any other symbol, the first default note will be "1) All dimensions in millimeters" or whatever the current lens units are. To use the default note, start the user supplied notes at note number 2. The line breaks and spacings in the note file will be replicated exactly on the element drawing.

Whenever a new element drawing is generated, or the "Reset" button is pressed, the default settings will be regenerated. The default tolerances are taken from the tolerance data editor. The maximum of the min/max tolerance range is used as the default. For example, if the TTHI thickness tolerance is -.03, +.05, the tolerance value will be 0.05. Only TTHI, TRAD, and TIRR tolerances are considered. If a suitable default cannot be generated, the tolerance is set to zero. Note all tolerance fields are text; and may be edited to suit any requirement.

A handy conversion between radius tolerance and the power tolerance in fringes for a Newton's rings type optical test against a test plate is given by

$$\#fringes = \frac{\Delta R \rho^2}{\lambda R^2},$$

where ΔR is the radius error, λ is the test wavelength, ρ is the radial aperture, and R is the radius of curvature. This formula is an approximation for shallow curvatures. For more information, see Malacara, Optical Shop Testing, J. Wiley & Sons, Inc.

Special characters

There are useful special characters that may be inserted and edited in the text fields of the tolerances. The symbols are inserted by pressing and holding Alt on the keyboard and typing a 4 digit number on the numeric keypad. Some common characters are ± - Alt 0177, μ - Alt 0181, ® - Alt 0174, © - Alt 0169.

ISO Element Drawing

Purpose:

This feature creates an ISO 10110 type drawing of surface, singlet, or doublet elements suitable for use in optical shop fabrication.

Settings:

Item	Description
First Surface	The first surface of the element to be drawn.
Show As	Choose surface, singlet, or doublet.
Other settings	There are many other settings available for this feature that correspond to the tolerances defined in the ISO 10110 specification. See the discussion for details.

Discussion:

The ISO 10110 Element Drawing is an interpretation of the drawing specification "ISO 10110 Optics and Optical Instruments -- Preparation of drawings for optical elements and systems: A User's Guide", by Ronald K. Kimmel and Robert E. Parks, eds., published by the Optical Society of America. For more information see OSA's web site at www.osa.org.

Although the ISO specification only covers the case of single elements, the ZEMAX ISO drawing also supports surfaces and doublets. However, not all the ISO drawing particulars are incorporated on a drawing of a cemented doublet. If the doublet drawing results are not satisfactory, draw the doublet as two separate singlet drawings instead. This feature does not consider the "Ignore Surface" setting and will draw all requested surfaces, whether they are ignored or not.

Summary of ISO 10110 symbols and codes

ZEMAX does not automatically include default values for all ISO tolerances; however text fields are provided for user specification of these tolerances. The text fields have limited width, and if very long strings are placed in the text fields the text may overrun the columns on the drawing. The following table summarizes the ISO 10110 symbols and codes used by ZEMAX.

Summary of ISO 10110 symbols and drawing codes

Symbol or Code Number	Description
R	Radius of curvature in lens units. CC indicates concave, CX indicates convex.
Ø	Diameter in lens units.
Øe	Effective diameter in lens units.
n	refractive index
v	Abbe number
(λ)	Coating specification
P; Rq; Lmin, Lmax	Indicates a polished surface on the lens drawing. The details of the polish specification are not provided on the drawing, but are specified in the surface data. Rq = maximum permissible rms surface roughness in micrometers, Lmin/Lmax = minimum and maximum length of sampling interval for rms calculation in mm.
0/A	Stress Birefringence; A = maximum optical path difference in nm/cm.
1/NxA	Bubbles & Inclusions; N = number of bubbles and/or inclusions, A = bubble grade number (size)
2/A; B	Inhomogeneity and Striae; A = homogeneity class number, B = striae class number
3/A(B/C) or 3/A(B/C) RMSx < D or 3/- RMSx < D where x is either t, i, or a	Surface Form Tolerance; A = maximum sagitta error in fringes or - if tolerance is part of radius tolerance, B = peak-to-valley irregularity in fringes or - if no tolerance is given, C = non-spherical rotationally symmetric error in fringes. If (B/C) is replaced by A(B), no tolerance is given. For RMS tolerances, D = maximum rms error in fringes, t = total rms deviation from nominal, i = rms irregularity, a = rms asymmetry remaining after asphericity is subtracted from irregularity.
4/S(L)	Centering Tolerance; S = surface tilt angle in minutes of arc, L = lateral displacement

Symbol or Code Number	Description
5/NxA; CN'xA'; LN" x A"; EA'''	Surface Imperfection Method 1; N = number of imperfections, A = grade number (square root of imperfection area), C = coating imperfection designation, N' = number of coating imperfections, A' = grade number, E = edge chips designation, A''' = chip protrusion from edge, L = long scratch designation, N" = number of scratches, A" = grade number (scratch width (mm))
5/TV or /TV; EA" 5/RV or RV; EA"	Surface Imperfection Method 2; T = transmissive test, V = visibility class number, R = reflective test, E,A" = edge chips, as in Method 1 above.
6/Hth;L;pdg;fp;nts x np or 6/Eth;L;nts	Laser Irradiation Damage Threshold; Hth = energy density threshold, L = laser wavelength (nm), pdg = pulse duration group, fp = pulse repetition rate (Hz), nts = number of test sites, np = number of pulses per test site, Eth = power density threshold in watt-cm ² .

Diameters and bevels

The semi-diameter value is used to determine the drawn radial aperture of each surface. A flat bevel is added to any surface whose semi-diameter is less than the largest semi-diameter being drawn. To place flat bevels on all surfaces, use the "Dia (flat) optional" setting available on the "L Surf Codes 3-4" tab. This optional value must exceed twice the semi-diameter of every surface drawn, otherwise the value is ignored.

Tolerance data

When the "Reset from TDE" button is pressed, radius, thickness, index, surface tilt, and surface decenter tolerances are taken from the tolerance data editor. The tolerances are automatically reset after changing the surface number or show as settings. Only TRAD, TCUR, TFRN, TTHI, TIND, TABB, TSTX, TSTY, TSDX, and TSDY tolerances are considered. If a tolerance is not provided, the tolerance is set to zero. The TFRN tolerance is applied to the surface form tolerance code 3/A. Note all tolerance fields may be edited to suit any requirement.

NSC 3D Layout

Purpose:

Draws 3D layout plots of the sources and objects in a single NSC group.

Settings:

Item	Description
Fletch Rays	If checked, small arrows are drawn on each ray to indicate the direction of propagation.
Split Rays	If checked, rays from NSC sources will be statistically split at ray-surface intercepts. Rays entering from the entry port are not affected by this setting.
Scatter Rays	If checked, rays from NSC sources will be statistically scattered at ray-surface intercepts. Rays entering from the entry port are not affected by this setting.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Suppress Frame	Suppresses drawing of the frame on the bottom of the window, which leaves more room for the layout plot itself. No scale bar, address block, or other data will be displayed.
Configuration	Select "All" to show all configurations, "Current" to show just the active configuration, or select any combination of other configurations.

Item	Description
Color Rays By	Select "Source #" to use color to distinguish rays traced from each source, unless a source has a user defined color assigned to that source, "Wave #" to distinguish between each wavelength, "Config #" to distinguish between configurations, "Wavelength" to approximate the color of wavelengths in the visible spectrum, or "Segment #" to distinguish between segment level number (number of object intersections since the ray left the source point).
Surface	Selects the NSC surface.
Scale Factor	If the scale factor is set to zero, then "Fill Frame" will be selected, which will scale the range of surfaces drawn to fill the graphic page. If a numeric value is entered, then the plot will be drawn in "real" scale, times the scale factor. For example, a scale factor of 1.0 will plot the lens actual size on the printer (not the display). A factor of 0.5 will plot the lens at half scale.
Rotation About X	The angle in degrees by which the lens appears to be rotated about the X axis.
Rotation About Y	The angle in degrees by which the lens appears to be rotated about the Y axis.
Rotation About Z	The angle in degrees by which the lens appears to be rotated about the Z axis.
Offset X, Y, Z	The X, Y, and Z direction offset between configurations in lens units. Only has an affect on the drawing if "All" configurations are being drawn.
Filter	If blank, all rays are drawn. Otherwise, only rays which meet the criteria defined by the filter string will be drawn. See "The filter string" on page 589 for information about the filter string syntax.
Ray Database	If "none" is selected, new rays will be traced and displayed. If a ZRD data base file is selected, then rays contained within the database will be displayed. In either case the filter, if any, is applied. Generally, reading rays from a large database is faster than retracing them. The other advantage to using a database of rays is that the rays are always the same, until the database is replaced. ZEMAX cannot tell if the selected database is for rays from the current lens being displayed; so care should be taken when selecting the ZRD file that it corresponds to the lens data being displayed. For more information on the ZRD file, see "Ray database (ZRD) files" on page 583.

Discussion:

The settings in the above table are very much like their counterparts on the 3D Layout feature. However, this feature only draws objects in a single NSC group, and only traces and draws rays from sources defined within the group.

To make surfaces partially transparent, see "Surface properties type tab" on page 77.

NSC Shaded Model

Purpose:

Draws a shaded solid model representation of the components in a non-sequential group.

Settings:

The options are very similar to those available for the NSC 3D Layout and Shaded Model features. One additional feature is the option to color detector objects by either the energy incident on the detector in the last analysis or by just the rays traced in the layout view. The detector shows false color or black and white displays, using either coherent intensity or any of the other options supported by the detector viewer. See "Detectors" on page 532 for details.

NSC Object Viewer

Purpose:

Draws a shaded solid model representation of a single NSC object.

Settings:

Item	Description
Object	Selects the object to draw.
Face	Selects the face of the object to highlight. All portions of the object that belong to the selected face will be shaded in red for identification. The highlighted face can be incremented or decremented while viewing the object using the keyboard command Shift+Page Up/Down, respectively.
Surface	Selects the NSC surface.
Brightness	This control varies the brightness of the image.
Opacity	The options are to ignore, which ignores the opacity values and draws all surfaces and objects as opaque; or consider, which draws partially transparent surfaces and objects. For information on setting the opacity of objects see "The object properties dialog box" on page 556.
Background	Selects the background color.
Rotation About X	The angle in degrees by which the lens appears to be rotated about the X axis.
Rotation About Y	The angle in degrees by which the lens appears to be rotated about the Y axis.
Rotation About Z	The angle in degrees by which the lens appears to be rotated about the Z axis.

Fans

Ray Aberration

Purpose:

Shows ray aberrations as a function of pupil coordinate.

Settings:

Item	Description
Plot Scale	Sets the maximum vertical scale for the plots. The maximum scale is in micrometers for ray fans, waves for OPD plots, or percent for entrance pupil aberration plots. This overrides the automatic selection of scale for the plots. Enter zero for automatic scaling.
Number of Rays	This is the number of rays traced on each side of the origin of the plot.
Wavelength	The wavelength number for which the calculation should be performed.
Field	The field number for which the calculation should be performed.
Tangential	Selects which aberration component to plot for the tangential fan. Since tangential fans are functions of the y pupil coordinate, the default is to plot the y component of the aberration.
Sagittal	Selects which aberration component to plot for the sagittal fan. Since sagittal fans are functions of the x pupil coordinate, the default is to plot the x component of the aberration.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.
Check Apertures	Specifies whether or not to check if rays pass all surface apertures. If this is selected, rays which do not pass surface apertures will not be drawn.

Item	Description
Vignetted Pupil	If checked, the pupil axis will be scaled to the unvignetted pupil, in which case the data will reflect the vignetting in the system. If unchecked, the pupil axis will be scaled to fit the vignetted pupil.
Surface	Selects the surface at which the fan is to be evaluated. This is useful for evaluating intermediate images. See "Evaluating results at intermediate surfaces" on page 128.

Discussion:

The tangential fans show either the x or the y component of the transverse ray aberration as a function of the y pupil coordinate of the ray. The default option is to plot the y component of the aberration. However, since transverse ray aberrations are vectors, this is an incomplete description of the aberration. When ZEMAX plots the y component, the plot is labeled EY, when plotting the x component of the aberration, the plot is labeled EX.

The vertical axis scale is given at the bottom of the graph. The data being plotted is the difference between the ray intercept coordinate and the chief ray intercept coordinate. The tangential fan is the plot of the difference between the ray x or y coordinate and the chief ray x or y coordinate at the primary wavelength, as a function of the y pupil coordinate. The sagittal plot is the difference between the ray x or y coordinate and the chief ray x or y coordinate as a function of the x pupil coordinate. The horizontal scale for each graph is the normalized entrance pupil coordinate, either PX or PY.

If "All" wavelengths are shown, then the plot is referenced to the primary wavelength chief ray. If monochromatic, then the chief ray for the selected wavelength is used as a reference. For this reason, the data for non-primary wavelengths will in general change when switching between monochromatic and polychromatic displays.

Because ray aberrations are vectors, with both an x and a y component, the ray aberration fan is an incomplete description of the aberrations, especially when the image surface is rotated or the system is otherwise non-rotationally symmetric. Also, the fans only indicate aberrations along two "slices" through the pupil, rather than over the entire entrance pupil. The primary purpose of the ray fan plot is to determine what aberrations are present in the system; it is not a complete description of the system performance, especially for systems without rotational symmetry.

Evaluating results at intermediate surfaces

ZEMAX can compute analysis results for surfaces other than the image surface by applying assumptions. These assumptions work in most cases, however, there may be cases where the methods described here for analyzing intermediate surfaces are not appropriate, including systems which require ray aiming. All of the changes made to the lens described are made on a copy of the original lens data made for analysis purposes so no changes are made to the original lens data.

If the field type is either real or paraxial image height, the field type is changed to angle or object height for infinite or finite conjugate systems, respectively. The angles and heights used correspond to the primary wavelength chief ray angles and heights as computed for the unaltered system.

If the selected surface precedes the stop surface, ZEMAX moves the stop surface to a (possibly virtual) dummy space prior to the existing surface 1. Unless the system aperture is object space numerical aperture or cone angle, the system aperture is changed to entrance pupil diameter, and the aperture value is set equal to the original paraxial entrance pupil diameter computed for the original stop position. Note this assumption might not be valid for systems that require ray aiming. If the selected surface follows the stop surface, no changes are made to the system aperture or stop definitions.

Surfaces which follow the selected surface are then deleted. The glass of the new image surface is set to be the same as the selected surface.

Most analysis features that compute results for focal mode systems make more sense if the rays are allowed to come to a focus after refraction from the desired surface. For example, the OPD plot is generally a useful diagnostic only when the OPD is measured at the (possibly virtual) focus of that surface. Other features which require a temporary image to be formed include PSF, MTF, and diffraction encircled energy. For these features requiring a temporary image, the new image surface is set to be a standard plane surface. A paraxial marginal ray height solve is placed on the selected surface thickness to place the image surface at paraxial focus for the selected surface. The analysis computation then proceeds at this newly created intermediate image surface. Note

the analysis occurs at the paraxial focus formed by the rays after refracting through the selected surface. This shift to paraxial focus is not performed if the intermediate system is in afocal mode. For more information on afocal mode see "Afocal Image Space" on page 103.

Some analysis features make more sense if the data is evaluated directly on the surface itself, without allowing the rays to focus to an image. These features include the various spot diagrams, footprint analysis, geometric encircled energy, line/edge spread function, and extended source encircled energy.

Optical Path

Purpose:

Shows optical path difference as a function of pupil coordinate.

Settings:

The options are identical to those for ray aberration fans, except the only option for "Tangential" and "Sagittal" is OPD, since OPD is a scalar quantity.

Discussion:

The vertical axis scale is given at the bottom of the graph. The data being plotted is the optical path difference, or OPD, which is the difference between the optical path length of the ray and the optical path length of the chief ray. Usually, the calculation is referenced back to the difference between the ray path lengths at the system exit pupil. The horizontal scale for each graph is the normalized entrance pupil coordinate.

If "All" wavelengths are shown, then the plot is referenced to the primary wavelength based reference sphere and chief ray. If monochromatic, then the reference sphere and chief ray for the selected wavelength is used as a reference. For this reason, the data for non-primary wavelengths will in general change when switching between monochromatic and polychromatic displays.

Pupil Aberration

Purpose:

Shows entrance pupil distortion as a function of pupil coordinate.

Settings:

The options are identical to those for ray aberration fans, except the only option for "Tangential" and "Sagittal" is pupil aberration, since pupil aberration is a scalar quantity. The "Surface" only allows "Image" because this data is always computed at the stop surface.

Discussion:

Entrance pupil aberration is defined as the difference between the real ray intercept on the stop surface and the on axis primary wavelength paraxial ray intercept as a percentage of the paraxial stop radius. If the maximum aberration exceeds a few percent, ray aiming (see "Ray Aiming" on page 106) may be required to get the rays in object space to correctly fill the stop surface. If ray aiming is switched on, the entrance pupil aberration will appear to be zero (or a very small residual value), because the distortion is accounted for by the ray tracing algorithm. This is used as a check that ray aiming is working correctly. The definition used here for pupil aberration is not intended to be complete or in agreement with other definitions. The sole purpose of this feature is to provide guidance as to whether or not ray aiming is required.

Spot Diagrams

Standard

Purpose:

Show spot diagrams.

Settings:

Item	Description
Pattern	The pattern may be either hexapolar, square, or dithered. These terms refer to the pattern of rays as they appear in a pupil plane. Defocus the lens significantly to see the pattern, if desired. Dithered spot diagrams are generated by pseudo-random rays which eliminate the symmetrical artifacts in the spot diagram typical of rectangular or hexapolar patterns. The pattern is distorted to give the correct distribution of rays if pupil apodization is specified. There is no "best" pattern to use, each shows a different character of the spot diagram.
Refer To	The spot diagrams by default are referenced to the real chief ray. The RMS and GEO spot radii listed at the bottom of the diagram (and defined in the discussion section) are calculated assuming the chief ray is the "zero aberration" point. However, this option allows selection of other reference points. The centroid is defined by the average position of the rays traced. The middle is defined so that the maximum ray errors are equal in the plus and minus x and y directions. The vertex is defined by the local coordinates 0,0 on the selected surface. If the system is in afocal mode, the chief ray reference will be used instead of the vertex reference.
Show Scale	Scale bar is the default. Selecting "Square" will draw a box, centered on the reference point, whose width is twice the distance from the reference point to the outermost ray. Selecting "Cross" will draw a cross through the reference point location. The "Circle" setting will draw a circle centered on the reference point.
Wavelength	The wavelength number for which the calculation should be performed.
Field	The field number for which the calculation should be performed.
Surface	Selects the surface at which the spot diagram is to be evaluated. This is useful for evaluating intermediate images or vignetting.
Plot Scale	Sets the maximum scale size of the display in micrometers. A setting of zero will generate a default scale.
Delta Focus	The delta focus option is only used if through-focus spot diagrams are selected. For focal systems, delta focus is the Z-axis spacing between the spot diagrams in micrometers. For afocal systems, delta focus is the optical power spacing between the spot diagrams in diopters. Five spot diagrams will be shown for each field position. The defocus will be -2, -1, 0, 1, and 2 times the delta focus provided.
Ray Density	The ray density specifies the number of hexapolar rings to be traced if a hexapolar or dithered pattern is selected, or the number of rays across the width and height if a rectangular pattern is selected. The more rays traced, the greater the accuracy of the RMS spot radius, although the computation time increases. There are 6 rays in the first hexapolar ring, 12 in the second, 18 in the third, and so on.
Use Symbols	If checked, this option will draw different symbols rather than dots for each wavelength. This helps distinguish the various wavelengths.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. See "Surface scattering settings" on page 84.
Airy Disk	If checked, an elliptical ring around each spot showing the size of the Airy ellipse will be drawn. The Airy disk radius is 1.22 times the wavelength (primary wavelength is used if polychromatic) multiplied by the F/# of the system; which in general depends upon field position, pupil orientation, and possibly vignetting of rays if the marginal rays cannot be traced.

Item	Description
Direction Cosines	If checked, the data presented will be the direction cosines of the rays rather than the spatial coordinates of the rays. The x direction data will be the x direction cosine of the ray, the y direction data will be the y direction cosine. The image coordinates will also be given as the reference point direction cosines. Direction cosines are dimensionless.
Configuration	Select "All" to show data for all configurations at once, or select the one configuration to show, or select "Current" to show the active configuration. If "All" is selected, and both focal and afocal mode configurations are defined, the Configuration setting will automatically be reset to "Current".
Color Rays By	Select "Field #" to use color to distinguish between each field position, or "Wave #" to distinguish between each wavelength, or "Config #" to distinguish between configurations, and "Wavelength" to approximate the color of wavelengths in the visible spectrum.

Discussion:

The ray density has a maximum value based upon the number of fields displayed, the number of wavelengths defined, and available memory. Through-focus spot diagrams will trace half of the maximum number of rays possible on standard spot diagrams.

The GEO spot radius listed on the plot for each field point is the distance from the reference point (which is either the chief ray at the primary wavelength, the centroid of all the rays traced, or the middle of the spot cluster) to the ray which is farthest away from the reference point. In other words, the GEO spot radius is the radius of the circle centered at the reference point which encloses all the rays.

The RMS spot radius is the root-mean-square radial size. The distance between each ray and the reference point is squared, and averaged over all the rays, and then the square root is taken. The RMS spot radius gives a rough idea of the spread of the rays, since it depends upon every ray. The GEO spot radius only gives information about the one ray which is farthest from the reference point.

For information on the X and Y RMS spot radii, see the "text" listing for the spot diagram.

The Airy disk radius is given by 1.22 times the wavelength (primary wavelength is used if polychromatic) times the F/# of the beam, which in general depends upon field position and pupil orientation. This is the radius to the first dark ring of the Airy disk for a circular, uniformly illuminated entrance pupil. The Airy disk may be optionally drawn to give an idea of the scale of the plot. For example, if all the rays are well within the Airy disk, then the system is often said to be "diffraction limited". If the RMS spot radius is significantly larger than the Airy disk radius, then the system is not diffraction limited. The threshold for diffraction limited performance depends upon which criterion is used. There is no absolute boundary at which the system becomes diffraction limited. The Airy disk shown is not an accurate representation of the diffraction dark ring shape or size if the system does not have uniform illumination or if vignetting is used to eliminate some of the rays. ZEMAX does not plot vignetted rays on spot diagrams, nor are they used in computing the RMS or GEO spot radii.

ZEMAX generates grids of rays based upon the wavelength weighting factors and the pupil apodization, if any. The wavelength with the largest weight uses the maximum grid size set by the "Ray Density" option. Wavelengths with lower weights use grids with fewer rays to maintain the correct representation in the diagram. Ray grids are also distorted to maintain the correct ray distribution, if apodization is specified. The RMS spot radius stated on the spot diagram considers the wavelength weighting and apodization factors. However, it is only an estimate of the RMS spot radius based on the rays actually traced. It is not a very accurate estimate for some systems.

The image surface intercept coordinates of the reference point are shown underneath each spot diagram. If a surface other than the image surface is specified, then the coordinates are the intercept coordinates of the reference point on that surface. Since the reference point may be selected to be the centroid, this provides a convenient way of determining the centroid coordinates. If the system is using afocal mode, the image coordinates are given in radians with respect to the local Z axis.

For graphic windows, the left/right cursor keys will change the surface number and recompute the data.

Through Focus

Purpose:

Show spot diagrams as they change through focal plane shifts.

Settings:

The settings are described in "Standard" on page 129.

Discussion:

The through focus spot diagrams are useful for estimating astigmatism, or for analyzing best focus or depth of focus.

Full Field

Purpose:

Shows spot diagram with all field points on a common scale.

Settings:

The settings are described in "Standard" on page 129.

Discussion:

The "Full Field" spot diagram type is similar to the "Standard" type, except all of the spots are plotted with respect to the same reference point, as opposed to a separate reference point for each field position. This provides some idea of how the spot would look relative to the other field points. For example, this can be used to determine if two closely spaced image points can be resolved. The "Full Field" spot diagram type is less useful if the spot radius is small compared to the total field size, because in this case the spots for each field will appear very small compared to the distance between the spots. To increase the detail of the spots, enter a factor for the "Exaggerate" setting. This dimensionless factor magnifies the transverse aberrations to make the spot structure more apparent, although clearly there is an inherent distortion in using this feature. If "chief ray" is selected as the reference point, then the chief ray for field position 1 will be used.

Matrix

Purpose:

Show spot diagram as a matrix of individual diagrams, with each field along a row and each wavelength down a column.

Settings:

The options for the matrix spot diagram are similar to those for the standard spot diagram, except for the addition of the following option.

Item	Description
Ignore Lateral Color	If checked, this option will reference each spot diagram to the reference point for each field and wavelength independently. This in effect ignores the effects of lateral color which can displace the reference points for each wavelength.

Discussion:

The matrix representation is a convenient way of distinguishing the components of wavelength dependent aberrations.

Configuration Matrix

Purpose:

Show spot diagram as a matrix of individual diagrams, with each field along a row and each configuration down a column.

Settings:

The options for the configuration matrix spot diagram are similar to those for the standard spot diagram.

Discussion:

The configuration matrix representation is a convenient way of distinguishing the components of configuration dependent aberrations. Along the left side of the plot the field positions are listed; only the field positions for the

last configuration are listed if more than one configuration is shown and the field definitions are part of the multi-configuration data which changes.

MTF

FFT MTF

Purpose:

Computes the diffraction modulation transfer function (MTF) data for all field positions using an FFT algorithm.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Show Diffraction Limit	Select whether or not the diffraction limited data should be displayed.
Max Frequency	Specify the maximum frequency (see "MTF Units" on page 106) plotted.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number to be used in the calculation.
Type	Select either modulation, real, imaginary, phase, or square wave response.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.
Surface	Selects the surface at which the MTF is to be evaluated. This is useful for evaluating intermediate images. See "Evaluating results at intermediate surfaces" on page 128.

Discussion:



See the discussion sections of the FFT and Huygens PSF. Those comments also apply to this feature.

The diffraction MTF computation is based upon an FFT of the pupil data. The resulting MTF is the modulation as a function of spatial frequency for a sine wave object, although optionally the real, imaginary, phase, or square wave response is available. The Square wave MTF is the modulation response for a square wave target of the specific frequency, as opposed to the response to a sine wave target for the other plots. The square wave response is computed from the MTF data using the following formula:

$$S(v) = \frac{4}{\pi} \left[\frac{M(v)}{1} - \frac{M(3v)}{3} + \frac{M(5v)}{5} - \frac{M(7v)}{7} + \dots \right]$$

where $S(v)$ is the square wave response, $M(v)$ is the sinusoidal modulation response, and v is the spatial frequency.

For focal systems, the cutoff frequency at any one wavelength is given by one over the wavelength times the working F/#. ZEMAX computes the working F/# at each wavelength for each field for the sagittal and tangential response separately. This yields accurate MTF data even for systems with anamorphic and chromatic distortion, such as those incorporating cylinders or gratings. For afocal systems, the cutoff frequency is the exit pupil diameter divided by the wavelength.

The diffraction calculations are more accurate as the sampling increases, the peak-to-valley and maximum slope of the OPD decrease, and the transverse ray aberrations decrease. If the peak-to-valley OPD in the pupil is too great, then the wavefront sampling is too coarse and aliasing occurs. Aliasing will result in inaccurate data. ZEMAX will attempt to detect when aliasing occurs, and issue an appropriate error message. However, ZEMAX cannot automatically detect when the sampling is too low in all cases, especially in the presence of very steep slopes on the wavefront phase.

The FFT based MTF assumes a (reasonably) uniform distribution of rays on the exit pupil in cosine space to be accurate. Some systems, such as extremely fast off-axis reflectors, have dramatic stretching of the exit pupil and the FFT based MTF will thus be inaccurate. For these systems, the Huygens MTF should be used instead. For more information, see “Huygens MTF” on page 137.

When the OPD in waves is large, such as more than 10 waves, it is probably a good idea to switch to geometric MTF instead of diffraction MTF. For these highly aberrated systems, the geometric MTF is very accurate, especially at low spatial frequencies (the higher frequency MTF falls off rapidly when aberrations are large). If shown, the diffraction limit curve is for the aberration free response at the reference field position; see “Diffraction limited” on page 52.

The spatial frequency scale of the MTF plot is in cycles per mm in image space, or cycles per milliradian in object space (see “MTF Units” on page 106). MTF in cycles per mm is computed in image space, so any magnification of the system needs to be considered when determining spatial frequency response for object space.

The nature of the FFT algorithm is that the computation is done in pupil space coordinates. For this reason, rotating the image surface will have no effect on the orientation of the computed MTF. The tangential response corresponds to a periodic target oriented with lines along the object space X axis, and the sagittal response corresponds to a periodic target oriented with lines along the object space Y axis. This is different from the conventions of the “Geometric MTF” on page 139 and the “Huygens MTF” on page 137.

FFT Through Focus MTF

Purpose:

Computes the FFT modulation transfer function data as a function of focus shift at a specific frequency.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Delta Focus	For focal systems, delta focus is the \pm Z-axis range of the plot in lens units. For afocal systems, delta focus is the \pm optical power range of the plot in diopters.
Frequency	The spatial frequency (see “MTF Units” on page 106) for which data is plotted.
# Steps	The number of focal planes at which the data is computed. A smooth curve is drawn through the computed points. More steps yield higher accuracy and longer computation times.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select either modulation, real, imaginary, phase, or square wave response.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.

Discussion:

See “FFT MTF” on page 133 for details.

FFT Surface MTF

Purpose:

Displays the FFT computed MTF data as a 3D surface, contour, grey scale or false color map. This plot is useful for visualizing the MTF response for object orientations other than purely sagittal or tangential.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Rotation	Rotation specifies how the surface plots are rotated for viewing; either 0, 90, 180, or 270 degrees.
Scale	The scale factor overrides the automatic vertical scaling set by the program on the surface plots. Generally, this value should be set to unity. The scale factor can be greater than unity to vertically stretch the plot, or less than unity to compress it.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number determines for which defined field position the calculation should be performed.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.
Type	Choose modulation, real part, or imaginary part of either the incoherent or coherent optical transfer function. If a text listing is generated and either the real or the imaginary part is requested, both the real and the imaginary parts are listed for completeness.

Discussion:

The regular MTF plot is just two orthogonal cross sections through the surface MTF plot. This plot is primarily qualitative. See “FFT MTF” on page 133 for details.

FFT MTF vs. Field

Purpose:

Computes the FFT MTF data as a function of field position, and displays the data in a graph.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Frequency 1-6	The spatial frequencies (see “MTF Units” on page 106) for which data is plotted.
Wavelength	The wavelength number to be used in the calculation.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.

Item	Description
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See the comments about vignetting factors in the discussion section below.
Field Density	The field density is the number of points between zero degrees and the maximum field at which the MTF is calculated, intermediate values are interpolated. A maximum of 100 field points is allowed.
Scan Type	Choose +y, +x, -y, or -x field scan direction.

Discussion:

See “FFT MTF” on page 133 for details. This feature plots MTF vs. field height up to the maximum defined radial field coordinate.

Comment about vignetting factors

Vignetting factors determine the size and shape of the pupil as seen from different field points (see “Vignetting factors” on page 115 for a full discussion). Because this analysis feature needs to trace rays at arbitrary intermediate field points where no specific vignetting factors are defined, the use of vignetting factors is not recommended. If “Remove Vignetting Factors” is checked on (the default), any defined vignetting factors will automatically be replaced with surface apertures for this computation. The surface aperture method is generally more accurate than the vignetting factor method when the pupil is overfilled with light. The resulting data may be different between the two methods. In some cases, particularly where the vignetting factors are being used to define the shape of the source beam rather than the apertures of the optics, it may be required to use the defined vignetting factors. In this case, check the “Remove Vignetting Factors” box off. ZEMAX will then use the closest defined field to determine the vignetting factors to use for an arbitrary field point.

FFT MTF Map

Purpose:

Computes the FFT MTF as a function of field position, and displays the data over a rectangular region of field.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
X or Y Field Width	The X or Y field width in field units. This is the total width or height, not the half width or height. Field units are degrees in object space if field angle is used, otherwise field units are the same as lens units.
Frequency	The spatial frequency (see “MTF Units” on page 106) for which data is plotted.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Wavelength	The wavelength number to be used in the calculation, or All for polychromatic MTF.
X or Y Pixels	The number of pixels at which to compute the MTF in each respective direction. Note the size of the pixels is determined by both the number of pixels and the width of the field; the pixels are not required to be square. The MTF is computed at the center of the pixel and the MTF is assumed to have that value over the entire region of the pixel for display purposes.
MTF Data	Choose Tangential, Sagittal, or average MTF to be displayed.
Reference Field	This control selects the field number that corresponds to the center of the map. If zero is selected the (0, 0) field coordinate is used as the center of the map.

Item	Description
Show As	Choose grey scale or false color map as the display option.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See “Comment about vignetting factors” on page 136.

Discussion:

See “FFT MTF” on page 133 for details.

This feature computes the MTF at each field point on a 2D grid. If the total number of points is large, the computation time may become quite large. See also the Geometric MTF Map feature on page 141.

Huygens MTF

Purpose:

Computes the diffraction modulation transfer function (MTF) data using a Huygens direct integration algorithm.

Settings:

Item	Description
Pupil Sampling	Selects the size of the grid of rays to trace to perform the computation. Higher sampling densities yield more accurate results at the expense of longer computation times.
Image Sampling	The size of the grid of points on which to compute the diffraction image intensity. This number, combined with the image delta, determine the size of the PSF used to compute the MTF. See also “Huygens PSF” on page 146.
Image Delta	The distance in micrometers between points in the image grid. Use zero for the default grid spacing.
Configuration	Select "All" to perform a calculation of the MTF based upon a coherent sum of the PSF at each wavelength across all configurations, or select the "Current" or any single configuration. Note the MTF is computed from the PSF, which is a coherent sum for the same wavelength in each configuration, followed by an incoherent sum of the resulting PSF's for different wavelengths. For this reason, each defined wavelength must be the same in all configurations. Wavelength and configuration weights may be used but the wavelength values must be identical. See CWGT and WLWT in “Summary of multi-configuration operands” on page 579. This coherent sum also assumes that the image surface is located in the identical position in all configurations. If "All" is selected, and both focal and afocal mode configurations are defined, the Configuration setting will automatically be reset to "Current".
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select the data to display, currently modulation is the only option.
Max Frequency	The maximum frequency in MTF units (see “MTF Units” on page 106) to display.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.

Discussion:



See the discussion sections of the Huygens PSF. Those comments also apply to this feature.

The Huygens MTF computes an FFT of the Huygens PSF. The settings for Image Sampling and Image Delta are the same as for the Huygens PSF, therefore it is instructive to do a Huygens PSF first (see “Huygens PSF” on page 146). Since the transform is done on the PSF in image space coordinates, the tangential response corresponds to spatial frequencies in the y direction in local image surface coordinates, and the sagittal response corresponds to spatial frequencies in the x direction. The Huygens MTF also has no dependence on the location of rays in the paraxial pupils. The MTF can therefore be computed for any system that the Huygens PSF can be computed for including many non-sequential systems using ports where reference rays required by other diffraction algorithms would not make it through, or for systems where pupils or images formed by multiple non-sequential sub-apertures are overlapped. Systems with extreme exit pupil distortion, such as very fast off-axis reflectors, are also handled correctly with the Huygens technique.

The nature of the Huygens algorithm is that the computation is done in image space coordinates. For this reason, rotating the image surface will affect the orientation of the computed MTF. The tangential response corresponds to the image of a periodic target oriented with lines along the image space X axis, and the sagittal response corresponds to the image of a periodic target oriented with lines along the image space Y axis. This is different from the conventions of the “FFT MTF” on page 133.

Huygens Through Focus MTF

Purpose:

Computes the diffraction modulation transfer function (MTF) data using a Huygens direct integration algorithm and displays the data as a function of delta focus.

Settings:

Item	Description
Delta Focus	For focal systems, delta focus is the \pm Z-axis range of the plot in lens units. For afocal systems, delta focus is the \pm optical power range of the plot in diopters.
Frequency	The spatial frequency (see “MTF Units” on page 106) for which data is plotted.
# Steps	The number of focal planes at which the data is computed. A smooth curve is drawn through the computed points. More steps yield higher accuracy and longer computation times.

The other settings are the same as for the Huygens MTF feature, as described on page 137.

Discussion:

This feature is very similar to the Huygens MTF feature, as described on page 137.

Huygens Surface MTF

Purpose:

Computes the diffraction modulation transfer function (MTF) data using a Huygens direct integration algorithm and displays the data as a grey scale or false color plot.

Settings:

Item	Description
Show As	Choose grey scale or false color map as the display option.
Type	Choose modulation, real part, or imaginary part of either the incoherent or coherent optical transfer function. If a text listing is generated and either the real or the imaginary part is requested, both the real and the imaginary parts are listed for completeness.

The other settings are the same as for the Huygens MTF feature, as described on page 137.

Discussion:

This feature is very similar to the Huygens MTF feature, as described on page 137.

Huygens MTF vs. Field

Purpose:

Computes the Huygens MTF data as a function of field position, and displays the data in a graph.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Frequency 1-6	The spatial frequencies (see “MTF Units” on page 106) for which data is plotted.
Wavelength	The wavelength number to be used in the calculation.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See the comments about vignetting factors in the discussion section below.
Field Density	The field density is the number of points between zero degrees and the maximum field at which the MTF is calculated, intermediate values are interpolated. A maximum of 100 field points is allowed.
Scan Type	Choose +y, +x, -y, or -x field scan direction.

Discussion:

This feature is very similar to the “FFT MTF vs. Field” on page 135, except the Huygens method is used to compute the MTF rather than the FFT method. See “Huygens MTF” on page 137 for details on the Huygens MTF method.

Geometric MTF

Purpose:

Computes the geometric MTF, which is an approximation to the diffraction MTF based upon ray aberration data.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Max Frequency	The maximum spatial frequency (see “MTF Units” on page 106) for which data is plotted.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Multiply by Diffraction Limit	When checked, will scale the geometric MTF by the diffraction limited MTF to yield a more realistic result for systems with small aberrations. Should always be used.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. See “Surface scattering settings” on page 84.

Discussion:

The geometric MTF is a useful approximation to the diffraction MTF if the system is not close to the diffraction limit. The primary advantage to using the geometric MTF is for systems which have too many waves of aberration to permit accurate calculation of the diffraction MTF. The geometric MTF is also very accurate at low spatial frequencies for systems with large aberrations.

The nature of the geometric algorithm is that the computation is done in image space coordinates. For this reason, rotating the image surface will affect the orientation of the computed MTF. The tangential response corresponds to the image of a periodic target oriented with lines along the image space X axis, and the sagittal response corresponds to the image of a periodic target oriented with lines along the image space Y axis. This is different from the conventions of the “FFT MTF” on page 133.

Geometric Through Focus MTF

Purpose:

Computes the geometric MTF data through focus at a specific spatial frequency.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Delta Focus	For focal systems, delta focus is the \pm Z-axis range of the plot in lens units. For afocal systems, delta focus is the \pm optical power range of the plot in diopters.
Frequency	The spatial frequency (see “MTF Units” on page 106) for which data is plotted.
# Steps	The number of focal planes at which the data is computed. A smooth curve is drawn through the computed points. More steps yield higher accuracy and longer computation times.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Multiply by Diffraction Limit	When checked, will scale the geometric MTF by the diffraction limited MTF to yield a more realistic result for systems with small aberrations. Should always be used.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. See “Surface scattering settings” on page 84.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.

Discussion:

See the Geometric Transfer Function section on page 139 for details.

Geometric MTF vs. Field

Purpose:

Computes the geometric modulation transfer function data as a function of field position.

Settings:

The settings are identical to those for the FFT MTF vs. Field feature, with the added ability to scatter rays, and a choice to scale by the diffraction limit or not.

Discussion:

This feature is nearly identical to the (diffraction) MTF vs. Field feature, except the geometric MTF is used rather than the diffraction based MTF.

Geometric MTF Map

Purpose:

Computes the geometric modulation transfer function data as a function of field position, and displays the data over a rectangular region of field.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
X or Y Field Width	The X or Y field width in field units. This is the total width or height, not the half width or height. Field units are degrees in object space if field angle is used, otherwise field units are the same as lens units.
Frequency	The spatial frequency (see "MTF Units" on page 106) for which data is plotted.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Wavelength	The wavelength number to be used in the calculation, or All for polychromatic GMTF.
X or Y Pixels	The number of pixels at which to compute the GMTF in each respective direction. Note the size of the pixels is determined by both the number of pixels and the width of the field; the pixels are not required to be square. The GMTF is computed at the center of the pixel and the GMTF is assumed to have that value over the entire region of the pixel for display purposes.
MTF Data	Choose Tangential, Sagittal, or average GMTF to be displayed.
Reference Field	This control selects the field number that corresponds to the center of the map. The value zero refers to the center of the object ($h_x = 0$, $h_y = 0$), even if no field point is defined there.
Show As	Choose grey scale or false color map as the display option.
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. See "Surface scattering settings" on page 84.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 136.

Discussion:

See the Geometric Transfer Function on page 139 for details. This feature computes the GMTF at each field point on a 2D grid. If the total number of points is large, the computation time may become quite large. See also the (Diffraction) MTF Map feature on page 136.

PSF

FFT PSF

Purpose:

Computes the diffraction point spread function (PSF) using the Fast Fourier Transform (FFT) method.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Display	The display size indicates what portion of the computed data will be drawn when a graphic display is generated. The display grid can be any size from 32 x 32 up to twice the sampling grid size. Smaller display sizes will show less data, but at higher magnification for better visibility. This control has no effect on the text display.
Rotation	Rotation specifies how the surface plots are rotated for viewing; either 0, 90, 180, or 270 degrees.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select linear (intensity), logarithmic (intensity), phase, real part (signed amplitude), or imaginary part (signed amplitude).
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Image Delta	The delta distance between points in image space, measured in micrometers. If zero, a default spacing is used. If negative, the image delta is set to the maximum allowed value and the full sampling grid is used. See the discussion for details.
Normalize	If checked, the peak intensity will be normalized to unity. Otherwise, the peak intensity is normalized to the peak of the unaberrated PSF (the Strehl ratio).
Surface	Selects the surface at which the PSF is to be evaluated. This is useful for evaluating intermediate images. See "Evaluating results at intermediate surfaces" on page 128.

Discussion:

The FFT method of computing the PSF is very fast, however, a few assumptions are made which are not always valid. The slower, but more general Huygens method makes fewer assumptions, and is described in "Huygens PSF" on page 146.

Assumptions used in the FFT PSF calculation

The FFT PSF computes the intensity of the diffraction image formed by the optical system for a single point source in the field. The intensity is computed on an imaginary plane which is centered on and lies perpendicular to the incident chief ray at the reference wavelength. The reference wavelength is the primary wavelength for polychromatic computations, or the wavelength being used for monochromatic calculations. Because the imaginary plane lies normal to the chief ray, and not the image surface, the FFT PSF computes overly optimistic (a smaller PSF) results when the chief ray angle of incidence is not zero. This is often the case for systems with tilted image surfaces, wide angle systems, systems with aberrated exit pupils, or systems far from the telecentric condition.

The other main assumption the FFT method makes is that the image surface lies in the far field of the optical beam. This means the computed PSF is only accurate if the image surface is fairly close to the geometric focus for all rays; or put another way, that the transverse ray aberrations are not too large. There is no hard and fast limit, however if the transverse aberrations exceed a few hundred wavelengths, the computation is likely not accurate. Note that even systems with very little wavefront aberration can have large transverse ray aberrations; for example, a cylinder lens which only focuses rays along one direction. In this case, the transverse aberrations along the unfocused direction will be on the order of the beam diameter. The Huygens PSF method may provide more accurate results in these cases as well.

For most lenses, a less important assumption is that scalar diffraction theory applies. The vectorial nature of the light is not accounted for. This is significant in systems that are very fast, around F/1.5 (in air) or faster. The scalar theory predicts overly optimistic (a smaller PSF) results when the F/# is very fast.

For systems where the chief ray is nearly normal (less than perhaps 20 degrees), the exit pupil aberrations are negligible, and the transverse ray aberrations are reasonable, then the FFT PSF is accurate and generally much faster than the Huygens PSF method.

When in doubt, both PSF methods should be employed for comparison. A solid understanding on the part of the user of these assumptions and the method of computation is essential to recognize cases where the accuracy may be compromised.

Discussion of the FFT method and sampling issues

The FFT PSF algorithm exploits the fact that the diffraction PSF is related to the Fourier transform of the complex amplitude of the wavefront in the exit pupil of the optical system. The amplitude and phase in the exit pupil are computed for a grid of rays, an FFT is performed, and the diffraction image intensity is computed.

There is a tradeoff between the sampling grid size in the pupil, and the sampling period in the diffraction image. For example, to decrease the sampling period in the diffraction image, the sampling period in the pupil must increase. This is done by "stretching" the pupil sampling grid so that it overfills the pupil. This process means fewer points actually lie within the pupil.

As the sampling grid size is increased, ZEMAX scales the grid on the pupil to yield an increase in the number of points that lie on the pupil, while simultaneously yielding closer sampling in the diffraction image. Each time the grid size is doubled, the pupil sampling period (the distance between points in the pupil) decreases by the square root of 2 in each dimension, the image surface sampling period also decreases by the square root of 2 in each dimension, and the width of the diffraction image grid increases by a factor of the square root of 2 (since there are twice as many points in each dimension). All ratios are approximate, and asymptotically correct for large grids.

The stretching is referenced to a grid size of 32 x 32. The 32 x 32 grid of points is placed over the pupil, and the points that lie within the pupil are actually traced. For this grid size, the default distance between points in the diffraction image surface is given by

$$\Delta = \lambda F \left(\frac{n-2}{2n} \right),$$

where F is the working F/# (not the same as the image space F/#), λ is the shortest defined wavelength, and n is the number of points across the grid. The -2 factor is due to the fact the pupil is not centered on the grid (since n is even), but is offset at $n/2 + 1$. The 2n in the denominator is due to the zero-padding described later.

For grids larger than 32 x 32, the grid is by default stretched in pupil space by a factor of $\sqrt{2}$ each time the sampling density doubles. The general formula for the sampling in image space is then

$$\Delta = \lambda F \left(\frac{n-2}{2n} \right) \left[\frac{32}{n} \right]^{1/2},$$

and the total width of the image data grid is

$$W = 2n\Delta.$$

Since the stretching of the pupil grid decreases the number of sample points in the pupil, the effective grid size (the size of the grid that actually represents traced rays) is smaller than the sampling grid. The effective grid size increases as the sampling increases, but not as quickly. The following table summarizes the approximate effective grid size for various sampling density values.

DEFAULT EFFECTIVE GRID SIZES FOR PSF CALCULATIONS

Sampling Grid Size	Approximate Effective Pupil Sampling
32 x 32	32 x 32
64 x 64	45 x 45
128 x 128	64 x 64
256 x 256	90 x 90
512 x 512	128 x 128
1024 x 1024	181 x 181
2048 x 2048	256 x 256
4096 x 4096	362 x 362
8192 x 8192	512 x 512

The sampling is also a function of wavelength. The discussion above is only valid for the shortest wavelength used in the calculation. If the computation is polychromatic, then the longer wavelengths will be scaled to have smaller effective grids. The scale factor used is the ratio of the wavelengths. This should be considered when selecting sampling grids for systems with broad wavelength bands. For polychromatic computations, the data for shorter wavelengths is more accurate than for longer wavelengths.

The image delta, Δ , can be selected manually if a different sampling distance is required. If the image delta is zero, ZEMAX uses the default spacing and sampling grids described above. If the image delta is greater than zero, then ZEMAX scales the pupil sampling to yield the desired image delta size. The actual amount of stretching depends upon the grid size, the image delta, the defined wavelengths, the F/#'s at each field and wavelength, and the aspect ratio of the exit pupil. If the image delta is set too small, then not enough points will be left to sample the pupil; if the image delta is too big, then the pupil grid will not extend over the full width of the exit pupil. Both of these cases are trapped by ZEMAX and an error message will be issued if they occur. If the image delta is less than zero, then ZEMAX does not stretch the pupil at all. This maximizes the extent of the PSF in image coordinates and uses the full grid of rays in the pupil, at the expense of not increasing the spatial resolution with sampling, as the image delta will be nearly constant with increased sampling.

Once the sampling is specified, ZEMAX doubles the array size in a process called "zero padding". This means for a 32 x 32 sampling, ZEMAX uses the center portion of a 64 x 64 grid. Therefore, the diffraction PSF will be distributed over a 64 x 64 size grid. The sampling in the image space is always twice the pupil sampling. Zero padding is performed to reduce aliasing.

The nature of the FFT algorithm is that the computation is done in pupil space coordinates. For this reason, rotating the image surface will have no effect on the orientation of the computed PSF. The X and Y orientation of the PSF corresponds to the X and Y orientation of rays in the entrance pupil, which is not always the same as the spatial X and Y orientation of the image. To compute the PSF in the image space coordinates, see "Huygens PSF" on page 146.

FFT PSF Cross Section

Purpose:

This feature plots cross sections through the diffraction PSF.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Row/Col	The row or column to display. For a sampling of 32 x 32, there are 64 rows and 64 columns (see the discussion section of the FFT PSF feature). Whether a row or a column is used depends upon the "Type" setting. The default "Center" displays the center of the PSF, which is not always the peak of the PSF.
Plot Scale	Sets the maximum horizontal scale in micrometers. Enter zero for automatic scaling.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select X or Y cross sections, either linear, logarithmic, phase, real part, or imaginary part. X cross sections are called rows, and Y cross sections are called columns.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Normalize	If checked, the peak intensity will be normalized to unity. Otherwise, the peak intensity is normalized to the peak of the unaberrated PSF (the Strehl ratio).

Discussion:



See the discussion section of the FFT PSF. Those comments also apply to this feature.

The cross sections are taken directly from the PSF data. Because the PSF is computed directly from the phase in the exit pupil, the orientation of the coordinate system may not be correct in all cases. What constitutes positive x or y may not agree with data presented in image space coordinates such as the spot diagram.

FFT Line/Edge Spread

Purpose:

This feature plots the edge or line spread functions based upon a computation and integration of the diffraction FFT PSF.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Spread	Choose either line or edge spread function.
Plot Scale	Sets the maximum horizontal scale in micrometers. Enter zero for automatic scaling.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select X or Y orientation, and either linear intensity, logarithmic intensity, phase, real part, or imaginary part. The phase, real part, and imaginary part computations are only available if the "Use Coherent PSF" option is selected.

Item	Description
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Use Coherent PSF	If checked, the computation is done assuming coherent illumination. If unchecked, the illumination is considered incoherent.

Discussion:



See the discussion section of the FFT PSF. Those comments also apply to this feature.

The line response function (or line spread function, LSF) is the cross section of the image of a line object. The edge spread function (ESF) is the cross section of the image of an edge (a semi-infinite plane). The X- or Y-orientation refers to the direction of the line or edge. The X-orientation means the line or edge is parallel to the X-axis. The Y-orientation means the line or edge is parallel to the Y-axis. The LSF and ESF are computed by integration of the FFT PSF. See also “Geometric Line/Edge Spread” on page 160.

Huygens PSF

Purpose:

Computes the diffraction PSF using direct integration of Huygens wavelets method. The Strehl ratio is also computed.

Settings:

Item	Description
Pupil Sampling	Selects the size of the grid of rays to trace to perform the computation. Higher sampling densities yield more accurate results at the expense of longer computation times.
Image Sampling	The size of the grid of points on which to compute the diffraction image intensity. This number, combined with the image delta, determine the size of the area displayed.
Image Delta	The distance in micrometers between points in the image grid. Use zero for the default grid spacing.
Rotation	Rotation specifies how the surface plots are rotated; either 0, 90, 180, or 270 degrees.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select linear (intensity), or logarithmic (intensity). The logarithmic scaling can range from 1 to 5 decades. Also available are real amplitude, imaginary amplitude, and phase in degrees.
Configuration	Select "All" to perform a coherent sum of the PSF at each wavelength across all configurations, or select the "Current" or any single configuration. Note this is a coherent sum for the same wavelength in each configuration, followed by an incoherent sum of the resulting PSF's for different wavelengths. For this reason, each defined wavelength must be the same in all configurations. Wavelength and configuration weights may be used but the wavelength values must be identical. See CWGT and WLWT in “Summary of multi-configuration operands” on page 579. This coherent sum also assumes that the image surface is located in the identical position in all configurations. If "All" is selected, and both focal and afocal mode configurations are defined, the Configuration setting will automatically be reset to "Current".
Normalize	If checked, the peak intensity will be normalized to unity. Otherwise, the peak intensity is normalized to the peak of the unaberrated PSF (the Strehl ratio).

Item	Description
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option. The True Color option creates an RGB color representation of the PSF by converting the wavelengths to the closest RGB equivalent and summing over all wavelengths. The accuracy of the True Color presentation is limited by the RGB method of rendering color on a computer display; and it is not possible to exactly represent monochromatic colors. The True Color option cannot be used if the Type is anything other than linear.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Use Centroid	If checked, the plot will be centered on the geometric image centroid. If unchecked, the plot will be centered on the chief ray.

Discussion:

One way of considering the effects of diffraction is to imagine each point on a wavefront as a perfect point source with an amplitude and phase. Each of these point sources radiates a spherical "wavelet", sometimes called a "Huygens wavelet" after Huygens, who first proposed the model. The diffraction of the wavefront as it propagates through space is given by the interference, or complex sum, of all the spherical wavelets radiated.

To compute the Huygens PSF, a grid of rays is launched through the optical system, and each ray represents a particular amplitude and phase wavelet. The diffraction intensity at any point on the image surface is the complex sum of all these wavelets, squared. The PSF is computed this way for every point on the image grid.

The Image Delta value determines the point spacing of the image space grid. If a value of zero is specified, a default grid spacing is used. The default Image Delta is given by:

$$\Delta = \frac{2\lambda F}{\sqrt{n}}$$

where n is the number of points in the image space grid, λ is the longest wavelength used in the calculation, and F is the working F/#. The exact value of the image Delta size is not critical, as long as the entire width of the PSF is included within the range of n * (Image Delta).

Unlike the FFT PSF, ZEMAX computes the Huygens PSF on an imaginary plane tangent to the image surface at the chief ray intercept. Note the imaginary plane is normal to the normal of the surface, not the chief ray. Therefore, the Huygens PSF accounts for any local tilt in the image surface caused by either the image surface slope, the chief ray incidence angle, or both.

The Huygens method accounts for the evolving shape of the diffraction image as the beam propagates along the image surface. This is an important effect if the image surface is tilted with respect to the incoming beam. Another advantage to the Huygens PSF method is that any grid size and spacing may be selected by the user. This allows direct comparison between PSF's from two different lenses, even if the F/#'s or wavelengths are different.

The only disadvantage of the Huygens PSF is speed. Direct integration is slow when compared to the FFT method (see the previous section for details). The computation time depends upon the pupil grid size squared times the image grid size squared, times the number of wavelengths. ZEMAX accounts for any symmetry the system has. The Huygens PSF automatically uses all available processors for maximum speed on multiple CPU computers.

Huygens PSF Cross Section

Purpose:

Computes the diffraction PSF using direct integration of Huygens wavelets method. The Strehl ratio is also computed. This feature is very similar to the "Huygens PSF" on page 146, with the difference being the data is plotted as a cross section.

Wavefront

Wavefront Map

Purpose:

Displays the wavefront aberration.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Rotation	Rotation specifies how the surface plots are rotated for viewing; either 0, 90, 180, or 270 degrees.
Scale	The scale factor is used to override the automatic vertical scaling set by the program on the surface plots. The scale factor can be greater than unity to vertically stretch the plot, or less than unity to compress it.
Polarization	If none is selected, polarization is ignored. If Ex, Ey, or Ez is selected, then the phase due to polarization effects for the specified component of the electric field is added to the optical path difference. If the polarization phase exceeds one wave, the wavefront map may display abrupt changes of 2π because no phase "unwrapping" is performed.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Reference To Primary	By default, the wavefront aberration is referenced to the reference sphere for the wavelength being used. If this box is checked, then the primary wavelength reference sphere will be used instead. In other words, checking this box will cause the data to exhibit the effects of lateral color.
Use Exit Pupil Shape	If checked, the shape of the pupil is distorted to show the approximate shape of the exit pupil as seen from the image point for the specified field. The shape is based upon the F/# of the beam in the X and Y pupil directions. If this box is unchecked, then instead the plot will be scaled to circular entrance pupil coordinates, no matter how distorted the exit pupil may actually be.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.
Surface	Selects the surface at which the data is to be evaluated. This is useful for evaluating intermediate images. See "Evaluating results at intermediate surfaces" on page 128.
Remove Tilt	If checked, the linear X- and Y- tilt is removed from the data. This is equivalent to referencing the OPD data to the centroid.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see "The Contour Format String" on page 151.
Subaperture Data: Sx, Sy, Sr	Defines a subaperture of the pupil to compute the wavefront data for. See "Subaperture computations" on page 197 for details.

Discussion:

See also the Interferogram feature described below.

Interferogram

Purpose:

Generates and displays interferograms.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Surface	Selects the surface at which the data is to be evaluated. This is useful for evaluating intermediate images. See "Evaluating results at intermediate surfaces" on page 128.
Scale Factor	Determines the number of fringes per wave of OPD. Useful for modeling double pass interferometers (i.e. use a scale factor of two).
Show As	Choose contour map, grey scale, or false color map as the display option.
Wavelength	The wavelength is the wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
X-Tilt	The number of fringes of tilt to add in the x-direction after applying the scale factor.
Y-Tilt	The number of fringes of tilt to add in the y-direction after applying the scale factor.
Beam 1	Selects the first beam for the interferogram. See the discussion.
Beam 2	Selects the second beam for the interferogram. See the discussion.
Ref Beam To Vertex	Normally, ZEMAX references the OPD to the chief ray; which in effect subtracts out tilt from the wavefront phase. For interferometry, it is sometimes desirable to retain the wavefront tilt. Checking this option on will add tilt to the beam based upon the deviation of the chief ray from the image surface vertex. This option is only useful for field positions whose chief ray is reasonably close to the surface vertex, where the assumption that tilt is described by the deviation of the chief ray is valid. This feature should also only be used when the scale factor is 1.0.
Use Exit Pupil Shape	If checked, the shape of the pupil is distorted to show the approximate shape of the exit pupil as seen from the image point for the specified field. The shape is based upon the F/# of the beam in the X and Y pupil directions. If this box is unchecked, then instead the plot will be scaled to circular entrance pupil coordinates, no matter how distorted the exit pupil may actually be. If Beam 1 is "reference" then the configuration of Beam 2 is used to determine the pupil shape, otherwise the configuration of Beam 1 is used. If both Beam 1 and Beam 2 are defined by different configurations, the shape of Beam 1 is used, and no attempt is made to confirm that the pupils are the same shape. If the pupil shapes are different for the two configurations, this feature cannot accurately predict the interferogram.
Consider Path Length Difference	If unchecked, the interferogram is computed assuming the phase at the center of each beam is zero, resulting in a "dark" fringe at the center. No attempt is made to determine the phases of the two beams with respect to each other. If checked, the difference between the total optical path length of the chief rays in each beam configuration is considered. This setting will alter the phase of the entire interferogram, but will not change the shape of the fringe pattern. For example, if the path difference is exactly one-half wave, dark fringes will become bright and bright will become dark. The optical path length of the chief rays from object to image surface is used, even if the fringes are not localized at the image surface. This assumption is only valid if the fringes are localized on a plane that is the same distance from the image surfaces in each configuration. If a reference beam is selected rather than a beam from a configuration, the optical path length for the reference beam is assumed to be zero.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see "The Contour Format String" on page 151.
Subaperture Data: Sx, Sy, Sr	Defines a subaperture of the pupil to compute the interferogram data for. See "Subaperture computations" on page 197 for details.

Discussion:

This feature works by computing two pupil maps, one each from beams 1 and 2. The phase (or OPD) of these two pupil maps is subtracted, and then optionally some linear phase is added as a function of the x and y pupil coordinate to simulate tilt fringes. The individual beams may be OPD as computed for any one configuration, or a "reference" beam which has identically zero OPD may be selected.

Interferometers may be simulated by modeling the two paths through the system using two configurations, and then computing the interferogram of the two resulting beams. The accuracy of this approach is limited by some simplifying assumptions:

-Any lateral shift or magnification difference between the two beams is ignored; it is assumed that the pupils perfectly overlap at the exit pupil.

-Any differences in transmission are ignored; so the two OPD values at any one point in the pupil are assumed to be of equal value in intensity and the phase can be subtracted to yield the net phase difference.

Foucault Analysis

Purpose:

Generates and displays Foucault knife-edge shadowgrams.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Type	Selects either linear or logarithmic display of the data.
Show As	Choose surface, contour map, grey scale, false color map, or cross section as the display option.
Wavelength	The wavelength is the wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Row/Column	When "Show As" is selected to be a cross section, this control defines the row or column number to display.
Knife	Choose Horizontal Above, Horizontal Below, Vertical Left, or Vertical Right. The Vertical Left knife blocks out all the light near focus from the knife position coordinate left; that is, towards negative x coordinates. The Vertical Right blocks all light from the knife position rightward. The Horizontal Above blocks all light from the knife position up, and Horizontal Below blocks all light from the knife position down. The terms left, right, up and down refer to the -x, +x, +y, and -y directions in the local coordinate system of the image surface.
Position	The position in micrometers relative to the chief ray of the knife. The coordinate is assumed to be in X or Y depending upon whether an X or Y knife is selected.
Data	The computed shadowgram, the reference shadowgram, or the difference between the two may be selected. See the discussion for details.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Reference	The name of the bitmap reference image file.

Item	Description
Decenter X/Y	The decenter in X or Y of the reference shadowgram image relative to the computed shadowgram image. The units are relative to the full width or height of the reference shadowgram image. For example, an X Decenter of 0.25 will shift the reference image relative to the computed image by 25% of the full width of the reference image.
Scale X/Y	The scale factor in X or Y of the reference shadowgram image pixels relative to the computed shadowgram image pixels.

Discussion:

This feature simulates the placement of either an X- or Y- oriented knife edge at the image surface; then computes the resulting shadowgram after propagating the vignetted beam back to the near field. The method of calculation involves computing the diffraction based complex amplitude PSF at focus via the FFT method; then a portion of the complex amplitude is vignetted by the simulated knife edge, and the remaining complex amplitude is propagated back to the near field. The shadowgram calculated this way is called the "computed" shadowgram for this feature.

This feature also allows the import of either a BMP or JPG bitmap file of a reference or measured shadowgram. The reference shadowgram may be displayed for convenient check of orientation.

The difference between the computed and reference shadowgram may be displayed. ZEMAX computes the RMS difference between the computed and reference shadowgrams, and this RMS difference may be optimized using the FOUC operand described in the chapter "Optimization". Optimizing the RMS difference permits quantitative determination of the aberrations present in the beam that created the measured shadowgram. When calculating the difference between the computed and reference shadowgram, the two images must be registered together to overlap correctly. The decenter x/y and scale x/y controls are used for registering the two images.

Surface

Surface Sag

Purpose:

Displays the sag of a surface as a 2D color or contour map, or as a 3D surface plot.

Settings:

Item	Description
Sampling	The size of the grid used. The sampling may be 33x33, 65x65, etc.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see "The Contour Format String" below. The contours are defined in lens units.
Surface	The surface number to compute the sag display for.
Data	Choose surface sag.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.

Discussion:

This feature accounts for the size and shape of any aperture present on the surface; even if the aperture is decentered. The sag is computed on a uniform grid of points in XY plane, and the Z value of the sag is the displayed data.

See also the Surface Phase feature described below.

The Contour Format String

The Contour Format string allows some control over the appearance of the contour map.

If left blank, the default contour spacing will be selected. If a single value is entered, then the contour spacing will be set to this value. For example, if 0.05 is entered, the contour interval will be 0.05. If multiple values

separated by spaces are entered, only those contours will be drawn. For example, if "0.8 0.5 0.2" is entered, only those three contours will be drawn. If the specified contour settings result in too many contours to be drawn properly, the default number of contours will be selected.

Surface Phase

Purpose:

Displays the phase of a surface as a 2D color or contour map, or as a 3D surface plot.

Settings:

Item	Description
Sampling	The size of the grid used. The sampling may be 33x33, 65x65, etc.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see "The Contour Format String" on page 151. The contours are defined in periods; each period is a phase change of 2π .
Surface	The surface number to compute the phase display for.
Data	Choose surface phase.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.

Discussion:

This feature accounts for the size and shape of any aperture present on the surface; even if the aperture is decentered. The phase is computed on a uniform grid of points in XY plane, and the phase value is the displayed data. This feature defines phase in units of periods; so one period represents a phase change of 2π . Surfaces which do not impart a phase change to the ray, such as the Standard surface, will display a phase of zero everywhere on the surface phase display. See also the Surface Sag feature described above.

Surface Curvature

Purpose:

Displays the tangential, sagittal, x, or y curvature of a surface as a 2D color or contour map, or as a 3D surface plot.

Settings:

Item	Description
Sampling	The size of the grid used. The sampling may be 33x33, 65x65, etc.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see "The Contour Format String" on page 151. The contours are defined in lens units.
Surface	The surface number to compute the sag display for.
Data	Choose tangential, sagittal, x, or y curvature.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.

Discussion:

This feature accounts for the size and shape of any aperture present on the surface; even if the aperture is decentered.

Surface Sag Cross Section

Purpose:

Displays the sag of a surface as a cross section.

Settings:

Item	Description
Sampling	The sampling may be 33, 65, etc.
Surface	The surface number to compute the sag display for.
Data	Choose surface sag.
Angle	Choose the orientation angle in degrees relative to the local x axis for the cross section.

Discussion:

This feature calculates the surface sag as a cross section across the full diameter of the surface. See also the Surface Phase Cross Section feature described below.

Surface Phase Cross Section

Purpose:

Displays the phase of a surface as a cross section.

Settings:

Item	Description
Sampling	The sampling may be 33, 65, etc.
Surface	The surface number to compute the phase display for.
Data	Choose surface phase.
Angle	Choose the orientation angle in degrees relative to the local x axis for the cross section.

Discussion:

This feature calculates the surface phase as a cross section across the full diameter of the surface. This feature defines phase in units of periods; so one period represents a phase change of 2π . Surfaces which do not impart a phase change to the ray, such as the Standard surface, will display a phase of zero everywhere on the surface phase display. See also the Surface Sag Cross Section feature described above.

Surface Curvature Cross Section

Purpose:

Displays the curvature of a surface as a cross section.

Settings:

Item	Description
Sampling	The sampling may be 33, 65, etc.
Surface	The surface number to compute the sag display for.
Data	Choose tangential, sagittal, x, y, tangential+sagittal, or x+y curvature.
Angle	Choose the orientation angle in degrees relative to the local x axis for the cross section.

Discussion:

This feature calculates the surface curvature cross section across the full diameter of the surface. See also the Surface Sag Cross Section feature described above.

RMS

RMS vs. Field

Purpose:

Plots RMS radial, x, and y spot radius, RMS wavefront error, or Strehl ratio as a function of field angle.

Settings:

Item	Description
Ray Density	If the method is Gaussian quadrature, then the ray density specifies the number of radial rays to be traced. The more rays traced, the greater the accuracy, although the computation time increases. The maximum number is 20 which is sufficient for pupil aberrations up to order 40. If the method is rectangular array, then the ray density indicates the grid size. Rays outside the circular entrance pupil are ignored. See the "Discussion" section for details.
Field Density	The field density is the number of points between zero degrees and the maximum field angle specified at which the RMS/Strehl ratio is calculated, intermediate values are interpolated. A maximum of 100 field points is allowed.
Plot Scale	Sets the maximum vertical scale for the plot. Zero results in automatic scaling.
Method	Select either Gaussian Quadrature (GQ) or Rectangular Array (RA). See "Comments about RMS computation methods" on page 155.
Data	Selects either wavefront error, spot radius, spot x-direction, spot y-direction, or Strehl ratio. Strehl ratio is only available for monochromatic calculations, and is computed using the method described in "Strehl ratio approximation" on page 198. For a polychromatic Strehl, see the "Huygens PSF" on page 146.
Refer To	Select either chief ray or centroid. For monochromatic calculations, the specified wavelength is used for reference. For polychromatic calculations, the primary wavelength is used for reference. Both reference points subtract out wavefront piston. The centroid reference mode also subtracts out the tilt of the wavefront, which yields smaller RMS values.
Orientation	Select +y, -y, +x, or -x field direction. Note the data will only be computed to the limits of the defined fields in the selected direction.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.
Wavelength	Select "All" to display data for each wavelength and a polychromatic computation, any one wavelength to display monochromatic data, "Poly Only" which only displays the polychromatic data, or "All With No Poly" to display all monochromatic data without the polychromatic data. Strehl ratio can only be computed for a single wavelength.
Show Diffraction Limit	If checked, then a horizontal line indicating the diffraction limited response will be drawn on the plot. For RMS radius, x, or y; the diffraction limit is assumed to be 1.22 times the working F/# on axis times the wavelength (primary wavelength if polychromatic). The change in the diffraction limit due to changes in F/# with field are ignored; a single value is used across the range of the plot. For Strehl ratio, 0.8 is used, and for RMS wavefront, 0.072 waves is used. These are all approximate indicators for convenience only; the actual meaning of "diffraction limited" may be open to interpretation.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 136.

Discussion:

This feature calculates the RMS error or Strehl ratio as a function of field angle for each wavelength.

Comments about RMS computation methods

Two different methods of calculation are available; either Gaussian Quadrature (GQ), or Rectangular Array (RA). For GQ, the rays traced are arranged in a radial pattern with an optimal weighting to estimate the RMS with a minimum number of rays. The method is described in G. W. Forbes, "Optical system assessment for design: numerical ray tracing in the Gaussian pupil", J. Opt. Soc. Am. A, Vol. 5, No. 11, p1943 (1988). Although the method is very efficient, the algorithm is not accurate if some of the rays are clipped due to surface apertures. To compute the RMS wavefront in systems with surface apertures requires the use of the RA method, and a larger number of rays for sufficient accuracy. For more information see "Selecting the pupil integration method" on page 485.

Comments about RMS wavefront computations

For RMS wavefront computations, ZEMAX always subtracts out the mean OPD. RMS wavefront, when referenced to the centroid, also requires computation of a shifted, tilted reference sphere. The raw, uncorrected OPD values do not account for the possible shift in image centroid location due to any coma that may be present. The correction method is described in M. Rimmer, "Analysis of Perturbed Lens Systems", Applied Optics Vol. 9, No. 3, p533 (1970). This same method is used when computing RMS wavefront referenced to the centroid by the various RMS wavefront merit function optimization operands. The method is more accurate than simple direct integration of the "raw" OPD values, which are referenced to a sphere centered on the chief ray - image surface intercept point.

RMS vs. Wavelength

Purpose:

Plots RMS radial, x, and y spot radius, RMS wavefront error, or Strehl ratio as a function of wavelength.

Settings:

Item	Description
Ray Density	If the method is Gaussian quadrature, then the ray density specifies the number of radial rays to be traced. The more rays traced, the greater the accuracy, although the computation time increases. The maximum number is 20 which is sufficient for pupil aberrations up to order 40. If the method is rectangular array, then the ray density indicates the grid size. Rays outside the circular entrance pupil are ignored. See the "Discussion" section for details.
Wave Density	The wave density is the number of points between the minimum and maximum defined wavelength at which the RMS/Strehl ratio is calculated, intermediate values are interpolated. A maximum of 100 points is allowed.
Plot Scale	Sets the maximum vertical scale for the plot. Zero results in automatic scaling.
Method	Select either Gaussian Quadrature (GQ) or Rectangular Array (RA). See "Comments about RMS computation methods" on page 155.
Data	Selects wavefront error, spot radius, spot x-direction, spot y-direction, or Strehl ratio.
Refer To	Select either chief ray or centroid. Both reference points subtract out wavefront piston. The centroid reference mode also subtracts out the tilt of the wavefront, which yields smaller RMS values.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.
Field	Select "All" to display data for all fields, select one field to display data for a single field.

Item	Description
Show Diffraction Limit	If checked, then a curve indicating the diffraction limited response will be drawn on the plot. For RMS radius, x, or y; the diffraction limit is assumed to be 1.22 times the working F/# on axis times the wavelength. The change in the diffraction limit due to changes in F/# with field are ignored. For Strehl ratio, 0.8 is used, and for RMS wavefront, 0.072 waves is used. These are all approximate indicators for convenience only; the actual meaning of "diffraction limited" may be open to interpretation.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.

Discussion:

This feature calculates the RMS error or Strehl ratio as a function of wavelength for each field position. The method of calculation is identical to that described in "RMS vs. Field" on page 154; see that section for a detailed discussion. The wavelength range is determined by the minimum and maximum wavelengths defined on the wavelengths dialog box.

RMS vs. Focus

Purpose:

Plots RMS radial, x, and y spot radius, RMS wavefront error, or Strehl ratio as a function of focus change.

Settings:

Item	Description
Ray Density	If the method is Gaussian quadrature, then the ray density specifies the number of radial rays to be traced. The more rays traced, the greater the accuracy, although the computation time increases. The maximum number is 20 which is sufficient for pupil aberrations up to order 40. If the method is rectangular array, then the ray density indicates the grid size. Rays outside the circular entrance pupil are ignored. See the "Discussion" section for details.
Focus Density	The focus density is the number of points between the minimum and maximum focus shift specified at which the RMS/Strehl ratio is calculated, intermediate values are interpolated. A maximum of 100 points is allowed.
Plot Scale	Sets the maximum vertical scale for the plot. Zero results in automatic scaling.
Method	Select either Gaussian Quadrature (GQ) or Rectangular Array (RA). See "Comments about RMS computation methods" on page 155.
Data	Selects either RMS wavefront error, RMS spot radius, RMS X-direction, RMS Y-direction, or Strehl ratio. Strehl ratio is only available for monochromatic calculations, and is computed using the method described in "Strehl ratio approximation" on page 198. For a polychromatic Strehl, see the "Huygens PSF" on page 146.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.
Refer To	Select either chief ray or centroid. For monochromatic calculations, the specified wavelength is used for reference. For polychromatic calculations, the primary wavelength is used for reference. Both reference points subtract out wavefront piston. The centroid reference mode also subtracts out the tilt of the wavefront, which yields smaller RMS values.
Wavelength	Select "All" to display data for a polychromatic computation, select any one wavelength to plot monochromatic data.
Min Focus	The minimum value of the defocus to plot. For focal systems, the units are lens units. For afocal systems, the units are diopters.

Item	Description
Max Focus	The maximum value of the defocus to plot. For focal systems, the units are lens units. For afocal systems, the units are diopters.
Show Diffraction Limit	If checked, then a horizontal line indicating the diffraction limited response will be drawn on the plot. For RMS radius, x, or y; the diffraction limit is assumed to be 1.22 times the working F/# times the wavelength (primary wavelength if polychromatic) on axis. The change in the diffraction limit due to changes in F/# with field are ignored; a single value is used across the range of the plot. For Strehl ratio, 0.8 is used, and for RMS wavefront, 0.072 waves is used. These are all approximate indicators for convenience only; the actual meaning of "diffraction limited" may be open to interpretation.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.

Discussion:

This feature calculates the RMS error or Strehl ratio as a function of a change in focus position for each field position. The method of calculation is identical to that described in "RMS vs. Field" on page 154; see that section for a detailed discussion. ZEMAX adds the specified focus shift to the value of the thickness of the surface prior to the image surface. If the optical system has an odd number of mirrors, this surface normally would have a negative thickness, and therefore negative focus values move the image surface farther away from the last component. For systems with an even number of mirrors, then negative focus values move the image surface closer to the last component.

RMS Field Map

Purpose:

Plots RMS radial, x, and y spot radius, RMS wavefront error, or Strehl ratio as a function of field position in a rectangular map.

Settings:

Item	Description
Ray Density	If the method is Gaussian quadrature, then the ray density specifies the number of radial rays to be traced. The more rays traced, the greater the accuracy, although the computation time increases. The maximum number is 20 which is sufficient for pupil aberrations up to order 40. If the method is rectangular array, then the ray density indicates the grid size. Rays outside the circular entrance pupil are ignored. See the "Discussion" section for details.
Data	Selects wavefront error, spot radius, spot x-direction, spot y-direction, or Strehl ratio. If the method selected is rectangular array, then a rough estimate of the PTV is also available as one of the Data options. Note that PTV depends upon just two points in the pupil, and for this reason it is difficult to obtain the exact PTV without resorting to high sampling and long computation times. Both Strehl Ratio and PTV require a single wavelength to be selected; polychromatic computations are not supported.
Method	Select either Gaussian Quadrature (GQ) or Rectangular Array (RA). See "Comments about RMS computation methods" on page 155.
Plot Scale	Sets the maximum vertical scale for the plot. Zero results in automatic scaling.
Wavelength	Select "All" to display polychromatic data, or select any one wavelength to display monochromatic data. Strehl Ratio and PTV Wavefront can only be computed for a single wavelength.
Field	Selects the reference field. The rectangular field of view will be centered on the selected field point.

Item	Description
Refer To	Select either chief ray or centroid. Both reference points subtract out wavefront piston. The centroid reference mode also subtracts out the tilt of the wavefront, which yields smaller RMS values.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.
Surface	Selects the surface at which the data is to be evaluated. This is useful for evaluating intermediate images. See “Evaluating results at intermediate surfaces” on page 128.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see “The Contour Format String” on page 151.
X/Y Field Size	The half-field width in the X or Y directions in field units. Note this field of view is centered on the reference field selected by the "Field" setting above.
X/Y Field Sampling	The number of points in the X or Y directions used to sample the Data.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See “Comment about vignetting factors” on page 136.

Discussion:

This feature calculates the RMS error or Strehl ratio as a function of field position. The method of calculation is identical to that described in “RMS vs. Field” on page 154; see that section for a detailed discussion.

Encircled Energy

Diffraction

Purpose:

Encircled energy diagram. This is the percentage of total energy enclosed as a function of distance from either the chief ray or the image centroid at the image of a point object.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Type	The analysis type option specifies how the encircled energy is calculated; either encircled (radial), X-only, Y-only, or ensquared.
Maximum Distance	This setting overrides the default scaling. The units are micrometers. To choose the default scaling option, enter zero.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.
Refer To	Select chief ray, centroid, or vertex as the reference point. Vertex refers to the coordinates (0, 0) on the image surface. This option will only return meaningful data if the diffraction image at all selected fields is within the maximum distance of the vertex. When vertex is selected, ZEMAX is unable to detect if the sampling is sufficient, so some care should be taken to set the sampling high enough for accurate results.
Surface	Selects the surface at which the data is to be evaluated. This is useful for evaluating intermediate images. See “Evaluating results at intermediate surfaces” on page 128.
Wavelength	The wavelength number to be used in the calculation.

Item	Description
Field	The field number for which the calculation should be performed.
Show Diffraction Limit	If checked, the diffraction limited results are computed and displayed. See the discussion below.
Use Huygens PSF	If checked, the more accurate but slower Huygens PSF method is used to compute the PSF. This option should always be used if the image surface is tilted, or if the chief ray is not close to normal to the image surface. The pupil and image sampling, and the other settings specific to the Huygens PSF, are automatically selected by ZEMAX to match the settings and properties used by the FFT PSF for this computation.
Image Sampling	If "Use Huygens PSF" is checked, this control allows a user selectable image sampling grid. The default setting will select a grid that is 2X larger in each direction than the pupil sampling grid.
Image Delta	If "Use Huygens PSF" is checked, this control selects the point spacing on the image grid in micrometers.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.

Discussion:



See the discussion sections of the FFT and Huygens PSF. Those comments also apply to this feature.

The accuracy of the diffraction encircled energy calculation is limited by the magnitude and slope of the OPD error and the sampling density used. If the sampling density is insufficient, ZEMAX will issue an error message indicating that the data is inaccurate. To increase the accuracy, increase the sampling density or decrease the OPD error. If shown, the diffraction limit curve is for the aberration free response at the reference field position (see "Diffraction Limited" in the chapter "Conventions and Definitions").

Geometric

Purpose:

Computes encircled energy using ray-image surface intercepts.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Type	The analysis type option specifies how the encircled energy is calculated; either encircled (radial), X-only, Y-only, or ensquared.
Max Distance	This setting overrides the default scaling. The units are micrometers. To choose the default scaling option, enter zero.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.
Refer To	Select either chief ray, centroid, vertex, or middle as the reference point. The middle is the coordinate upon which the smallest circle enclosing all the ray intercepts is centered.
Surface	Selects the surface at which the data is to be evaluated. This is useful for evaluating intermediate images. See "Evaluating results at intermediate surfaces" on page 128.

Item	Description
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Multiply by Diffraction Limit	If checked, ZEMAX approximates the diffraction encircled energy by scaling the geometric data by the theoretical diffraction limit curve computed for a rotationally symmetric Airy disk. The only way to compute an obscured or asymmetric pupil diffraction limit function would be to perform an exact diffraction calculation, in which case the diffraction encircled energy feature should be used instead. The diffraction limit approximation is only useful for systems with unobscured pupils, reasonably rotationally symmetric images, and modest field angles since the approximation ignores the change in F/# with field.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. See “Surface scattering settings” on page 84.

Discussion:

The X- and Y-only options will compute the fraction of rays which are contained with plus or minus the specified distance from either the chief ray or the image centroid. If a scale of 10 micrometers is shown, then the region enclosed is 20 micrometers across (and infinite in the other direction). The geometric encircled energy is not a good indicator of performance if the system is close to diffraction limited.

Geometric Line/Edge Spread

Purpose:

Computes the geometric response to a line object and an edge object.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Max Radius	The maximum radius setting overrides the default scaling. The units are micrometers. To choose the default scaling option, enter zero.
Surface	Selects the surface at which the data is to be evaluated. This is useful for evaluating intermediate images. See “Evaluating results at intermediate surfaces” on page 128.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number to be used in the calculation.
Type	The type option specifies which data is displayed on the graph; line and edge, line only, or edge only.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.

Discussion:

The line response function (or line spread function, LSF) is the cross section of the intensity pattern of the image of a line object. The edge spread function (ESF) is the cross section of the intensity pattern of the image of an edge (a semi-infinite plane). The X- or Y-orientation refers to the direction of the line or edge. The X-orientation means the line or edge is parallel to the X-axis. The Y-orientation means the line or edge is parallel to

the Y-axis. The zero coordinate is referenced to the chief ray. This is a geometric calculation. See also “FFT Line/Edge Spread” on page 145.

Extended Source

Purpose:

Computes encircled energy using an extended source similar to the geometric image analysis feature.

Settings:

Item	Description
Field Size	This value defines the full width of the square image file in field coordinates, which may be either lens units or degrees, depending upon the current field definition (heights or angles, respectively).
Rays x 1000	This setting determines approximately how many rays will be traced. The number of rays traced is approximately 1000 times the specified value. The reason the number of rays is only approximate is because the distribution of rays over the pixels in the image must be uniform. For example, if there are 1500 pixels in an image file, then at least 1500 rays will be traced, even if a value of 1 is selected. The distribution of rays at each wavelength is in proportion to the wavelength weights.
Type	The analysis type option specifies how the encircled energy is calculated; either encircled (radial), X-only, Y-only, or ensquared. The X-only and Y-only options are sometimes called "enslitted" and correspond to the total fraction of energy contained within an expanding slit. There are also options for X- or Y-distributions, which show the energy distribution in either x or y directions. These latter two options also report the geometric full width at half max. The X- or Y- distributions are the amount of energy falling on a pixel which is narrow in one direction and infinite in the other direction.
Refer To	Select either chief ray, centroid, or surface vertex as the reference point.
Surface	Selects the surface at which the data is to be evaluated. This is useful for evaluating intermediate images. See “Evaluating results at intermediate surfaces” on page 128.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Multiply by Diffraction Limit	If checked, ZEMAX approximates the diffraction encircled energy by scaling the geometric data by the theoretical diffraction limit curve. The diffraction limit curve ZEMAX uses is based upon the unobscured circular pupil. The only way to compute the obscured pupil diffraction limit function would be to perform the exact diffraction calculation, in which case the diffraction encircled energy feature should be used instead. The diffraction limit approximation is only useful for systems with unobscured pupils and modest field angles, since the approximation ignores the change in F/# with field. This option is only available for Type = encircled.
Wavelength	The wavelength number to be used in the calculation.
Field	The image file may be centered on any defined field position. This permits a small target such as a bar chart to be moved to any location in the field of view.
File	The name of the .IMA image file. This file must reside in the <images> folder (see “Folders” on page 68). See the discussion section in the Geometric Image Analysis feature for a full description of the IMA file format.
Max Distance	This setting overrides the default scaling. The units are micrometers. To choose the default scaling option, enter zero.

Item	Description
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See “Comment about vignetting factors” on page 136.

Discussion:

The X- and Y-only options will compute the fraction of rays which are contained with plus or minus the specified distance from either the chief ray or the image centroid. If a scale of 10 micrometers is shown, then the region enclosed is 20 micrometers across (and infinite in the other direction). The geometric encircled energy is not a good indicator of performance if the system is close to diffraction limited.

See the Geometric Image Analysis feature discussion for details about extended source modeling and the IMA file format.

Image Simulation

Image Simulation

Purpose:

This feature simulates the formation of images by convolving a source bitmap file with an array of Point Spread Functions. The effects considered include diffraction, aberrations, distortion, relative illumination, image orientation, and polarization.

Settings:

Item	Description
Source Bitmap Settings	
Input File	The name of the file for the source bitmap. The file may be in BMP, JPG, IMA, or BIM file formats. For a description of the IMA and BIM file formats see “The IMA format” on page 167 and “The BIM format” on page 168. The input files are stored in the <images> folder (see “Folders” on page 68). IMA and BIM files must have more than 1 but no more than 8,000 pixels in each direction, while JPG and BMP files must have more than 1 and no more than 16,000 pixels in each direction.
Field Height	This value defines the full height of the source bitmap in field coordinates, which may be either lens units or degrees, depending upon the current field definition (heights or angles, respectively). If the field type is Real Image Height, the field type is automatically changed to Paraxial Image Height for this analysis. This avoids the problem of Real Image Height masking the distortion in the image. The Field Height applies to the resulting bitmap <u>after</u> any oversampling, guard band, or rotation is applied.
Oversampling	Oversampling increases the pixel resolution of the source bitmap by copying one pixel into 2, 4, or more identical adjacent pixels. The purpose of this feature is to increase the number of pixels per field unit. Oversampling will be applied a factor of 2 at a time until the specified oversampling is achieved, as long as the maximum number of pixels in any direction will not exceed 16,000. This limit only applies to the oversampling feature, and not to the input file. Oversampling is applied to the image before any effects of the optical system are considered.
Flip Source	Flips the source bitmap left-right, top-bottom, or both. The source bitmap is flipped before any effects of the optical system are considered.

Item	Description
Guard Band	This feature increases the pixel resolution of the source bitmap by repeatedly doubling the number of pixels, but the original pixel size of the source bitmap is not changed. This results in a black "guard band" all around the original image. The purpose of this feature is to increase the number of pixels per field unit while adding a region around the desired image. This is particularly useful if the point spread function(s) are large compared to the source bitmap field size. The guard band will be applied a factor of 2 at a time until the specified size is achieved, as long as the maximum number of pixels in any direction will not exceed 16,000. This limit only applies to the guard band feature, and not to the input file. The guard band is applied to the source bitmap before any effects of the optical system are considered.
Rotate Source	Rotates the source bitmap. The rotation is applied to the source bitmap before any effects of the optical system are considered.
Wavelength	<p>If "RGB" is selected, then 3 wavelengths will be defined, 0.656, 0.587, and 0.486 micrometers for red green, and blue, respectively; no matter what the current wavelength definitions are. If "1+2+3" is selected, then wavelengths 1, 2, and 3 as currently defined on the wavelength data box will be used. The red channel of the source bitmap will be used for wavelength 3, green channel for wavelength 2, and the blue channel for wavelength 1. The displayed image will be in RGB format no matter what wavelengths are defined using this option. For selection of a specific single wavelength, such as 6, the B channel image will be used for wavelengths 4, 7, 10, 13, etc. The G channel image will be used for wavelengths 5, 8, 11, 14, etc. The R channel image will be used for wavelengths 6, 9, 12, 15, etc.</p> <p>Wavelength weighting is usually ignored. The color weighting that results in the final image depends only upon the color amplitude in the source bitmap and the relative transmission of the optical system. The exception is when using monochromatic IMA files, the selected wavelength is "1+2+3", and 3 wavelengths are defined. In this case only, the wavelength weighting is used to define the relative color weighting.</p>
Field	The source bitmap may be centered on any defined field position. The resulting image is then centered on this field position.
Convolution Grid Settings	
Pupil Sampling	The grid sampling to use in the pupil space for computing the PSF grid.
Image Sampling	The grid sampling to use in the source bitmap space for computing the PSF grid.
PSF X/Y Points	The number of field points in the X/Y direction at which to compute the PSF. For field points in between these grid points, an interpolated PSF is used.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Apply Fixed Apertures	If checked, all surfaces with optical power that have no aperture defined are modified for this computation to have a circular aperture at the current semi-diameter value. Without this change in the aperture definition, rays may pass the surface beyond the listed semi-diameter, especially if the Field Height exceeds the field of view defined by the field points. This leads to misleading illumination, typically at the edges of the image.
Aberrations	Choose None to ignore aberrations, Geometric to consider only ray aberrations, or Diffraction to use the Huygens PSF to model the aberrations. If Diffraction is selected, the analysis may automatically switch to Geometric if the aberrations are so severe that the diffraction PSF cannot accurately be computed. See the discussion below.
Detector and Display Settings	
Show As	Choose Simulated Image, Source Bitmap to see the input bitmap (including the effects of oversampling, guard band, and rotation), or PSF Grid to see all of the PSF functions computed over the field of view. See the discussion below.

Item	Description
Reference	Selects the reference coordinate for the center of the plot: chief ray, surface vertex, or primary wavelength chief ray. The latter option chooses the primary wavelength chief ray even if some other wavelength is the only one selected.
Suppress Frame	If checked, the frame is not drawn and the entire window is used to display the image.
X/Y Pixels	The number of pixels to use in the detected image. Use 0 for the default value, which is the number of pixels in the source bitmap.
Pixel Size	The detected image pixel size (square). For focal systems, the units are lens units. For afocal systems, the units are cosines. Use 0 for the default value, which is computed from the magnification of the optical system for the center pixel in the source bitmap.
Output File	If a file name ending in the BMP or JPG extension is provided, then the output image will be saved in the specified file and be placed in the <images> folder (see "Folders" on page 68).

Discussion:

The Image Simulation algorithm consists of the following steps to compute the appearance of the image.

- The source bitmap is oversampled, rotated, and a guard band is applied, if these options are selected.
- A "grid" of Point Spread Functions (PSFs) are computed. The grid spans the field size, and describes the aberrations at selected points in the field of view defined by the bitmap and field size settings. The PSF grid also includes the effects of polarization and relative illumination.
- The PSF grid is interpolated for every pixel in the modified source bitmap. At each pixel, the effective PSF is convolved with the modified source bitmap to determine the aberrated bitmap image.
- The resulting image bitmap is then scaled and stretched to account for the detected image pixel size, geometric distortion, and lateral color aberrations.

The most important part of the algorithm is the computation of the PSF grid. The PSF X/Y Points settings determine how many PSFs are computed in each direction over the field of view. There are three different methods for modeling the aberrations using the PSF: Diffraction, Geometric, and None. Diffraction uses the Huygens PSF (for details see "Huygens PSF" on page 146). The Huygens PSF accurately models diffraction effects for almost any optical system, but is slower than the Geometric PSF, which is based upon an integration of the Spot Diagram. The Huygens PSF also cannot be accurately computed if the aberrations are severe (for this feature, the threshold is 20 times the diffraction limit). The Image Simulation algorithm detects when the Huygens PSF cannot be computed, and will automatically switch to the Geometric PSF if required. This switching is done independently for each field and wavelength, and so imaging systems that are diffraction limited in one part of the field but severely aberrated in other parts of the field can still be accurately modeled if Diffraction is selected. If None is selected for the aberration type, the PSF grid is an array of delta functions and aberrations other than lateral color and distortion are not considered.

Relative illumination and (optionally) polarization are considered in the PSF computation. The PSF Grid is smoothly interpolated between computed field points to estimate the PSF at each pixel in the modified source bitmap. The resulting PSF is then convolved with the source bitmap to yield the aberrated image. The more PSF grid points that are used, the more accurate the simulation is, but the longer the computation time.

If the image surface is not a plane, the aberrations and distortion are all computed on the curved image surface, and the Simulated Image is projected onto the XY plane, ignoring the z coordinate of the image surface. The pixels are assumed to be square on the projected plane.

The brightness of the output image is determined by normalizing the image to have the same peak brightness as the input image.

Suggestions for use

The accuracy of the simulation is always ultimately limited by the resolution of the input image. If the resolution of the optical system is sufficient, the discrete pixel nature of the source bitmap may be apparent. This can be

manifested by stair-case like edges in the simulated image. The oversampling feature can reduce these effects at the expense of longer computation time.

If the field of view is large, it is often the case that the PSFs are small compared to the area of a single pixel. In this case, a lot of computation of the PSF may yield what is effectively a delta function, with only a single pixel representing the entire PSF. The computation can be done much faster in these cases by ignoring aberrations, or at least by switching from diffraction to geometric PSF computations.

If the magnification or field of view are small enough to yield an image that is comparable in size to the PSF, or perhaps even smaller than the PSF, then the Guard Band feature is likely required. The reason is that the PSF convolution is done over the source image pixels. If the PSF is larger than the source bitmap in image space, the convolution loses the "spread" outside the original bitmap caused by the large PSF. The guard band adds a black region around the original image which can be used to show the spread due to the PSF.

By far the most important diagnostic is the "PSF Grid" under the "Show As" setting. It is always a good idea to look at the PSF grid, and confirm it shows well sampled data, before computing the Simulated Image.

Comments on using field angles

The Field Height determines the physical size of the source bitmap as seen by the optical system. For example, if a 30 x 30 pixel source bitmap is used, and the Field Height is 2.0 mm (this assumes the fields are defined in terms of object heights), then each pixel represents a 66.67 micrometer region. If the same source bitmap is later used with a system with a 40 degree full field of view (using field angles), then the Field Height can be set to 40 to cover the entire field. Each pixel will now represent a 1.33 degree square. The difficulty in using field angle for defining the object field of view is that field angle units are inherently anamorphic. X-direction angles represent a different subtended angle at a Y angle of 80 degrees than at an a Y angle of 10 degrees. If field angles are being used, and the field of view is fairly wide (more than about 40 degrees in any direction) then great care should be taken in interpreting the results for an extended object. For a precise definition of the field angles ZEMAX uses, see "Field angles and heights" on page 53.

Comments on using paraxial and real image height

Note that if fields are defined by image height, then Field Height determines the size of the object in image space, not object space. The Field Height is always in whatever units the fields are defined in, and so when using image height as a field type the Field Height determines the source bitmap height in image space. The size of the source in object space is then determined by the Field Height divided by the magnification of the lens. When using paraxial image height, the anamorphic magnification of the system (if any) will be masked in the simulated image. When using real image height, all forms of distortion will be masked. For this reason, this feature will use paraxial image height for the field definition even if real image height is selected. However, it is recommended that neither real nor paraxial image height be used. Instead, use object height for finite conjugate systems or field angle for infinite conjugate systems, as these field types unambiguously define the size and orientation of the object in image space.

Comments on JPG and BMP files

The BMP files used as source bitmaps must be standard Windows format uncompressed bitmaps. Most 24 bit color JPG files will work. Some JPG file formats are not compatible with ZEMAX and do not read properly. If a source file does not read properly, try opening the file with the Windows Paint accessory and then save the file to a new name or format.

Geometric Image Analysis



For a feature better suited to simulating images, see "Image Simulation" on page 162.

Purpose:

The geometric image analysis feature has many applications. It can be used to model extended sources, analyze useful resolution, represent the appearance of imaged objects, and provide intuition as to image rotation. Image analysis is also useful for estimating multi-mode fiber coupling efficiency.

Settings:

Item	Description
Field Size	This value defines the full width of the square image file in field coordinates, which may be either lens units or degrees, depending upon the current field definition (heights or angles, respectively). Note that when using "field angle" as the field type the square IMA file is inherently distorted; see the discussion below for details.
Image Size	If "Show" is selected as a spot diagram, this value sets the size of the scale bar which is superimposed on the image. It has no effect on the actual size of the image. The image size is set by the object scale and the magnification and aberrations of the system. The default may not be acceptable to see the desired portion of the image. If "Show" is selected as any other option, this value sets the size of the detector used to capture rays. Rays landing outside of the image size are ignored, and are not included in the total detected rays, which will decrease the computed efficiency.
Parity	The "Even" setting leaves the object as it would appear when viewed looking down the negative Z axis in object space. The parity can be set to "Odd" which reverses the object from top to bottom.
Rotation	The rotation can be set to any angle in degrees. The algorithm actually rotates the object before tracing the rays, so this feature can be used to switch from tangential to sagittal orientation of bar targets, for example.
Rays x 1000	This setting determines approximately how many rays will be traced. The number of rays traced is approximately 1000 times the specified value. The reason the number of rays is only approximate is because the distribution of rays over the pixels in the image must be uniform. For example, if there are 1500 pixels in an image file, then at least 1500 rays will be traced, even if a value of 1 is selected. The distribution of rays at each wavelength is in proportion to the wavelength weights.
Show	Choose surface plot, contour map, grey scale, false color map, spot diagram, cross x, or cross y as the display option.
Source	The source may be uniform or Lambertian. The uniform setting weights rays equally. Lambertian weights rays by the cosine of the angle the ray makes with local normal of the object surface.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. See "Surface scattering settings" on page 84.
Wavelength	The wavelength number to be used in the calculation.
Field	The image file may be centered on any defined field position. This permits a small target such as a bar chart to be moved to any location in the field of view. The resulting image is then centered on the chief ray coordinate of this field position.
File	The name of the .IMA or .BIM image file. This file must reside in the <images> folder (see "Folders" on page 68). See the discussion section for a full description of the IMA and BIM file formats.
Edit IMA File	Pressing this button will invoke the Windows Notepad editor which allows modification of the currently selected IMA file. This button is disabled if the file type is BIM.
Surface	The surface number at which to evaluate the rays. The default is the image surface. Other surfaces may be selected, for example, to visualize the beam footprint on an optical surface.
# Pixels	The number of pixels across the width of the selected image size. This value is not used if "spot diagram" is the method of displaying the image data.

Item	Description
NA	The numerical aperture (NA) cut-off. If zero, this feature is ignored. If a number greater than 0 is entered, then all rays with a numerical aperture greater than the specified number are ignored.
Total Watts	The total power in watts radiated by the source into the entrance pupil of the optical system. This flux is then used to normalize the detected power according to the relative pixel values and the total efficiency.
Plot Scale	For false color and gray scale plots, this value defines the maximum scale range.
Use Symbols	If checked, this option will draw different symbols rather than dots for each wavelength. This helps distinguish the various wavelengths. This value is only used if "spot diagram" is the method of displaying the image data.
Configuration	Select "All" to show data for all configurations at once, or select the one configuration to show, or select "Current" to show the active configuration. If "All" is selected, and both focal and afocal mode configurations are defined, the Configuration setting will automatically be reset to "Current". If "All" is selected all coordinates are referenced to the chief ray image coordinate in configuration 1, at either the primary wavelength (if all wavelengths are selected) or the selected wavelength.
Reference	Selects the reference coordinate for the center of the plot: chief ray, surface vertex, or primary wavelength chief ray. The latter option chooses the primary wavelength chief ray even if some other wavelength is the only one selected.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 136.
Delete Vignetted	If checked, rays are not drawn if they will be vignetted by any surface.
Use Pixel Interpolation	If checked, ray energy is distributed among adjacent pixels according to the distance from where the ray lands to the center of the pixel. This yields smoother, more accurate results for images where the pixels are small compared to the spot size. If unchecked, all the energy for a ray is placed only in the pixel struck. Pixel Interpolation should not be used if the pixels are large compared to the spot size. This control has no affect if "Show" is set to "spot diagram".
Save As BIM File	If a file name ending in the BIM extension is provided, and "Show" is not set to "Spot Diagram", then the output image will be saved in the specified file and be placed in the <images> folder (see "Folders" on page 68).

Discussion:

ZEMAX supports three different file formats. Two of these formats end in the IMA extension, one in the BIM extension.

The IMA format

There are two different IMA file formats, one text and one binary. Whichever file format is used, the file must end in the extension IMA. ZEMAX will distinguish between the two types of file formats automatically.

The text image file is a text file which ends in the extension .IMA. At the top of the file is one number which indicates the size of the file in pixels. The remaining rows and columns contain the pixel data, with one character to each pixel. All IMA files must be square, with n x n pixels defined. For example, a 7 x 7 representation of the letter "F" could be described by the following IMA file:

```
7
0111110
0100000
0100000
0111100
0100000
0100000
0100000
```

Note that the single entry "7" starts the file, and it is followed by a carriage return. Then there are 7 rows of 7 columns each, and each row is terminated with a carriage return. The columns are not separated by a space or any other character. The image file must be square. ZEMAX will attempt to allocate enough memory to hold the image file and will report an error if there is not enough memory.

The "intensity" at each pixel can be any digit between 0 and 9. The number of rays each pixel will generate is proportional to this value. Pixels with a value of 0 do not radiate any rays.

The binary IMA file format is more complicated than the text format, and binary IMA files cannot be edited with a text editor. However, the binary IMA files are dramatically more powerful. Each pixel in the binary IMA file is represented by an unsigned byte, which means there are 256 "gray-scale" levels of intensity. Furthermore, each wavelength can be assigned a separate pixel map. Therefore, very realistic photograph like extended sources can be modeled.

The binary IMA file format requires 3 16-bit header values. The first 16-bit value is a signed integer that must be equal to zero. The second 16-bit signed integer is the width of the pixel map in pixels, which can be any number from 1 to 8000. The third 16-bit signed integer is the number of pixel maps, which correspond to the number of colors (or wavelengths) represented in the file.

For example, a 3-color binary pixel map of a 50 by 50 image would have 6 bytes of header (0, 50, and 3), followed by 2500 bytes for color 1, then 2500 bytes for color 2, then 2500 bytes for color 3, for a total of 7506 bytes. The data for each color is stored by columns for each row (the column index changes faster than the row index).

The BIM format

The drawback to the IMA format is that a maximum of 256 grey scale levels are supported. The BIM format is a binary double precision floating point file format which effectively makes the number of grey scales many trillions. The BIM format consists of the following binary values:

1 32 bit integer representing the number of x pixels, nx.

1 32 bit integer representing the number of y pixels, ny.

followed by $n \times n$ 64 bit double precision floating point values representing the relative intensity. The first pixel is the bottom left corner, and the remaining pixels are listed by rows along x.

Currently, the nx and ny values must be identical or an error message will be issued.

Comments on using field angles

The comments regarding field angles for the Image Simulation also apply to this feature. For a full discussion, see "Comments on using field angles" on page 165.

Comments on using paraxial and real image height

The comments regarding paraxial and real image height for the Image Simulation also apply to this feature. For a full discussion, see "Comments on using paraxial and real image height" on page 165.

How rays are chosen for analysis

The rays generated by each pixel are chosen randomly from coordinates within the pixel cell. The entrance pupil coordinates are also chosen randomly for each ray. The distribution of rays is uniform over the pixel and over the circular paraxial entrance pupil (if ray aiming is used, then there may be some pupil distortion). For the text IMA files, the number of rays generated by each pixel is equal to the pixel intensity times the number of wavelengths times the ray density. The wavelength used for each ray is selected randomly in proportion to the wavelength weights. For the binary IMA files, the number of rays generated from each pixel is proportional to the ray density times the fractional intensity relative to 256.

By separating the form of the object from the scale, the same image file can be used for many applications. For example, the sample image file "letterf.ima" contains a 7x7 grid of pixels defining the capital letter F. The object scale can be set to 1 mm, then 0.1 mm, then 0.01 mm to get a feel for how small a character F the optical system can resolve, without the need to change the IMA file.

Note that if fields are defined by image height, then field size determines the size of the object in image space, not object space. The field size is always in whatever units the fields are defined in, and so for image height the

field size determines image height. The size of the object is then determined by the field size divided by the magnification of the lens.

The choice of field position also permits great flexibility in analyzing image quality. For example, the letter F image file can be tested at several field points to see if the resolution is strongly affected by field aberrations. The object scale is set to the height of the letter, but the image will be centered about the chief ray intercept of the selected field point.

The source is by default a uniform radiator of rays. Uniform here means uniform in the entrance pupil. All rays generated fall within the entrance pupil, and they are all weighted equally. Since ray wavelengths are selected randomly in proportion to the wavelength weights, no explicit wavelength weighting is required. The uniform setting is usually preferred for large object distance systems with small fields of view. The source may also be defined to be Lambertian, which weights all rays by a cosine factor. Once the input ray weight is determined, any pupil apodization (see "Apodization Type" on page 102) that is defined will further modify the input ray weight.

The percent efficiency is defined by

$$\%E = \frac{\sum W_i}{\sum W_j}$$

where the sum i is over all rays which were unvignetted, and the sum over j is over all rays which were launched. The efficiency calculation considers vignetting, source distribution, wavelength weights, and reflection and transmission losses in the optical system if the "Use Polarization" checkbox is selected. If "spot diagram" is selected for the "show" option, then the percent efficiency includes all unvignetted rays. The other display options vignette rays that are beyond the extent of the detector size. Therefore, the percent efficiency will be different between spot diagram displays and other displays if rays fall outside the region of the detector.

Calculating efficiency of multi-mode fibers

ZEMAX has an algorithm for accurately computing fiber coupling into single-mode fibers; for details see "Fiber Coupling Efficiency" on page 206. Also see "Computing Fiber Coupling" on page 645.

To estimate the coupling efficiency for multi-mode fibers, a geometric approach may be used. Place a circular aperture at or just before the image surface with the appropriate maximum radial aperture representing the core size. Then set the NA (see the table above) to the maximum acceptance NA of the fiber. The percent efficiency will then be calculated by summing all the rays that pass the core aperture within the specified NA. The NA of a typical multi-mode fiber with an inner core of index n_i and an outer cladding of index n_o is given by

$$NA = \sqrt{n_i^2 - n_o^2}.$$

Text output

Selecting the "Text" option on the image analysis window menu bar will generate and display a text file listing the ray data. If the "Show" option is set to "Spot Diagram", the file will have 9 columns. The first column is the sequential ray number. The second and third columns are the x and y field coordinates (either degrees or object height). The fourth and fifth columns are the normalized pupil coordinates, Px and Py. The sixth column is the integer wavelength number. The seventh column is the weight of the ray, which depends upon the source properties. The eighth and ninth columns are the image coordinates in lens units, relative to the reference ray.

If the "Show" option is not set to "Spot Diagram", then the text display will list the weighted ray count in each pixel. The pixel listing starts at the -x, -y (bottom left) corner pixel, and proceeds up the y columns. Use the "Escape" key to terminate a lengthy image analysis computation.

Geometric Bitmap Image Analysis



The feature "Image Simulation" on page 162 is generally superior to this feature, unless the simulated image is desired on an intermediate surface away from the image.

Purpose:

This feature creates an RGB color image based upon ray tracing data using an RGB bitmap file as the source. This feature has many applications. It can be used to model extended sources, analyze useful resolution, display distortion, represent the appearance of imaged objects, provide intuition as to image rotation, display beam footprints, and show surface plots of illumination on any surface, to name just a few.

This feature is based strictly upon geometrical ray tracing. The bitmap image analysis feature uses standard Windows BMP and JPG files as the source image, see the discussion for details.

Settings:

Item	Description
Field Y Size	This value defines the full y direction size of the source bitmap in field coordinates, which may be either lens units or degrees, depending upon the current field definition (heights or angles, respectively).
Parity	The "Even" setting leaves the object as it would appear when viewed looking down the negative Z axis in object space. The parity can be set to "Odd" which reverses the object from top to bottom.
Rotation	The rotation can be set to any angle in degrees. The algorithm actually rotates the object before tracing the rays, so this feature can be used to switch from tangential to sagittal orientation of bar targets, for example.
Rays/Pixel	This setting determines how many rays per source image pixel per color channel will be traced. Selecting 10 rays/pixel on a 640 x 480 3-color image will in general trace 9,216,000 rays.
X-Pixels	The number of pixels across the X direction on the detector.
Y-Pixels	The number of pixels across the Y direction on the detector.
X-Pixel Size	The size in lens units of each detector pixel measured in the X direction.
Y-Pixel Size	The size in lens units of each detector pixel measured in the Y direction.
Source	The source may be uniform or Lambertian. The uniform setting weights rays equally. Lambertian weights rays by the cosine of the angle the ray makes with local normal of the object surface.
Normalize	The image may be normalized by either the peak pixel intensity or by the average pixel intensity. Using the peak pixel intensity scales all pixels equally so the pixel with the highest intensity determines the overall brightness of the image. This method generally produces fairly dark images if the signal-to-noise ratio is low. Using the average pixel intensity scales all pixels equally until the average brightness equals the original bitmap image brightness. The disadvantage to this method is that some pixels become oversaturated, and are thus clipped at the maximum brightness that can be displayed.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Grey Scale	If checked, the RGB intensities will be averaged at each detector pixel to yield a grey scale detected image. Rays will still be traced according to the relative RGB intensities of the source bitmap, but all color information will be lost when the detected image is displayed.

Item	Description
Wavelength	If "RGB" is selected, then 3 wavelengths will be defined, 0.656, 0.587, and 0.486 micrometers for red, green, and blue, respectively; no matter what the current wavelength definitions are. If "1+2+3" is selected, then wavelengths 1, 2, and 3 as currently defined on the wavelength data box will be used. The red channel of the source bitmap will be used for wavelength 3, green channel for wavelength 2, and the blue channel for wavelength 1. The displayed image will be in RGB format no matter what wavelengths are defined using this option. For selection of a specific wavelength, such as 1, 2, 3, etc... the B, G, or R channel image will be used; for wavelengths higher than 3 the B channel is always used, for wavelengths selected which are not defined the highest defined wavelength will be used.
Field	The source image may be centered on any defined field position. This permits a small target such as a bar chart to be moved to any location in the field of view. The resulting image is then centered on the chief ray coordinate of this field position.
Input	The name of the .BMP or .JPG source image file. This file must reside in the <images> folder (see "Folders" on page 68).
Surface	The surface number at which to evaluate the rays.
Show Source Bitmap	If checked, then the source bitmap will be drawn and no rays will be traced. The number of rays, pixels, and pixel sizes are ignored when drawing the source bitmap. This feature is used to check the appearance of the source bitmap and to verify ZEMAX properly reads the bitmap file.
Output	The name of the .BMP or .JPG to write the detected bitmap to. The detected bitmap size is determined by the number of x and y pixels defined; but the pixels size must be the same in x and y for the aspect ratio to be correct in the output bitmap file. The file name must end in either a BMP or JPG extension, with no path name supplied. This file will be created or overwritten without warning and will be placed in the <images> folder (see "Folders" on page 68).
Reference	Selects the reference coordinate for the center of the plot, either the chief ray or the surface vertex.
Delete Vignetted	If checked, rays are not drawn if they will be vignetted by any surface.
Suppress Frame	If checked, the frame is not drawn and the entire window is used to display the image.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 136.
Apply Fixed Apertures	If checked, all surfaces with optical power that have no aperture defined are modified to have a circular aperture at the current semi-diameter value. Without this change in the aperture definition, rays may pass the surface beyond the listed semi-diameter, especially if the Field Y Size exceeds the field of view defined by the field points. This leads to misleading illumination, typically at the edges of the image.

Discussion:

See also "Geometric Image Analysis" on page 165, as this feature is very similar. See also "Comments on JPG and BMP files" on page 165.

Comments on field units

The Field Y Height determines the physical size of the source file as seen by the optical system. For example, if a 50H x 100W pixel source file is used, and the Field Y Height is 2.0 mm (this assumes the fields are defined in terms of object or image heights), then each pixel represents a 0.040mm x 0.040mm region, and the source bitmap covers a 2.0mm high x 4.0 mm wide area.

Note that if fields are defined by image height, then Field Y Height determines the size of the object in image space, not object space. The Field Y Height is always in whatever units the fields are defined in, and so when using image height as a field type the Field Y Height determines the source bitmap height in image space. The size of the source in object space is then determined by the Field Y Height divided by the magnification of the

lens. When using paraxial image height, the anamorphic magnification of the system (if any) will be masked in the simulated image. When using real image height, all forms of distortion will be masked. For this reason, this feature will automatically switch to paraxial image height for the field definition if real image height is selected. However, it is recommended that neither real nor paraxial image height be used. Instead, use object height or field angle, as these field types unambiguously define the size and orientation of the object in image space.

Selection and tracing of rays

The rays generated by each pixel are chosen randomly from coordinates within the source pixel cell. The entrance pupil coordinates are also chosen randomly for each ray. The distribution of rays is uniform over the pixel, the ray distribution over the entrance pupil is uniform unless pupil apodization is defined.

Once a ray is generated, it is traced through to the selected surface. If the ray is vignetted or an error occurs, the ray is ignored. Otherwise, the pixel on the receiving detector that the ray struck is determined, and the intensity of the ray is added to the detector's bin count in the appropriate color channel.

After all the rays are traced, an RGB image is created from the normalized counts in each detector pixel.

The percent efficiency is defined by

$$\%E = \frac{\sum W_i}{\sum W_j}$$

where the sum i is over all rays which were unvignetted, and the sum over j is over all rays which were launched. The efficiency calculation also includes consideration of reflection and transmission losses in the optical system if the "Use Polarization" checkbox is selected, as well as ray errors. Rays that are beyond the extent of the detector are considered to be vignetted.

This feature automatically divides the computation over all available processors for maximum execution speed. Any symmetry in the lens is also used to speed up the ray tracing computation. Use the "Escape" key to terminate a lengthy image analysis computation.

Partially Coherent Image Analysis

Purpose:

This feature considers diffraction and aberrations of the optical system, as well as partial coherence of the illumination, to compute the image appearance. This method accounts for the finite pass band and other diffraction related effects real optical systems have on image formation. This feature can compute the coherent, incoherent, or partially coherent diffraction image of scene defined by a bitmap image. For purely incoherent analysis the Image Simulation feature described on page 162 is generally superior.

The image analysis feature uses IMA/BIM files to describe the object to be imaged. See "The IMA format" on page 167 and "The BIM format" on page 168.

Settings:

Item	Description
File Size	The full width in lens units of the region defined by the file measured in image space. Note IMA and BIM files are always square. This feature also supports a subset of ZBF files, see the discussion for details.
Oversampling	Sets the factor by which the file pixels are oversampled. This increases the effective resolution of the file without the need to define a new file. If the original file has an odd number of pixels, oversampling will make the number of pixels even because the oversampling values are all even.

Item	Description
Zero Padding	<p>Determines the actual size of the region to compute the diffraction image of by adding zero intensity values around the file pixels. This increases the size of the displayed diffraction image without changing the size of the unaberrated image; this allows study of the energy diffracted well away from the perfect image location.</p> <p>To keep the file data centered, zero padding will yield an odd number of pixels if the original IMA file has an odd number of pixels. The number of pixels will be $2*n-1$, where n is the original number of pixels. If the original file has an odd number of pixels, the modified image will have $2*n$ pixels.</p> <p>Zero Padding only affects IMA and ZBF, and not BIM format files.</p>
OTF Sampling	The grid size of the sampling in the pupil; larger grids yield a more accurate representation of the system OTF. This has no affect on the size of the diffracted image, just the accuracy of the predicted frequency response.
Show As	Choose surface plot, contour map, grey scale, false color, cross section x, or cross section y as the display option. For information on cross sections, see "Computing MTF using Cross Section X/Y" on page 177.
Data Type	<p>Choose either the incoherent image, coherent image, raw image, incoherent transfer function, coherent transfer function, or the transform of the raw image.</p> <p>For partial coherence computations, choose partially coherent image (either mostly coherent or mostly incoherent), or partially coherent PSF or gamma. For a full discussion on partially coherent image computations, see "Computing partially coherent images" on page 176.</p> <p>See the discussion below for information about the coherent transfer function and the limitations of this computation.</p>
Gamma	The form of the complex degree of coherence function.
Alpha	The scaling parameter used in the complex degree of coherence function.
Fraction	The fraction of encircled energy in the PSF and/or Gamma functions to consider when computing partially coherent images.
Diffraction Limited	If checked, aberrations are ignored. Apertures are still considered.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number indicates for which field position the optical transfer function is computed.
File	The name of the IMA, BIM, or ZBF image file. The IMA and BIM files must reside in the <images> folder, while ZBF files must reside in the <pop> (see "Folders" on page 68). See the discussion section for a full description of the supported file formats.
Edit IMA File	Pressing this button will invoke the Windows Notepad editor which allows modification of the currently selected IMA file. BIM and ZBF files may not be edited directly.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see "The Contour Format String" on page 151. The contours are defined in units of relative intensity.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Resample Image	If checked, the final image will be resampled to simulate the image recorded by a detector with finite resolution. The controls below are used to define the resampling. This feature is only available for coherent, incoherent, raw, and partially coherent spatial domain images.
Number X/Y	The number of X and Y direction pixels to use for resampling the image.

Item	Description
Delta X/Y	The width/height of the pixels used to resample the image. In focal mode, the Delta values are measured in Lens Units. In afocal mode, the Delta values are measured in Afocal Mode Units.
Decenter X/Y	The X/Y decenter of the resampled detector relative to the final image. The units are resampled delta pixels. For example, a Decenter X of +0.5 will decenter the resampling detector one half the Delta X value to the right relative to the final image.
Output File	If the input file is a ZBF file, and either a fully coherent or incoherent image is selected for the Data Type, then the resulting complex amplitude may be saved in a ZBF file of the name specified in this control. If the Output File name is left blank, no file is saved. No path name should be provided, as the file will be saved in the <pop> folder (see "Folders" on page 68). To save ZBF files, the zero pad must be set to yield a total number of pixels that is a power of 2, as this is required for the ZBF file format.

Discussion:

This feature can compute complex diffraction image properties from extended sources. The method involved in the computation is based upon Fourier Optics, which is described in clear and insightful detail in "Introduction to Fourier Optics" by Joseph Goodman, McGraw-Hill 1968. A more detailed treatment of coherent imaging theory may be found in "Linear Systems, Fourier Transforms, and Optics" by Jack Gaskill, John Wiley 1978. See those references for more information on coherent vs. incoherent imaging and other Fourier optics theory. There are several important assumptions in the method that must be understood by the user before this feature may be used to draw important conclusions.

The file size parameter determines the full width of the input file in *the image space of the optical system*. For details on the IMA format see "The IMA format" on page 167. For details on the BIM file format see "The BIM format" on page 168. For details on the ZBF file format see "ZEMAX Beam File (ZBF) binary format" on page 640 and "ZEMAX Beam File (ZBF) text format" on page 641.

Note this feature is different from the geometric image analysis feature, where the file size defines the size and shape in "field" space, which may be either object or image space. For this feature, the input file defines the ideal image in image space. The data values in the IMA or BIM file are interpreted to be in units of relative irradiance, not amplitude. The ZBF file is in units of relative complex amplitude.

Comment about using the IMA file format

Although the IMA file is convenient for defining simple shapes, the resolution of files created by hand is generally too low to see sufficient detail in the diffracted image. The oversampling option remedies this problem by increasing the resolution; the number of pixels is increased, and the data from each pixel is replicated as required to yield the same shape at higher resolution. With an oversampling of 2X; the letter "F" file IMA image becomes a 14 x 14 grid:

```
00111111111100
00111111111100
00110000000000
00110000000000
00110000000000
00110000000000
00110000000000
00111111110000
00111111110000
00110000000000
00110000000000
00110000000000
00110000000000
00110000000000
00110000000000
00110000000000
```

Note the shape is the same, there is just twice the sampling in each direction. The width of the image is unaffected; the pixels are half the width and there are twice as many in each direction now.

Because the effects of diffraction tend to blur and extend the ideal image, it is desirable to increase the size of the displayed area beyond the limits of the IMA file. This can be done with the zero padding option. This option

increases the IMA file size by adding zero intensity values around the defined IMA file pixels. The letter "F" file would look like this with a zero padding of 2X:

```
00000000000000
00000000000000
00000000000000
00001111100000
00001000000000
00001000000000
00001111000000
00001000000000
00001000000000
00001000000000
00001000000000
00000000000000
00000000000000
00000000000000
```

The size and shape of the image is the same, but additional area has been defined that some energy may now diffract into. Note the width of the new image has increased by the zero padding factor, but the width of the "F" part of the image is the same. If the original file has an odd number of pixels (as this sample letter F does) then the number of pixels after zero padding will remain odd to preserve the notion of a "center" pixel. The number of total pixels is thus $2*n-1$ if n is odd, or $2*n$ if n is even, where n is the number of original pixels. Zero padding is not available when using the binary IMA file format.

The oversampling and zero padding may be used together, however the array sizes quickly become quite large. For an oversampling of 8X, a zero padding of 4X, and an original pixel size of 12 x 12, the resulting array becomes 384 x 384, which ZEMAX internally further zero pads up to 512 x 512 for transform purposes.

Comment about using the ZBF file format

When using the ZBF file format for this feature, note that only the complex amplitude of the unpolarized beam is considered. None of the other data in the ZBF file is relevant, including the pilot beam data, wavelength, and pixel size. If the ZBF file has an unequal number of pixels in each direction, the file size is the larger of the two dimensions, and zero-padding is added to make the file square. The size of the pixels as defined in the ZBF file are not used. The ZBF file format is only used here for convenience in specifying complex amplitude, and most of the features specific to physical optics analysis are not used.

When performing an incoherent image analysis using a ZBF file as input, the intensity of the data is computed by summing the squares of the real and imaginary values at each point. The phase information is thus lost, as is correct for incoherent image formation. The resulting image is a real valued intensity image. When performing coherent image analysis, the complex amplitude as defined in the ZBF file is used.

The optical transfer function algorithm

Once the input image is defined, the image is transformed into frequency space, multiplied by the OTF, and transformed back into position space. The primary assumption made by this implementation of the Fourier method is that the OTF does not change over the extent of the image region (to avoid this assumption see "Image Simulation" on page 162). This means the field of view defined by the size of the image is small enough so that the OTF is the same over all points in the image. The user must take care to be sure the image region defined is small with respect to the rate of change of the field aberrations. ZEMAX computes the OTF for the field point selected, and assumes this OTF is valid over the entire region covered by the image.

Because of this assumption, distortion will not be visible in the predicted image, since only the variation in OTF over the field will introduce distortion. To see the effects of distortion or other "large field" effects, see "Image Simulation" on page 162.

Note that this feature is good at computing detailed image data for small images under partially coherent light, while the Image Simulation feature is good at computing image data for larger scale images or incoherent light.

Comment about the coherent optical transfer function

The other assumption this calculation makes is in the method used to compute the coherent optical transfer function. The coherent optical transfer function is assumed to be the complex pupil function:

$$H(f_x, f_y) = P(\lambda d_i f_x, \lambda d_i f_y) e^{iKW(\lambda d_i f_x, \lambda d_i f_y)}$$

where H is the complex OTF, d_i is the pupil to image distance, f_x and f_y are the spatial frequencies, P is the pupil function (which determines the relative transmission over the pupil, and is zero outside the pupil), and W is the wavefront aberration function. This approximation is valid for a wide range of optical systems. For more information on the development of coherent imaging theory, see the references by Gaskill or Goodman.

This feature also accounts for any vignetting at the image surface due to surface apertures (apertures at other surfaces are accounted for via the effects of those apertures on the optical transfer function). After the diffraction image is computed, any energy that lies outside of the aperture defined on the image surface is eliminated. The fraction of energy within the unvignetted portion of the image surface aperture is reported on the text version of the analysis, if the image surface has a defined aperture.

Computing partially coherent images



This feature is only available in the EE edition of ZEMAX.

The prior discussion applies to fully incoherent or fully coherent illumination. Partially coherent imaging is more complex. The method ZEMAX uses to model partially coherent illumination is described in B. Saleh and M. Rabbani, "Simulation of partially coherent imagery in the space and frequency domains and by modal expansion", Applied Optics Vol. **21**, No. 15, p2770 (1982).

The diffraction image for the partially coherent case can be described as:

$$I_i(\vec{r}) = \iint \sqrt{I_o(\vec{r}_1)} \sqrt{I_o(\vec{r}_2)} \gamma(\vec{r}_1 - \vec{r}_2) h^*(\vec{r} - \vec{r}_1) h(\vec{r} - \vec{r}_2) d\vec{r}_1 d\vec{r}_2,$$

where $I_i(\vec{r})$ is the irradiance of the image, $I_o(\vec{r})$ is the irradiance of the unaberrated, undiffracted image of the object in image space, $\gamma(\vec{r})$ is the complex degree of coherence in image space, $h(\vec{r})$ is the complex point spread function in image space, and the symbol * indicates the complex conjugate. The position vector \vec{r} has components x and y in image space. The function $h(\vec{r})$ is assumed to be invariant over the entire region of the image computed, which is the same assumption made in the fully coherent and incoherent cases.

The key difference between the fully coherent or fully incoherent illumination and partially coherent illumination is contained with the function $\gamma(\vec{r})$, which defines the degree of coherence between any two points in the object illumination. The function $\gamma(\vec{r})$ is symmetric, with a magnitude between 0 and 1 for all \vec{r} . For fully incoherent imaging, $\gamma(\vec{r}) = \delta(\vec{r})$, where $\delta(\vec{r})$ is the Dirac delta function. For fully coherent imaging, $\gamma(\vec{r}) = 1$ for all \vec{r} . In either of these two extremes, the general imaging equation above may be substantially simplified. For the partially coherent case, no easy simplification is possible, and for this reason partially coherent computations are generally much slower than fully coherent or incoherent cases.

To speed up the partially coherent image computation, two distinct algorithms may be employed to solve the imaging integral above. The two algorithms are user selectable. If the Mostly Incoherent method is selected, the algorithm assumes the function $\gamma(\vec{r})$ has a spatial extent that is small compared to the width of the image being computed. If the Mostly Coherent method is selected, the algorithm assumes the spatial extent of the function $\gamma(\vec{r})$ is large, or at least comparable to the width of the image. The exact computation speed depends upon the size and structure of both $h(\vec{r})$ and $\gamma(\vec{r})$, as well as the function $I_o(\vec{r})$. A good rule of thumb is that if $\gamma(\vec{r})$ is wider than about 20% of the image, the Mostly Coherent method is faster. Some experimenting may be required

to determine the faster method for any specific case. Both of these algorithms are multi-threaded and will automatically use parallel execution on any number of available CPU's for maximum computation speed.

Two different functional forms for $\gamma(\vec{r})$ are currently supported. The Gaussian form is:

$$\gamma(\vec{r}) = e^{-\left(\frac{\vec{r} \cdot \vec{r}}{\alpha^2}\right)}.$$

The Sinc form is:

$$\gamma(\vec{r}) = \text{SINC}\left(\frac{x}{\alpha}\right)\text{SINC}\left(\frac{y}{\alpha}\right), \text{ where } \text{SINC}(x) = \frac{\sin \pi x}{\pi x}.$$

Both the Mostly Coherent and Mostly Incoherent methods accelerate the computation by assuming the various functions have a finite width that has a significant effect on the final image. This finite width is defined by the function and by the user defined parameter called "Fraction". This is the fraction of total energy within the functions $h(\vec{r})$ and $\gamma(\vec{r})$ that is considered significant. ZEMAX will determine the spatial dimension of significance of these functions using the defined fraction. The closer the fraction is to unity, the slower the computations will be. Generally, values between 0.90 and 0.96 are recommended. Values larger than 0.96 may substantially increase the computation with little effect on the final image.

Two additional "Data Type" options are available for use with partially coherent computations. The "Partially Coherent Test: PSF" will display the diffraction PSF (this is the intensity of the function $h(\vec{r})$) used for the partially coherent image computation. The various sampling and image width settings affect the sampling used for the PSF. Before any partially coherent image can be considered accurate, the PSF should be computed to confirm adequate sampling of the PSF. At least 10 points should be displayed over the width of the non-zero portion of the PSF for adequate sampling. The data type "Partially Coherent Test: Gamma" displays the magnitude of the function $\gamma(\vec{r})$.

Computing MTF using Cross Section X/Y

The "Show As" option supports "Cross Section X" and "Cross Section Y" options. These options compute the same complete diffraction image, but then plot only a single cross section across either the columns or rows of the image, for X and Y, respectively. The intended purpose of these cross sections is to visualize the effective MTF of the image of a bar target. The MTF can be estimated by determining the maximum and minimum relative intensity across the cross section. To reduce the effect of edges, the analysis parameters should be set to provide at least 5 well defined peaks across the cross section. The MTF is computed by looking for the minimum and maximum intensity at all points between the second and second-to-last local peaks in the intensity data. By considering only data within these two peaks, the effects of the edges is somewhat reduced. The estimated MTF is then given by the usual computation of $(I_{\max} - I_{\min}) / (I_{\max} + I_{\min})$. Finally, note that if a bar target is used the resulting MTF is the square-wave, not sine-wave modulation.

Extended Diffraction Image Analysis



The feature "Image Simulation" on page 162 is generally superior to this feature.

Purpose:

The extended diffraction image analysis feature is similar to "Partially Coherent Image Analysis" on page 172, except the Optical Transfer Function (OTF) may vary over the field of view of the image and the illumination must be fully coherent or incoherent. This feature uses IMA/BIM files to describe the object to be imaged. See "The IMA format" on page 167 and "The BIM format" on page 168.

Settings:

Item	Description
File Size	The full width in lens units of the region defined by the IMA file. Note IMA files are always square. This control does not set the size of the resulting image; that is defined by the Display Size.
Oversampling	Sets the factor by which the IMA file pixels are oversampled. This increases the effective resolution of the IMA file without the need to define a new IMA file. If the original IMA file has an odd number of pixels, oversampling will make the number of pixels even because the oversampling values are all even.
Display Size	This defines the width of the displayed image in lens units.
OTF Sampling	The grid size of the sampling in the pupil; larger grids yield a more accurate representation of the system OTF. Larger grids will also yield a larger maximum display size (see the discussion below).
OTF Grid	The grid size of the OTF computation. A denser OTF grid yields a more accurate computation of the variation of OTF over the field of view of the image at the expense of more memory usage and slower computation.
Resolution	This multiplier yields more points in the final image while keeping the display size fixed, at the expense of more memory usage.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.
Data Type	Choose either the incoherent image or coherent image. See “Comment about the coherent optical transfer function” on page 176 for information about the coherent transfer function and the limitations of this computation.
Diffraction Limited	If checked, aberrations are ignored. Apertures are still considered.
Use Delta Functions	If checked, each pixel in the IMA file will be assumed to represent a delta function. This is useful for checking the imaging of point sources like stars. If unchecked, the whole pixel is assumed to be a luminous square area. If Use Delta Functions is checked, it is very difficult to visually see the effects of aliasing in the computation. For this reason, it is highly recommended that this option be <u>unchecked</u> while setting up the file size, display size, and sampling parameters. Once the proper values for these settings are determined, then Use Delta Functions may be checked if desired.
File	The name of the image file. This file must reside in the <images> folder (see “Folders” on page 68). See “The IMA format” on page 167 and “The BIM format” on page 168 for a description of file formats.
Edit IMA File	Pressing this button will invoke the Windows Notepad editor which allows modification of the currently selected IMA file. BIM files cannot be edited this way.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number indicates for which field position the image will be centered on.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see “The Contour Format String” on page 151. The contours are defined in units of relative intensity.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Consider Distortion	If checked, the real vs. paraxial ray distortion will be considered in forming the image appearance.

Item	Description
Output File	If a name is provided, the complex amplitude of the resulting image will be saved in a ZBF file with the specified name. The file will be saved in the <pop> folder (see "Folders" on page 68). See "ZEMAX Beam File (ZBF) binary format" on page 640.
Use Relative Illumination	If checked, then the computation described in "Relative Illumination" on page 187 is used to weight the OTF over the field of view to accurately account for the effects of exit pupil radiance and solid angle.

Discussion:

This feature can compute complex diffraction image properties from extended sources while accounting for the variation in the optical transfer function (OTF) over the field of view. Most of the discussion in "Partially Coherent Image Analysis" on page 172 also applies to this feature. The differences are described below.

Diffraction image formation can be thought of as a filtering or as a convolution process. Suppose the ideal, unaberrated, undiffracted image is described by a function "A" which describes image amplitude as a function of spatial coordinates in the image space of an optical system. Convoluting this function with the system PSF (see "FFT PSF" on page 141) here denoted by "P" yields the final image "I":

$$I(x, y) = A(x, y) \otimes P(x, y),$$

where the notation $A \otimes P$ denotes the convolution of A and P. Taking the Fourier transform of this equation yields the spatial filtering perspective of the image forming process:

$$i(f_x, f_y) = a(f_x, f_y) \times o(f_x, f_y),$$

where i, a, and o are the transforms of I, A, and P into the spatial frequency domain. The function o is called the optical transfer function (OTF); which acts as a filter scaling the amplitude and phase of the spatial frequency components of the image.

The Extended Diffraction Image Analysis eliminates one major assumption of the Partially Coherent Image Analysis feature: that the OTF is constant over the field of view represented by the function A. This is accomplished by considering the source IMA file one pixel at a time, and computing the Fourier transform of the one pixel. The one pixel transform is multiplied by the OTF corresponding to that pixel. The sum over all pixels is then computed in spatial frequency space, and finally the sum of the filtered pixel transforms is Fourier transformed back to form the final image.

Computing the OTF for every pixel in the IMA file is slow and not practically required, as the OTF over any reasonably small field of view does not change rapidly. As an alternative, this feature computes a grid of OTF which spans over the field of view, and then interpolation is used to compute the effective OTF at any single pixel. The greater the size of the OTF grid, the more accurate the results will be, at the expense of longer computation times and more memory usage.

This feature can require large amounts of RAM. For an OTF grid size of n x n, ZEMAX will actually store n*n+1 OTF's: the n*n OTF's on the grid plus one to use as a running sum of the image frequency components. The OTF itself is a complex value array of double precision points. ZEMAX automatically zero pads the OTF grid to achieve accurate sampling. If the OTF sampling is 64 x 64, a single OTF will be 128 x 128, using up 256k of RAM (128 x 128 x 2 x 8 bytes). For an OTF grid of 5 x 5, the total comes to 26*256k = 6.5 Mb. For 128 x 128 sampling and a 9 x 9 grid, the RAM required is 82 Mb.

The file size parameter determines how big each pixel in the IMA file is in *the image space of the optical system*. Note this is different from the geometric image analysis feature, where the IMA file defines the size and shape in "field" space. For this analysis, the IMA file defines the ideal image shape in image space. If the file size is 0.1 mm, then each pixel is 14.286 micrometers wide for this 7 x 7 pixel image.

IMA/BIM File Viewer

Purpose:

This feature displays IMA/BIM files. See “Geometric Image Analysis” on page 165 for a discussion of IMA/BIM files.

Settings:

Item	Description
File	The name of the .IMA or .BIM image file. This file must reside in the <images> folder (see “Folders” on page 68). See the discussion section for a full description of the IMA and BIM file formats.
Show	Choose surface plot, contour map, grey scale, or false color map as the display option.
Color	If the selected file is a binary IMA, this control will select the color channel to display.
Edit IMA File	Pressing this button will invoke the Windows Notepad editor which allows modification of the currently selected IMA file. This button is disabled if the file type is BIM.

Discussion:

All input values are normalized to a relative intensity of 1.0 for display purposes. See “Geometric Image Analysis” on page 165 for a discussion of IMA/BIM files.

JPG/BMP File Viewer

Purpose:

This feature displays JPG/BMP files without any processing.

Settings:

Item	Description
File	The name of the .JPG or .BMP image file. This file must reside in the <images> folder (see “Folders” on page 68).

Biocular Analysis

Field of View

Purpose:

Displays the field of view for up to four configurations. See the discussion for important assumptions before using this feature.

Settings:

Item	Description
Config 1, 2, 3, 4	Selects which configurations to display the field of view data for.
X Points, Y Points	Defines the number of points to use in each direction.
Left/Right X	The left and right limits in field of view units.
Top/Bottom Y	The upper and lower limits in field of view units.
Use Angles	If checked, the field of view units are degrees, otherwise, the units are direction cosines.
Wavelength	The wavelength number to be used in the calculation.

Discussion:

This feature is intended to show the field of view for up to four configurations. Field of view in this context means the angle of rays emanating from the stop surface, *not the object surface*, which are unvignetted all the

way to the image surface. The primary usage is the analysis of biocular systems, where both eyes look through the same optical path to see a projected image. This feature makes all of the following assumptions:

- The field of view is either angles in degrees or in direction cosines directly. The angles/cosines are measured along the chief ray from the stop position relative to the local Z axis at the stop surface. Surfaces prior to the stop are ignored.
- The image surface is assumed to be the location of the image the eyeball is looking at. Each selected configuration should be represented by a single (usually decentered) eyeball position.
- The eyeball decenter should be set up so that the X, Y coordinates on the image surface represent the same point on the image source in all configurations. For example, if the image source is a CRT, then the point with image coordinates (x = 1, y = 2) should correspond to the same physical location on the CRT in all configurations.
- All surfaces should have fixed apertures for the purposes of vignetting rays that are outside of the field of view.

The field of view is computed by tracing chief rays at various angles from each configuration. If the ray passes all apertures, then the point on the plot corresponding to that angle will be inside of the closed curve for that configuration. Each configuration will generate a closed curve representing the limits of the field of view for that configuration. The ray tracing is done iteratively to determine the exact angles at which the chief rays become vignetted.

If the field of view units are cosines, then the plot is linear in cosine space, with the Z direction cosine determined from the X and Y direction cosines. If the field of view units are angles, then the direction cosines of the chief rays are determined from the X and Y field angles using these formulas:

$$\tan(\theta_x) = \frac{l}{n}, \tan(\theta_y) = \frac{m}{n}, l^2 + m^2 + n^2 = 1,$$

where l, m, and n are the X, Y, and Z direction cosines, respectively.

Dipvergence/Convergence

Purpose:

Displays the dipvergence and convergence for biocular analysis.

Settings:

Item	Description
Left Eye Config	Defines the configuration to use as the left reference eye.
Right Eye Config	Defines the configuration to use as the right eye, whose dipvergence and convergence is determined relative to the left eye.
Do X Scan	If checked, the plot scan is across the X direction field of view, otherwise, the scan is across the Y direction field of view.
Use Angles	If checked, the field of view units are degrees, otherwise, the units are direction cosines.
Min/Max Angle/ Cosine	Sets the limits of the scan in X or Y, in either cosine space or in degrees. Note the limits do not need to be symmetric, and the min value need not be less than the max value.
X/Y Angle/Cosine 1-6	If "Do X Scan" is checked, then up to 6 Y angles/cosines offsets may be selected for display. If "Do X Scan" is unchecked, then up to 6 X angles/cosines offsets may be selected for display. To turn off one or more of the offset values, set the value to -99.
# Points	The number of points to use across the full scan.

Item	Description
Wavelength	The wavelength number to be used in the calculation.
Convergence Offset	This value is subtracted from the computed convergence value. This allows only the deviation from the desired value to be displayed.

Discussion:

The assumptions made for this feature are the same as for “Field of View” on page 180.

When two eyes are used to look through a biocular lens, there is usually a small angular difference between the direction the two eyes must look to see the same image point. The vertical (up/down) angle is called dipvergence. The horizontal left-right angle is called convergence if the eyes gaze toward a common point in front of the viewer’s head, so the chief rays from the two eye positions would converge as they move toward the lens and away from the viewer’s head. If the two chief rays diverge as the rays propagate toward the lens, as if to see a virtual image point behind the head, then the angle is called divergence. Convergence and divergence are really the same thing from a computational standpoint. ZEMAX uses the common convention that convergence is a positive value and divergence is a negative value; for this reason only the term convergence will be used in the subsequent discussion, with the understanding that if the convergence is negative then it is properly called divergence. Usually, convergence is more tolerable than divergence and the two aberrations may have different specifications. Dipvergence and convergence are both measured in milliradians, and typical limits are on the order of 1.0 milliradians for visual systems.

The computation proceeds for a given point in the field of view by tracing a reference ray in the left eye configuration. The same angle chief ray is then traced in the right eye configuration. In general, the right eye ray will not land on the exact same X, Y coordinates on the image surface as the left eye reference ray did. ZEMAX iterates the right eye trace until the chief ray is found that matches the left eye reference ray intercept coordinates. The resulting right eye chief ray will in general make some angle with respect to the left eye in both the vertical and horizontal directions, and it is these angles that are the dipvergence and convergence, respectively.

Note that both the left eye and the right eye chief rays must pass all surfaces without errors and without vignetting for the computation to be valid. If both rays do not trace, no data will be returned for that field of view. The field of view overlap between the configurations is very useful for setting appropriate min/max scan values, see the “Field of View” on page 180 for determining the field of view overlap.

It is possible for the iteration required for the right eye to fail. This typically happens if either no solution exists or if the dipvergence/convergence is so large that algorithm becomes unstable. Failed iteration will usually show up as gaps in the plot.

Miscellaneous

Field Curvature/Distortion

Purpose:

Displays the field curvature and distortion plots.

Settings:

Item	Description
Max Curvature	Defines the maximum scale for the field curvature plot in lens units. Enter zero for automatic.
Max Distortion	Defines the maximum scale for the distortion plot. Enter zero for automatic.
Wavelength	The wavelength number to be used in the calculation.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.
Ignore Vignetting Factors	See the discussion section.

Item	Description
Distortion	Select F-Tan(theta), F-Theta, Calibrated F-Theta, or Calibrated F-Tan(theta) distortion. See the discussion section for details.
Display As	Selects whether to display distortion as a percentage or as an absolute length.
Scan Type	Choose +y, +x, -y, or -x field scan direction.

Discussion:

The field curvature plot shows the distance from the image surface to the paraxial image surface as a function of field coordinate. The tangential data are the distances measured along the Z-axis from the image surface to the paraxial image surface measured in the tangential (YZ) plane. The sagittal data are the distances measured in the plane orthogonal to the tangential plane. The base of the plot is on axis, and the top of the plot represents the maximum field (angle or height). There are no units on the vertical scale because the plot is always normalized to the maximum field coordinate along the scan direction.

The field scan is along the +Y field by default. If +X or -X is selected, then the field scan is along the X direction, in which case the tangential curve is for the XZ plane, and the sagittal curve is in the YZ plane.

The field curvature plot does not always start at zero for zero field. The reason is that the plot shows the distance from the image surface to the paraxial image surface, and the image surface need not be coincident with the paraxial image surface. If there is any defocus, then the two surfaces are offset, and so is the field curvature data.

Distortion in percent is defined as the real chief ray height, minus the reference ray height, divided by the reference ray height, times 100:

$$Distortion = 100 \times \frac{y_{chief} - y_{ref}}{y_{ref}},$$

where all heights are taken to be the image surface radial coordinate, at whatever image surface is defined (the data is not referred to the paraxial image surface). The reference ray height is computed by tracing a real ray from a very small field height, and then scaling the results as required. This generalization permits the computation of reasonable distortion values even for systems not well described by paraxial ray tracing. For F-Tan(theta) distortion, the reference height for the undistorted ray in a rotationally symmetric system at paraxial focus is given by

$$y_{ref} = f \tan \theta,$$

where f is the focal length and θ is the angle in object space. Note that this reference is meaningless for optical systems that use field angles at or greater than 90 degrees. If a system with field angles over 89 degrees is analyzed with this feature, the maximum computed field angle will be limited to 89 degrees.

F-Theta distortion does not use the tangent relationship, but instead uses the height given by the focal length of the lens multiplied by the angle the chief ray makes in object space. This so-called F-Theta height is only meaningful in systems with the object at infinity, when field heights are measured in angles. The reference height for the undistorted ray at the paraxial focal plane is given by

$$y_{ref} = f \theta,$$

where f is the focal length and θ is the angle in object space. F-Theta is used in scanning systems where the image height needs to be linear with scan angle.

Calibrated distortion, whether F-Tan(theta) or F-Theta, is similar to the other distortion definitions, except that the "best-fit" focal length is used rather than the system focal length. Calibrated distortion measures the deviation

between the reference and actual image height and selects the best f value to minimize the distortion, without the restriction that the proportion of linearity be defined by the focal length of the system. A focal length is chosen that best fits the data, rather than the system focal length, although in general the best fit focal length is close to the system focal length. The calibrated focal length used is given on the text listing for this feature.

One effect of this definition of calibrated distortion is the non-zero distortion at zero field angle. The reason for this non-zero distortion is best explained by looking at the limiting behavior of the definition of calibrated distortion. Calibrated distortion in percent is defined as

$$Distortion = 100 \times \frac{y_{chief} - y_{ref}}{y_{ref}}$$

For small angles, the real y chief ray coordinate in any reasonable optical system is well described by $y_{chief} = f\theta$, and the reference ray coordinate by $y_{ref} = f'\theta$, where f' is the calibrated focal length, so the distortion near zero field is

$$Distortion = 100 \times \frac{y_{chief} - y_{ref}}{y_{ref}} = 100 \times \frac{f\theta - f'\theta}{f'\theta} = 100 \times \frac{f - f'}{f'}$$

which is not zero, unless $f = f'$. Therefore, the calibrated distortion is not generally zero on axis. This does not mean the image height is not zero, it is an artifact of the different choice of focal lengths for the reference and actual ray coordinates near the optical axis. Note that a percentage distortion near the axis is not significant because the field itself approaches zero at zero field angle.

For non-rotationally symmetric systems, distortion is poorly defined and the data presented is probably meaningless. The reason is that no single number adequately describes distortion at a single field point if the system is not rotationally symmetric. Instead, see "Grid Distortion" on page 184.

Strictly speaking, the field curvature and distortion plots are only valid for rotationally symmetric systems with plane object and image surfaces. However, ZEMAX uses a generalization of the field curvature and distortion concepts to give reasonable results for some, but not all, non-rotationally symmetric systems. Caution should be used when interpreting these data for non-rotationally symmetric systems or systems with object and/or image surfaces that are not planes.

By default, ZEMAX ignores vignetting factors when computing the field curvature and distortion plots. Vignetting factors can change the chief ray location on the stop surface, such that the chief ray no longer goes through the center of the stop.

Computing distortion when using real image height as the field type

The distortion cannot be computed in the manner described here if the field type is real image height. The reason is that when using real image height, ZEMAX iterates each chief ray trace to find the exact object space angle to hit the desired image coordinate. Because the desired image coordinate is always reached, the image height is linear with field coordinate. The iteration is thus implicitly removing the distortion. Instead, ZEMAX automatically changes the field type from real image height to field angle for the purposes of this computation. The field angles at each defined field point are then chosen to yield the identical real image height coordinates. Using this method, the deviation from linearity is now retained and the distortion may be computed accurately. If the optical system is not axial symmetric, the conversion from real image height to field angle may yield unexpected results as the field angle scan direction and the real image coordinate direction may no longer be coplanar, and care should be taken in interpreting the results. It is recommended that the system be converted to use field angle or object height rather than real image height for non-axial systems when evaluating distortion.

Grid Distortion

Purpose:

Displays a grid of chief ray intercept points to indicate distortion.

Settings:

Item	Description
Display	Select either "Cross" to mark each chief ray intercept with a cross, or "Vector" to plot a vector from the ideal image point to the actual chief ray image point.
Grid Size	The size of the grid.
Wavelength	The wavelength number to be used in the calculation.
Ref. Field	The reference field position. See the "Discussion".
Scale	If the scale is different than 1.0, then the "x" points on the distortion grid will be exaggerated by the selected scale factor. The scale factor may be negative to change positive to negative distortion on the plot.
H/W Aspect	If unity, then a square field will be selected. The output image may not be square if the system is not symmetric, but the object field will be square. If the H/W aspect is greater than 1, then the "height" or "y" field will be expanded by the aspect ratio. If the aspect is less than 1, then the "height" or "y" field will be compressed by the aspect ratio. The aspect ratio is the "y" field height divided by the "x" field width. The aspect ratio only affects the input field; the image aspect ratio is determined by the optical system imaging properties.
Symmetric Magnification	If checked, then the X magnification is required to be identical to the Y magnification. This causes distortion to be referenced to a symmetric predicted grid rather than an anamorphic predicted grid.
Field Width	The full width, in field units, of the "x" field of view. If the current field type is real image height, ZEMAX will temporarily change the field height to field angle, and the "Field Width" value will be in angular units. See "Computing distortion when using real image height as the field type" on page 184.

Discussion:

This feature displays or computes the coordinates of a grid of chief rays. In a system without distortion, the chief ray coordinates on the image surface follow a linear relationship with the field coordinate:

$$\begin{bmatrix} x_p \\ y_p \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} f_x \\ f_y \end{bmatrix},$$

where x_p and y_p are the predicted image coordinates relative to a reference image point and f_x and f_y are the linear coordinates on the object surface relative to a reference point. For optical systems using angles as a field definition, f_x and f_y are the tangents of the field angles (the field coordinates must be linear, therefore, tangents of angles rather than angles are used). To compute the ABCD matrix, ZEMAX traces rays over a very small region centered upon the reference field position. Usually, this is the center of the field of view. ZEMAX allows selection of which field position to use for reference.

By default, ZEMAX sets the corner of the field grid in object space to be at the maximum radial field distance. Because object height is linear with the tangent of the field angle, the full width of the field when angles are used to define the field is given by

$$\theta_{wide} = 2 \tan^{-1} \left[\frac{\sqrt{2}}{2} \tan \theta_r \right],$$

where θ_r is the maximum radial field angle at the corner of the field.

The ray coordinates in image space for the very small field of view are used to determine the ABCD matrix components. The use of an ABCD matrix allows for coordinate rotations. If the image surface is rotated, such that a y object coordinate images to both an x and a y image coordinate, the ABCD matrix will automatically account for the rotation. The grid distortion plot shows the linear grid, and then marks the actual chief ray intercept for a ray with the same linear field coordinates with an "X" for each point on the grid.

If the optical system is not rotationally symmetric, then the distortion is not generally radial. The distortion is a vector, and the magnitude of the distorted vector must be used to compute the total distortion.

The text listing available tabulates the predicted image coordinate, the actual image coordinate, and the "percent distortion" defined by


$$P = 100\% \frac{R_{\text{distorted}}}{R_{\text{predicted}}}, \text{ where}$$

$$R_{\text{real}} = \sqrt{(x_r)^2 + (y_r)^2},$$

$$R_{\text{predicted}} = \sqrt{(x_p)^2 + (y_p)^2}, \text{ and}$$

$$R_{\text{distorted}} = \sqrt{(x_p - x_r)^2 + (y_p - y_r)^2},$$

and the subscripts r and p refer to the real and predicted coordinates on the image surface relative to the reference field position image location, respectively. Because the values for R_{real} and $R_{\text{predicted}}$ are always positive, this definition will always yield a positive value for P . However, it is still a frequently useful concept to distinguish between "positive" and "negative" distortion. To support this, if R_{real} is less than $R_{\text{predicted}}$, then the sign of P is changed to be negative.


 ***This definition of generalized distortion may not be applicable in all cases, and the results should be used with caution.***

SMIA-TV Distortion

The text listing of this feature also computes an alternate measure of distortion called Standard Mobile Imaging Architecture TV (SMIA-TV) Distortion. SMIA-TV Distortion is defined as

$$P_{TV} = 100\% \frac{A - B}{B},$$

where $A = 0.5 \cdot (A1 + A2)$ and $A1$, $A2$, and B are distances measured between three pairs of points in a rectangular field of view as shown in the adjacent image. The A distances are measured between the top and bottom corners of the field of view at the left and right sides of the image, and the B distance is measured between the top and bottom of the image at the x-field center of the image. If any one of the 6 required rays cannot be traced, the SMIA-TV distortion will be listed as zero,



Limitations of Grid Distortion

ZEMAX cannot compute the distortion in the manner described here if the field type is real image height. Instead, ZEMAX automatically changes the field type from real image height to field angle for the purposes of this

computation. For more information see “Computing distortion when using real image height as the field type” on page 184.

Grid distortion cannot be calculated if the field units are angles and the maximum angle equals or exceeds 90 degrees. This limitation is due to the assumption that the predicted image height is proportional to the tangent of the field angle in object space. When the field angle exceeds 90 degrees, the tangent does not predict the linear image height correctly.

The grid distortion plot is oriented in field units, so that the +y direction corresponds to the +y field coordinate in field units (see “Fields” on page 115). This may or may not be different than the orientation of a spot diagram or other plot plotted in image space units.

Relative Illumination

Purpose:

Computes the relative illumination as a function of radial field coordinate for a uniform Lambertian scene. This feature also computes the Effective F/#.

Settings:

Item	Description
Ray Density	The number of rays on one side of an array of rays used to integrate the illumination of the exit pupil. A value of 10 will trace about $10 \times 10 \times \pi / 4$ or 78 rays. Higher ray densities yield more accurate results at the expense of longer computation times.
Field Density	The number of points along the radial field coordinate to compute the relative illumination for. Larger field densities yield smoother curves.
Wavelength	Selects the wavelength for computation. Relative illumination is a monochromatic entity.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Log Scale	If checked, a logarithmic rather than linear scale will be shown.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See “Comment about vignetting factors” on page 136.
Scan Type	Choose +y, +x, -y, or -x field scan direction.

Discussion:

This feature computes the relative illumination (RI) as a function of radial y field coordinate. RI is defined as the intensity of illumination per unit area of image surface normalized to the illumination at the point in the field that has maximum illumination (which may not be on axis). The computation considers apodization, vignetting, apertures, aberrations of both the image and pupils, variations in F/#, chromatic aberrations, image surface shape, angle of incidence, and optionally, polarization effects assuming unpolarized light. The method is based upon one described in M. Rimmer, "Relative illumination calculations", Proc. SPIE Vol. **655**, p99 (1986). The published method was extended to include apodization, transmission, polarization, and non-planar image surface effects. The computation method assumes the following are all true:

1. The object scene is plane, uniform, and Lambertian.
2. The image surface is a reasonably good conjugate (that is, an image) of the object surface, so that rays coming from small patches of light on the object surface are imaged to patches of light on the image surface. Aberrations are fine, but the rays should be reasonably localized on the image surface.
3. The exit pupil is not too close to the image surface. This condition will be satisfied if the F/# is larger than about 0.1 and the ray aberrations are small compared to the exit pupil distance.
4. The cosine space aberrations are not so severe as to form caustics in angle space. A caustic in angle space means that rays in different parts of the entrance pupil have the same angle in image space. To check this, use the "Direction Cosines" option on the spot diagram feature (See “Standard” on page 129).

The relative illumination is computed by integration of the effective area of the exit pupil as seen from the image point(s). The integration is carried out in direction cosine space using a uniform grid in image cosine space.

Note that the RI computation will not in general yield a cosine-fourth curve, because the so-called cosine-fourth "law" is actually an approximation based upon a thin, slow, aberration free lens with the stop at the lens. For more general lenses, including telecentric, aberrated, or vignetted lenses, the RI can be computed using an integration of the projected solid angle or effective area of the exit pupil as seen from the image location, and this computation will not yield a simple cosine-fourth curve. If a system violates the assumptions of the computation an error message will be displayed and the RI will not be computed.

Effective F/#

The text listing of the relative illumination data also includes data for the Effective F/#. The Effective F/# is the F/# required for a perfect optical system with 100% transmission and a circular exit pupil to have the same image illumination as the system being evaluated. The Effective F/# is computed by:

$$E = \sqrt{\frac{\pi}{4A}},$$

where A is the area of the projected solid angle of the pupil in cosine space weighted for system transmission. The Effective F/# is a useful metric for comparing the brightness of the image formed by different optical systems, because it accounts for RI and is independent of the aperture shape. For more information see "F-Number and the radiometry of image forming optical systems with non-circular aperture stops," R. Siew, Proc. of SPIE Vol. 5867 (2005).

Vignetting Plot

Purpose:

Calculates fractional vignetting as a function of field angle.

Settings:

Item	Description
Ray Density	The ray density specifies the number of rays to be traced. The more rays traced, the greater the accuracy, although the computation time increases. For a ray density of n, ZEMAX traces a grid of (2n+1) x (2n + 1) rays at each field point.
Field Density	The field density is the number of points between zero degrees and the maximum field angle specified at which vignetting is calculated, intermediate values are interpolated.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 136.

Discussion:

Fractional vignetting is the percentage of rays incident upon the entrance pupil which pass all obscurations and apertures in the system and survive to the image surface, normalized to relative pupil area. The graphic generated by this function shows fractional vignetting as a function of field position. If too few rays are used, the results may be inaccurate. This is especially true in systems with many apertures and large field angles.

Only the primary wavelength is used in this calculation. This is a geometric calculation. Only positive y field positions are used, therefore this feature is only appropriate for rotationally symmetric lenses and fields. Rays which cause errors such as missing a surface or those which are TIR are considered vignetted.

See also "Relative Illumination" on page 187.

Footprint Diagram

Purpose:

Displays the footprint of the beam superimposed on any surface. Used for showing the effects of vignetting and for checking surface apertures.

Settings:

Item	Description
Ray Density	Determines the number of rays traced across the half pupil; a setting of 10 will trace a grid of 21 x 21 rays. The "Ring" option will trace 360 marginal rays around the edge of the pupil at each field and wavelength. The ring option will automatically determine the radial coordinate of the marginal ray that is not vignetted to approximate the shape of the beam on any surface, but will yield incorrect results if the beam is in a caustic.
Surface	The surface to show the beam footprint on.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number to be used in the calculation.
Delete Vignetted	If checked, then rays which are vignetted by subsequent surfaces will not be drawn. Rays which are vignetted by prior surfaces are never drawn.
Use Symbols	If checked, this option will draw different symbols rather than dots for each wavelength. This helps distinguish the various wavelengths.
Configuration	Select "All" to show all configurations, "Current" to show just the active configuration, or select any combination of other configurations.
Color Rays By	Select "Field #" to use color to distinguish between each field position, or "Wave #" to distinguish between each wavelength, "Config #" to distinguish between configurations, and "Wavelength" to approximate the color of wavelengths in the visible spectrum.

Discussion:

This feature will draw the shape of the surface, and then superimpose on that surface a grid of rays. If there is no aperture on the surface, then a circular surface shape with a radial clear aperture of the semi-diameter value is shown. Otherwise, the shape of the aperture is shown. The surface aperture is always shown as centered in the frame; even if the aperture is decentered on the actual surface. If there is an obscuration on the surface, then the obscuration will be drawn along with the circular aperture defined by the semi-diameter. If the surface aperture on the selected surface is modified between different configurations, and more than one configuration is selected, the aperture for the first configuration only will be drawn.

The ray grid size is specified by the ray density parameter, and rays may be from any or all fields, at any or all wavelengths. Rays which are vignetted by surfaces prior to the surface shown are not drawn. Rays which are vignetted by the surface or subsequent surfaces are not drawn if "delete vignetted" is checked, otherwise, they are drawn. The ray set is apodized if any system pupil apodization is selected. The number of rays shown divided by the total number of rays launched is listed as the percentage of rays through if the Ray Density is not "ring".

For graphic windows, the left/right cursor keys will change the surface number and recompute the data.

Longitudinal Aberration

Purpose:

Displays the longitudinal aberration as a function of pupil height at each wavelength.

Settings:

Item	Description
Plot Scale	Defines the maximum scale for the plot in lens units. Enter zero for automatic.
Wavelength	The wavelength number to be used in the calculation.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.

Discussion:

This feature computes the distance from the image surface to where a zonal marginal ray "focuses", or crosses the optical axis. The computation is performed only for the on axis field point, and only for zonal marginal tangential rays as a function of pupil zone. The base of the plot is on axis, and the top of the plot represents the maximum entrance pupil radius. There are no units on the vertical scale because the plot is always normalized to the maximum entrance pupil radius. The horizontal scale is in lens units, and represents the distance from the image surface to the point where the ray crosses the optical axis. For afocal systems, the horizontal axis is diopters of power required to bring the rays to collimation.

Because longitudinal aberration is defined in terms of the distance to the ray-axis crossing point, this feature may produce meaningless data for non-rotationally symmetric systems. Great care should be exercised in interpreting this plot for non-rotationally symmetric systems.

Lateral Color

Purpose:

Displays the lateral color as a function of field height.

Settings:

Item	Description
Plot Scale	Defines the maximum scale for the plot in lens units. Enter zero for automatic.
Use Real Rays	If checked, real rays are used, otherwise paraxial rays are used.
All Wavelengths	If checked, then data for all defined wavelengths will be displayed. Each wavelength will be referenced to the primary wavelength. If not checked, then the difference between the shortest and longest wavelength rays will be used. See the discussion.
Show Airy Disk	If checked, then the Airy disk radius at the primary wavelength will be plotted on either side of the reference line to indicate the extent of the Airy disk.

Discussion:

This feature computes the lateral color either of two ways:

If "All wavelengths" is checked: The data computed is the distance on the image surface from the intercept of the chief ray at each wavelength to the primary wavelength chief ray intercept.

If "All Wavelengths" is not checked: The data computed is the distance on the image surface from the shortest wavelength chief ray intercept to the longest wavelength chief ray intercept.

The base of the plot is on axis, and the top of the plot represents the maximum field radius. Only positive field angles or heights in the Y direction are used. The vertical scale is always normalized to the maximum field angle or height. The horizontal scale is in micrometers. Either real or paraxial rays may be used.

If "Show Airy Disk" is check on, then the Airy disk radius will be approximated as the square root of the sum of the squares of the X and Y direction Airy disk radii.

This feature may produce meaningless data for non-rotationally symmetric systems. Great care should be exercised in interpreting this plot for non-rotationally symmetric systems.

Y-Ybar Diagram

Purpose:

Y-Ybar diagram.

Settings:

Item	Description
First Surface	The first surface for which data will be plotted.
Last Surface	The last surface for which data will be plotted.
Wavelength	The wavelength number to be used in the calculation.
Plot Scale	Sets the maximum scale for the plot. The plot is always shown in a square box; the default scale is the maximum transverse ray coordinate. Enter zero for automatic.

Discussion:

The Y-Ybar diagram is a plot of marginal ray height as a function of chief ray height for a paraxial skew ray at every surface in the lens.

Chromatic Focal Shift

Purpose:

Chromatic focal shift plot.

Settings:

Item	Description
Maximum Shift	The maximum extent in lens units for the horizontal axis. The vertical axis scale is set by the range of wavelengths defined. Enter zero for automatic.
Pupil Zone	The radial zone in the pupil used to compute the back focus. The default value of zero means a paraxial ray will be used. Values between 0 and 1 mean real marginal rays are used in the appropriate zone in the entrance pupil. A 1 is at the edge of the pupil, or full aperture.

Discussion:

This is a plot of the shift in back focal length with respect to the primary wavelength. At each plotted wavelength, the shift in image space required to reach focus for that color marginal ray is computed. The shift distance is computed in the same media (glass or air) as the surface prior to the image surface. This plot may not be meaningful for non-paraxial systems.

The maximum shift setting overrides the default scaling. Units are in lens units. The entire plot is always referenced to the primary wavelength paraxial focus. The diffraction limited depth of focus listed is given by the formula $4\lambda F^2$, where F is the working F/# and λ is the primary wavelength.

System Summary Graphic

Purpose:

Displays in a graphic window a summary of the system data, similar to the text based system data report.

Settings:

None.

Discussion:

This graphic is primarily used to display a summary of system data within the printed page of the report graphics 4 or 6 feature; described in "Report Graphics 4/6" on page 273.

Power Field Map

Purpose:

This feature computes optical power or effective focal length as a function of field position. A common application of this feature is analysis of the spherical and cylinder power of a Progressive Addition Lens (PAL). See also "Power Pupil Map" on page 193.

Settings:

Item	Description
Sampling	The rectangular grid size to use for the analysis.
Field Width	The full width of the field in field units.
Offset	A user defined offset value that is added to the computed data.
Wavelength	The wavelength number to use for the analysis.
Surface	The last surface with optical power. The power and EFL data is computed for the system as a whole up to and including refraction from this surface. See the discussion below.
Use Tangent Of Field Angle	If checked, and the field type is angle in degrees, the displayed spatial distribution of the data is proportional to the tangent of the angle rather than angle directly.
Data	The data to be computed. The options are spherical, cylinder, maximum, minimum, tangential, sagittal, Y direction, X direction, and Astigmatic (X minus Y) optical power in diopters or effective focal length (EFL) in lens units.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option. A scan across the X- or Y-field direction is also available.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see "The Contour Format String" on page 151.

Discussion:

This feature computes optical power or focal length as a function of field coordinate. The power or focal length is determined for the optical system as a whole up to and including refraction from any surface. The method used is to trace a ring of real rays around the chief ray at each point in the field. The ray data are used to determine the focal length for each field position. This focal length can then be used to compute the optical power in units of diopters (inverse meters). In the general case, the focal length is a function of orientation in the entrance pupil. By tracing a ring of rays, the average, maximum, and minimum optical power or focal length around the pupil can be determined. From this data, several types of optical power can be computed. Note that the definitions used by this feature may vary from other uses of these terms.

The spherical power is given by the average power over the entrance pupil. The maximum and minimum powers are the extreme values over the pupil. The cylinder power is defined as the difference between the maximum and minimum powers. Note using this definition the amount of cylinder power is always positive. For cylinder, the EFL value reported is the difference between the maximum and minimum focal lengths. For astigmatic EFL, the value reported is the difference between the X and Y direction focal lengths.

This feature can also compute the power along the tangential or sagittal directions. In this context, the tangential direction is along the plane that contains both the chief ray and the local z axis in object space. The sagittal direction is orthogonal to the tangential direction. Also available are the Y and X direction optical powers. The directions are defined by the local coordinate orientation on the entrance pupil. Astigmatic power is defined as the X direction power minus the Y direction power.

The surface setting should define the last surface of the portion of the lens being evaluated that has optical power. It is recommended that this be set manually to the correct value rather than use the default "Image" surface setting. The reason is that exact ray tracing is used to compute the optical power, and the most accurate data is the ray data immediately after refracting from the last powered surface. ZEMAX may not be able to correctly resolve some sign ambiguities if the rays propagate a long way past the lens.

The "offset" value is a user defined value added to all the computed data. This feature is useful for evaluating the deviation relative to some user defined reference value.

To optimize a lens for specific power characteristics, see the POWF optimization operand.

Power Pupil Map

Purpose:

This feature computes optical power or effective focal length as a function of pupil position. See also "Power Field Map" on page 191.

Settings:

Item	Description
Sampling	The rectangular grid size to use for the analysis.
Field	The field number for which the calculation should be performed.
Offset	A user defined offset value that is added to the computed data.
Wavelength	The wavelength number to use for the analysis.
Surface	The last surface with optical power. The power and EFL data is computed for the system as a whole up to and including refraction from this surface. See the discussion below.
Data	The data to be computed. The options are spherical, cylinder, maximum, minimum, tangential, sagittal, Y direction, X direction, and Astigmatic (X minus Y) optical power in diopters or effective focal length (EFL) in lens units.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option. A scan across the X- or Y-field direction is also available.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see "The Contour Format String" on page 151.

Discussion:

This feature computes optical power or focal length as a function of pupil coordinate. The power or focal length is determined for the optical system as a whole up to and including refraction from any surface. The method used is to trace a ring of real rays around a central reference ray at each point in the pupil. The ray data are used to determine the focal length for each pupil position. This focal length can then be used to compute the optical power in units of diopters (inverse meters). In the general case, the focal length is a function of orientation. By tracing a ring of rays, the average, maximum, and minimum optical power or focal length around the reference ray can be determined. From this data, several types of optical power can be computed. Note that the definitions used by this feature may vary from other uses of these terms.

The spherical power is given by the average power over the ring of rays. The maximum and minimum powers are the extreme values over the ring of rays. The cylinder power is defined as the difference between the maximum and minimum powers. Note using this definition the amount of cylinder power is always positive. For cylinder, the EFL value reported is the difference between the maximum and minimum focal lengths. For astigmatic EFL, the value reported is the difference between the X and Y direction focal lengths.

This feature can also compute the power along the tangential or sagittal directions. In this context, the tangential direction is along the plane that contains both the chief ray and the local z axis in object space. The sagittal direction is orthogonal to the tangential direction. Also available are the Y and X direction optical powers. The directions are defined by the local coordinate orientation on the entrance pupil. Astigmatic power is defined as the X direction power minus the Y direction power.

The surface setting should define the last surface of the portion of the lens being evaluated that has optical power. It is recommended that this be set manually to the correct value rather than use the default "Image" surface setting. The reason is that exact ray tracing is used to compute the optical power, and the most accurate data is the ray data immediately after refracting from the last powered surface. ZEMAX may not be able to correctly resolve some sign ambiguities if the rays propagate a long way past the lens.

The "offset" value is a user defined value added to all the computed data. This feature is useful for evaluating the deviation relative to some user defined reference value.

Incident Angle vs. Image Height

Purpose:

This feature computes the incident angle on the image surface of the lower marginal, chief, and upper marginal rays as a function of image height.

Settings:

Item	Description
Field Density	The number of points in the field, from 0 to the maximum +y field coordinate.
Wavelength	The wavelength number to use for the analysis.

Discussion:

This feature traces three rays: one at the bottom of the entrance pupil, the chief ray, and another at the top of the entrance pupil. The (signed) angle of incidence on the image surface (which is assumed to be flat) is computed for each of these three rays, as the field coordinate changes from 0 to the maximum +y field. The image height based upon the chief ray y-intercept coordinate is also computed. The feature then plots (or tabulates as text) the three resulting angles as a function of the resulting image height.

This feature is only intended to be used for rotationally symmetric systems with no vignetting, and may not produce meaningful results for other types of systems.

Aberration coefficients

Seidel Coefficients

Purpose:

Displays Seidel (unconverted, transverse, and longitudinal), and wavefront aberration coefficients.

Settings:

Item	Description
Wavelength	The wavelength number to use for the calculation.

Discussion:

ZEMAX will compute the unconverted Seidel, transverse, longitudinal, and some wavefront coefficients. The Seidel coefficients are listed surface by surface, as well as a sum for the entire system. The coefficients listed are for spherical aberration (SPHA, S1), coma (COMA, S2), astigmatism (ASTI, S3), field curvature (FCUR, S4), distortion (DIST, S5), longitudinal color (CLA, CL), and transverse color (CTR, CT). The units are always the same as the system lens units, except of course for the coefficients measured in waves.

These calculations are only valid and accurate for systems consisting entirely of surfaces that are axial symmetric spheres, conic, second order, or fourth order aspheres. Any systems which contain coordinate breaks, gratings, paraxial, or other surfaces that are not radially symmetric are not adequately described by the paraxial rays which are used to compute the coefficients. Specifically, the supported surface types are the Standard, Even Asphere, Odd Asphere, Extended Asphere, and Extended Odd Asphere.

Transverse aberration coefficients are listed for each surface and for the system as a whole. The coefficients given are transverse spherical (TSPH), transverse sagittal coma (TSCO), transverse tangential coma (TTCO), transverse astigmatism (TAST), transverse Petzval field curvature (TPFC), transverse sagittal field curvature (TSFC), transverse tangential field curvature (TTFC), transverse distortion (TDIS), transverse axial color (TAXC), and transverse lateral color (TLAC). The transverse aberrations are in the system lens units. The transverse aberration coefficients may be very large in optical spaces where the light is nearly collimated, and have little meaning in these optical spaces.

Longitudinal aberration coefficients are computed for longitudinal spherical aberration (LSPH), longitudinal sagittal coma (LSCO), longitudinal tangential coma (LTCO), longitudinal astigmatism (LAST), longitudinal Petzval field curvature (LPFC), longitudinal sagittal field curvature (LSFC), longitudinal tangential field curvature (LTFC),

longitudinal distortion (LDIS), longitudinal axial color (LAXC), and longitudinal lateral color (LLAC). The longitudinal aberrations are in the system lens units. The longitudinal aberration coefficients may be very large in optical spaces where the light is nearly collimated, and have little meaning in these optical spaces.

The wavefront coefficients given are spherical aberration (W040), coma (W131), astigmatism (W222), field curvature Petzval (W220P), distortion (W311), axial color defocus term (W020), lateral color tilt term (W111), field curvature sagittal (W220S), field curvature medial (W220M), and field curvature tangential (W220T). All the wavefront coefficients are in units of wavelengths at the edge of the exit pupil. The various aberration coefficients are interrelated according to the following table. The symbols n and u refer to the index of refraction and the paraxial marginal ray angle in the object space side of each surface. The primes above the n and u symbols indicate that these are the values on the image space side of the surface. For a discussion of the meaning and derivation of the Seidel aberration coefficients, see Welford, Aberrations of Optical Systems, Smith, Modern Lens Design, or O' Shea, Elements of Modern Optical Design. A list of good references can be found in "REFERENCES ON LENS DESIGN" on page 37.

INTERRELATIONSHIP OF ABERRATION COEFFICIENTS

Name	Seidel	Wave	Description	Transverse	Longitudinal
Spherical	S_1	$\frac{S_1}{8}$	Spherical	$-\frac{S_1}{2n'u'}$	$\frac{S_1}{2n'u'^2}$
Coma	S_2	$\frac{S_2}{2}$	Sagittal	$-\frac{S_2}{2n'u'}$	$\frac{S_2}{2n'u'^2}$
			Tangential	$-\frac{3S_2}{2n'u'}$	$\frac{3S_2}{2n'u'^2}$
Astigmatism	S_3	$\frac{S_3}{2}$	From tangential to sagittal foci	$-\frac{S_3}{n'u'}$	$\frac{S_3}{n'u'^2}$
Field Curvature	S_4	$\frac{S_4}{4}$	Gaussian to Petzval	$-\frac{S_4}{2n'u'}$	$\frac{S_4}{2n'u'^2}$
	$S_3 + S_4$	$\frac{S_3 + S_4}{4}$	Gaussian to sagittal	$-\frac{(S_3 + S_4)}{2n'u'}$	$\frac{S_3 + S_4}{2n'u'^2}$
	$2S_3 + S_4$	$\frac{2S_3 + S_4}{4}$	Gaussian to medial	$-\frac{(2S_3 + S_4)}{2n'u'}$	$\frac{2S_3 + S_4}{2n'u'^2}$
	$3S_3 + S_4$	$\frac{3S_3 + S_4}{4}$	Gaussian to tangential	$-\frac{(3S_3 + S_4)}{2n'u'}$	$\frac{3S_3 + S_4}{2n'u'^2}$
Distortion	S_5	$\frac{S_5}{2}$	Distortion	$-\frac{S_5}{2n'u'}$	$\frac{S_5}{2n'u'^2}$

Name	Seidel	Wave	Description	Transverse	Longitudinal
Axial Color	C_L	$\frac{C_L}{2}$	Chromatic aberrations are measured between the extreme defined wavelengths, referenced to the selected wavelength.	$\frac{C_L}{n'u'}$	$\frac{C_L}{n'u'^2}$
Lateral Color	C_T	C_T		$\frac{C_T}{n'u'}$	$\frac{C_T}{n'u'^2}$

The Petzval radius of curvature is in system lens units, as is the Lagrange (or optical) invariant.

Seidel Diagram

Purpose:

Displays Seidel unconverted aberration coefficients as a bar chart.

Settings:

Item	Description
First/Last Surface	The surface range to display. The sum is for the selected range of surfaces. The image surface is not displayed.
Wavelength	The wavelength number to use for the calculation.
Plot Scale	Sets the maximum scale in lens units for the plot. Enter zero for automatic.
Ignore Distortion	If checked, distortion data will not be displayed.
Ignore Chromatic	If checked, axial and lateral color data will not be displayed.

Discussion:

The unconverted Seidel coefficients are displayed surface by surface, as well as a sum for the selected range of surfaces. For details on the Seidel coefficients, see "Seidel Coefficients" on page 194.

Zernike Fringe Coefficients

Purpose:

Calculates the Zernike coefficients using the Fringe (also called the "University of Arizona") polynomials.

Settings:

Item	Description
Sampling	Specify the density in the pupil to use for coefficient fitting. Larger grid sizes are more accurate, although the computation time increases.
Maximum Term	Specify the maximum Zernike coefficient to compute. Any value up to 37 may be specified.
Wavelength	The wavelength number to use for the calculation.
Field	The field number to use for the calculation.
Ref OPD To Vertex	Normally, ZEMAX references the OPD to the chief ray; which in effect subtracts out tilt from the wavefront phase. For interferometry, it is sometimes desirable to retain the wavefront tilt. Checking this option on will add tilt to the beam based upon the deviation of the chief ray from the image surface vertex. This option is only useful for field positions whose chief ray is reasonably close to the surface vertex, where the assumption that tilt is described by the deviation of the chief ray is valid.

Item	Description
Surface	Selects the surface at which the data is to be evaluated. This is useful for evaluating intermediate images. See "Evaluating results at intermediate surfaces" on page 128.
Subaperture Data: Sx, Sy, Sr	Defines a subaperture of the pupil to compute the Zernike data for. See "Subaperture computations" on page 197 for details.

Discussion:

This feature displays the individual Zernike coefficients as well as the peak-to-valley, RMS, variance, Strehl ratio, residual RMS fit error, and maximum fit error.

The RMS of the wavefront error, σ , when referenced to the mean, is defined as

$$\sigma^2 = \overline{W^2} - \bar{W}^2$$

where W is the wavefront error, $\overline{W^2}$ is the mean square error, and \bar{W} is the average wavefront error. The RMS can actually be computed several different ways. If the mean wavefront term is ignored, then the RMS "referenced to zero" results. This computation yields the square root of $\overline{W^2}$ directly, and is rarely used.

If the mean wavefront is subtracted from all the wavefront phase values (the absolute phase reference has no physical meaning), then the RMS is "referenced to the mean".

Typically, the RMS is further referenced to the tilted and shifted reference sphere which minimizes the RMS. This is equivalent to subtracting out not only the mean (which is piston) but the average tilt in x and y as well. This is justified because tilt shifts the location of the diffraction image centroid, but otherwise has no affect on image quality. For brevity, ZEMAX calls this reference point the "centroid", although it is a reference point which is usually close to but not exactly at the diffraction image centroid. Most of the time, the RMS is taken to mean the RMS referenced to the centroid, which is always the lowest of the three numbers.

The W values used by the Zernike RMS computations are the "raw" OPD values measured by computing the phase of each ray as it intercepts the reference sphere. The reference sphere is centered on the chief ray - image surface intercept point, and has a radius equal to the "exit pupil position" length (see "Exit pupil position" on page 53). This method does not consider a subtle effect due to the change in the location of the center of the reference sphere in the presence of coma. For this reason, the Zernike based estimate of the RMS wavefront referenced to the centroid may differ slightly from that computed by the RMS analysis feature; see "Comments about RMS wavefront computations" on page 155 for more information.


Subaperture computations

By default, the Zernike polynomials are determined for the entire pupil. The subaperture values S_x , S_y , and S_r allow computation of the Zernike data only over a circular subaperture portion of the pupil. The subaperture is a circular region of normalized radius S_r , decentered from the usual pupil by the normalized coordinates S_x and S_y . All three subaperture data values (S_x , S_y , and S_r) are normalized dimensionless parameters relative to the full pupil unit circle.

The transformation between full aperture and subaperture normalized pupil coordinates is given by:

$$P_x = S_x + S_r \times F_x, \text{ and } P_y = S_y + S_r \times F_y,$$

where F_x and F_y denote the subaperture coordinates and P_x and P_y denote the full aperture coordinates.



Strehl ratio approximation

The Strehl ratio is computed using the RMS referenced to the centroid by the following approximation:

$$S = e^{-(2\pi\sigma)^2}.$$

This approximation is valid for monochromatic calculations resulting in a Strehl ratio higher than about 0.10.

This feature computes a maximum of 37 Zernike terms. The particular Zernike terms used are not orthonormal, but are instead all normalized to have unity magnitude at the edge of the pupil. Some of the higher order terms in the expansion were dropped to keep the total number of terms small, and the terms remaining were selected to favor accurate fitting of higher order spherical aberration. This particular set of Zernike polynomials is sometimes called the "Fringe" or "University of Arizona" notation. The more formal, and more general polynomial set is the Standard notation, sometimes called the "Born & Wolf" or the similar "Noll" notation, which is described under the "Zernike Standard Coefficients" feature.

The Zernike Fringe polynomials are defined in the following table. The angle φ is measured counter clockwise from the local +x axis. The radial coordinate is the normalized dimensionless parameter ρ .

ZERNIKE FRINGE POLYNOMIALS

Term	$Z(\rho, \varphi)$
1	1
2	$\rho \cos \varphi$
3	$\rho \sin \varphi$
4	$2\rho^2 - 1$
5	$\rho^2 \cos 2\varphi$
6	$\rho^2 \sin 2\varphi$
7	$(3\rho^2 - 2)\rho \cos \varphi$
8	$(3\rho^2 - 2)\rho \sin \varphi$
9	$6\rho^4 - 6\rho^2 + 1$
10	$\rho^3 \cos 3\varphi$
11	$\rho^3 \sin 3\varphi$
12	$(4\rho^2 - 3)\rho^2 \cos 2\varphi$
13	$(4\rho^2 - 3)\rho^2 \sin 2\varphi$
14	$(10\rho^4 - 12\rho^2 + 3)\rho \cos \varphi$
15	$(10\rho^4 - 12\rho^2 + 3)\rho \sin \varphi$

Term	$Z(\rho, \varphi)$
16	$20\rho^6 - 30\rho^4 + 12\rho^2 - 1$
17	$\rho^4 \cos 4\varphi$
18	$\rho^4 \sin 4\varphi$
19	$(5\rho^2 - 4)\rho^3 \cos 3\varphi$
20	$(5\rho^2 - 4)\rho^3 \sin 3\varphi$
21	$(15\rho^4 - 20\rho^2 + 6)\rho^2 \cos 2\varphi$
22	$(15\rho^4 - 20\rho^2 + 6)\rho^2 \sin 2\varphi$
23	$(35\rho^6 - 60\rho^4 + 30\rho^2 - 4)\rho \cos \varphi$
24	$(35\rho^6 - 60\rho^4 + 30\rho^2 - 4)\rho \sin \varphi$
25	$70\rho^8 - 140\rho^6 + 90\rho^4 - 20\rho^2 + 1$
26	$\rho^5 \cos 5\varphi$
27	$\rho^5 \sin 5\varphi$
28	$(6\rho^2 - 5)\rho^4 \cos 4\varphi$
29	$(6\rho^2 - 5)\rho^4 \sin 4\varphi$
30	$(21\rho^4 - 30\rho^2 + 10)\rho^3 \cos 3\varphi$
31	$(21\rho^4 - 30\rho^2 + 10)\rho^3 \sin 3\varphi$
32	$(56\rho^6 - 105\rho^4 + 60\rho^2 - 10)\rho^2 \cos 2\varphi$
33	$(56\rho^6 - 105\rho^4 + 60\rho^2 - 10)\rho^2 \sin 2\varphi$
34	$(126\rho^8 - 280\rho^6 + 210\rho^4 - 60\rho^2 + 5)\rho \cos \varphi$
35	$(126\rho^8 - 280\rho^6 + 210\rho^4 - 60\rho^2 + 5)\rho \sin \varphi$
36	$252\rho^{10} - 630\rho^8 + 560\rho^6 - 210\rho^4 + 30\rho^2 - 1$
37	$924\rho^{12} - 2772\rho^{10} + 3150\rho^8 - 1680\rho^6 + 420\rho^4 - 42\rho^2 + 1$

Zernike Standard Coefficients

Purpose:

Calculates the orthonormal Zernike coefficients computed using the notation defined in R. Noll, "Zernike polynomials and atmospheric turbulence", J. Opt. Soc. Am., Vol. 66, No. 3, p207 (1976). This notation allows terms to be consistently defined up to (almost) arbitrary order, rather than the limited 37 term Fringe polynomials.

Settings:

Item	Description
Sampling	Specify the density in the pupil to use for coefficient fitting. Larger grid sizes are more accurate, although the computation time increases.
Maximum Term	Specify the maximum Zernike coefficient to compute. Any value up to 231 may be specified.
Wavelength	The wavelength number to use for the calculation.
Field	The field number to use for the calculation.
Ref OPD To Vertex	Normally, ZEMAX references the OPD to the chief ray; which in effect subtracts out tilt from the wavefront phase. For interferometry, it is sometimes desirable to retain the wavefront tilt. Checking this option on will add tilt to the beam based upon the deviation of the chief ray from the image surface vertex. This option is only useful for field positions whose chief ray is reasonably close to the surface vertex, where the assumption that tilt is described by the deviation of the chief ray is valid.
Surface	Selects the surface at which the data is to be evaluated. This is useful for evaluating intermediate images. See "Evaluating results at intermediate surfaces" on page 128.
Subaperture Data: Sx, Sy, Sr	Defines a subaperture of the pupil to compute the Zernike data for. See "Subaperture computations" on page 197 for details.

Discussion:

This feature is similar to the "Zernike Fringe Coefficients" feature, with a somewhat different numbering scheme is used, more terms in the expansion are available for fitting, the terms are orthogonal and normalized, and none of the terms are skipped. See the description of "Zernike Fringe Coefficients" on page 196 for some important details.

The most general way to express the Zernike polynomials is in the form

$$R_n^m(\rho)e^{im\theta} = \begin{cases} R_n^m(\rho) \cos m\theta \\ R_n^m(\rho) \sin m\theta \end{cases}$$

The radial portion of the polynomial is defined by two indices, n and m. The n index defines the order of the radial power; so an n value of 5 would indicate all polynomials whose maximum radial power was ρ^5 . Only certain values for m are allowed once n is chosen; n + m must be even, and $0 \leq m \leq n$. For details on Zernike Standard polynomials, see "Principles of Optics", by Born & Wolf, referenced in Chapter 1, or the Noll reference listed above. The terms are orthonormal such that the magnitude of the coefficient of each term is the RMS contribution of the term.

This feature lists the formulas next to the fitted coefficients; the entire 231 term table is too long to include here. The first 28 terms are given below. The angle φ is measured counter clockwise from the local +x axis. The radial coordinate is the normalized dimensionless parameter ρ .

ZERNIKE STANDARD POLYNOMIALS

Term	$Z(\rho, \varphi)$
1	1
2	$\sqrt{4}\rho \cos \varphi$
3	$\sqrt{4}\rho \sin \varphi$
4	$\sqrt{3}(2\rho^2 - 1)$
5	$\sqrt{6}(\rho^2 \sin 2\varphi)$
6	$\sqrt{6}(\rho^2 \cos 2\varphi)$
7	$\sqrt{8}(3\rho^3 - 2\rho) \sin \varphi$
8	$\sqrt{8}(3\rho^3 - 2\rho) \cos \varphi$
9	$\sqrt{8}\rho^3 \sin 3\varphi$
10	$\sqrt{8}\rho^3 \cos 3\varphi$
11	$\sqrt{5}(6\rho^4 - 6\rho^2 + 1)$
12	$\sqrt{10}(4\rho^4 - 3\rho^2) \cos 2\varphi$
13	$\sqrt{10}(4\rho^4 - 3\rho^2) \sin 2\varphi$
14	$\sqrt{10}\rho^4 \cos 4\varphi$
15	$\sqrt{10}\rho^4 \sin 4\varphi$
16	$\sqrt{12}(10\rho^5 - 12\rho^3 + 3\rho) \cos \varphi$
17	$\sqrt{12}(10\rho^5 - 12\rho^3 + 3\rho) \sin \varphi$
18	$\sqrt{12}(5\rho^5 - 4\rho^3) \cos 3\varphi$
19	$\sqrt{12}(5\rho^5 - 4\rho^3) \sin 3\varphi$
20	$\sqrt{12}\rho^5 \cos 5\varphi$
21	$\sqrt{12}\rho^5 \sin 5\varphi$
22	$\sqrt{7}(20\rho^6 - 30\rho^4 + 12\rho^2 - 1)$

Term	$Z(\rho, \varphi)$
23	$\sqrt{14}(15\rho^6 - 20\rho^4 + 6\rho^2)\sin 2\varphi$
24	$\sqrt{14}(15\rho^6 - 20\rho^4 + 6\rho^2)\cos 2\varphi$
25	$\sqrt{14}(6\rho^6 - 5\rho^4)\sin 4\varphi$
26	$\sqrt{14}(6\rho^6 - 5\rho^4)\cos 4\varphi$
27	$\sqrt{14}\rho^6 \sin 6\varphi$
28	$\sqrt{14}\rho^6 \cos 6\varphi$

Zernike Annular Coefficients

Purpose:

Calculates the orthonormal Zernike coefficients computed using the notation defined in R. Noll, "Zernike polynomials and atmospheric turbulence", J. Opt. Soc. Am., Vol. **66**, No. 3, p207 (1976), with extensions for annular pupils, as described in V. N. Mahajan, "Zernike annular polynomials for imaging systems with annular pupils", J. Opt. Soc. Am., Vol. **71**, No. 1, p75 (1981). Thanks are due to Brett Patterson of the UK Astronomy Technology Centre who assisted in the development of this feature.

Settings:

Item	Description
Sampling	Specify the density in the pupil to use for coefficient fitting. Larger grid sizes are more accurate, although the computation time increases.
Maximum Term	Specify the maximum Zernike coefficient to compute. Any value up to 231 may be specified.
Wavelength	The wavelength number to use for the calculation.
Field	The field number to use for the calculation.
Ref OPD To Vertex	Normally, ZEMAX references the OPD to the chief ray; which in effect subtracts out tilt from the wavefront phase. For interferometry, it is sometimes desirable to retain the wavefront tilt. Checking this option on will add tilt to the beam based upon the deviation of the chief ray from the image surface vertex. This option is only useful for field positions whose chief ray is reasonably close to the surface vertex, where the assumption that tilt is described by the deviation of the chief ray is valid.
Obscuration	The obscuration ratio, ε of the annular pupil.
Surface	Selects the surface at which the data is to be evaluated. This is useful for evaluating intermediate images. See "Evaluating results at intermediate surfaces" on page 128.
Subaperture Data: Sx, Sy, Sr	Defines a subaperture of the pupil to compute the Zernike data for. See "Subaperture computations" on page 197 for details.

Discussion:

This feature is similar to the "Zernike Standard Coefficients" feature, except the pupil may be annular rather than circular. See that discussion, and the references above, for important information. The annular nature of the pupil is defined by the single parameter ε , defined as

$$\varepsilon = \frac{R_{inner}}{R_{outer}},$$

where R_{inner} and R_{outer} are the inner and outer radii of the annular pupil. When ε is zero, the annulus is a circle, and the Annular and Standard Zernike polynomials are identical. When ε is greater than zero, the annular Zernike polynomials are distinct from the Standard Zernike polynomials. ZEMAX only traces rays in the annular region when fitting annular polynomials, and unvignetted wavefront outside the annular region is ignored.

The annular polynomials are too complex in general to list on the text output window, so the first 22 terms are listed here. The angle φ is measured counter clockwise from the local +x axis. The radial coordinate is the normalized dimensionless parameter ρ .

ZERNIKE ANNULAR POLYNOMIALS

Term	$Z(\rho, \varphi, \varepsilon)$
1	1
2	$\frac{2\rho \cos \varphi}{\sqrt{1 + \varepsilon^2}}$
3	$\frac{2\rho \sin \varphi}{\sqrt{1 + \varepsilon^2}}$
4	$\frac{\sqrt{3}(1 + \varepsilon^2 - 2\rho^2)}{\varepsilon^2 - 1}$
5	$\frac{\sqrt{6}\rho^2 \sin 2\varphi}{\sqrt{1 + \varepsilon^2 + \varepsilon^4}}$
6	$\frac{\sqrt{6}\rho^2 \cos 2\varphi}{\sqrt{1 + \varepsilon^2 + \varepsilon^4}}$
7	$\frac{2\sqrt{2}\rho(-2 - 2\varepsilon^4 + 3\rho^2 + \varepsilon^2(-2 + 3\rho^2)) \sin \varphi}{\sqrt{(\varepsilon^2 - 1)^2(1 + \varepsilon^2)(1 + 4\varepsilon^2 + \varepsilon^4)}}$
8	$\frac{2\sqrt{2}\rho(-2 - 2\varepsilon^4 + 3\rho^2 + \varepsilon^2(-2 + 3\rho^2)) \cos \varphi}{\sqrt{(\varepsilon^2 - 1)^2(1 + \varepsilon^2)(1 + 4\varepsilon^2 + \varepsilon^4)}}$
9	$\frac{2\sqrt{2}\rho^3 \sin 3\varphi}{\sqrt{1 + \varepsilon^2 + \varepsilon^4 + \varepsilon^6}}$

Term	$Z(\rho, \varphi, \varepsilon)$
10	$\frac{2\sqrt{2}\rho^3 \cos 3\varphi}{\sqrt{1 + \varepsilon^2 + \varepsilon^4 + \varepsilon^6}}$
11	$\frac{\sqrt{5}(1 + \varepsilon^4 - 6\rho^2 + 6\rho^4 + \varepsilon^2(4 - 6\rho^2))}{(\varepsilon^2 - 1)^2}$
12	$\frac{\sqrt{10} \sqrt{\frac{(\varepsilon^2 - 1)^4 (1 + \varepsilon^2 + \varepsilon^4)}{1 + 4\varepsilon^2 + 10\varepsilon^4 + 4\varepsilon^6 + \varepsilon^8}} \rho^2 (3 + 3\varepsilon^6 - 4\rho^2 + \varepsilon^2(3 - 4\rho^2) + \varepsilon^4(3 - 4\rho^2)) \cos 2\varphi}{(\varepsilon^2 - 1)^3 (1 + \varepsilon^2 + \varepsilon^4)}$
13	$\frac{\sqrt{10} \sqrt{\frac{(\varepsilon^2 - 1)^4 (1 + \varepsilon^2 + \varepsilon^4)}{1 + 4\varepsilon^2 + 10\varepsilon^4 + 4\varepsilon^6 + \varepsilon^8}} \rho^2 (3 + 3\varepsilon^6 - 4\rho^2 + \varepsilon^2(3 - 4\rho^2) + \varepsilon^4(3 - 4\rho^2)) \sin 2\varphi}{(\varepsilon^2 - 1)^3 (1 + \varepsilon^2 + \varepsilon^4)}$
14	$\frac{\sqrt{10}\rho^4 \cos 4\varphi}{\sqrt{1 + \varepsilon^2 + \varepsilon^4 + \varepsilon^6 + \varepsilon^8}}$
15	$\frac{\sqrt{10}\rho^4 \sin 4\varphi}{\sqrt{1 + \varepsilon^2 + \varepsilon^4 + \varepsilon^6 + \varepsilon^8}}$
16	$\frac{(-2\sqrt{3})EQ\rho \cos \varphi}{(\varepsilon^2 - 1)^3 (1 + 4\varepsilon^2 + \varepsilon^4)}, \text{ where}$ $E = \sqrt{\frac{(\varepsilon^2 - 1)^2 (1 + 4\varepsilon^2 + \varepsilon^4)}{1 + 9\varepsilon^2 + 9\varepsilon^4 + \varepsilon^6}}, \text{ and}$ $Q = (3 + 3\varepsilon^8 - 12\rho^2 + 10\rho^4 - 12\varepsilon^6(\rho^2 - 1) + 2\varepsilon^4(15 - 24\rho^2 + 5\rho^4) + 4\varepsilon^2(3 - 12\rho^2 + 10\rho^4)).$
17	$\frac{(-2\sqrt{3})EQ\rho \sin \varphi}{(\varepsilon^2 - 1)^3 (1 + 4\varepsilon^2 + \varepsilon^4)}, \text{ where E and Q are defined as in term 16 above.}$
18	$\frac{(-2\sqrt{3})EQ\rho^3 \cos 3\varphi}{(\varepsilon^2 - 1)^4 (1 + \varepsilon^2)(1 + \varepsilon^4)}, \text{ where}$ $E = \sqrt{\frac{(\varepsilon^2 - 1)^6 (1 + \varepsilon^2)(1 + \varepsilon^4)}{1 + 4\varepsilon^2 + 10\varepsilon^4 + 20\varepsilon^6 + 10\varepsilon^8 + 4\varepsilon^{10} + \varepsilon^{12}}}, \text{ and}$ $Q = 4 + 4\varepsilon^8 - 5\rho^2 + \varepsilon^2(4 - 5\rho^2) + \varepsilon^4(4 - 5\rho^2) + \varepsilon^6(4 - 5\rho^2).$
19	$\frac{(-2\sqrt{3})EQ\rho^3 \sin 3\varphi}{(\varepsilon^2 - 1)^4 (1 + \varepsilon^2)(1 + \varepsilon^4)}, \text{ where E and Q are defined as in term 18 above.}$

Term	$Z(\rho, \varphi, \varepsilon)$
20	$\frac{2\sqrt{3}\rho^5 \cos 5\varphi}{\sqrt{1 + \varepsilon^2 + \varepsilon^4 + \varepsilon^6 + \varepsilon^8 + \varepsilon^{10}}}$
21	$\frac{2\sqrt{3}\rho^5 \sin 5\varphi}{\sqrt{1 + \varepsilon^2 + \varepsilon^4 + \varepsilon^6 + \varepsilon^8 + \varepsilon^{10}}}$
22	$\frac{\sqrt{7}(1 + \varepsilon^6 - 12\rho^2 + 30\rho^4 - 20\rho^6 - 3\varepsilon^4(-3 + 4\rho^2) + 3\varepsilon^2(3 - 12\rho^2 + 10\rho^4))}{(\varepsilon^2 - 1)^3}$

Calculations

Ray Trace

Purpose:

Paraxial and real trace of a single ray.

Settings:

Item	Description
Hx	Normalized x-field coordinate. The value should be between -1 and 1.
Hy	Normalized y-field coordinate. The value should be between -1 and 1.
Field	Select either a specific field or "Arbitrary" to enter in Hx and Hy. If a specific field is selected; the Hx and Hy controls are grayed out.
Wavelength	The wavelength number of the ray to trace.
Px	Normalized x-pupil coordinate. The value should be between -1 and 1.
Py	Normalized y-pupil coordinate. The value should be between -1 and 1.
Global Coordinates	If checked, all ray trace data is given in global coordinates rather than local coordinates, except tangent angles.
Type	Select "Dir Cosines" to display the ray direction cosines at each surface, "Tangent Ang" to display the tangent of the angle the ray makes at each surface, or "Ym,Um,Yc,Uc" to display the paraxial marginal and chief ray intercept/tangent angle display. The tangent angle is the ratio of the x (or y) direction cosine to the z direction cosine.

Discussion:

The Hx, Hy, Px, Py, and Global Coordinates settings are ignored if the "Ym, Um, Yc, Uc" option is selected. For a description of normalized coordinates, see the chapter "Conventions and Definitions".

For the other options, this feature allows the user to specify the normalized object coordinates, the normalized pupil coordinates, and wavelength number and then view the real and paraxial ray coordinates at every surface.

The first set of data presented is for a real ray. The values presented are the coordinates (in the surface local or the global coordinate system) of the ray intercept point. The direction cosines (or tangent angles) are the values of the ray in the surface after refraction. The direction cosine value is the cosine of the angle the ray makes with respect to the specified axis (the x-direction cosine is the cosine of the angle the ray makes with respect to the x-axis, etc.) The second set of data is just like the first set, except it is for a paraxial ray. Tangent angles are always computed relative to the local Z axis; regardless of the Global Coordinates setting.

Fiber Coupling Efficiency

Purpose:

This feature computes the coupling efficiency for single mode fiber coupling systems. For multi-mode fiber coupling, see “Calculating efficiency of multi-mode fibers” on page 169. Also see “Computing Fiber Coupling” on page 645.

Settings:

Item	Description
Source Fiber	
NA x/y	Numerical Aperture of the source fiber in object space in the xz and yz planes respectively. This is n times the sin of the half angle to the 1 over e squared intensity point.
X Angle, Y Angle	The angular rotation of the source fiber in object space, in degrees, measured from the nominal orientation of the source fiber. The X angle is the angle measured in the XZ plane; the Y angle is measured in the YZ plane.
Field	The field number to use for the calculation.
Wavelength	The wavelength number to use for the calculation.
Sampling	The grid size to use for the numerical integration.
Align Source to Chief Ray	If unchecked, the fiber is oriented along the object z axis, centered on the field point. If checked, the center of the fiber points along the chief ray for that field point.
Ignore Source Fiber	If checked, then illumination of the pupil as specified by the System, General, Apodization setting will be used. All computations are then referenced to the total energy incident upon the entrance pupil.
Receiving Fiber	
NA x/y	Numerical Aperture of the receiving fiber in image space in the xz and yz planes respectively. This is n times the sin of the half angle to the 1 over e squared intensity point.
Tilt About X, Y	The angular rotation of the receiving fiber in image space, in degrees, measured from the nominal orientation of the receiving fiber. If both X and Y rotations are specified, the rotation is first around X, then around Y. Note the rotations are about the specified axes; so Tilt About X rotates the fiber toward or away from the Y axis. Rotations are done <u>after</u> the XYZ decenters below.
Decenter XYZ	The linear position shift in the x/y/z direction in lens units of the receiving fiber, measured from the nominal position of the receiving fiber. Shifts are done <u>before</u> the XY tilts above.
Align Receiver to Chief Ray	If unchecked, the fiber is positioned at XY coordinates (0.0, 0.0) on the image surface. If checked, the fiber is repositioned in x and y to be centered at the point the chief ray intercepts the image surface. In either case the receiving fiber is oriented along the image space z axis, unless rotated using the controls above.
Use Polarization	If checked, polarization is considered. See “Polarization (ZEMAX-EE only)” on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Use Huygens Integral	If checked, the Huygens PSF is used for the beam mode at the receiving fiber. Computing the Huygens PSF is much slower than the direct ray-based integration used if this option is unchecked. However, there are cases where the full Huygens PSF is required. These cases generally have severe caustics or discontinuities in the wavefront. The recommended approach is to leave this option unchecked. If the fast ray-based integration cannot be computed accurately, the fiber coupling will be zero and a warning will be printed which will recommend switching to the Huygens algorithm.

Discussion:

This feature computes fiber coupling for single-mode fibers with a Gaussian shaped mode. For multi-mode fiber coupling, see “Calculating efficiency of multi-mode fibers” on page 169. For a more detailed and flexible physical optics treatment of fiber coupling see “Computing Fiber Coupling” on page 645.

Fiber coupling efficiency is computed based upon a two fiber or a one fiber model. In the two fiber model, light emerges from a source fiber to fill or overfill the entrance pupil of an optical system. Energy not collected by the entrance pupil is lost, reducing the overall efficiency. If desired, the source fiber may be ignored for a one fiber model; and in this case the efficiency is computed relative to the energy entering the entrance pupil; which in turn is a function of the system apodization (see “Apodization Type” on page 102).

The system efficiency (S) is the sum of the energy collected by the entrance pupil which passes through the optical system, accounting for both the vignetting and transmission of the optics (if polarization is used), divided by the sum of all the energy which radiates from the source fiber:

$$S = \frac{\left[\iint t(x, y) F_s(x, y) dx dy \right]^2}{\iint F_s(x, y) dx dy},$$

where F_s is the source fiber amplitude function and the integral in the numerator is only done over the entrance pupil of the optical system, and $t(x, y)$ is the amplitude transmission function of the optics. The transmission is affected by bulk absorption and optical coatings (if use polarization is checked on). If the source fiber is ignored, then the integral in the denominator is only done over the entrance pupil, and the F_s function is determined by the system apodization, if any.

Aberrations in the optical system introduce phase errors which will affect the coupling into the fiber. Maximum coupling efficiency is achieved when the mode of the wavefront converging towards the receiving fiber perfectly matches the mode of the fiber in both amplitude and phase at all points in the wavefront. This is defined mathematically as a normalized overlap integral between the fiber and wavefront amplitude:

$$T = \frac{\left| \iint F_r(x, y) W'(x, y) dx dy \right|^2}{\iint F_r(x, y) F_r'(x, y) dx dy \iint W(x, y) W'(x, y) dx dy},$$

where $F_r(x, y)$ is the function describing the receiving fiber complex amplitude, $W(x, y)$ is the function describing the complex amplitude of the wavefront from the exit pupil of the optical system, and the ' symbol represents complex conjugate. Note that these functions are complex valued, so this is a coherent overlap integral.

T has a maximum possible value of 1.0, and will decrease if there is any mismatch between the fiber amplitude and phase and the wavefront amplitude and phase.

ZEMAX computes the values S and T. The total power coupling efficiency is the product of these numbers. A theoretical maximum coupling efficiency is also computed; this value is based upon ignoring the aberrations but accounting for all vignetting, transmission, and other amplitude mismatches between the modes.

To model the reflection loss from the front face of the receiving fiber, set the image space glass to the material the fiber is made of, and check the box "use polarization". The reflectivity of the receiving fiber (and other surfaces) will be accounted for.

YNI Contributions

Purpose:

This feature lists for each surface the paraxial YNI value, which is proportional to the Narcissus contribution of that surface.

Settings:

Item	Description
Wavelength	The wavelength number to use for the calculation.

Discussion:

See “Narcissus: reflections on retroreflections in thermal imaging systems”, Applied Optics, Vol. 21, No. 18, p3393 (1982) for a discussion.

Sag Table

Purpose:

This feature lists for the selected surface the sag (z-coordinate) at various Y-coordinate distances from the vertex. The best fit spherical (BFS) radius is computed, and the sag of the BFS and the difference between the BFS and the surface is also tabulated. The depth of material that would need to be removed from the BFS to yield the asphere is also listed. The dimensionless slope (dz/dr) of the actual surface and the BFS are given, and the difference between the two.

Only the Y-coordinate of the surface is considered, therefore this feature should not be used if the surface is non-rotationally symmetric.

Settings:

Item	Description
Surface	The surface number to use for the calculation.
Step Size	The distance between steps measured from the vertex at which the sag is computed. The default value of zero will automatically choose a reasonable step size.
Reverse Direction	See "Discussion".
MinR, MaxR	The minimum and maximum radial aperture over which the fit is performed. This feature is intended to support radial symmetric annular optics. If both of these values are zero, the MinR is 0.0 and the MaxR is the semi-diameter of the surface.

Discussion:

This feature can be used to determine the maximum aspheric deviation of a surface. The sag table is also useful for lens fabrication.

This feature lists out eight columns of data:

Y-Coord: The y coordinate of the point on the surface being computed.

Sag: The sag of the surface at the y coordinate.

BFS Sag: The sag of the best fit sphere. The BFS is determined by the radius of the sphere that minimizes the volume of material that would need to be removed from a spherical surface to yield the aspheric surface.

Deviation: The difference between the BFS Sag and the Sag, ignoring any vertex offset (see Remove below).

Remove: The depth of material at the specified y coordinate to remove assuming that the surface was first generated to the best fit sphere. Note that for some aspheric surfaces, material may need to be removed from the vertex; for other types of aspheres the minimum volume to remove is achieved by removing more material from the central annular zones of the BFS.

Slope: The slope (dz/dr) of the surface.

BFS Slope: The slope of the BFS surface.

Slope Delta: The difference between the Slope and the BFS Slope values.

The BFS radius and position (vertex offset) relative to the aspheric surface depends on the properties of the aspheric surface, and whether the surface defines a boundary between air and glass, or glass and air, and whether the surface is concave or convex. ZEMAX always attempts to choose a BFS that will "add" material on the air side of the surface. The vertex offset and BFS radius will be selected to minimize the volume of material that needs to be removed from the BFS to yield the asphere surface. For some surfaces, such as those between glass in a cemented doublet, ZEMAX cannot determine which side is "air". For these surfaces, if the desired data is not generated by default, selecting "Reverse Direction" will choose the other side of the surface as the air surface. The user must check the data and sign convention carefully before any fabrication decisions are made.

Cardinal Points

Purpose:

This feature lists for the selected range of surfaces and wavelengths the locations of the focal, principal, anti-principal, nodal, anti-nodal planes. The calculation is done for any defined wavelength and either the X-Z or Y-Z orientation.

Settings:

Item	Description
First Surface	The starting surface number of the group to compute the cardinal points for.
Last Surface	The ending surface number of the group to compute the cardinal points for.
Wavelength	The wavelength number to use for the computation.
Orientation	The orientation to use for computing the cardinal plane locations.

Discussion:

This feature may not return reliable results if coordinate breaks or non-centered optics are included within the specified surface range.

Glass and Gradient Index

Dispersion Diagram

Purpose:

Plots index of refraction as a function of wavelength for any material in the glass catalog.

Settings:

Item	Description
Min Wavelength	Defines the left X- axis of the dispersion plot.
Max Wavelength	Defines the right X- axis of the dispersion plot.
Minimum Index	Defines the bottom Y- axis of the dispersion plot. Enter zero for automatic scaling.
Maximum Index	Defines the top Y- axis of the dispersion plot. Enter zero for automatic scaling.
Glass	The name of the material.
Use Temperature, Pressure	If checked, then the change in index due to temperature and pressure effects will be considered.

Discussion:

This is useful for checking if constants of dispersion or other formula data were entered correctly. Table glass files (see "Using table glasses" on page 601) are listed at the bottom of the "Glass" control.

Glass Map

Purpose:

Plots names of glasses on map according to the d-light index of refraction and Abbe V-number. The index and Abbe number are computed from the data in the glass catalog. All currently loaded glass catalogs are searched, and glasses within the boundary values specified in the following table are plotted.

Settings:

Item	Description
Min Abbe	Defines the left X- axis of the plot.
Max Abbe	Defines the right X- axis of the plot.
Min Index	Defines the bottom Y- axis of the plot. Enter zero for automatic scaling.
Max Index	Defines the top Y- axis of the plot. Enter zero for automatic scaling.
Apply Glass Substitution Template	If checked, only glasses that meet the glass substitution template will be shown. See "Glass Substitution Template" on page 240.

Discussion:

This is useful for locating a glass with particular refractive and dispersive properties. By convention, the glass map is shown with the Abbe number decreasing while going from left to right, which is why the default min and max Abbe numbers seem to be reversed.

Athermal Glass Map

Purpose:

Plots names of glasses on a map according to the chromatic and thermal power. The chromatic and thermal power are computed based upon the currently defined wavelengths and the data in the glass catalog. All currently loaded glass catalogs are searched, and glasses within the boundary values specified in the following table are plotted.

Settings:

Item	Description
Min/Max Chrom. Power	Minimum/Maximum Chromatic power to plot. If both the Min and Max Chromatic power are zero, a default scaling is selected.
Min/Max Therm. Power	Minimum/Maximum Thermal power to plot. If both the Min and Max Thermal power are zero, a default scaling is selected.
Apply Glass Substitution Template	If checked, only glasses that meet the glass substitution template will be shown. See "Glass Substitution Template" on page 240.
Reference Glass	If a glass is selected, a reference line is drawn from the origin of the plot to the selected glass. Glasses along this line may be used to make athermalized achromatic doublets. See the discussion below.

Discussion:

The Chromatic power, which is the inverse of a generalized Abbe V-number, is given by

$$\varphi = \frac{n_{min} - n_{max}}{n_{cen} - 1},$$

where the subscripts on the index values refer to the relative refractive index measured at the minimum, center, and maximum wavelengths, respectively. The wavelength values are determined by the defined wavelengths.

The Thermal power is given by

$$\psi = \frac{\frac{\partial n}{\partial t}}{n_{cen} - 1} - \alpha,$$

where $\frac{\partial n}{\partial t}$ is the derivative of the relative refractive index with respect to temperature measured at the center wavelength and the reference temperature of the glass, and α is the linear expansion coefficient of the glass.

All these values are computed using the data for the glass defined in the glass catalog. The air pressure is assumed to be 1.0.

To make a doublet that is both achromatic and athermalized requires the solution of three simultaneous equations:

$$\Phi = \Phi_1 + \Phi_2,$$

$$\Delta\Phi = \varphi_1\Phi_1 + \varphi_2\Phi_2, \text{ and}$$

$$\frac{d\Phi}{dT} = \psi_1\Phi_1 + \psi_2\Phi_2.$$

To solve these three equations simultaneously with two glasses requires selection of glasses that lie along a line that also crosses the origin of the athermal glass map. This reference line may be drawn on the map by selecting a reference glass, which defines one point on the line (the other point is the origin of the map).

Internal Transmittance vs. Wavelength

Purpose:

Plots the internal transmittance for any thickness as a function of wavelength for any material in the glass catalog.

Settings:

Item	Description
Min Wavelength	Defines the left X- axis of the plot.
Max Wavelength	Defines the right X- axis of the plot.
Minimum Transmission	Defines the bottom Y- axis of the plot.
Maximum Transmission	Defines the top Y- axis of the plot. Enter zero for automatic scaling.
Glass	The name of the material.
Thickness	The thickness of the glass in millimeters.

Discussion:

This is useful for checking the transmittance of a particular glass. See the Chapter “Polarization Analysis”.

Grin Profile

Purpose:

Plots the index of refraction along any axis for gradient index surfaces.

Settings:

Item	Description
Surface	The surface number of the gradient index surface.
Wavelength	The wavelength number to use.
Show Index vs.	Choose X, Y, or Z as the independent variable.
Min X, Y, or Z	Defines the left X- axis of the plot, the minimum value of the independent variable.
Max X, Y, or Z	Defines the right X- axis of the plot, the maximum value of the independent variable.
X, Y, or Z Value	The value of the two fixed coordinates.
Minimum Index	Defines the bottom Y- axis of the plot. Use zero for the default value.
Maximum Index	Defines the top Y- axis of the plot. Enter zero for the default value.

Discussion:

This plot shows index of refraction along either the X, Y, or Z axis for gradient index surfaces.

Gradium™ Profile

Purpose:

Plots the axial index profile of a Gradium surface.

Settings:

Item	Description
Profile	The name of the profile. See the chapter on surface types for details.
Surface	The surface number to plot.
Wavelength	The wavelength number to use for the calculation.

Discussion:

If the surface is selected to "None", then the reference wavelength for the profile glass family is used, no matter what wavelength is selected. If a Gradium surface is selected, then any defined wavelength, or the reference wavelength may be selected. Also, if a surface number is selected, then the starting and ending points for the glass blank will be indicated using an "X" on the plot. The starting and ending positions include consideration of the sag of the surface at the defined semi-diameter.

Universal Plot



This feature is only available in the EE edition of ZEMAX.

New Universal Plot 1D...

Purpose:

Displays as either a plot or as a text listing the value of any optimization operand as a function of one other parameter. See also "New Universal Plot 2D..." on page 214.

Settings:

Item	Description
Independent Variable X Settings	
	Select either surface, system, multi-configuration, or NSC data as the independent variable.
	Select the parameter to use as the independent variable. The available selections include radius, curvature, thickness, conic, parameter, and extra data values if surface parameters are chosen. For system parameters, the selections include the system aperture, field and wavelength data, apodization factor, temperature, and pressure. For configuration data, all multi-configuration operands are listed. For NSC data, object position and parameter data are listed.
	The surface number for the independent variable parameter if the parameter is associated with a surface. This field is used for the configuration number if configuration data is selected. If NSC data is selected, this field is used for the object number.
Start/Stop Value	The beginning and ending range of the independent variable.
# of Steps	The number of values between the start and stop values, inclusive, to compute the dependent variable function value.
Dependent Variable Y Settings	
Operand	The optimization operand to use as the dependent variable function. The current system merit function is also included as "Merit". If the system merit function is selected, then either the entire merit function or any specific operand's value may be selected using the "Line" control below.
Line	The optimization operand # to use as the dependent variable function if "Merit" is selected for the Operand.
Int1/Int2	The Int1 and Int2 values for the selected operand. The dialog box will display the short name for the Int1 and Int2 values to aid in identifying the data that needs to be supplied. If the edit control is grayed out; then the selected operand does not use that value.
Data1-Data4	The Data values for the selected operand. The dialog box will display the column heading name for these values to aid in identifying the data that needs to be supplied. If the edit control is grayed out; then the selected operand does not use that value.
Min/Max Plot Value	The min and max value to plot for the dependent variable (Y) function. If both of these values are zero, then a default scale will be chosen. If either one is not zero, then the y range of the plot will extend from the minimum to the maximum. If the data does not fit within the specified range, then the data may plot beyond the borders of the data box.
Plot Title	Any text may be entered which will be used as the title for the plot or text listing.
Save As New Universal Plot	This will save the current settings in a Universal Plot 1D (UPL) file specified by the user. Any name which may be used as a valid file name may be entered. Although the "browse" window allows the file to be saved in any folder, ZEMAX will only recognize the file if it resides in the <data>\Miscellaneous folder (see "Folders" on page 68). Once the settings are saved, the file name will appear as a menu option under the Universal Plot 1D menu option for quick regeneration of the plot.
Load From	This will load the settings from a previously saved Universal Plot 1D (UPL) file specified by the drop down list to the right of the button.

Discussion:

This feature will either plot as a graph or create a text listing of the value of any optimization operand as a function of a single system, surface, or multi-configuration value.

For example, suppose a plot of the sagittal MTF at 30 lp/mm was needed as a function of the decenter of a lens group (which may be a useful diagnostic for tolerancing analysis). Since the MTF operand computes the sagittal MTF, the Universal Plot can generate such a graphic or text listing. See "Optimization operands" on page 489 for a list of available optimization operands. A decenter of a lens group is defined by Parameter 1 or 2 on the relevant coordinate break surface, and Parameter 1 and 2 are both listed in the available surface group parameters.

Because of the number of different plots this feature can generate, there are no intelligent defaults for either the independent or dependent settings. These values need to be carefully provided in the settings dialog box. If the optimization operand cannot be computed, an error message will be displayed and the plot will not be generated.

Because many optimization operands accept Hx and Hy values to define the field point for the calculation, these operands may require that the number of fields be set to 1, then set Hx = 0 and Hy = 1, then finally choose Y Field 1 as the independent variable for a plot of the operand as a function of field. A similar trick works for getting plots as a function of wavelength.

Some individual operands, like DIFF and SUMM, have no meaning if selected from the "Operand" column because these operands are only defined when used as part of a larger merit function. If results from a calculation like DIFF are required, then the merit function should be selected from the Operand list and the individual DIFF's within the merit function selected for the "Line".

When selecting independent variables that are controlled by Multi-Configuration operands, note that the Multi-Configuration operand will override the changes made to the independent variable. For this reason ZEMAX automatically alters the corresponding data on the Multi-Configuration operand as well.

The escape key will terminate the analysis if the computation is taking too long.

New Universal Plot 2D...

Purpose:

Displays as either a plot or as a text listing the value of any optimization operand as a function of two other parameters. See also "New Universal Plot 1D..." on page 212.

Settings:

Item	Description
Independent Variable X or Y Settings	
	Select either surface, system, multi-configuration, or NSC data as the independent variable.
	Select the parameter to use as the independent variable. The available selections include radius, curvature, thickness, conic, parameter, and extra data values if surface parameters are chosen. For system parameters, the selections include the system aperture, field and wavelength data, apodization factor, temperature, and pressure. For configuration data, all multi-configuration operands are listed. For NSC data, object position and parameter data are listed.
	The surface number for the independent variable parameter if the parameter is associated with a surface. This field is used for the configuration number if configuration data is selected. If NSC data is selected, this field is used for the object number.
Start/Stop Value	The beginning and ending range of the independent variable.
# of Steps	The number of values between the start and stop values, inclusive, to compute the dependent variable function value. The minimum value is 2.
Dependent Variable Settings	

Item	Description
Operand	The optimization operand to use as the dependent variable function. The current system merit function is also included as "Merit". If the system merit function is selected, then either the entire merit function or any specific operand's value may be selected using the "Line" control below.
Line	The optimization operand # to use as the dependent variable function if "Merit" is selected for the Operand.
Int1/Int2	The Int1 and Int2 values for the selected operand. The dialog box will display the short name for the Int1 and Int2 values to aid in identifying the data that needs to be supplied. If the edit control is grayed out; then the selected operand does not use that value.
Data1-Data4	The Data values for the selected operand. The dialog box will display the column heading name for these values to aid in identifying the data that needs to be supplied. If the edit control is grayed out; then the selected operand does not use that value.
Display Settings	
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see "The Contour Format String" on page 151.
Min/Max Plot Scale	These controls define the min and max range of the Z-direction scale for the display. If the value is set to exactly zero, a default value will be selected. To define a value of zero, use a very small number such as 1.0E-20. If the defined plot scale values do not span the resulting data, the plot may clip, saturate, or otherwise distort the data display. These settings have no affect on the surface or contour plots.
Plot Title	Any text may be entered which will be used as the title for the plot or text listing.
Save As New Universal Plot	This will save the current settings in a Universal Plot 2D (UP2) file specified by the user. Any name which may be used as a valid file name may be entered. Although the "browse" window allows the file to be saved in any folder, ZEMAX will only recognize the file if it resides in the <data>Miscellaneous folder (see "Folders" on page 68). Once the settings are saved, the file name will appear as a menu option under the Universal Plot menu option for quick regeneration of the plot.
Load From	This will load the settings from a previously saved Universal Plot 2D (UP2) file specified by the drop down list to the right of the button.

Discussion:

This feature will either plot as a graph or create a text listing of the value of any optimization operand as a function of any two system, surface, or multi-configuration values.

For example, suppose a plot of the sagittal MTF at 30 lp/mm was needed as a function of the X and Y decenter of a lens group (which may be a useful diagnostic for tolerancing analysis). Since the MTF operand computes the sagittal MTF, the Universal Plot 2D can generate such a graphic or text listing. See "Optimization operands" on page 489 for a list of available optimization operands. The X and Y decenter of a lens group is defined by Parameters 1 and 2 on the relevant coordinate break surface, and Parameters 1 and 2 are both listed in the available surface group parameters.

Because of the number of different plots this feature can generate, there are no intelligent defaults for either the independent or dependent settings. These values need to be carefully provided in the settings dialog box. If the optimization operand cannot be computed, an error message will be displayed and the plot will not be generated.

Because many optimization operands accept Hx and Hy values to define the field point for the calculation, these operands may require that the number of fields be set to 1, then set Hx = 0 and Hy = 1, then finally choose Y Field 1 as the independent variable for a plot of the operand as a function of field. A similar trick works for getting plots as a function of wavelength.

Some individual operands, like DIFF and SUMM, have no meaning if selected from the "Operand" column because these operands are only defined when used as part of a larger merit function. If results from a calculation like DIFF are required, then the merit function should be selected from the Operand list and the individual DIFF's within the merit function selected for the "Line".

When selecting independent variables that are controlled by Multi-Configuration operands, note that the Multi-Configuration operand will override the changes made to the independent variable. For this reason ZEMAX automatically alters the corresponding data on the Multi-Configuration operand as well.

The escape key will terminate the analysis if the computation is taking too long.

Polarization



These features are only available in the EE editions of ZEMAX.

For more information on polarization calculations, see "POLARIZATION ANALYSIS" on page 609.

Polarization Ray Trace

Purpose:

The polarization ray trace feature generates a text window which displays all the polarization data for a single ray.

Settings:

Item	Description
Jx, Jy	Jones electric field. see "Defining the initial polarization" on page 622.
X, Y Phase	Phase of the X, Y component of the Jones electric field in degrees.
Hx	Normalized x-field coordinate. The value should be between -1 and 1.
Hy	Normalized y-field coordinate. The value should be between -1 and 1.
Px	Normalized x-pupil coordinate. The value should be between -1 and 1.
Py	Normalized y-pupil coordinate. The value should be between -1 and 1.
Wavelength	The wavelength number of the ray to trace.
Global Coordinates	If checked, all coordinate and cosine data are given in global rather than local coordinates.

Discussion:

This feature tabulates all of the data computed by ZEMAX to perform polarization ray tracing. See "POLARIZATION ANALYSIS" on page 609.

Polarization Pupil Map

Purpose:

Generates a graph of the polarization ellipse as a function of pupil position. This aids in visualizing the change in polarization over the pupil. This feature may also be used to compute transmission in systems modeled using interfering paths in separate configurations, such as birefringent polarizing beam splitters and interferometers.

Settings:

Item	Description
Jx, Jy	Jones electric field. see "Defining the initial polarization" on page 622.
X, Y Phase	Phase of the X, Y component of the Jones electric field in degrees.
Wavelength	The wavelength number of the rays to trace.
Field	The field position number of the rays to trace.
Surface	The surface at which the data is shown. Data is after refraction through the specified surface.
Sampling	The grid size of the sampling in the pupil.
Add Configs Sub Configs	If the "Add Configs" and "Sub Configs" settings are left blank, the polarization ellipse and transmission is computed for a single configuration (defined by the Window, Configuration menu option; see "Graphic windows operations" on page 42) and no interference effects are considered. To model interference between configurations see the discussion below.

Discussion:

The polarization ellipse is a representation of the figure traced out by the electric field vector as the wave propagates during one cycle. The magnitude of the ellipse is determined by the transmission of the ray, which is generally a function of pupil position. See "POLARIZATION ANALYSIS" on page 609.

The text of the data computed by this feature lists the Px, Py, Ex, Ey, Intensity, and Phase values. These are the normalized pupil x and y coordinates, electric field magnitude x and y components, the electric field intensity ($E_x^2 + E_y^2$), and the phase *difference* between the Ex and Ey phase. Note Ex and Ey are both complex and the values listed here are the magnitude of the amplitude, or $\sqrt{E \cdot E}$.

Modeling interference between configurations

Some optical systems require more than one configuration to model the complete optical beam. The most common example is a polarizing beamsplitter that uses uniaxial crystals to separate orthogonal polarizations of the incident electric field. Typically, these beamsplitters implement 2 crystals with orthogonal axis orientations, each requiring an ordinary and an extraordinary ray trace configuration, and therefore these systems require a total of 4 configurations to model all possible ray paths.

Computing transmission through these systems can be complicated. If a single incident ray is decomposed into 4 individual rays, one for each configuration, each ray will carry a portion of the electric field and energy. It is a fundamental property of rays that the ray trace does not depend upon other rays in other configurations. Thus, if each ray is propagated individually, each will carry some portion of the electric field. If in two or more configurations the ray paths are superimposed, then the coherent sum of the individual ray paths must be considered. It is possible that two or more configurations create electric fields of comparable magnitude but opposing phase, resulting in destructive interference. This destructive interference results in a lower total transmission for these rays. From a physical optics viewpoint, this destructive interference means the energy could never have propagated along these rays in the first place, and must have propagated elsewhere. There is no practical, general way for ZEMAX to compute where this energy actually went.

To address this difficulty, the "Add Configs" and "Sub Configs" allow user specification of which configurations should be coherently summed. The "Add" configurations are all ray traced, and the electric fields are coherently summed. The "Sub" configurations are also all ray traced, and these are separately coherently summed. As part of the coherent sum, ZEMAX can determine how much energy was lost due to destructive interference between the rays in the "Sub" configurations. This energy is then added to the "Add" configurations, and the total transmission is then determined. To list multiple configurations, separate the configuration numbers with a blank space. For example, to define configurations 1 and 2 as the "Add" configurations enter "1 2".

This feature is particularly useful for analysis of birefringent systems, where two or more configurations (ordinary and extraordinary) are required to trace all possible beam paths. The computations may be applicable

to other systems as well, but great care should be taken before any computed results are trusted. This feature may yield meaningless or misleading results if the "Add" and "Sub" configurations are defined in a way that does not reflect how the beams actually interfere with one another. If the same configuration number is listed in both the "Add" and the "Sub" configuration lists, the duplicate configurations in the "Sub" list are ignored.

Transmission

Purpose:

Computes the integrated and surface by surface transmission through the optical system considering polarization effects.

Settings:

Item	Description
Jx, Jy	Jones electric field. see "Defining the initial polarization" on page 622.
X, Y Phase	Phase of the X, Y component of the Jones electric field in degrees.
Sampling	Specify the density in the pupil. Larger grid sizes are more accurate, although the computation time increases.
Unpolarized	If checked, the electric field definition is ignored and the unpolarized computation is performed.

Discussion:

This feature tabulates for each field position and wavelength the integrated transmission of the optical system for the specified polarization. The transmission is computed as a fraction of 100%, with 100% being that transmitted if there were no absorption, reflection, or vignetting losses. The transmission calculation accounts for vignetting factors, vignetting due to apertures or obscurations, ray clipping due to ray trace errors, surface Fresnel or coating losses, and bulk internal transmittance due to absorption.

Also tabulated for each field and wavelength is the relative and total transmission of the chief ray. This allows identification of where significant surface losses originate.

See also "Relative Illumination" on page 187.

Phase Aberration

Purpose:

Computes the polarization induced phase aberration of an optical system.

Settings:

Item	Description
Jx, Jy	Jones electric field. see "Defining the initial polarization" on page 622.
X, Y Phase	Phase of the X, Y component of the Jones electric field in degrees.
Wavelength	The wavelength number of the rays to trace.
Field	The field position number of the rays to trace.

Discussion:

Polarization phase aberrations are induced by the effects of refraction through dielectric media, and by reflection from metallic or dielectric mirrors. This feature computes for the specified field position and wavelength the polarization phase aberration in the image space for the X and Y orientations of the electric field vector as a function of entrance pupil coordinate. The aberration is defined as the variation in the phase of the electric field in one direction, such as X or Y, as a function of the pupil position. For example, if the electric field in the X direction is $(0.7 + 0.7i)$ for the chief ray (note the E field values are complex), and at some other point in the pupil the electric field in the X direction is $(-0.7 + 0.7i)$ then the polarization phase aberration between these points is one-quarter

wave (the phase of E_x changed from 45 degrees to 135 degrees, or one-quarter wave). Note this is completely distinct from the phase differences between the E_x and E_y fields, which describe the polarization state (linear or circular for example). For phase differences between E_x and E_y see “Power Pupil Map” on page 193.

Like the usual OPD plots, the polarization phase aberration is referenced to the chief ray. However, there are cases where the chief ray phase cannot be determined in both orientations. For example, in an axial symmetric system, if the incident polarization is linear in the Y direction, there is zero intensity in the X direction for the chief ray, and therefore the X phase is indeterminate. For other rays in the pupil, there is generally a slight rotation of the polarization, and therefore the resulting electric field in the X direction yields a valid phase angle. To avoid this discontinuity in the phase, ZEMAX averages two rays on either side of the chief ray to interpolate the chief ray phase. The problem may still appear under certain circumstances, even with this averaging technique. In all cases, the phase data is still valid, because the phase aberration has no effect on image quality if the intensity is zero.

Transmission Fan

Purpose:

Generates a graph of the transmitted intensity for each field and wavelength as a function of either tangential or sagittal pupil fans.

Settings:

Item	Description
Jx, Jy	Jones electric field. see “Defining the initial polarization” on page 622.
X, Y Phase	Phase of the X, Y component of the Jones electric field in degrees.
Wavelength	The wavelength number of the rays to trace.
Field	The field position number of the rays to trace.
Unpolarized	If checked, the electric field definition is ignored and the unpolarized computation is performed.

Discussion:

The polarization transmission fan is useful for determining the transmission variation over the pupil as a function of field and wavelength. See “POLARIZATION ANALYSIS” on page 609.

Coatings



These features are only available in the EE editions of ZEMAX.

For more information on coating calculations, see “POLARIZATION ANALYSIS” on page 609.

Reflection vs. Angle

Purpose:

Computes the S, P, and average polarization intensity coefficients for reflection for the specified surface as a function of incident angle.

Settings:

Item	Description
Min Angle	The minimum angle of incidence to plot. This defines the left edge of the plot.
Max Angle	The maximum angle of incidence to plot. This defines the right edge of the plot.
Min Y	The minimum y value to plot. This defines the bottom edge of the plot.

Item	Description
Max Y	The maximum y value to plot. This defines the top edge of the plot.
Surface	The surface number which defines the interface to use.
Object	The object number if the surface is a Non-Sequential Components surface.
Face	The face number which defines the interface to use, if the surface is a Non-Sequential Components surface.
Wavelength	The wavelength number to use.
Reverse Direction	If the selected surface is a Non-Sequential Components surface, choosing this option will reverse the incident media and substrate.

Discussion:

The incident angle is measured in the medium prior to the specified surface.

Transmission vs. Angle

Purpose:

Computes the S, P and average polarization intensity coefficients for transmission for the specified surface as a function of incident angle.

Settings:

See "Reflectivity vs. Angle" above.

Discussion:

The incident angle is measured in the medium prior to the specified surface.

Absorption vs. Angle

Purpose:

Computes the S, P, and average polarization intensity coefficients for absorption for the specified surface as a function of incident angle.

Settings:

See "Reflectivity vs. Angle" above.

Discussion:

The incident angle is measured in the medium prior to the specified surface.

Diattenuation vs. Angle

Purpose:

Computes the R (reflected) and T (transmitted) diattenuation for the specified surface as a function of incident angle.

Settings:

See "Reflectivity vs. Angle" above.

Discussion:

The incident angle is measured in the medium prior to the specified surface.

Phase vs. Angle

Purpose:

For sequential surfaces, this feature computes the S and P polarization phase for reflection (if the glass is MIRROR) or for transmission (if the glass is not MIRROR) as a function of incident angle. For non-sequential objects, this feature computes the S and P polarization phase for reflection (if the material is MIRROR) or for transmission (if the material is not MIRROR) as a function of incident angle.

Settings:

See "Reflectivity vs. Angle" above.

Discussion:

The incident angle is measured in the medium prior to the specified surface.

Retardance vs. Angle

Purpose:

Computes the retardance for the specified surface as a function of incident angle.

Settings:

See "Reflectivity vs. Angle" above.

Discussion:

The incident angle is measured in the medium prior to the specified surface.

Reflection vs. Wavelength

Purpose:

Computes the S, P, and average polarization intensity coefficients for reflection for the specified surface as a function of incident wavelength.

Settings:

Item	Description
Min Wave	The minimum wavelength to plot. This defines the left edge of the plot.
Max Wave	The maximum wavelength to plot. This defines the right edge of the plot.
Min Y	The minimum y value to plot. This defines the bottom edge of the plot.
Max Y	The maximum y value to plot. This defines the top edge of the plot.
Surface	The surface number which defines the interface to use.
Object	The object number if the surface is a Non-Sequential Components surface.
Face	The face number which defines the interface to use, if the surface is a Non-Sequential Components surface.
Angle	The angle of incidence, in degrees, to use.
Reverse Direction	If the selected surface is a Non-Sequential Components surface, choosing this option will reverse the incident media and substrate.

Discussion:

See "POLARIZATION ANALYSIS" on page 609.

Transmission vs. Wavelength

Purpose:

Computes the S, P, and average polarization intensity coefficients for transmission for the specified surface as a function of incident wavelength.

Settings:

See "Reflectivity vs. Wavelength" above.

Discussion:

See "POLARIZATION ANALYSIS" on page 609.

Absorption vs. Wavelength

Purpose:

Computes the S, P, and average polarization intensity coefficients for absorption for the specified surface as a function of incident wavelength.

Settings:

See "Reflectivity vs. Wavelength" above.

Diattenuation vs. Wavelength

Purpose:

Computes the R (reflected) and T (transmitted) diattenuation for the specified surface as a function of incident wavelength.

Settings:

See "Reflectivity vs. Wavelength" above.

Phase vs. Wavelength

Purpose:

For sequential surfaces, this feature computes the S and P polarization phase for reflection (if the glass is MIRROR) or for transmission (if the glass is not MIRROR) as a function of incident wavelength. For non-sequential objects, this feature computes the S and P polarization phase for reflection (if the material is MIRROR) or for transmission (if the material is not MIRROR) as a function of incident wavelength.

Settings:

See "Reflectivity vs. Wavelength" above.

Retardance vs. Wavelength

Purpose:

Computes the retardance for the specified surface as a function of incident wavelength.

Settings:

See "Reflectivity vs. Wavelength" above.

Physical Optics

Paraxial Gaussian Beam

Purpose:

Computes Paraxial Gaussian Beam parameters. Paraxial Gaussian Beams are limited to axial analysis of rotationally symmetric systems. See also "Skew Gaussian Beam" on page 225 and "Physical Optics Propagation" on page 225.

Settings:

Item	Description
Wavelength	The wavelength number to use for the calculation.
M ² Factor	The M-squared quality factor used to simulate mixed mode beams. See the Discussion.
Waist Size	The radial size of the embedded (perfect TEM00 mode) beam waist in object space in lens units.
Surf 1 to Waist	The distance from surface 1 (NOT the object surface) to the beam waist location. This parameter will be negative if the waist lies to the left of surface 1.

Item	Description
Update	See below.
Orient	
Surface	

Discussion:

This feature computes ideal and mixed mode M-squared Gaussian beam data, such as beam size, beam divergence, and waist locations, as a given input beam propagates through the lens system. This discussion is not meant to be a complete tutorial on laser beam propagation theory. For more information on Gaussian beam propagation, see one of the following references: A. E. Siegman, *Lasers*, University Science Books (1986), R. Herloski, S. Marshall, and R. Antos, "Gaussian beam ray-equivalent modeling and optical design", *Applied Optics* Vol. **22**, No. 8 p1168 (1983), M. W. Sasnett and T. F. Johnston, Jr., "Beam characterization and measurement of propagation attributes", *Proc. SPIE* Vol. **1414**, p21 (1991), and A. E. Siegman, "New developments in laser resonators", *Proc. SPIE* Vol. **1224**, p2 (1990).

Limitations of the analysis

Gaussian beams are idealized, "perfect" beams, and this limits the range of systems for which Gaussian beam analysis is useful. Because the calculation of Gaussian beam parameters is based upon paraxial ray data the results cannot be trusted for systems which have large aberrations, or those poorly described by paraxial optics, such as non-rotationally symmetric systems. Flat fold mirrors will not invalidate the results, but decentered or tilted components generally will. This feature ignores all apertures, and assumes the Gaussian beam propagates well within the apertures of all the lenses in the system.

For a more general beam propagation analysis feature, see "Physical Optics Propagation" on page 225.

Overview of Gaussian beams

A Gaussian laser beam is described by a beam waist size, a wavelength, and a location in object space. The Gaussian beam is an idealization that can be approached but never attained in practice. However, real laser beams can be well described by an "embedded" Gaussian beam with ideal characteristics, and a quality factor, called M^2 , which defines the relative beam size and divergence with respect to the embedded Gaussian mode. The ideal M^2 value is unity, but real lasers will always have an M^2 value greater than unity.

This feature requires the definition of the initial input embedded beam properties, and the M^2 value. The input embedded beam is defined by the location of the input beam waist relative to surface 1 (note this is not the object surface, but the first surface after the object) and the waist radial size at this location. ZEMAX then propagates this embedded beam through the lens system, and at each surface the beam data is computed and displayed in the output window. ZEMAX computes the Gaussian beam parameters for both X and Y orientations.

Default beam parameters

ZEMAX defaults to a waist size of 0.05 lens units (no matter what the lens units are) and a surface 1 to waist distance of zero; this of course means the waist is at surface 1. The default values may be reset by clicking on the "Reset" button. After the default values are computed and displayed, any alternate beam waist size and location may be entered and that Gaussian beam will be traced instead.

Propagating the embedded beam

Once the initial beam waist and location parameters are established, ZEMAX traces the embedded beam through the system and computes the radial beam size, the narrowest radial waist, the surface coordinate relative to the beam waist, the phase radius of curvature of the beam, the semi divergence angle, and the Rayleigh range for every surface in the system. ZEMAX calls these parameters the Size, Waist, Waist Z, Radius, Divergence, and Rayleigh, respectively, on the text listing that is generated.

The dimensions for all parameters are lens units except for the semi divergence angle, which is in units of radians. The embedded Gaussian beam parameters can be computed using the following standard formulas. The Rayleigh range, z_r , is defined as

$$z_r = \frac{\pi \omega_0^2}{\lambda}.$$

The phase radius of curvature is given by

$$R(z) = z + \frac{z_r^2}{z},$$

where z is the distance from the beam waist. The radial beam size at any z is computed from

$$\omega(z) = \omega_0 \left[1 + \left(\frac{z}{z_r} \right)^2 \right]^{1/2},$$

where ω_0 is the beam waist. The divergence angle of the beam is given by

$$\theta = \tan^{-1} \frac{\omega_0}{z_r}.$$

Finally, the Gouy shift is the phase of the center of the Gaussian beam:

$$\varphi = \tan^{-1} \frac{z}{z_r}$$

The quality factor

All of the preceding results are correct for the ideal embedded Gaussian beam. For aberrated, mixed-mode beams, an extension to the fundamental Gaussian beam model has been developed by Siegman. The method uses a term called the beam quality factor, usually denoted by M^2 . The factor M^2 can be thought of as "times diffraction limited" number, and is always greater than unity. The M^2 factor determines the size of the real, aberrated Gaussian beam by scaling the size and divergence of the embedded Gaussian mode by M . It is common practice to specify M^2 for a laser beam, rather than M , although the factor M is used to scale the beam size. The M^2 factor must be measured to be determined correctly. If the M^2 factor is set to unity, the default value, ZEMAX computes the TEM00 data described above. If M^2 is greater than unity, then ZEMAX computes both the embedded Gaussian beam parameters as well as the scaled data.

Interactive analysis

The Settings dialog box for this feature also supports an interactive mode. After defining the various input beam parameters, clicking on "Update" will immediately trace the specified Gaussian beam, and display the usual results in the dialog box. The parameters for any surface may be viewed, and the surface number selected from the drop down list. The orientation may also be selected using the provided control.

The interactive feature does not in any way modify the lens or the system data; it is a handy "calculator" for displaying paraxial Gaussian beam data.

Skew Gaussian Beam

Purpose:

Computes Skew Gaussian Beam parameters. Skew beams may enter an optical system at any surface from any field position, and may travel through the optical system off axis. The Skew Gaussian Beam parameters are computed using real rays and account for astigmatism but not higher order aberrations. See also "Paraxial Gaussian Beam" on page 222 and "Physical Optics Propagation" on page 225.

Settings:

Item	Description
X Waist Size	The x direction radial size of the beam waist in the space prior to the "start surface" in lens units.
Y Waist Size	The y direction radial size of the beam waist in the space prior to the "start surface" in lens units.
Surf to Waist	The distance from the "start surface" to the beam waist location. This parameter will be negative if the waist lies to the left of the start surface.
Wavelength	The wavelength number to use for the calculation.
Field	The field number to use for the calculation.
Start Surface	The surface to begin the beam propagation.
End Surface	The surface to end the beam propagation.

Discussion:

For important background information on Gaussian beams, see "Paraxial Gaussian Beam" on page 222.

This feature computes the ideal skew Gaussian beam data, such as beam size, beam divergence, and waist locations, as a given input beam propagates through the lens system. The beam is not limited to being on axis in a symmetric optical system, and beam at any angle may be traced anywhere in the optical system.

The beam is initially defined by a waist size in the x and y directions in the optical space prior to the starting surface. The field and wavelength settings are used to define a chief ray which is traced through the optical system. This chief ray becomes the axis for the skew Gaussian beam throughout the propagation.

The x and y directions are initially defined by the entrance pupil orientation. The +y direction relative to the chief ray is defined as a vector pointing from the chief ray coordinate at the entrance pupil to the coordinate of the $p_x = 0$, $p_y = 1$ marginal ray coordinate at the entrance pupil. The $p_x = 1$, $p_y = 0$ marginal ray defines the +x direction. For a discussion of the normalized pupil coordinates, see "Normalized pupil coordinates" on page 57.

As the chief ray propagates through the optical system, the two marginal rays are propagated as well, and are used to define the +x and +y directions in each optical space. All the data listed for x and y directions are measured in this moving coordinate system.

The Skew Gaussian Beam feature results are generally not accurate for systems including non-sequential component groups. For the most accurate results, replace non-sequential component groups with sequential equivalents.

Physical Optics Propagation



This feature is only available in the EE editions of ZEMAX.

Purpose:

This feature propagates an arbitrary coherent optical beam through the optical system. For a complete description of this feature see "PHYSICAL OPTICS PROPAGATION" on page 627. The Physical Optics Propagation feature results are generally not accurate for systems including non-sequential component groups; see

“Propagating through non-sequential surfaces” on page 637 for more information. For the most accurate results, replace non-sequential component groups with sequential equivalents.

Settings:

General Tab:

Item	Description
Start Surface	The starting surface for the initial beam. The beam will begin in the optical space just prior to the starting surface. The entrance pupil may be used as the starting surface. If the system is telecentric in object space, then surface 1 will be used as the starting surface even if the entrance pupil is selected.
End Surface	The surface at which to terminate the propagation. The beam will stop just after entering the optical space of the end surface.
Wavelength	The wavelength number to use for the beam.
Field	The field number of the chief ray used to align the initial beam.
Surf to Beam	The distance from the starting surface to the starting beam position in lens units. This value is negative if the beam starts to the left of the surface.
Use Polarization	If checked, 2 beam arrays will be propagated, one each for the X and Y polarizations of the beam.
Separate X, Y	If checked, the X and Y directions are propagated independently. This allows much more efficient sampling if the beam is astigmatic or the optics are cylindrical or toroidal.
Use Disk Storage To Save Memory	See “Memory requirements” on page 637 for a discussion.

Beam Definition Tab:

Item	Description
X- Y- Sampling	The number of points used to sample the beam.
X- Y- Width	The initial width in lens units of the region represented by the array. See Auto below.
Auto	Pressing this button will calculate the optimum X- and Y- Width values to maintain the approximately the same number of pixels across the beam both within and outside the Rayleigh range. See “Comments about point spacing and sampling” on page 636.
Type	See the discussion and “Defining the initial beam” on page 638.
File	See the discussion and “Defining the initial beam” on page 638.
Peak Irradiance	The initial beam peak irradiance in power per area. This is an alternative to Total Power.
Total Power	The initial beam total power. This is an alternative to Peak Irradiance.
Read power from file	If the beam type is File, pressing this button will read the current power and peak irradiance values from the selected file.
Parameters	Different initial beam types require various parameters to define the distribution. These parameter names and values will change depending upon the beam type setting. For details, see “Defining the initial beam” on page 638.

Display Tab:

Item	Description
Show As	Choose surface, contour, false color, grey scale, cross sections, encircled, ensquared, or enslitted displays. Note that the encircled energy is a function of radius; and the ensquared and enslitted widths are the half widths.
Row/Column	If cross section is chosen for Show As, this control selects the row or column to view.
Data	Choose irradiance (beam power per area), phase (of the beam), transfer function intensity, or transfer function phase. The Ex and Ey refer to the X and Y polarization components of the beam. The transfer function values are the intensity and phase of transmission for the final surface only; these are primarily used for debugging and analysis of the POP results. For text POP windows only, a propagation report is also available. The propagation report lists various data about the propagation algorithm parameters used, the orientation matrix for the pop results, and may contain warnings that indicate if a possible inaccuracy is detected in the results.
Scale	Choose linear or logarithmic scaling of the data. Logarithmic scaling only applies to irradiance data on graphical (not text) displays.
Project	The beam may be viewed from any one of these perspectives: along the beam (this is along the chief ray used by the pilot beam), along the surface normal, or along the local Z axis of the surface. Note ZBF beam file data is always "along the beam". See the comments about beam projection below.
Save Output Beam To	If checked and a file name is provided, the complex amplitude of the beam at the end surface will be written to a file. This file can be read back in as a starting beam using the Beam Type control. ZEMAX will add the extension "ZBF" to the provided file name. ZBF files are stored in the <pop> folder (see "Folders" on page 68).
Save Beam At All Surfaces	If checked, the beam file will be saved at every surface from the start surface to the end surface that uses the physical optics propagator. The "Save Output Beam To" option must also be checked and a file name provided. The individual surface beam files will be named according to the provided file name, with a suffix indicating the surface number. For example, if the output file name is "MyBeamData.ZBF", then the beam file name for surface 12 will be "MyBeamData_12.ZBF". ZBF files are stored in the <pop> folder (see "Folders" on page 68).
Zero Phase For Relative Irradiance Below	Phase values are meaningless if the irradiance is extremely low. Computing the phase angle of data points with nearly zero irradiance (relative to the peak irradiance in the beam) will result in noisy plots and text listing of meaningless data. This value sets the lower limit on the relative irradiance of data points for which the phase is computed. Data points with relative irradiance lower than this threshold will have a phase value of zero.
Plot Scale	If zero, the vertical scale on false color, grey scale, and cross section plots will be set by the maximum data value, or for log plots, the next largest power of 10. Otherwise, the maximum vertical scale is set by the normalization value. If a log plot is generated, the maximum scale is set by the next largest power of 10 greater than the plot scale value. Phase plots always scale from -pi to pi.
Contour Format	The contour format string. For a discussion of the contour format string syntax, see "The Contour Format String" on page 151. The contours are defined in the same units as the data being displayed.
Zoom In	Magnifies the displayed area of the beam. This is useful for eliminating some of the guard band from the displayed beam data.

Fiber Data Tab:

Item	Description
Compute Fiber Coupling Integral	If checked, the fiber data on this tab will be used and the fiber coupling computed; otherwise, no fiber coupling computations will be performed. See "Computing Fiber Coupling" on page 645.
Ignore Polarization	If checked, the fiber mode will be unpolarized even if the beam itself is polarized. If unchecked and the beam is polarized, then the fiber coupling overlap integral is computed for the two orthogonal polarizations independently and the results combined to yield the overall coupling. See "Defining the fiber mode" on page 645.
Tilt About X/Y (deg)	The tilt about each axis in degrees of the fiber mode with respect to the beam. The phase of the fiber mode will be modified with a linear tilt proportional to the tilt angles.
Fiber Type	Selects the mode for the fiber. See "Computing Fiber Coupling" on page 645.
File	Selects the name of the DLL or data file describing the fiber mode. See "Computing Fiber Coupling" on page 645.
Parameters	Different fiber mode types require various parameters to define the mode. These parameter names and values will change depending upon the fiber type setting.
Fiber Position	The center of the receiving fiber may be referenced to the surface vertex or to the chief ray. The incoming POP beam is referenced to the chief ray position and angle.

Discussion:

For a full discussion of physical optics propagation, see "PHYSICAL OPTICS PROPAGATION" on page 627.

Comments about beam projection

The analysis computes the beam irradiance or phase on a plane tangent to the chief ray at the point where the chief ray intercepts the surface. The orientation matrix of the plane may be determined from the propagation report. The irradiance or phase data is shown after the beam refracts into, or reflects from, the end surface. The chief ray representing the center of the beam will generally intercept the surface at some angle of incidence other than zero, and the surface normal will generally be at some angle to the local Z axis. Using the "project" option on the "display" tab, the beam data may be viewed along the beam, along the surface normal, or along the local Z axis. The latter two options are implemented by elongating the beam by the cosine of the angle between the beam and the desired projection.

The beam itself is always represented by the data array in a coordinate system normal to the chief ray. For this reason, the beam projected onto a surface by this feature may be elongated compared to the same beam stored as a file and viewed with the Beam File Viewer, see "Beam File Viewer" on page 229.

Beams in ZEMAX are always centered on the chief ray for the selected field and wavelength. Therefore, the data in the beam file is positioned relative to the chief ray. The center point in the beam file is at the coordinate (nx/2+1, ny/2+1).

The analysis window may include some of the following data:

Display X Width/ Y Height: The width and height of the beam array in lens units at the end surface.

Peak Irradiance: The maximum irradiance in power per area. To set the power and area units see "Units" on page 104.

Total Power: The total power in the beam. To set the power units see "Units" on page 104.

Pilot Beam Data: The Pilot Beam data includes the radial beam size, beam waist, position, and Rayleigh range. For detailed information on the Pilot Beam see "The pilot beam" on page 632.

Fiber Coupling: If "Compute Fiber Coupling Integral" is checked, then the fiber coupling is computed and displayed. The system efficiency is the fraction of beam power that transmits from the starting surface to the surface at which the fiber coupling is computed. The receiver efficiency is the fraction of the power in

the beam mode that couples into the fiber mode. The total coupling efficiency is the product of these two numbers. For more information see “Computing Fiber Coupling” on page 645.

Beam File Viewer



This feature is only available in the EE editions of ZEMAX.

Purpose:

This feature allows viewing and analysis on previously stored ZEMAX Beam Files (ZBF). These files may be user defined or may be generated by the Physical Optics Propagation feature, see the feature “Physical Optics Propagation” on page 225. See also the Chapter “PHYSICAL OPTICS PROPAGATION” on page 627. For information on the ZBF file format, see “ZEMAX Beam File (ZBF) binary format” on page 640.

Settings:

The available settings are a subset of what is available for the Physical Optics Propagation feature described above.

Description:

The primary advantage to using this feature is that the beam may be propagated once through the optical system, using the Physical Optics Propagation feature described above, and then the beam files may be analyzed and viewed without the need to re-propagate the beam. Beam files for every surface may be generated by choosing “Save Beam At All Surfaces” on the Physical Optics Propagation feature settings.



For rapid viewing through a range of files that differ only by the surface number, use the left and right cursor keys.

See the “Comments about beam projection” on page 228.

ZBF files are stored in the <pop> folder (see “Folders” on page 68).

Non-Sequential

These features are only used for non-sequential analysis. For a full description, see “NSC - OVERVIEW” on page 345.

Ray Trace Control

See “The Ray Trace Control” on page 377.

Detector Viewer

See “The Detector Viewer” on page 384.

Ray Database Viewer

“The Ray Database Viewer” on page 382.

Ray Database Analysis

These features perform analysis on previously saved ZRD files. See “Ray database (ZRD) files” on page 380.

Flux vs. Wavelength

Purpose:

This feature selects all rays from a ZRD file, and sums the intensity of rays striking a specific object as a function of wavelength.

Settings:

Item	Description
Ray Database	Selects the existing ZRD file to analyze.
First/Last Ray	Choose the ray range. To use all rays, set the First ray to 1 and the last ray to -1.
Filter	The filter string. See "The filter string" on page 386.
Object	The object the rays strike. A specific object must be selected, and the intensity for any segment that strikes the object will be added to the sum for the associated wavelength.
Bins	Because the ray wavelengths are discrete, the sum for each wavelength must be over a finite bandwidth. This setting determines how many wavelength bins are used to span the range from minimum to maximum wavelength.
Minimum/Maximum Wavelength	The lowest and highest wavelength. Rays outside this wavelength range are ignored.
Plot Scale	The maximum scale for the plot. Use zero to choose a default scale.

Discussion:

Because the wavelengths for each segment are discrete and are only determined by the rays in the ZRD file, this plot can take on a noisy look with lots of data "spikes". To minimize the noisy nature of this computation, the number of rays traced should be large, and the number of bins kept reasonably small.

Obsolete

These features are considered obsolete. Generally, this is because newer features that provide similar and superior capabilities exist. Although these features are still available for back compatibility, they are no longer supported and will likely be removed entirely from ZEMAX in a future release.

Wireframe



This feature has been replaced by the "3D Layout" on page 118.

Purpose:

Draws a representation of the lens.

Settings:

The options are similar to those available for the 3D layout feature. The "Hide Lens Edges" and "Hide X Bars" checkboxes are not available, and several new controls are added as described below.

Item	Description
Draw Section	Select "Full" to draw each lens element completely. The 3/4, 1/2, and 1/4 options draw just that much of the element, yielding a cut away perspective of the lens interior.
Radial Segments	The number of radial segments used to approximate the lens shapes. Larger numbers require more processing time.
Angular Segments	The number of angular segments used to approximate the lens shapes. Larger numbers require more processing time.

Discussion:

The wireframe algorithm describes the lens as a collection of polygons. The number of polygons used to display the lens elements can be modified using the radial and angular segment options. The wireframe model is identical to the solid model, except hidden lines are not removed. This representation may cause the drawing

to become cluttered with lines. The "Hide Lens Faces" option can be used to clean up the display. The advantage to this display method is speed; it is faster than the solid model.

Pressing the left, right, up, down, Page Up, or Page Down keys will rotate the displayed image for a different perspective.

If rays miss a surface, then the rays will not be drawn to the surface where the error occurred. If the ray is total internal reflected, then the ray will be drawn up to but not past the surface where the error occurred. Ray failures can be evaluated in detail, see "Ray Trace" on page 205.

See also "The orientation indicator" on page 119.

Solid Model



This feature has been replaced by the "Shaded Model" on page 120.

Purpose:

Draws a hidden-line representation of the lens.

Settings:

The options are similar to those available for the 3D layout feature. The "Hide Lens Edges" and "Hide X Bars" checkboxes are not available, and several new controls are added as described below.

Item	Description
Draw Section	Select "Full" to draw each lens element completely. The 3/4, 1/2, and 1/4 options draw just that much of the element, yielding a cut away perspective of the lens interior.
Radial Segments	The number of radial segments used to approximate the lens shapes. Larger numbers require more processing time.
Angular Segments	The number of angular segments used to approximate the lens shapes. Larger numbers require more processing time.

Discussion:

The solid model algorithm describes the lens as a collection of polygons. The lines and polygons which are hidden from view are removed, which gives the lens a solid appearance. This algorithm is slower than the other layout plots. The number of polygons used to display the lens elements can be modified using the radial and angular segment options.

Pressing the left, right, up, down, Page Up, or Page Down keys will rotate the displayed image for a different perspective.

If rays miss a surface, then the rays will not be drawn to the surface where the error occurred. If the ray is total internal reflected, then the ray will be drawn up to but not past the surface where the error occurred. Ray failures can be evaluated in detail, see "Ray Trace" on page 205.

See also "The orientation indicator" on page 119.

Illumination XY Scan



This feature has been replaced by "Image Simulation" on page 162.

Purpose:

Computes relative illumination for the image of an extended source along a line across any surface. This feature should only be used for systems that form a reasonably good image of the object source. For non-imaging systems, see "Geometric Image Analysis" on page 165.

Settings:

Item	Description
Sampling	Selects the grid size used to sum the illumination. The sampling determines how many pixels are used to collect the ray data.
Rays x 1000	Determines the approximate total number of rays to be traced in computing the illumination due to the extended source.
File	The name of the IMA file used to define the shape of the extended source. See the Geometric Image Analysis feature for details.
Source Size	The full width in field units of the extended source.
Rotation	The angle to rotate the extended source in object space about the normal to the center of the extended source.
Show As	Select either X or Y direction scan.
Smoothing	Smoothing helps remove jagged lines due to low ray sampling by applying a smoothing operator to average the data in adjacent pixels.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number indicates which field position is used as the center reference point for the extended source.
Surface	The scan may be computed at any surface, however, the relative illumination computation is only used at the image surface.
Detector Size	The total width of the detector in lens units. The detector size is divided into pixels according to the "Sampling" setting.
Use Relative Illumination	If checked, then the computation described in "Relative Illumination" on page 187 is used to weight the rays from various points in the field of view to accurately account for the effects of exit pupil radiance and solid angle. The computation is generally more accurate, but slower, if this feature is used. RI may only be used if the selected surface is the image surface, and ZEMAX assumes the image surface is a reasonable conjugate of the object surface. See the RI feature for details. If unchecked, the RI is estimated from the system distortion. RI is only considered if the surface is the image surface.
Assume Lambertian Source	If checked, rays are weighted according to the cosine of the angle they make with respect to the object surface. This cosine factor is not appropriate in all cases. This option is not available if "Use Relative Illumination" is checked, since the RI calculation already assumes the object is Lambertian.
Use Polarization	If checked, polarization is considered. See "Polarization (ZEMAX-EE only)" on page 108 for information on defining the polarization state and how polarization is used by analysis features.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 136.

Discussion:

The illumination XY scan is similar to the relative illumination (RI) feature, with the added capability to estimate the RI for non-uniform extended sources. For uniform extended Lambertian sources, the RI feature is faster and more accurate. However, for systems with complex source properties, the illumination XY scan can estimate illumination by Monte Carlo ray tracing combined with the conventional RI computation.

The extended sources are defined in the same way as described in the Geometric Image Analysis feature, see "Geometric Image Analysis" on page 165 for details.

Illumination 2D Surface



This feature has been replaced by "Image Simulation" on page 162.

Purpose:

Computes relative illumination for an extended source over a 2D surface. This feature should only be used for systems that form a reasonably good image of the object source. For non-imaging systems, see "Geometric Image Analysis" on page 165.

Settings:

The options are identical to the XY scan, except that a 2D surface is drawn as either an isometric surface or contour plot, or as a grey scale or false color map.

Discussion:

See "Illumination XY Scan" on page 231.

Export IGES Line Work



This feature has been replaced by "Export IGES/SAT/STEP Solid" on page 257.

Purpose:

Exports the current lens data as an IGES line work format file, with various options. See also the Export IGES/STEP Solid feature.

Settings:

Item	Description
First/Last Surface	The range of surfaces to include in the exported data.
Wavelength	The wavelength number to use for any traced rays to be exported.
Field	The field number to use for any traced rays to be exported.
Number of Rays	The number of rays to trace; the exact meaning depends upon the "Ray Pattern".
Ray Pattern	Selects the type of ray set to export. This control is very similar to that defined for the 3D layout plot; see "3D Layout" on page 118 for details.
Lens/Ray Layer	Selects the layers of the output file to place lens and ray data.
Delete Vignetted	If checked, rays are not included in the data export if they will be vignetted by any surface.
Spline Segments	The number of segments to use when spline entities are exported.

Discussion:

For information on the accuracy of exported surfaces see "Comments about imported objects" on page 513.

The Initial Graphics Exchange Specification (IGES) is an American National Standard whose intended purpose is to facilitate transfer of data between CAD programs. ZEMAX currently supports version 5.2 of the IGES standard. For more information on IGES, contact U.S. Product Data Association, P. O. Box 3310, Gaithersburg, MD 20885-3310.

ZEMAX exports lines, spherical arcs, and splines to represent the shape and position of each surface. ZEMAX does not export the "edge" of the lens, and makes no distinction between dummy surfaces, glass surfaces, and mirrors. The type and number of each IGES entity exported for a given surface depends upon the surface aperture, if any, and the symmetry, if any of the surface. The following table lists how ZEMAX determines the best representation of the surface. ZEMAX exports all surfaces and rays in a 3D coordinate system referenced to the global coordinate reference surface which is described in the System Menu chapter.

This feature will prompt for the file name and path to save the exported data to. The default file is EXPORT.IGS in the <objects>\CAD Files folder (see “Folders” on page 68). The output will be in IGES format no matter what extension is chosen.

ZEMAX makes no attempt to export objects defined as part of a Non-Sequential Components surface. See the Export IGES/STEP Solid feature if this capability is required.

IGES ENTITIES EXPORTED BY SURFACE APERTURE AND SHAPE

Surface Aperture	Surface Shape	Entities Used
None, Circular Aperture or Obscuration	Plano	One line along the X axis, one line along the Y axis, one circle in the XY plane at the semi-diameter radius.
	Spherical	One arc in the XZ plane, one arc in the YZ plane, one circle in the XY plane at the semi-diameter radius.
	Other	One spline in the XZ plane, one spline in the YZ plane, one circle in the XY plane at the semi-diameter radius.
Rectangular Aperture or Obscuration	Plano	Three lines parallel to the X axis, three lines parallel to the Y axis, arranged at the edges and center of the surface to form a grid.
	Other	Three splines parallel to the X axis, three splines parallel to the Y axis, arranged at the edges and center of the surface to form a grid when projected upon the XY plane.
Elliptical Aperture or Obscuration	Plano	One line along the X axis, one line along the Y axis, one spline defining the edge of the surface in the XY plane.
	Other	One spline along the X axis, one spline along the Y axis, one spline defining the edge of the surface.
User Defined Aperture or Obscuration	Any	One line for each user defined segment.

The line entity is IGES entity 110. The circular arc entity is IGES entity 100. The spline entity is IGES entity 112. Rays are exported as line entities. Within GRIN media, rays are exported as a series of line entities. A warning will be issued if not all surfaces within the specified range were exported.

For surfaces where both an obscuration and another shape need to be exported, use a dummy surface collocated with the surface that has the second aperture on it. For example, to export a rectangular lens with a rectangular obscuration on the surface, use two surfaces with a zero thickness between them; one with the aperture and one with the obscuration. Make sure the surfaces have the same shape by setting the surface types the same and using pickup solves where appropriate.

Export 2D DXF File



This feature has been replaced by “Export IGES/SAT/STEP Solid” on page 257.

Purpose:

Exports the current lens data as a 2D DXF format file. Only available for spherical symmetric lenses.

Settings:

Item	Description
First/Last Surface	The range of surfaces to include in the exported data.
Wavelength	The wavelength number to use for any traced rays to be exported.
Field	The field number to use for any traced rays to be exported.
Number of Rays	The number of rays to trace.
Delete Vignetted	If checked, rays are not included in the data export if they will be vignetted by any surface.
Marginal and Chief Only	Draws only the marginal and chief rays, overriding the other ray settings.

Discussion:

This feature will prompt for the file name and path to save the exported data to. The default file is EXPORT.DXF in the <objects>\CAD Files folder (see “Folders” on page 68). The output will be in DXF format no matter what extension is chosen.

Export 3D DXF File



This feature has been replaced by “Export IGES/SAT/STEP Solid” on page 257.

Purpose:

Exports the current lens data as a 3D DXF format file.

Settings:

Item	Description
First/Last Surface	The range of surfaces to include in the exported data.
Wavelength	The wavelength number to use for any traced rays to be exported.
Field	The field number to use for any traced rays to be exported.
Radial Segments	The number of radial segments used to approximate the lens shapes. Larger numbers require more processing time.
Angular Segments	The number of angular segments used to approximate the lens shapes. Larger numbers require more processing time.
Number of Rays	The number of rays to trace.
Ray Pattern	Choose XY Fan, X Fan, Y Fan, Ring, List, Random, or Grid to indicate what the pattern of rays to be traced should be. The List option indicates that the rays to be traced are user defined and listed in a file, see “Raylist file format” on page 120 for information on the ray list format. If List is selected the Number of Rays setting is ignored.
Delete Vignetted	If checked, rays are not included in the data export if they will be vignetted by any surface.

Discussion:

This feature is similar to the 2D DXF export feature described in the previous section. The key difference is that the 3D DXF supports systems that are not rotationally symmetric.

Conjugate Surface Analysis



This feature has been replaced by the multi-configuration tool "make conjugate", see "Tools" on page 90.

Purpose:

This feature computes the RMS spot radius formed on a surface when a specific point on a previous surface is used as the object. The feature works by deleting surfaces before and after the conjugate surfaces, and then evaluating the image quality for a specific point on the new object surface.

Settings:

Item	Description
Obj, Stop, Ima	The object, stop, and image surface numbers. These numbers refer to existing surfaces in the current system. The object and image surface numbers are assumed to be conjugate surfaces, which means an image of the object surface is formed at the image surface location. There is however no requirement on the image quality between these surfaces. The stop surface is the surface to be used as the stop for the conjugate surfaces. This surface is in general not the same stop surface used for the whole optical system to which the conjugate surfaces belong.
Data	The data to be computed for the subset system. Currently only RMS Spot Radius referenced to the centroid is supported.
Wavelength	The wavelength number to use for the analysis.
Sampling	The rectangular grid size to use for the analysis.
X/Y Coordinate	The coordinates in lens units of the point on the new object surface to use as the source field point. If the object surface is zero and the resulting system is at infinite conjugates, then the X/Y Coordinates are assumed to be angles in degrees.
Save System As	If a valid name is provided, then the system used to evaluate the data is saved. The file name should be a valid name with a ZMX extension and a full path.

Discussion:

This feature allows a simple analysis of the RMS spot radius of the image formed by some range of surfaces using a single point on another surface as the source. The feature works using the following procedure on a temporary copy of the lens being analyzed:

The new stop surface semi-diameter is fixed to the current value.

All solves are removed.

All surfaces prior to the new object surface are deleted.

All surfaces following the new image surface are deleted.

The stop surface is reassigned to the new stop surface.

Ray aiming is turned on.

The system aperture type is set to float by stop size.

Two fields are defined, one at (0, 0) and the other at the requested object field coordinates.

The field type is set to either object height or field angle for finite and infinite conjugates, respectively.

The analysis then proceeds to compute the RMS spot radius at the requested field. Note that if the resulting system is at infinite conjugates, the field coordinates are angles in degrees, otherwise the field coordinates are in lens units.

A much more general way of analyzing conjugate surfaces is to use the "ignore surface" feature to select a subset of the optical system to consider as an alternate configuration. For more information see "Ignore This Surface" on page 78.

Optimization**Optimization*****Purpose:***

Optimization uses a Damped Least-Squares or Orthogonal Descent algorithm to improve or modify a design to meet specific conditions.

Discussion:

See "OPTIMIZATION" on page 483 for a full discussion of building merit functions, setting variables, and using the optimization feature.

Global Search***Purpose:***

This feature initiates a search for the global optimum, which is the best possible design for a given merit function and set of variables.

Discussion:

See "GLOBAL OPTIMIZATION" on page 539.

Hammer Optimization***Purpose:***

This feature automates the repetitive optimization of a design to escape local minima in the merit function.

Discussion:

See "GLOBAL OPTIMIZATION" on page 539.

Find Best Asphere***Purpose:***

This feature replaces spherical surfaces with aspherical surfaces, and reoptimizes the design, to determine which surface would benefit the most from aspherization.

Settings:

Item	Description
First/Last Surface	These controls set the range of surfaces to consider adding an asphere to.
Asphere Type	Choose an asphere defined by a conic or an even polynomial asphere to the specified maximum order.
Start/Run Again	Begins the analysis. The first time the analysis is run, the button is labelled "Start". After the first analysis, the button changes to "Run Again". Selecting "Run Again" will discard the previous results and search for the best asphere using the current surface range and asphere type selections.
Terminate	Terminates a running analysis.
Keep and Exit	Retains the best asphere found and exits.
Cancel	Discards the best asphere found, if any, and exits.

Discussion:

This feature automates the process of determining which spherical surface, if re-optimized to be an aspheric surface, would provide the lowest merit function. Each surface is evaluated to see if it is a candidate asphere. To be considered, the surface must be of type Standard, have no conic value, define a boundary between air and glass (cemented surfaces usually make poor aspheres), have a curvature that is either variable or controlled by a marginal ray angle or F/# solve, and have a curvature that is not under control of the Multi-configuration Editor. Surfaces that do not meet this test are ignored.

When a candidate surface is identified, the surface is converted into an asphere of the user-selected type. The aspheric terms are set as variables for optimization. The local damped-least squares optimizer is then called to optimize the modified system. If the resulting system has the lowest merit function yet found, the system is retained. The procedure repeats until all surfaces have been tested. Finally, this feature reports which surface, when converted to an asphere, provided the lowest merit function. Note the currently defined merit function is used, and all data that are variable are re-optimized during this process. The current merit function should be appropriate for an aspheric design, which may require higher sampling than a non-aspheric design for good optimization. For a discussion of sampling in the merit function see "Sampling considerations for Rings and Arms" on page 486.

Note that like all local optimization results, there is no way to know if the solution found is the optimum "global minimum" for that combination of merit function, variables, and design parameters. For this reason, once the best candidate asphere is determined, it is usually a good idea to run the Hammer Optimization (see "Hammer Optimization" on page 239) on the resulting design to see if any further gains are possible.

No attempt is made to determine whether the resulting asphere is practical to fabricate, or is more or less costly to manufacture as compared to making other surfaces aspheric. Some care must be taken in using this feature, as the results may not be useful for multi-configuration systems, especially those using ignored surfaces in some configurations but not in others. To prevent any particular surface from being selected as a possible asphere, either remove the variable status from the radius of curvature, or temporarily enter a small conic value (such as 1E-20). This feature will then not include that surface in the possible candidate aspheres.

Merit Function Listing

Purpose:

This function generates a text listing of the merit function which can be saved or printed. The overall merit function numerical value is listed, as well as the numerical merit function of all user added operands (defined as those operands placed before the DMFS operand) and the default merit function (all operands after the DMFS operand). See also "Merit function definition" on page 489.

Remove All Variables

Purpose:

This tool provides a quick way to remove all of the variables defined in the current lens.

Discussion:

Variables are removed by setting each of the values to "fixed" at the current value.

Glass Substitution Template

Purpose:

The template is used to define the maximum limits of the glass cost as well as the mechanical resistance properties of glasses. The template is used to select acceptable glasses during global optimization, when converting from model to real glasses (see "Using model glasses" on page 601), and when computing the value of the RGLA operand ("RGLA" on page 518).

Settings:

Item	Description
Use Glass Substitution Template	If checked, the defined parameters are used to restrict the glasses available for substitution by the global optimizer.
Exclude Glasses With Incomplete Data	If checked, glasses with incomplete template data will not be selected by the global optimizer. Only data corresponding to non-zero template values are tested before excluding a glass.
Standard, Preferred, Obsolete, Special	These checkboxes indicate the acceptable status settings of glasses. These options are additive. For example, if both Standard and Obsolete are checked, then either Standard or Obsolete glasses are acceptable. Note a single glass may have only one status. See "Description of catalog data" on page 588.
Maximum Relative Cost	The maximum relative cost. The cost of Schott BK7 is 1.0, all other glasses are more expensive, and have larger relative costs. If zero the relative cost is ignored.
Maximum Climatic Resistance (CR)	The maximum allowed value of CR. If zero the CR is ignored.
Maximum Stain Resistance (FR)	The maximum allowed value of FR. If zero the FR is ignored.
Maximum Acid Resistance (SR)	The maximum allowed value of SR. If zero the SR is ignored.
Maximum Alkali Resistance (AR)	The maximum allowed value of AR. If zero the AR is ignored.
Maximum Phosphate Resistance (PR)	The maximum allowed value of PR. If zero the PR is ignored.
Save	Saves the current settings to a *.GST file.
Load	Loads the current settings from a previously saved *.GST file.
Save As New Glass Catalog	Saves all the glasses that meets the current criteria into a new glass catalog.
OK	Accepts the current settings.
Cancel	Reverts to the prior settings.
Reset	Restores the default settings.

Discussion:

If "Use Glass Substitution Template" is unchecked, any glass defined in the current catalogs may be selected if not explicitly excluded (by having "exclude substitution" checked in the glass catalog) and if the wavelength range is suitable for the current lens. For glasses that do not include the cost or resistance codes, a "?" will be displayed in the glass catalog. If no cost or resistance data is provided for a glass, the glass will not be selected if "Exclude Glasses With No Data" is checked. Template data is stored along with the lens file; so the template parameters are lens file specific. If no glass meets the specified template parameters, the glass substitution status is automatically removed.

The glass substitution template is also used when converting from model to real glasses. When converting model glasses to real glasses, it is possible that no glass can be found that meets the template specification. If this happens, ZEMAX will first ignore the "Exclude Glasses With No Data" selection and then try again to find a suitable glass. If still no glass can be found that meets the criteria, then ZEMAX ignores the template entirely and tries again. If still no suitable glass can be found, glass number 1 in the catalog, whatever properties it may have, will be used, even if the glass does not have a suitable wavelength range or is far from the model glass parameters.

Individual tests for Relative Cost, CR, FR, SR, AR, and PR may be activated and deactivated by checking the box next to each description.

Tolerancing

Tolerancing

Purpose:

Tolerancing.

Discussion:

See "TOLERANCING" on page 545 for a complete discussion.

Tolerance Listing

Purpose:

This function generates a text listing of the tolerances which can be saved or printed.

Tolerance Summary

Purpose:

This function generates a text listing of the tolerances which can be saved or printed. The format is somewhat easier to read than the Tolerance Listing, and it uses no mnemonics which makes the tolerances easier to understand for fabricators and others not familiar with ZEMAX terminology.

Test Plates

Test Plate Fitting

Purpose:

Provides automatic fitting of radii to a vendor test plate list.

Settings:

Item	Description
File Name	Selects different test plate lists
Method of Fit	Chooses the fitting order, which may affect the quality of the fit.
# Opt Cycles	Chooses the number of optimization cycles to run, or Automatic.

Discussion:

This feature automates the task of matching radii of lens elements to the existing tooling of a particular vendor. The current merit function is used as a figure of merit during the fitting process.

To match a particular radius, make the radii variable on the lens data, or make the CRVT operand corresponding to the radii variable on the multi-configuration editor. As many radii as desired can be matched at one time. ZEMAX will attempt to fit all radii with variable status, even if the radii correspond to surfaces that are not spheres, such as toroids, cylinders, or other aspheric surfaces. To prevent fitting of the radii on these surfaces, remove the variable flag before performing the fitting.

Now select the test plate fitting tool. Select which vendor test plate file is to be used. The fitting method may be selected as follows:

Try All Methods: Try all of the following methods, and uses whichever one yields the lowest merit function.

Best To Worst: Fits the radii with the closest test plate match (measured in fringes) first.

Worst To Best: Fits the worst fitting radii first.

Long To Short: Fits the longest radii first.

Short To Long: Fits the shortest radii first.

Press OK to begin the fitting. ZEMAX starts by searching the test plate list for the closest match (in fringes) between all of the radii and all of the test plates. Each test plate must be of the correct shape (convex or concave or both, as required) and must have sufficient diameter to test the clear aperture of the lens surface (as determined

by the semi-diameter value on the Lens Data Editor). The test plate has sufficient diameter if the diameter is at least 3/4 of the clear aperture diameter of the lens surface (if the test plate diameter is defined as zero, this test is not applied and any semi-diameter value is accepted). The surface being fitted must also have an edge sag no more than 0.2 mm different than the test plate at the semi-diameter of the surface. In practice this is hundreds of fringes of power and test plates that do not meet this test would likely increase the merit function unacceptably.

The radius of the test plate that best fits one of the radii is then substituted in for the actual radius. The variability of the chosen radius is removed, and the lens is reoptimized. For this reason it is important to allow compensators such as spacings to be variable as well as the radii to be fit. The reoptimization will adjust all remaining variables, including the remaining radii. Note the optimization will use the current merit function. After optimization, if there are any variable radii remaining, the procedure repeats. Note that the radii are not fit in general in the order they appear in the lens data or multi-configuration editors.

During the fitting process, the number of radii remaining and the current merit function will be presented. After all of the radii have been fit, a report will be presented. The report will provide the vendor identification information, and a list of the radii that were changed and which vendor test plate ID numbers were selected.

There is no way to know if the test plates selected are optimal. If there are a large number of test plates in the test plate list, and the values are reasonably continuous without large gaps, the fitting will usually be quite good. If the merit function increases unacceptably during the fitting process, either a different vendor test plate list should be used, the design needs to be modified, or some of the lenses may have to have custom test plates made. Usually, the last radii to be fit are the likely candidates for custom test plates. The report shows the order in which the radii were fit.

Using the "Try All Methods" option will yield the lowest overall merit function, since all 4 methods are tried and the only the one yielding the lowest merit function is retained. However, there may be some algorithm other than the 4 listed that might yield a better fit.

There are occasions where the algorithm is unable to find any suitable test plate match for a given radius. This can occur when no test plate with a large enough diameter is sufficiently close to the desired radii. If this occurs, a message stating "NO MATCH FOUND!" will be printed in the report file, and that radii will subsequently be ignored. Usually this means test plates for this radii will have to be custom made.

All test plate data is provided by the respective vendors, and no warranty is provided as to the accuracy or completeness of the data. For details on the latest vendor test plate lists, contact the vendor. New test plate files may be added to ZEMAX. The test plate data files end in the extension .TPD, and must be placed in the <data>\Testplate folder (see "Folders" on page 68). The text files use the following format:

```
! Header line 1
! Header line 2
! etc... up to 15 lines of header info is supported
partname radius diameter code
partname radius diameter code
partname radius diameter code
partname radius diameter code
etc...
```

where partname is an ID number or name for the test plate, radius is the radius of curvature in millimeters, diameter is the actual plate diameter in millimeters, and code is an integer value. The code should be -1 for concave only, 0 for both concave and convex, and 1 for convex only. All four values must be on one line, separated by spaces or tabs. No spaces or tabs are allowed in the partname parameter. Each test plate (or pair) should be on one line, and the lines separated by carriage returns. The numeric values can be in free format, but must be in millimeters. The maximum number of test plates allowed per file is 30,000. The header lines can be used for any purpose, and typically indicate the vendor's name, address, phone, e-mail, and other contact data. The "!" is replaced by a space when ZEMAX lists the data out on test plate lists and fitting reports.

Test Plate Lists

Purpose:

Displays in a text window a list of test plates from a specific vendor.

Settings:

Item	Description
File Name	The name of the test plate file. The file must be in the <data>\Testplate folder (see "Folders" on page 68).

Discussion:

All units on the report are in mm. The CC and CX columns indicate the availability of concave and convex test plates, respectively.

Catalogs

Glass Catalogs

Purpose:

Provides access to the glass catalogs.

Discussion:

See the chapter "Using Glass Catalogs" for an extended discussion of this feature.

Glass Compare

Purpose:

Compares 2 or 3 glass catalogs to assist in finding glasses used in visible light with similar index and dispersion characteristics.

Settings:

Item	Description
Primary	The name of the primary glass catalog. This catalog is used as the reference catalog.
Secondary	The name of the secondary glass catalog. Only glasses which are considered equivalent to those in the primary catalog will be listed.
Tertiary	The name of the tertiary glass catalog. Only glasses which are equivalent to those in the primary catalog will be listed.
Index Tolerance	The maximum allowed difference in index at d-light for a glass from the secondary or tertiary catalog to be considered equivalent to a glass from the primary catalog.
Abbe Tolerance	The maximum allowed difference in Abbe number at d-light for a glass from the secondary or tertiary catalog to be considered equivalent to a glass from the primary catalog.
Show	Either all glasses in the primary catalog may be listed, or just those used in the current lens file and configuration.
Sort	Results may be sorted by index at d-light, name, or Abbe number at d-light.

Discussion:

This feature only lists glasses that are used in the F-C wavelength range. If "Show Used Glasses" is selected, the name of the glass displayed in the glass column of the lens data editor is used to determine if the glass from the primary catalog is used by the current lens and configuration. However, there is no check made that the glass from the primary catalog is actually used by the lens. For example, if the current lens uses glass A1 from catalog X, the primary catalog is Y, and Y also contains a glass called A1, then the data for A1 from catalog Y will be displayed, even though the lens actually uses the A1 as defined in catalog X. If more than one catalog uses the same name for different glasses, great care should be taken to verify all data is for the intended glass.

Glass Fitting

Purpose:

Fits glass dispersion data at either a single temperature or thermal dispersion over a range of temperatures.

Settings:

Item	Description
Fit Type	Choose either Index Data or Thermal Index Data.
Formula	Choose the name of the desired glass dispersion formula. For a complete description of the formulas, see "The glass dispersion formulas" on page 590.
Data File	The name of the data file that contains the dispersion data. When fitting index data, the file extension is IND. For thermal index data, the file extension is TID. See the discussion below for information on the file formats.
Add To Catalog	If "none" is selected the data will be fit and a report generated, but the data will not be added to a catalog. If a glass catalog name is selected, then the fit data will be added to the specified catalog. Note that adding glass data to a catalog that is supplied with ZEMAX is not allowed, therefore only the names of catalogs not supplied with ZEMAX are listed. If no catalog is listed, then a new catalog must first be created. For information on doing this, see "Creating a new catalog" on page 589.
Glass Name	The name of the glass to add to the selected glass catalog.

Discussion:

There are two different file formats and fitting algorithms supported by this feature. For index data, the text file should end in the extension IND. The first line of the file should contain the word TEMPERATURE followed by the temperature of the glass in degrees Celsius at which the data was measured. The next line should contain the word PRESSURE followed by the ambient air pressure in atmospheres. If one or both of these lines are omitted the default values are 20 degrees and 1.0 atmospheres. The remaining lines consist of two columns of data: wavelength and index. The wavelength values are in micrometers, measured in air at the specified temperature and pressure. The index must be relative to air at the temperature and pressured defined by the first two lines. ZEMAX will automatically convert the wavelength and index data to be relative to air at a pressure of 1.0 atmospheres at the reference temperature before fitting. A sample file might look like this:

```
TEMPERATURE 20
PRESSURE 1.0
 0.3500000 1.5391663
 0.4500000 1.5253195
 0.5500000 1.5185224
```

etc...

For index data, the maximum number of points allowed to be fit is 1000.

For thermal index data, the text file should end in TID. The first line should contain the word PRESSURE followed by the ambient air pressure in atmospheres. If this value is omitted the default value is 1.0 atmosphere. The remaining lines consist of three columns of data: temperature, wavelength, and index. **The temperature of the first data point defines the reference temperature for the glass.** It is important that the reference temperature data be listed first. Data for other temperatures should be listed in subsequent lines. All temperatures are measured in degrees Celsius. The wavelength values are in micrometers, measured in air at the reference temperature and pressure. The index values must be relative to air at the temperature of each data point at the pressure defined by the PRESSURE argument. ZEMAX will automatically convert the wavelength and index data to be relative to air at a pressure of 1.0 atmospheres before fitting. A sample file might look like this:

```
PRESSURE 0.0
20.0000000 0.3500000 1.5391663
20.0000000 0.4500000 1.5253195
20.0000000 0.5500000 1.5185224
20.0000000 0.6500000 1.5145203
```

etc...

```

-20.0000000  0.3500000  1.5391663
-20.0000000  0.4500000  1.5253195
-20.0000000  0.5500000  1.5185224
-20.0000000  0.6500000  1.5145203
etc...
 40.0000000  0.3500000  1.5391663
 40.0000000  0.4500000  1.5253195
 40.0000000  0.5500000  1.5185224
 40.0000000  0.6500000  1.5145203
etc...

```

The maximum number of points for the reference temperature fitting is 1,000. The total number of points for all temperatures and wavelengths that may be fit is limited only by available memory. Note that the fit will only be as good as the data provided. At least 6 points should be used at each wavelength, and preferably more than 12. At least 6 different temperature values should be provided. Using more than about 30 points at any one temperature tends to only increase the fit time and does not improve accuracy of the fit.

As with all fitting of data, the resulting fit should be carefully checked for suitable accuracy before using the data. It is particularly important to look at the fit index data at temperatures and wavelengths not included in the fitted data, because dispersion formulas can diverge somewhat from reasonable values between data points if the data points are sparse.

Both the IND and TID files should be placed in the folder <glass> (see "Folders" on page 68).

Lens Catalogs

Purpose:

Used to search for a particular lens from or browse through the stock lens catalogs.

Settings:

Item	Description
Vendor	Lists the available stock lens vendors. Each vendor name is the name of the file which contains the stock lenses available from that vendor. The vendor files must be placed in the <data>\Stockcat folder (see "Folders" on page 68).
Use Effective Focal Length	If checked, then the specified effective focal length range is used as part of the search criteria, otherwise, any value will be accepted.
Focal Length Min/Max	Used to define the range of acceptable effective focal lengths in millimeters.
Use Entrance Pupil Diameter	If checked, then the specified entrance pupil diameter range is used as part of the search criteria, otherwise, any value will be accepted.
Diameter Min/Max	Used to define the range of acceptable entrance pupil diameters in millimeters.
Equi-, Bi-, Plano-, Meniscus	If any of these options are checked, then the search will be limited to lenses meeting at least one of the selected criteria. For example, if both Equi- and Bi- are selected, then both equiconvex/equiconcave and biconvex/biconcave lenses will be included.
Spherical, GRIN, Aspheric, Toroidal	If any of these options are checked, then the search will be limited to lenses meeting at least one of the selected criteria. The "Spherical" category is also used as "other" for lenses which are not gradient index, aspheric, or toroidal.
# Elements	Selects the number of elements allowed in lenses that meet the search criteria. Choose "Any #" to include all lenses in the search.
Search Results	Lists all lens files in the specified catalog which match the current search criteria.

Search	Causes a new search to be conducted for lenses in the specified vendor file based upon the current search criteria.
Load	Loads the currently selected lens file into the Lens Data Editor, and updates all windows to display data for the newly loaded lens.
Insert	<p>Inserts the currently selected lens file to the existing lens.</p> <p>In sequential mode, the lens will be inserted into the Lens Data Editor. A prompt will be issued for the insert file options, such as at which surface the lens should be inserted.</p> <p>In non-sequential mode, the lens will be inserted into the Non-Sequential Components Editor at the bottom of the list of current objects. If the inserted lens consists of more than one component, the second and subsequent component will use the first as its reference object, so only the first component needs to be properly positioned.</p>
Prescription	Displays a prescription report of the currently selected lens file.
Layout	Displays a layout plot (or 3D Layout for systems without axial symmetry) of the currently selected lens file.
Close	Closes the dialog box.

Discussion:

The "Search Results" table will list the part name, focal length, and diameter of all the lenses from the selected vendor which meet the specified search criteria. After the diameter, a 3 part code is listed, such as (P,S, 1). The first entry is the shape code, which is E, B, P, M, or "?", for Equi-, Bi-, Plano-, Meniscus, or other. Other is used if the lens is multi-element. The second entry is S, G, A, or T, for Spherical, GRIN, Aspheric, or Toroidal. The third entry is the number of elements.

Upon installation, ZEMAX creates a <data>\Stockcat folder (see "Folders" on page 68). Within this folder, many individual files are stored with the extension .ZMF. Each of these files contains a large number of individual ZMX lens files, each of which represent stock lenses available from various vendors. The following table lists the represented vendors and contact data.

STOCK LENS VENDORS

File Name	Website
ANTERYON.ZMF	www.anteryon.com
ARCHEROPTX.ZMF	www.archeroptx.com
ASPHERICON.ZMF	www.asphericon.com
BEFORT-WETZLAR.ZMF	www.befort-optic.com
COMAR.ZMF	www.comaroptics.com
CVI MELLES GRIOT.ZMF	www.cvimellesgriot.com
DAHENG OPTICS.ZMF	www.dahengoptics.com
DIAS INFRARED.ZMF	www.dias-infrared.de
EALING.ZMF	www.ealingcatalog.com
EDMUND OPTICS.ZMF	www.edmundoptics.com
ESCO.ZMF	www.escoproducts.com
GELTECH.ZMF	www.lightpath.com

File Name	Website
ISP OPTICS.ZMF	www.ispoptics.com
JML.ZMF	www.jmloptical.com
LIGHTPATH.ZMF	www.lightpath.com
LINOS PHOTONICS.ZMF	www.linos-photonics.com
MIDWEST OPTICAL.ZMF	www.midwestopticalsystems.com
NEWPORT CORP.ZMF	www.newport.com
NSG.ZMF	www.nsgamerica.com
OPTICS FOR RESEARCH.ZMF	www.ofr.com
OPTOSIGMA.ZMF	www.optosigma.com
QIOPTIC_POLYMER.ZMF	www.qiopticpolymer.com
ROLYN OPTICS.ZMF	www.rolyn.com
ROSS OPTICAL.ZMF	www.rossoptical.com
RPO.ZMF	www.rpoptics.com
SIGMA KOKI.ZMF	www.sigma-koki.com
SPECIAL OPTICS.ZMF	www.specialoptics.com
THORLABS.ZMF	www.thorlabs.com

To find a catalog lens that closely matches the properties of one or more surfaces in the Lens Data Editor, place the cursor on the first surface of the surface group in the Lens Data Editor, then select Tools, Lens Catalogs. The default focal length and diameter search ranges will be set to appropriate values for that lens. ZEMAX computes the lens focal length and diameter and adds plus/minus 5% to define the default search ranges.

The Lens Search tool provides the capability to search through the available lenses to find an appropriate choice. Once selected, the lens file may be loaded or appended to an existing design.

Make Private Lens Catalog

Purpose:

Used to convert a folder full of ZMX files to a stock lens catalog file.

Settings:

Item	Description
Source Folder	Enter the name of the folder where the source ZMX files reside. All ZMX files in this folder will be placed in the ZMF file.
ZMF File Name	The name of the ZMF file to be created - this file will be placed in the Source Folder.
Convert	Begins the conversion process.
Cancel	Closes the dialog box.

Discussion:

This feature is used to convert multiple ZMX files into a single ZMF file that can be used by the Lens Catalogs feature. See "Lens Catalogs" on page 246 for a description of this feature.

There are limitations as to what kind of lenses can sensibly be placed in a private lens catalog. The intended purpose is to include in the catalog only simple lenses that are reasonably described as singlets, doublets, or triplets. Lenses with complex surface shapes, DLLs, non-sequential groups, etc. may not properly convert to the simplified stock lens format used to make these files searchable. File names should use simple ANSI names, as the first space delimited word in the name of the lens file in the source folder is used as the "part name" in the stock lens catalog. For example, a file with the name "Lens number 5.ZMX" will be assigned a part name of "Lens". A better name would be L005.ZMX so that the name "L005" would be used as the part name. Unicode file names are not supported, and will be converted to the nearest equivalent ANSI string, which in some cases may result in odd looking part names.

Coatings



These features are only available in the EE edition of ZEMAX.

Edit Coating File

Purpose:

Invokes the Windows NOTEPAD editor to edit the coating file for the current lens. This file contains the material and coating definitions.

Discussion:

See the chapter "Polarization Analysis". If the coating file is edited, the file must be reloaded manually by selecting "Reload Coating File" as described below, or the file may be saved and reloaded to update the new coating data.

Reload Coating File

Purpose:

Reloads the coating file for the current lens. This is required if the file has been edited since the lens was loaded.

Add Coatings to All Surfaces

Purpose:

Adds any specified coating to all surfaces with air-glass boundaries.

Discussion:

When selected, this tool will prompt for the name of the coating to apply. Any defined coating name may be selected, or choose "none" to remove any existing coatings from surfaces. All surfaces which define a transition from glass to air will have the coating applied (or removed); so this feature is primarily for applying anti-reflection coatings.

Coating Listing

Purpose:

This function generates a text listing of the definitions of materials and coatings contained in the coating file for the current lens.

Export Encrypted Coating

Purpose:

This feature exports a single coating to a ZEMAX Encrypted Coating (ZEC) file. To export an entire lens, see "Export ZEMAX Black Box Data" on page 260.

Settings:

Item	Description
Coating	The name of the coating to export.

File Name	The name of the file to create. The file name should not include any path or extension. The file will be created in the coating file folder (see “Folders” on page 68) and the extension "ZEC" will be automatically appended.
Wavelength	The wavelength, in micrometers, measured in air at system temperature and pressure, used to define the layer thicknesses. This value is only used when converting layers from optical thickness to absolute thickness. See below.
Convert To Absolute Thickness	Converts optical thickness defined relative to the primary wavelength (for example, a quarter wave thickness) in the coating to absolute thickness. For maximum security, this option should be selected. If the relative optical thickness is left in the coating definition, then the end user of the encrypted coating may be able to determine some properties of the coating by changing the design wavelength. The conversion from thickness in primary waves to absolute thickness is done using the wavelength value from the above setting.

Discussion:

Optical coating designers and manufacturers are often hesitant to divulge the details of their coating designs. This feature allows a coating defined in ZEMAX format to be exported to a single file that is encrypted. The encrypted file may be distributed to users that need to trace rays and perform analysis on the coating as applied to a specific optical system, without the details of the coating design being revealed.

The coating data that is exported includes all the material, index, extinction, dispersion, and layer data required to exactly reproduce the coating performance. No approximations or tables are used. However, no coating tapers, ideal coatings, or table coatings are supported. A maximum of 12 different materials and 400 layers may be used in defining the coating to be exported.

To use an encrypted coating file, see “The ENCRYPTED data section” on page 618.

The encryption is based upon a 256-bit algorithm that provides good, but not unbreakable security. This feature is provided as is and the user should determine if the encryption implemented provides adequate security for the intended application before distributing sensitive data in the encrypted format. Radiant ZEMAX LLC provides no warranty and assumes no liability for the use of this feature.

Scattering

ABg Scatter Data Catalogs

Purpose:

Provides access to the ABg scatter data catalogs.

Settings:

Item	Description
File	The current ABg scatter data file name. This file must be placed in the <data>\ABg_Data folder (see “Folders” on page 68); the file name may be selected on the System, General dialog box on the Files tab, as described on page 109.
Name	The name of the ABg data being edited.
Wave	The wavelength in micrometers at which the ABg data is defined (or was measured). If this value is zero, no scaling with wavelength will be performed. Otherwise, the A and B values will be scaled with wavelength as defined in the discussion below.
Use	If checked, then the displayed row will be used. Up to 10 data rows may be defined for each ABg data name.
Angle	The angle of incidence in degrees. The value should be between 0.0 and 90.0, inclusive. The BSDF for intermediate angles of incidence are linearly interpolated in cosine space. If data is required at angles less than the smallest angle defined, data for the smallest angle defined will be used. Data for the largest angle defined will be used for any incidence angles greater than the largest angle defined.

Item	Description
A, B, g	The A, B, and g values at the specified angle of incidence.
TIS	Pressing this button will compute the total integrated scatter (TIS) for the displayed data. This value must be less than 1.0 for accurate scattering results. See the discussion.
Save	Saves the modified data to the catalog.
Save As	Saves the modified data to a catalog of a different name.
Exit	Closes the dialog box.
Insert	Inserts a new ABg data entry into the catalog; and prompts for the name of the new entry.
Delete	Deletes the displayed ABg data entry.
Reload	Loads from the last saved data file the ABg data; eliminating any changes made since the last save.
Sort	Sorts in ascending order by angle the displayed data. The data is automatically sorted whenever the catalog is saved or loaded.
Report	Sends the data to a text window which may then be printed.
Rename	Prompts for a new name for the ABg data.

Discussion:

ABg data values are used to define scattering properties of both sequential and non-sequential surfaces. Each ZEMAX lens files may use a separate ABg data catalog if desired. The ABg data file used is specified on the System, General dialog box on the Files tab, as described on page 109. The selected ABg catalog only may be edited from within ZEMAX using the ABg Scatter Data Catalog dialog box.

For a theoretical discussion of ABg scattering; see “ABg model scattering” on page 570.

The integral of the ABg distribution function over the unit vector circle that represents physically possible scattering ray directions determines the total integrated scatter, or fraction to scatter. This integral is automatically computed by ZEMAX when required; to view this integral press the TIS button. The TIS must be less than 1.0; otherwise conservation of energy may not be preserved and the scattering will not be correct.

Wavelength scaling of ABg data

If a reference wavelength other than zero is defined, then the A and B coefficients are scaled according to the following relations:

$$A' = A \left(\frac{\lambda}{\lambda_{ref}} \right)^{g-4}, \text{ and}$$

$$B' = B \left(\frac{\lambda}{\lambda_{ref}} \right)^g,$$

where λ is the wavelength being traced and the unsubscripted A, B, and g parameters are measured at the reference wavelength, λ_{ref} .

When using the ABg scatter models, ZEMAX can only scatter light using ABg parameters scaled for wavelengths defined in the System Wavelengths Dialog box. If you wish to define a broadband source via a Source Color spectrum (e.g. Tristimulus XYZ), and the ABg data varies strongly with wavelength, then you need to defined a reasonable number of System Wavelengths to span the bandwidth of the source. For an arbitrary wavelength being traced, ZEMAX will use the closest System Wavelength to determine the ABg coefficients and

scatter distribution. This is the only situation in which System Wavelengths need to be defined when using a Source Color model.

Scatter Function Viewer

Purpose:

Displays BSDF as a function of the magnitude of the scatter vector x .

Settings:

Item	Description
Lambertian	Selects Lambertian scattering.
Gaussian	Selects Gaussian scattering.
ABg	Selects ABg scattering.
User Defined	Selects User Defined scattering.
Wavelength	The wavelength number to use. ABg scattering uses the specified wavelength, along with the reference wavelength defined in the ABg catalog to scale the ABg coefficients.
Fraction	The fraction of total energy to scatter. This value must be between 0 and 1. The fraction to scatter linearly scales the BSDF. This fraction is ignored for ABg scatter functions.
Sigma	The sigma value to use for Gaussian scattering.
ABg Name	The name of the ABg scattering data to use. ABg scatter data names are defined in the ABg Scatter Data Catalogs; see "ABg Scatter Data Catalogs" on page 250.
DLL Name	The name of the user defined scattering DLL to use. See "User defined scattering" on page 572.
File Name	The name of the data file, if any, used by the user defined scattering DLL.
(parameter names)	These parameters are for user defined scattering DLL.
Bins	The number of histogram bins to use in the computation and plot of the BSDF for user defined scattering. Using a larger number of bins increases the resolution of the plot, at the expense of requiring more rays to generate smooth data.
Rays	The number of rays to scatter to determine the shape of the BSDF curve for user defined scattering.
Alpha X, Y	The angle in the X-Z and Y-Z planes respectively, of the specular ray. See the discussion for details.
X, Y Minimum	The minimum value of the scatter vector x , and BSDF (y), to plot, respectively. Note these are powers of 10 as the BSDF plot is a log-log plot.
X, Y Decades	The number of log scale decades to plot in the x and BSDF (y) directions, respectively.

Discussion:

This feature plots or lists BSDF data for any of the surface scattering models supported by ZEMAX. The BSDF plot is a log - log plot, with BSDF on the vertical axis, and the magnitude of the scattering vector x on the horizontal axis. The scattering vector x is a 2D vector that defines the change in the x and y direction cosines between the specular ray and the scattered ray, when the local surface normal is oriented along the z direction. The vector x is therefore the difference between the specular and scattered ray when these rays are projected onto the plane of incidence. For a detailed discussion and drawing illustrating the scattering conventions used by ZEMAX, see "Scatter models" on page 568.

For User Defined and ABg scattering, the Alpha X and Y values are used to define the specular ray direction cosines. The defining equations which convert angles to cosines are as follows:

$$n = 1 / (\sqrt{1 + (\tan \alpha_x)^2 + (\tan \alpha_y)^2}),$$

$$l = n \tan(\alpha_x), \text{ and}$$

$$m = n \tan(\alpha_y).$$

The comments below apply to each of the available scattering models:

Lambertian: Lambertian scattering is independent of incident angle and wavelength. The total integrated scatter (TIS) is 1.0. The resulting BSDF "curve" is a flat line with a value of $1/\pi$.

Gaussian: Gaussian scattering is symmetric about the specular ray projection. The TIS is always 1.0.

ABg: The ABg values are defined in the ABg Scatter Data Catalogs; see "ABg Scatter Data Catalogs" on page 250 for more information. The TIS is not generally 1.0; the actual value is listed in the output plot or listing. The ABg data may be scaled with wavelength; if a reference wavelength is defined in the catalog. ABg data is also generally a function of incident angle. ZEMAX interpolates the BSDF data between the two points with the angles of incidence closest to that of the defined specular ray.

User Defined: ZEMAX cannot determine the exact curve of the BSDF for user defined scattering. As an alternative, a number of identical specular rays are scattered by the DLL, and the resulting scatter vector x values are binned in a histogram, normalized, and plotted. BSDF plots made in this way are typically noisy in a complex way that depends upon the nature of the scattering function defined by the DLL. Note that if the BSDF as defined by the DLL is not symmetric about the specular ray in the projected tangent plane, then this plot is not an accurate description of the BSDF.

Apertures

Convert Semi-Diameters to Circular Apertures

Purpose:

Converts all surfaces without surface apertures to fixed semi-diameters with circular apertures set to the semi-diameter value.

Discussion:

The primary purpose of this feature is to simplify the analysis of vignetting effects. For most optical designs, it is best to use vignetting factors (see "Vignetting factors" on page 115) during optimization. However, vignetting factors are an approximation. This feature converts all semi-diameters to surface apertures. The vignetting factors may then be deleted (this is not done automatically by this feature) and the pupil can be overfilled to see where the rays really make it through the system.

Convert Semi-Diameters to Floating Apertures

Purpose:

Converts all surfaces without surface apertures to floating apertures which vignette at the semi-diameter value.

Discussion:

This feature is very similar to "Convert Semi-Diameters to Circular Apertures" except floating apertures are used instead of fixed circular apertures. Floating apertures leave the semi-diameter value of the surface in "automatic" mode, and dynamically adjust the vignetting aperture to match the semi-diameter value. Note if any semi-diameters are "fixed" they remain fixed, and all vignetting will occur at the specified semi-diameter on each surface.

Convert Semi-Diameters to Maximum Apertures

Purpose:

Converts all surfaces without surface apertures to semi-diameters set to the maximum required value to pass all rays without vignetting in all configurations.

Discussion:

The primary purpose of this feature is to set adequate apertures on all surfaces in zoom lens systems. See "Semi-Diameter: Maximum" on page 478.

Remove All Apertures

Purpose:

Turns off all surface apertures.

Discussion:

All surface aperture types are set to "none" by this tool. Note floating apertures are added automatically to surfaces with power that have user defined semi-diameter values, and so these apertures cannot be deleted; see "Semi-diameters" on page 59 for more information.

Replace Vignetting With Apertures

Purpose:

Turns off vignetting factors and places surface apertures.

Discussion:

This tool uses the identical algorithm as many ZEMAX analysis features do to replace vignetting factors with surface apertures. Analysis features that need to trace rays from field points in between defined field points need to do this because vignetting is only defined for discrete field points. This tool is useful as a check that ZEMAX does this conversion properly; and is also handy for removing vignetting factors.

Coordinates

These tools are used to modify tilts, decenters, and other coordinate break related data.

Tilt/Decenter Elements

Purpose:

Tilts or decenters a single surface or a group of surfaces forming an element.

Settings:

Item	Description
First Surface	This is the first surface of the lens group to be tilted/decentered.
Last Surface	The last surface of the lens group to be tilted/decentered.
Decenter X, Y	The X and Y decenters in lens units.
Tilt X, Y, Z	The rotation about X, Y, and Z in degrees.
Order	Either tilt then decenter, or decenter then tilt.
Coord Brk Color	Selects the spreadsheet display row color for the inserted coordinate breaks. See "Row Color" on page 77 for a description of Row color.
Coord Comment	The comment string for the inserted coordinate break surfaces.
Hide Trailing Dummy Surface	If checked, any trailing dummy surface inserted as part of the tilt and decenter tool modifications will be checked with the options "Skip Rays To This Surface" and "Do Not Draw This Surface". See "Surface properties draw tab" on page 78.

Discussion:

The feature inserts coordinate break and dummy surfaces as required to tilt and/or decenter a single surface or a range of surfaces. For more information on coordinate breaks see “Coordinate Break” on page 293.

Add Fold Mirror

Purpose:

Inserts a fold mirror, including the required coordinate breaks, to bend a beam.

Settings:

Item	Description
Fold surface	Selects which surface will become the fold mirror. The selected surface should be a dummy surface already located at the desired fold position.
Reflect angle	The angle between the incident and reflected beam.
Tilt type	Select either tilt about the local x or y axis.

Discussion:

This feature inserts two dummy surfaces, one before and one after the selected fold surface. The fold surface then becomes a mirror, and the two newly inserted adjacent surfaces are set to be coordinate breaks with the appropriate tilt angles. The second tilt angle is set as a pickup from the first tilt angle. Finally, all subsequent surface thicknesses and curvatures are changed sign to account for the new mirror.

This feature may not provide useful or even correct alterations if the selected fold surface is not a plane, standard type dummy surface in air. The dummy surface should be located at the desired fold mirror position before using this feature. For example, to insert a fold mirror midway between two lenses 100 mm apart, a dummy surface should be inserted between the two lenses, with the thickness before and after the dummy being set to 50 mm. The dummy surface then should be used as the fold surface.

Limitations of reversing surfaces following the fold mirror

After inserting a mirror, all subsequent surfaces must be reversed in the sense that +z becomes -z for the rays to trace the same. For most ZEMAX surfaces, this means changing the sign of the radius of curvature, or any other parameters that determine the sign of the surface sag. For user defined surfaces, ZEMAX cannot automatically make this adjustment. For non-sequential component (NSC) surfaces, it may not be possible. ZEMAX attempts to reverse NSC surfaces by changing the sign of the Z coordinate of the exit port, and rotating all components around the y axis 180 degrees. However, this method fails if the objects inside the NSC are not symmetric in their local XZ and YZ planes. Because the z parity of the objects must be modified, there is no general way to do this for all NSC objects. This feature may also not work for lenses with multi-configuration data that alters the thicknesses or glass types of any surfaces from the fold surface onward.

Delete Fold Mirror

Purpose:

Deletes an existing fold mirror, including any neighboring coordinate breaks.

Settings:

Item	Description
Fold surface	Selects which mirror surface will be deleted.

Discussion:

This feature deletes a fold mirror. If the mirror surface is preceded by a coordinate break with zero thickness, the coordinate break will be deleted. If the mirror surface is followed by a coordinate break, and the mirror surface has zero thickness, the following coordinate break will be deleted. The thicknesses of all deleted surfaces will be added to preceding surfaces as required. All following surfaces and thicknesses will be automatically reversed

as required to follow the correct sign convention. See “Limitations of reversing surfaces following the fold mirror” on page 255 for details. This feature may not provide useful or even correct alterations if the selected surface is not a plane, standard type mirror surface in air.

Local To Global

Purpose:

This feature combines, modifies, and deletes existing coordinate break surfaces as required to position groups of surfaces using a single decenter and rotation relative to a reference surface. The original position of the surfaces is not modified, only the coordinate decenters and rotations are modified. See also “Global To Local” on page 256

Settings:

Item	Description
First/Last Surface	The range of surfaces to modify. The first surface must precede the last surface. Neither of these surfaces may be a coordinate break.
Reference Surface	The surface number to which the coordinates and rotations are referenced. This surface must precede the First Surface.

Discussion:

This feature automates the process of placing surfaces in a common reference coordinate system. A group of surfaces is defined by any number of surfaces bounded by a coordinate break or the image surface. Each group is then positioned relative to the reference surface by three coordinate break surfaces. The first coordinate break returns the coordinate system to the reference surface. The second coordinate break positions the group using the decenter x, decenter y, and thickness (z coordinate) relative to the reference surface. Finally, the third coordinate break rotates the group into the correct orientation. Each set of three coordinate breaks precedes each group of surfaces in the system, within the defined surface range. The primary application of this tool is to place all groups of surfaces relative to a common reference for tolerancing or perturbation analysis.

This feature is not supported if any surface thickness, ignored surfaces, or any coordinate break properties within the range are controlled by the multi-configuration editor.

Global To Local

Purpose:

This feature combines all consecutive coordinate breaks into single coordinate breaks. The original position of the surfaces is not modified, only the coordinate decenters and rotations are modified. Each group of surfaces will be referenced to the previous group’s local coordinates. The combination of consecutive coordinate breaks can also be used to reverse the process performed by the Local To Global tool. See also “Local To Global” on page 256.

Settings:

Item	Description
First/Last Surface	The range of surfaces to modify. The first surface must precede the last surface. Neither of these surfaces may be a coordinate break.
Order	Select "Forward" to use coordinate breaks that decenter then tilt, or "Reverse" to use coordinate breaks that tilt and then decenter. For more information, see “The order flag” on page 294

Discussion:

This feature combines all consecutive coordinate breaks into single coordinate breaks. When multiple coordinate breaks are placed on consecutive surfaces, they can be combined into a single coordinate break with the same result.

This feature is not supported if any coordinate break properties within the range are controlled by the multi-configuration editor.

Export Data

Export IGES/SAT/STEP Solid

Purpose:

Exports the current lens data to a file as a collection of 3D solids in IGES, SAT, STEP, or STL format.

Settings:

Item	Description
First/Last Surface or Object	The range of surfaces to include in the exported data. In Non-Sequential mode, these options are First and Last Object.
Wavelength	The wavelength number to use for any traced rays to be exported.
Field	The field number to use for any traced rays to be exported.
Number of Rays	The number of rays to trace; the exact meaning depends upon the "Ray Pattern". When exporting to STL format, no rays are exported as STL does not support a line entity.
Split NSC Rays, Scatter NSC Rays, Use Polarization	These settings only affect rays from NSC sources. For more information, see the discussion on the similar controls described in "NSC 3D Layout" on page 125.
Ray Pattern	Selects the type of ray set to export. This control is very similar to that defined for the 3D layout plot; see "3D Layout" on page 118. The solid beam option exports a solid volume representing the envelope of the rays. The solid beam option only works properly if the marginal rays are all traceable and unvignetted and the marginal rays do not form a caustic at any surface. Beams are not generated for non-sequential surfaces. The number of rays must be at least 8 to export a solid beam, and a smoother solid results if more rays are used.
Lens/Ray Layer	Selects the layers of the output file to place lens and ray data.
Config	One or more configurations may be exported at once using this control. Current will export only data from the current configuration. Selecting a configuration number, such as 1, 2, or 3, will export data only from that configuration. "All By File" will export all configurations, but each to a separate file name. The file name for each configuration is the specified save file name (the save file name is selected after OK is pressed) with the name "_config000x" appended to indicate the data is for configuration "x". "All By Layer" will export all configurations, but each to a separate layer within the same file. The first configuration will be placed in the layer specified by the lens layer and ray layer settings. Subsequent configurations are placed in the lens and ray layers incremented by 1 for each configuration. For example, if the lens layer is 1 and the ray layer is 10 and there are 3 configurations, then configuration 1 will have lens data in layer 1 and ray data in layer 10. Configuration 2 will have lens and ray data in layers 2 and 11, and configuration 3 will have lens and ray data in layers 3 and 12, respectively. "All At Once" will export all configurations at once in a single layer. No attempt is made by ZEMAX to avoid overlapping or redundant data export.
File Type	Selects the output file format type.
Delete Vignetted	If checked, rays are not included in the data export if they will be vignetted by any surface.

Item	Description
Surfaces As Solids	This control only applies to sequential surfaces, and not NSC surface objects. If checked, surfaces on either side of a glass material will be combined to form a single solid object, where possible. Not all ZEMAX surface types may be combined to form a single solid. Both surfaces must have the same aperture shape to export surfaces as solids. For example, both surfaces must have circular apertures, or both surfaces must have rectangular apertures. If the aperture types do not match then the surfaces are exported as single surfaces.
Export Dummy Surfaces	If checked, dummy surfaces will be exported. A dummy surface is a surface that does not reflect or refract rays.
Dummy Thickness	The distance in lens units for the thickness of dummy surfaces if they are automatically converted to solids. Both "Export Dummy Surfaces" and "Surfaces As Solids" need to be checked for this control to have any affect.
Tolerance	The approximate accuracy of the exported objects in lens units. Smaller tolerances yield more accurate representations at the expense of larger exported files and longer computation times.
Spline Segments	The number of segments to use when spline entities are exported.

Discussion:

For information on the accuracy of exported surfaces see "Comments about imported objects" on page 513.

The Initial Graphics Exchange Specification (IGES) is an American National Standard whose intended purpose is to facilitate transfer of data between CAD programs. ZEMAX currently supports version 5.3 of the IGES standard. For more information on IGES, contact U.S. Product Data Association, P. O. Box 3310, Gaithersburg, MD 20885-3310. The SAT format is used by the ACIS geometry modeling engine developed by Spatial Technologies. The STEP format is AP203 as defined by ANS US PRO/IPO-200-203-1994, also available from the U.S. Product Data Association.

ZEMAX exports five types of data:

Lines: Lines are used to indicate the paths of rays traveling through the optical system. Within GRIN media, rays are exported as a series of line entities. The STL format does not support lines, so rays are not exported.

Surfaces: Surfaces may be of arbitrary shape, including user defined surfaces; with any aperture shape ZEMAX supports, including user defined apertures. Data for some types of surfaces (particularly aspheres and toroids) which are exported as splines have an accuracy that depends on the number of spline points used; more points yields more accuracy at the expense of larger, slower exported data files. For more information on the limitations of the accuracy of the exported objects see the discussion below.

Lens Solid: Lens Solids include a non-zero volume enclosed by two surfaces, with the edge region formed by the extrusion of the edges of the two surfaces. Most surfaces surrounding glass with similar shaped apertures (e.g. both surfaces have rectangular apertures) may be exported as solids.

Faceted Solid: A Faceted Solid is defined by a collection of triangles that form a volume, such as a prism or Fresnel lens. Non-sequential STL and POB objects, among others, are exported as Faceted Solids.

Parametric Solid: Some NSC objects, such as the Torus Volume, are exported as exact solid NURBS objects.

ZEMAX can export the surface or solid data for all sequential surfaces and non-sequential objects. ZEMAX exports all surfaces and rays in a 3D coordinate system referenced to the global coordinate reference surface which is described in "Global Coordinate Reference Surface" on page 111.

If exporting surfaces as solids, ZEMAX will export sequential dummy surfaces as thin solids. The front and back face will be of the shape of the surface, with a thickness equal to the dummy thickness in lens units. Mirror surfaces are exported according to the settings for the mirror substrate; for more information see "Mirror Substrate and Thickness" on page 79.

Most NSC objects are exported exactly, or by using one or more spline based surfaces similar to those used for sequential surface export. However, some objects, such as user defined objects, are exported using a faceted approximation to the smooth surfaces. Increasing the drawing resolution for these objects will also increase the

resolution of the exported representation, see “Draw tab” on page 565. Because STL is inherently a faceted format, all STL objects are exported as faceted approximations.

This feature will prompt for the file name and path to save the exported data to. The default file name and folder are the same as the current lens file, with an extension that depends upon the CAD file type format selected.

To import IGES, SAT, and STEP files see “Imported” on page 512. To Import STL files see “STL Object” on page 523.

Limitations of exported data

ZEMAX supports many types of complex surface and object shapes. Within ZEMAX, these shapes are modeled exactly, yielding the very high numerical precision required for optical ray trace analysis. However, most of the complex optical shapes ZEMAX supports have no analog in commonly used CAD programs, and no exact representation in common CAD file formats such as IGES, STEP, and SAT. For this reason, ZEMAX must approximate the exact surface shape with the closest available analog within the CAD format. This approximation is usually very good, but for some highly aspheric surface shapes, the approximation is not sufficiently accurate for ray tracing, and may not even be good enough to be a reasonable mechanical description of the surface shape. This is not a limitation of ZEMAX. This is a limitation of the CAD data exchange format and the CAD programs ZEMAX is being asked to export data to. In these cases, it may be necessary to recreate the complex geometry within the CAD program directly rather than use the ZEMAX export capability.

There is always the possibility that the conversion from the ZEMAX representation of surfaces and objects to the CAD file format representation will cause a loss of precision. For more information, see “Ray tracing accuracy for imported objects” on page 514.

Export Source Data

Purpose:

This features exports intensity data recorded by the Detector Polar to a text file. For detailed information on the Detector Polar, see “Detector Polar object:” on page 534. The exported file format may be in EULUMDAT *.LDT or IESNA *.IES format. These are commonly used file formats for lamp data. These are the same format ZEMAX uses for the EULUMDAT Source described in “Source EULUMDAT File” on page 546 and for the IESNA source described in “Source IESNA File” on page 549.

Settings:

Item	Description
Format	The file format to use.
Detector	The Detector Polar to use. Only Detector Polar objects are supported.
File Name	The name of the file to use. The file name may include a full path, otherwise, the file is placed in either the folder <objects>\Sources\EULUMDAT or the folder <objects>\Sources\IESNA depending upon the file format used (see “Folders” on page 68).
Smoothing	This control selects the number of times to apply the smoothing operator, which averages adjacent detector pixels to smooth out sharp discontinuities in the detected data. The smoothing operator is exactly the same as described in “The Detector Viewer” on page 587.

Discussion:

The EULUMDAT file format includes a table of luminous intensity data measured along polar arcs at various azimuthal angles. This arrangement of data closely matches the data recorded by the Detector Polar object. The same angular sampling in both the polar and azimuthal directions that the Detector Polar uses to record the data are used for export to the data file.

Most of the data fields in the exported file, other than the luminous intensity values, are not used by ZEMAX and may contain blanks, zeroes, or other “dummy” data. These fields may need to be edited to include valid data, depending upon what program is used to import the resulting files.

Export ZEMAX Black Box Data

Purpose:

This feature exports data from one or more surfaces in the Lens Data Editor into a "black box" file. The file can then be later applied to a Black Box surface.

Settings:

Item	Description
First Surface	The first surface in the range of surfaces to export.
Last Surface	The last surface in the range of surfaces to export.
Create and load...	If this option is checked, the current file will be modified to use the newly created Black Box file. Use of this option is highly recommended to confirm the Black Box reproduces the desired analysis data.

Discussion:

Optical designers may under certain circumstances be hesitant to divulge the details of their optical designs. This is especially true when significant time and cost is involved in creating the design, and the designer may not be compensated for the work until the end customer has verified the design will perform as specified. This creates a dilemma, because sending the complete ZMX file to a customer allows the customer to verify the design performance, but also reveals in full detail the design itself. This feature is designed to satisfy the needs of both the designer and the end customer in these circumstances.

A ZEMAX Black Box is an encrypted file that contains the prescription data for a range of surfaces. The file contains most all of the data ZEMAX normally associates with a surface, including the radius, thickness, aperture, coordinate, dispersion, parameter, and extra data. The file is created by the original designer that holds the intellectual property rights. The file is then sent, usually as part of a complete "playback" ZEMAX lens file, to the end customer. The encrypted data appears as a "black box" in the Lens Data Editor. Rays go in, and the rays come out, exactly as if the original surfaces and design data were present in the Lens Data Editor. However, the detailed data is actually hidden inside the Black Box, and the user cannot review or modify any of the data contained within the hidden Black Box section.

To create a Black Box file, highlight the desired range of surfaces in the Lens Data Editor, and then select Tools -> Export Data -> Export ZEMAX Black Box Data. The surface range may also be selected from the Export Black Box Data dialog control directly. The only other option is a check box to "Create and load test file after black box export". Use of this latter option is highly recommended, as it allows confirmation that the Black Box was created properly and yields the identical ray tracing and analysis results. The "Create" option will create a test lens file with the same name as the original file, except the letters "_BB" are appended to the end of the name. If a file with the same name as the test file already exists, it will be automatically overwritten. The black box data itself is created in a file with the extension ZBB, and this file is placed in the folder <data>\BlackBoxes. This newly created file will be referenced on the Black Box Lens surface in the Lens Data Editor, and will replace the original surface data in the editors.

Once the Black Box and test playback file are created, it is recommended that the archive feature be used to make a single archival file that can be sent to the end user. For information on the archival feature, see "Backup To Archive File" on page 65.

For more information on the Black Box surface, see "Black Box Lens" on page 292.

The Black Box data structure is designed so that no ZEMAX feature can read or query the prescription data of the surfaces within the Black Box. Rays enter the Black Box, and the rays are traced using the data defined within, but no feature is able to access the curvature, thickness, glass type, aspheric constants, or other data. The refractive lenses in the Black Box are modeled using the full dispersion data from the original file, and thus can be "played back" at any wavelength. If a playback wavelength exceeds the limits of the dispersion data, an error is issued. If any ray is vignetted or TIRs inside the Black Box, it is treated as a "ray miss" error. Note that some features, such as those that compute the weight of a lens, will not be able to access the Black Box data either and will thus return meaningless data for Black Box Lens surfaces.

The black box is drawn on layout plots as a cylinder with a radial aperture defined by the semi-diameter of the first surface in the selected range of surfaces when the black box was created, or by the largest radial aperture of any surfaces within the selected range, whichever is larger. The length of the black box as drawn is equal to the sum of the lengths of the surfaces within the selected range.

There are some restrictions on the type of data that may be exported to a Black Box file:

- Neither the first nor the last surface of the specified range may be ignored.

- If the first surface follows an even number of mirrors, so must the last surface. If the first surface follows an odd number of mirrors, so must the last surface. Neither the first nor the last surface may be mirrors. The range of surfaces may be intended to be used in a mirror space or not, however, the playback will only work if the file is properly set up after either an even or odd number of mirrors as was the set up for the original file.

- Coordinate Break surfaces and surface tilts and decenters are permitted, however the last surface must be oriented in the same coordinate system as the first. This means the local X, Y, and Z axis of the first and last surfaces must all be parallel, and the X and Y decenters must be zero (both limits are tested to a fine tolerance set internally by ZEMAX).

- Not all surface types are supported. Nonsequential, Birefringent, Paraxial, and surfaces that use external data, such as user defined DLL surfaces and grid sag surfaces, and a few other special surface types are not supported. If a need arises for a surface type that is not supported, contact technical support for assistance.

- The stop cannot be placed in between the first and the last surfaces, although the stop may be either of these surfaces. To export lens data on either side of the stop surface, use two (or more) Black Boxes. Any number of Black Box Lens surfaces may be used in the playback system.

- If any coatings are applied to the surfaces within the Black Box, they must be either standard coatings with no tapers, or encrypted coatings (see "Export Encrypted Coating" on page 249). No ideal or table coatings are supported. The export to Black Box feature will automatically convert all non-encrypted coatings to encrypted coatings. The coating name will be changed by prepending "ZEC_" to the existing name. For example, the coating named BB_AR_5 will be exported as an encrypted coating file with the name ZEC_BB_AR_5. If this file already exists, it will be overwritten without warning. The encrypted coating data is not stored within the ZBB file, the data must be included separately when sent to the end user. By far the easiest and most reliable way to do this is to send the resulting Black Box lens system with the coatings in a ZEMAX Archive (ZAR) file (see "Backup To Archive File" on page 65 and "Restore From Archive File" on page 66).

- Apertures must be placed on all surfaces with power, so that when used the black-box file will correctly clip rays outside of the aperture of the lens. For a tool that converts semi-diameters to apertures see "Convert Semi-Diameters to Floating Apertures" on page 253, or apertures can be set by hand using the Aperture tab of the Surface Properties dialog described in "Surface properties aperture tab" on page 79. User-defined aperture data files (.UDA) are not supported. Note any rays vignetted by the black box surfaces are treated as ray miss errors in the playback system, and this can cause some features to not work properly with certain optical systems, such as telescopes with central obscurations.

- Not all analysis features can reproduce the same results from a black box file as the original file can. For more information see "Features not supported when using Black Box surfaces" on page 292.

- Only data for the current configuration is exported. In a multi-configuration system, each configuration must be exported separately. Upon loading into the playback system, the file name of the ZBB files used can be changed for each configuration using the MCOM operand (see "MCOM" on page 581) in the multi-configuration editor if multiple-configuration control of the black-box system is required.

- Any data controlled by solves are replaced by their current fixed values. No solves are supported inside the Black Box surface. If "Create and load" is selected any solves for surfaces outside the newly created blackbox that reference data now within the blackbox are also turned off.

- Model and Table glasses are not allowed - all dispersion data must be defined in the current glass catalogs.

The dispersion data will be extracted from the glass catalog and placed in the black box file. The exception is for gradient index surfaces as described below.

-Index data for gradient index surfaces is generally defined in parameter data that is encrypted within the black box. Gradient 5 surfaces, when converted to a black box, store the gradient polynomial parameters in encrypted form, but the dispersion data remains as defined in the unencrypted SGRIN.DAT file normally installed with ZEMAX. See "Gradient 5" on page 307 for more information. Gradient 6 surfaces, which use dispersion data defined in the GLC.DAT file, also use the unencrypted data in that file normally installed with ZEMAX.

The encryption is based upon a 256-bit algorithm that provides good, but not unbreakable security. This feature is provided as is and the user should determine if the encryption implemented provides adequate security for the intended application before distributing sensitive data in the encrypted format. Radiant ZEMAX LLC provides no warranty and assumes no liability for the use of this feature.

Miscellaneous

Reverse Elements

Purpose:

Reverses a lens element or group.

Settings:

Item	Description
First Surface	This is the first surface of the lens group to be reversed.
Last Surface	The last surface of the lens group to be reversed.

Discussion:

The feature may not work correctly in all cases if mirrors, coordinate breaks, multi-configuration controlled data, surface tilts and decenters, or non-standard surfaces are included in the range of surfaces. Solves on data affected by the reversal are removed by this feature. Check the reversed system results carefully to verify that the desired result was achieved.

Scale Lens

Purpose:

Scale will scale the entire lens by the specified factor. This is useful for scaling an existing design to a new focal length, for example. Wavelengths are not scaled. The scale lens feature may also be used to change the units from mm to inches, or other combinations of unit types.

Settings:

Item	Description
Scale by factor	If selected, then a scale factor may be entered directly.
Scale by units	If selected, then the lens will be converted by the selected units.

Discussion:

Scaling of data is performed based upon the units of measure for the data. If the scale factor is X, then data measured in lens units of length will be scaled by the factor X. Data measured in units of lens units squared (such as millimeters squared) will be scaled by X squared. Some polynomial coefficients, such as those on the Even Aspheric surface, have units that change from term to term, and ZEMAX accounts for this when scaling the data. Other parameters, such as the conic constant, are dimensionless, and are therefore not scaled.

ZEMAX will generally scale all lens data in the Lens Data Editor, the Extra Data Editor, and the Non-Sequential Components Editor correctly. ZEMAX also will scale almost all of the data in the Merit Function Editor, the Multi-Configuration Editor, and the Tolerance Data Editor correctly as well.

ZEMAX will not attempt to scale some types of data, even though scaling would be appropriate for those types of data. These exceptions are for data items such as the multi-configuration operand PAR3, that modify general parameter data of surfaces or objects, or the parameter data of user-defined DLL surfaces or objects. These data items should be scaled manually, if required.

Data items that are not scaled properly should be reported to ZEMAX technical support.

Make Focal

Purpose:

Make focal length is similar to scale lens, except the desired focal length is typed in directly. The entire lens is then scaled to make the focal length the specified value.

Ghost Focus Generator

Purpose:

Ghost focus analysis. For a description of ghosts, see “Ghost reflections” on page 54.

Settings:

Item	Description
Bounces	Select either single-bounce or double-bounce analysis.
First Surface	The first surface to consider reflections from.
Last Surface	The last surface to consider reflections from.
Save Files	If checked, the files used to compute the ghost ray trace are saved to a file.
Image Surface Only	If checked, only data for the image surface will be shown when computing double bounce ghosts.
Ghost Reflector Coating	The name of the coating, if any, to apply to surfaces that will be changed from refractive to reflective to simulate ghosting. For example, if a 1% reflective coating is applied to all optical surfaces; then enter a coating name of "I.99". The result will be a surface that reflects 1% of the energy; and this will aid in accurate computation of total ghost energy using the polarization transmission analysis, if desired. This example uses an ideal coating, but real coatings may also be used as ghost coatings if the substrates are properly defined in the coating model. This feature is intended for computing detailed transmission data when looking at the generated and saved ghost files using polarization ray tracing and the transmission feature, see “Transmission” on page 218.

Discussion:



This feature may not work correctly for systems which contain coordinate breaks or some special surface types such as the non-sequential components surface. Only the current configuration is considered.

This feature generates lens files which are derived from the current lens prescription data. The generated files are set up to reflect light at a given surface, rather than refract it. The portion of the optical system prior to the new reflective surface is then duplicated so the rays can be traced back through. The purpose of this analysis is to check whether or not rays reflected from any optical surface can form "ghost" images on other components or near the focal plane. These effects are significant in high power laser systems, where focused reflections can damage optics. Ghost images also reduce contrast. Both single and double reflections are supported.



An alternative to this type of ghost analysis is to use the non-sequential capabilities of ZEMAX. For details, see “NON-SEQUENTIAL COMPONENTS” on page 477.

For each ghost system, the paraxial marginal ray height, paraxial ray F/#, and real ray RMS spot radius on axis are listed. Glass surfaces which may have an internal focus are also indicated. For the image surface, when doing double-bounce ghost analysis, the paraxial marginal and chief ray heights, distance from the image surface to the ghost, and the EFL of the ghost system are also provided.

The lens files generated are stored in the file GHfffsss.ZMX, and these lens files can be opened for more analysis. The fff is the number of the first ghost surface, and sss is the number of the second ghost surface. For example GH007002.ZMX is a ghost focus file of the double bounce off surface 7 then 2. For single bounce ghosts, only the fff number is non-zero.

Dummy surfaces (those that do not define an index boundary) and mirror surfaces are ignored as candidate bounce surfaces. The only exception is the image surface, which is considered a potential first bounce surface even if there is no index change on that surface.

Not all ghost reflection configurations can be traced; there are occasionally problems with total internal reflection or rays completely missing surfaces. The GHfffsss.ZMX file may need to be opened and modified to analyze it in detail.

If the first reflected surface is before the stop, then the entrance pupil position is incorrectly calculated as well. This problem is easily overcome (for the purposes of analyzing the ghosts only) by using the following procedure before you begin the ghost focus generation:

- 1) Record the entrance pupil position and diameter.
- 2) Create a dummy surface at surface one.
- 3) Give this new dummy surface a thickness equal to the negative of the entrance pupil position value recorded.
- 4) Make the dummy surface the stop, and use an entrance pupil diameter equal to the entrance pupil diameter recorded.
- 5) For finite conjugates only, increase the object surface thickness by an amount equal to the thickness of the dummy surface.

These steps will give you a real entrance pupil in object space, and the rays will be correctly traced when reflected off surfaces before the stop. Ghost focus analysis can be very complex for moderately complicated systems, and care is required in interpreting the results of this analysis.

Make Double Pass

Purpose:

Creates a double pass system by reflecting the system about any surface, copying all preceding surface data with pick-up solves as appropriate.

Settings:

Item	Description
Reflect At Surface	The surface at which to insert a mirror and create the reflected system.

Discussion:



This feature may not work correctly for systems which contain coordinate breaks or some special surface types such as the non-sequential components surface. Multiple configurations are not supported.

This feature creates a double pass system. The reflect at surface becomes a mirror. All surface data beyond this surface is lost, all data preceding this surface is reflected and copied with pick-up solves to create a double pass system.

Performance Test

Purpose:

System ray trace and update performance test.

Discussion:

Performance test runs a check on the number of ray-surfaces per second and the number of system updates per second the computer hardware/lens combination is capable of.

The ray surfaces per second performance number is measured by tracing a large number of random skew rays through the current optical system, and then dividing the number of rays times the number of surfaces traced through by the elapsed time in seconds.

The system updates per second is calculated by performing many system updates and then dividing the number of system updates performed by the elapsed time in seconds. System updating includes recomputing the pupil positions, field data (such as ray aiming coordinates), lens apertures, index of refraction, solves, and other fundamental checks on the lens that must be performed prior to any ray tracing.

The speed will vary tremendously depending upon the system processor, clock speed, and lens complexity.

Lock/Unlock All Windows

Purpose:

Locks or Unlocks all open analysis windows at once.

Discussion:

This tool locks or unlocks all windows. For a description of the lock function, see "Graphic windows operations" on page 42.

Quick Focus

Purpose:

Adjusts back focal distance for best focus. See also "Quick Adjust" on page 265.

Settings:

Item	Description
Spot Size Radial	Focus to best RMS spot radius image surface.
Spot Size X Only	Focus to best RMS x direction spot size image surface.
Spot Size Y Only	Focus to best RMS y direction spot size image surface.
Wavefront Error	Focus to best RMS wavefront error image surface.
Use Centroid	Reference all calculations to the image centroid rather than the chief ray. This option is slightly slower, but more appropriate for systems dominated by coma.

Discussion:

This feature adjusts the thickness of the surface prior to the image surface. The thickness chosen is that which minimizes the RMS aberration. There are several different RMS computations that may be performed as described in the table above. The "best focus" position depends upon which criterion is selected. The RMS is always computed as a polychromatic average over the field, using the defined fields, wavelengths, and weights.

Quick Adjust

Purpose:

Adjusts any radius or thickness to achieve best transverse or angular ray focus evaluated at any subsequent surface. See also "Quick Focus" on page 265.

Settings:

Item	Description
Adjust Surface	Selects the surface to adjust.
Radius/Thickness	Selects either Radius or Thickness to adjust.
Criterion	Selects the best focus criterion. All criterion use the image centroid for a reference.
Evaluate Surface	Selects the surface at which to evaluate the criterion. Note the criterion is evaluated after refraction into the specified surface for angular criterion.

Discussion:

This feature adjusts any radius or thickness to minimize the RMS aberration at any subsequent surface. The RMS is always computed as a polychromatic average over the field, using the defined fields, wavelengths, and weights. Note RMS angular data is computed after refraction into the specified evaluation surface. This tool can remain open on the desktop while other editing or work is done. Click on the "Adjust" button to recompute the best focus data.

Slider

Purpose:

The slider control is used to interactively adjust any system or surface parameter while viewing any analysis window.

Settings:

Item	Description
The data to be modified is defined by the following three settings:	
Group (item name not shown)	Select either the surface, system, configuration, or NSC data group.
Data (item name not shown)	Select the data to modify. The available selections include radius, curvature, thickness, conic, parameter, and extra data values if surface data are chosen. For system parameters, the selections include the system aperture, field and wavelength data, apodization factor, temperature, and pressure. For configuration data, all multi-configuration operands are listed. For NSC data, object position and parameter data are listed.
Number (item name not shown)	The surface number for the data to be modified if the data is associated with a surface. This field is used for the configuration number if configuration data is selected. If NSC data is selected, this field is used for the object number.
Window	Select either "All" or any specific analysis window to update when the slider is adjusted. See the discussion.
Start/Stop Value	The beginning and ending range of the data corresponding to the extreme limits of the slider control.
Animate	Automatically increments the data over the defined range, and updates the selected windows in a continuous loop. Press "Stop" (the animate button changes to a "Stop" when animation is running) to terminate the animation.
Save	Saves the current data, but leaves the slider control box open.
Exit	Restores the original (or last saved) data for the modified parameter and closes the slider control.

Discussion:

The slider control can be used to adjust any surface or system parameter while monitoring how the changed value affects the data displayed in any open analysis window or in all windows.

When the slider control is activated, a copy is made of the data to be modified. If a new type of data is selected (for example, the surface number or parameter or surface/system setting is altered) the old data is automatically restored, unless the "Save" button has been pressed. The Save button replaces the saved copy of the original data with the newly modified data.

If any or all of the analysis windows require a long time to update, the slider control may act erratically. This is due to the Windows operating system handling of the slider control requiring a very fast update while dragging or scrolling the slider bar control. If this occurs, select a single window for updating rather than "all" or choose analysis options to reduce computation time.

Convert to NSC Group

Purpose:

Converts a range of surfaces in the Lens Data Editor into a group of components in a Non-Sequential Components surface or system.

Settings:

Item	Description
First Surface	This is the first surface of the lens group to be converted to non-sequential components.
Last Surface	The last surface of the lens group to be converted.
Ignore errors and convert as much as possible	If checked, the conversion will not stop even if errors or limitations in the conversion process are detected. This switch is useful because often the majority of a system may be converted properly even though a small portion of the system cannot be converted.
Convert file to non-sequential mode	If checked, the selected range of surfaces will be converted into a "pure" non-sequential file. In this case, the field positions, pupil data, and stop location are not considered.

Discussion:

This feature converts a range of sequential surfaces in the Lens Data Editor into a group of non-sequential components.



Don't assume this feature perfectly converts your sequential surface data; carefully check the conversion results before doing any important analysis.

There are some restrictions on the ability of ZEMAX to convert sequential surfaces to NSC objects:

The surfaces must all be of type Standard, Even Aspheric, Toroidal, Biconic, or Coordinate Break. This restriction will be relaxed as more NSC object types and more conversion logic is added to ZEMAX.

All coordinate break surfaces must be between elements, in "air" and not imbedded within the lenses.

The surface prior to the first surface and the last surface in the selected range must both be made of "air" with unity index.

The stop must precede the first surface. This restriction is removed if the "Convert file to non-sequential mode" box is checked.

ZEMAX can only convert rectangular, circular, and elliptical apertures on standard surface mirrors. Most other combinations of non-circular apertures and surface types may not convert properly. Non-circular apertures on standard surfaces are modeled using user defined aperture (UDA) files that ZEMAX automatically creates. For more information on UDA's see "User defined apertures and obscurations" on page 80.

There are times when the ray tracing after conversion will not yield exactly the same results as before. The possible explanations for this include the following:

Some rays may not trace the same inside a non-sequential group if some of the surfaces have central obscurations or annular apertures. This is especially likely to occur in telescopes with obscured apertures. In the non-sequential representation, the chief ray will either not make it through the group or will take a radically different path.

Rays that are exactly at the edge of a sequential surface may miss the surface, or strike the outer cylinder wall of a lens, in the non-sequential trace. For this reason, ZEMAX adds a very small increase in the diameter of each surface when converting to a NSC group.

For systems with concave surfaces at the ends of the selected range, the edges of the surface may fall outside of the entry and/or exit ports, which will cause ray tracing errors. The solution is to manually edit the entry/exit port locations, or, insert dummy surfaces before and after the selected range prior to conversion.

The feature works by the following algorithm:

First a Non-Sequential Component surface and a dummy exit port surface are inserted.

The exit port is positioned to coincide with the last surface in the defined range.

The range of surfaces is converted into either NSC surface objects (for mirrors) or NSC lens objects (for pairs of surfaces with glass in between).

Coordinate breaks are replaced with null objects with equivalent rotations; and subsequent objects will reference the null objects.

The original surfaces are then deleted.

If the "Convert file to non-sequential mode" box is checked, then the resulting NSC surface is the only surface retained and the file is placed in NSC mode.

The feature may not work correctly in all cases. Some additional dummy surfaces with small spacings may need to be inserted prior to and after the range of lens surfaces to be converted. The feature may also not work for lenses with multi-configuration data that alters the thicknesses or glass types of any surfaces within the range of surfaces to convert.

The most important point:



Don't assume this feature perfectly converts your sequential surface data; carefully check the conversion results before doing any important analysis.

Convert File Formats

Purpose:

This feature is used to convert data files from other programs into the equivalent ZEMAX format.

Settings:

Item	Description
Type	This selects the type of file conversion to perform. The selected type also controls which file extensions ZEMAX looks for when populating the "File" list, and the suggested extension of the converted file.
File	The file name to be converted. Only files in the most recently read folder with the proper extension are listed. However, ZEMAX does not confirm that the files are of the proper format until the conversion is attempted. Use the "Browse" button to select a file in any folder.

Discussion:

The file selected for conversion must be in the proper format. The default output file will have the same name, but a modified extension as appropriate to ZEMAX. This feature may not support all possible variations of file formats. All converted data should be checked carefully for accuracy.



Errors may easily be introduced when converting data files from different programs, especially with regards to orientation, scaling, and polynomial ordering.

For a detailed discussion of each format see below.

***.INT Zernike files**

The proper file format for .INT Zernike files is as follows:

```
! comment (optional)
! comment (optional)
Title line. A title line is required.
<zerntype> <numterms> <datatype>
<value of Zernike 1 term>
<value of Zernike 2 term>
```

There may be any number of comment lines at the beginning of the file. All comment lines must start with the "!" symbol. The first line in the file that does not start with the "!" symbol is the title line. All .INT files must have a title line (which can consist of up to 80 characters). The next line in the file consists of various parameters which define the format of the file. <zerntype> defines the Zernike format (either "ZRN" for Zernike Standard or "ZFR" for Zernike Fringe). At this time, only the Zernike Fringe format (ZFR) is supported. <numterms> specifies the number of Zernike terms that will be provided in the file. <datatype> refers to what type of data is stored in the file (either "SUR" for surface deformation data, "WFR" for wavefront deformation data or "FIL" for intensity apodization filter data). ZEMAX does not differentiate between the three different types of data when importing the data to a Zernike surface.

The remaining lines in the file contain the Zernike coefficients of the corresponding polynomial, either Zernike Standard or Zernike Fringe.

***.INT Grid Files**

The proper file format for .INT Grid files is as follows:

```
! comment (optional)
! comment (optional)
Title line. A title line is required.
GRD <xsize> <ysize> <datatype> WVL <wavelength> SSZ <scalesize> NDA <nodatavalue>
XSC <xscale>
<data value 1> <data value 2> <data value 3> . . . . . <data value n>
<data value n+1> . . .
```

There may be any number of comment lines at the beginning of the file. All comment lines must start with the "!" symbol. The first line in the file that does not start with the "!" symbol is the title line. All .INT files must have a title line (which can consist of up to 80 characters). The next line in the file consists of various parameters which define the format of the file. For Grid Files, this line must start with a "GRD" statement. <xsize> specifies the number of grid points in the x-direction while <ysize> refers to the number of grid points in the y-direction. <datatype> refers to what type of data is stored in the file (either "SUR" for surface deformation data, "WFR" for wavefront deformation data or "FIL" for intensity apodization filter data). ZEMAX does not differentiate between the three different types of data, the interpretation is determined by whether the data is imported to a Grid Sag or Grid Phase surface. <wavelength> is the wavelength, in microns, at which the data was measured. The <wavelength> value must be prefaced a "WVL" statement. <scalesize> specifies the value for the input data values which corresponds to one wave of deformation. The <scalesize> value must be prefaced by an "SSZ" statement. <nodatavalue> specifies the value of the input data that should be interpreted as missing data. Any rays that are incident in these regions of the grid are vignetted. This is an optional parameter. If provided, the "nodatavalue" parameter must be prefaced by an "NDA" statement. <xscale> is the ratio of the size of the grid in the x-direction to the size of the grid in the y-direction. This is an optional parameter. If the value is omitted, a value of 1.0 (which corresponds to a square grid) is used. If provided, the <xscale> parameter must be prefaced by an "XSC" statement.

The remaining lines in the file contain the individual grid point values. There should be a total of <xsize> * <ysize> values specified. Note that this data is in "free format". As such, each line of data does not necessarily correspond to a single row of data on the actual grid. The data is specified from -x to +x starting at the +y edge of the grid.

OptiWave *F3D files

OptiWave F3D files are used to describe either linear polarized or generally polarized electric fields. These files are converted to the ZEMAX Beam File (ZBF) format. For information about ZBF files see “ZEMAX Beam File (ZBF) binary format” on page 640.

The F3D file format consists of 3 lines of header information of the form

```
BCF3DCX 3.0  
Nx Ny  
xmin xmax ymin ymax
```

The identifier BCF3DCX identifies a linear polarized beam file, polarized beams use the header BCF3DCXV. Only file format version 3.0 is currently supported.

The values Nx and Ny are the integer number of points in each direction. ZEMAX ZBF files only support powers of 2, while F3D files may have any integer number of points. If Nx or Ny are not powers of two, the ZBF representation will be zero-padded to make the beam dimensions equal to the nearest larger power of 2. For example, if there are 115 points in the F3D file, the ZBF will contain 128 points, with zero values added to center the original data.

The values for xmin, xmax, ymin, and ymax define the physical size of the beam. All these values must all be in units of micrometers. ZEMAX will scale these values to millimeters, and use millimeters for the exported beam size. ZBF files only contain the x- and y-direction widths of the beam, and not any absolute position.

The F3D format does not include any wavelength data, and all E-Field phase data is referenced to a local plane. ZEMAX will fit a Gaussian beam to the E-Field data to estimate the pilot beam properties required for subsequent propagation in ZEMAX, and store this data in the ZBF file.

Introduction

This chapter provides detailed descriptions of each of the report features ZEMAX supports. The Settings menu option allows the parameters for the calculation to be changed using a dialog box. The dialog box typically has five buttons: OK, Cancel, Save, Load, and Reset. OK causes the window to recalculate and redisplay the data with the currently selected options. Cancel reverts the options to those selected before the dialog box appeared, and does not update the data in the window. Save stores the currently selected options as defaults for future use, and then causes the window to recalculate and redisplay the data. Load loads the most recently saved default options, but does not exit the dialog box. Reset returns the options to the “factory default” values, but does not exit the dialog box. A report window can be updated by double clicking with the left mouse button within the window. The settings dialog box is invoked by clicking with the right mouse button.

Surface Data*Purpose:*

Display surface data.

Settings:

Item	Description
Surface	The surface number to display the data for.

Discussion:

This feature generates a text window which displays surface-specific data. The data includes surface and element powers and focal lengths, edge thickness, index of refraction, and other data for the surface.

If the glass type of the surface is a “model” glass, then ZEMAX will list out the index of refraction at each defined wavelength calculated from the model glass parameters. Also listed is the name of the “Best Fit Glass”, which is the name of the glass in the currently loaded catalog(s) which has the closest index of refraction to the model glass in an RMS sense. Specifically, ZEMAX computes the index error as the sum of the squares of the difference between the model index and the actual glass index using the dispersion formulas. The sum is over the defined wavelengths. The index error is computed for every glass in the current catalogs, and the glass with the lowest RMS index deviation is designated as the best fit glass.

Note that the best fit glass may have a different V number than the model glass, however, this is due to the approximations made in the model glass dispersion. Since index of refraction is the physically significant parameter, only the index is used in making the glass selection. When changing from model glass to a real “fixed” glass, the same algorithm is used to make the selection. For more information on model glasses, see “Using model glasses” on page 601.

System Data*Purpose:*

Display system data.

Settings:

None.

Discussion:

This feature generates a text window which lists many system-related parameters, such as pupil positions and sizes, magnification, F/#, etc.

Prescription Data*Purpose:*

This function generates a list of all surface data, and summarizes the lens system. This is the feature to use for printing the contents of the Lens Data Editor.

Settings:

Item	Description
General Data	Includes F/#, pupil positions, magnification, etc.
Surface Data	Surface type, radii, thickness, glass, semi-diameter, conic.
Surface Detail	Surface Parameter data. All NSC object data is also listed in this section for NSC groups.
Edge Thickness	The X and Y edge thickness of each surface.
Multi-Config Data	A table of multi-configuration operand data.
Solves/Variables	Solve types and data, variables.
Index Data	Index of refraction and TCE data for each wavelength/surface.
Global Vertex	Global coordinates of the vertex and the rotation matrix for each surface. See "Global Coordinate Reference Surface" on page 111 for a discussion of this data.
Center Of Curvature	Global coordinates of the center of curvature of rotationally symmetric surfaces. The radius, position, and orientation of the surface are used to determine the point at the vertex center of curvature. The location of this point is useful for alignment.
Element Volume	Volume in cc, density, and mass. See below.
F/ Numbers	Lists working F/# for each field position and wavelength combination.
Cardinal Points	Lists the locations of focal, principal, anti-principal, nodal, and anti-nodal points.
POP Settings	Lists the surface by surface settings used by the Physical Optics Propagation feature.
Files Used	Lists all the files used by the optical system model, including glass data, macros, extensions, CAD files, DLL files, and others.

Discussion:

This report lists specification data, indices of refraction, global coordinates, element volumes, and more. It is suitable for describing the lens prescription.

Comments on computing element volumes

When ZEMAX computes element volumes for spherical or plane standard surfaces with circular edges, the edge of the surface with the smaller diameter is assumed to be "squared up" to the semi-diameter of the larger surface. The volume are exact for this common special case.

For other elements, with possibly aspheric faces, elliptical, annular, or rectangular apertures and/or decentered apertures, a numerical integration technique is used that yields approximate volumes to an accuracy of about 0.1%. If zero volume is reported than the algorithm is unable to compute the volume to sufficient accuracy. This happens when the aperture types are not the same, or have different sizes or decenters.

When computing the density of elements, the density in grams per cubic centimeter for a catalog glass is retrieved from the glass catalog. For gradient index surfaces, ZEMAX assumes the density is 3.6 g/cc, which may not be a good estimate.

System Check

This report tests the current lens system for common errors and suspicious settings that may cause errors in analysis. The feature does not check for every possible error. Each detected potential problem is reported as either an error or a warning. Errors should be considered serious flaws in the system setup that should be corrected. Warnings are less severe, and indicate conditions that should be reviewed and the relevant settings verified for correctness.

Report Graphics 4/6

Purpose:

This function generates a graphical window that simultaneously displays either 4 or 6 analysis graphics. The primary advantage of this feature is that multiple analysis graphs may be printed on a single page, making a suitable summary for reports, archival documentation, or promotional literature.

Settings:

The report graphics window works somewhat differently from other analysis windows. If the "Settings" option is selected from the window's menu bar, then a dialog box will appear which allows selection of the type of graphic to be displayed in each position within the window. The selected graphs to display can be saved as default selections like any other window. The configuration number for each window may also be selected.

Additionally, the settings may be saved as a custom report for later re-use. The Save As New Report button will save the current settings in a file. Once the settings are saved, the name will appear as a menu option for quick regeneration of the plot. To change the settings for an individual graph within the window, the right mouse button must be used. First, unzoom the display (if the graph has been zoomed), then right mouse click anywhere in the window within the region of the plot whose settings need changing.

Edit/Run ZPL Macros*Purpose:*

Runs ZEMAX Programming Language macros. This option invokes a dialog box which allows editing, viewing, and executing of macros.

Discussion:

See “ZEMAX PROGRAMMING LANGUAGE” on page 651 for details on the ZPL macro language.

This dialog box is most useful for developing and debugging new macros. It also displays a terminate button, which allows the macro execution to be terminated.

Refresh Macro List*Purpose:*

Updates the macro list.

Discussion:

This feature updates the macro list; which may be required if any macros have been added or deleted since ZEMAX was started or the last time the list was refreshed.

Macro Names*Purpose:*

A list of all ZPL macros in the folder for macros is shown (see “Folders” on page 68). By clicking on the macro name, it immediately executes.

Extensions

This feature is only available in the EE edition of ZEMAX.

Purpose:

Runs ZEMAX extensions.

Discussion:

See “ZEMAX EXTENSIONS” on page 749 for details on creating ZEMAX extensions. The extensions supplied with ZEMAX are described in the following table.

EXTENSIONS INCLUDED WITH ZEMAX

Extension Name	Description
ArrayDemo	Creates a table of rays traced for the current lens. This example program includes source code and illustrates the use of array ray tracing commands.
DDE_DEMO	Creates a text listing of ray and system data. This example includes source code and illustrates the basic use of DDE features.
PhasePlot	Creates a graph showing the phase and the inverse of the phase slope for Binary 2 surfaces. The plot indicates the phase over the semi-diameter of the surface.
Pupil Map	Maps an x-y grid of rays from the entrance pupil or a ray-aimed stop surface to another surface in the system.
SagCalculation	Creates an extended listing of sag data for non-radial symmetric surfaces.
TransmissionPlot	Computes transmission vs. Wavelength.

Refresh Extensions List

Purpose:

Updates the list of extensions.

Discussion:

This feature updates the list of extensions that appears in the Extensions menu; which may be required if any extensions have been added or deleted since the last time the list was refreshed. Any new extensions must be placed in the <program>\Extend folder (see “Folders” on page 68).

Extension Names

Purpose:

A list of all ZEMAX extensions in the extensions folder is shown. Clicking on the extension name will execute that extension.

Discussion:

See “ZEMAX EXTENSIONS” on page 749 for details on creating and executing ZEMAX extensions.

Introduction

ZEMAX models many types of optical components. These include conventional spherical glass surfaces, plus aspheres, toroids, cylinders, and others, ZEMAX can also model components such as diffraction gratings, binary optics, Fresnel lenses, holograms, and others.

Because of the large number of surface types ZEMAX supports, a conventional spreadsheet arrangement for the user interface would be difficult to use. For example, there is no need to have a column for diffraction order for a surface that has no diffractive capability. To make the user interface as uncluttered as possible, ZEMAX uses different surface types to indicate what kinds of data are needed to define that type of surface.

Parameter data

A Standard surface can be a plane, spherical, or conic aspheric surface which is followed by a homogeneous material (such as air, mirror, or glass). The only parameters required are a radius (which may be infinity to yield a plane), a thickness, a conic constant (the default zero value indicates a sphere), and the name of the glass type.

Other surface types use these same basic data, as well as other values. For example, the Even Asphere surface uses all the Standard surface column data plus additional values which describe coefficients on a polynomial. These values are called parameters. The most important property of the parameter values to understand is that their meaning changes depending upon the surface type selected. For example, the Even Asphere surface uses parameter 1 to specify the coefficient on the parabolic aspheric term. However, the Paraxial surface uses parameter 1 to specify the focal length of the surface. Both surface types use parameter 1, but for different purposes, since both surface types are never used on the same surface at the same time.

This sharing of data storage simplifies the ZEMAX interface, as well as reduces the total memory required to run the program. After you change a surface from Standard to some other surface type, ZEMAX will automatically change the column headings on the parameter columns to reflect what each parameter does on that surface. As you move the cursor from cell to cell, the column headings will always show you what that cell is used for. If the current surface does not use the parameter column, the column heading will display "Unused". For more information see "Lens Data" on page 75.

Extra data

Some surfaces cannot be described by only a few parameter values. For example, the Binary Optic 1 surface type requires nine parameters as well as up to several hundred additional numbers. This would make for a very large spreadsheet, and so a separate editor is used for the extra data values. However, the concept is the same. The extra data values are shared by each surface type, and change meaning depending upon the surface type selected. The column headings in the Extra Data Editor also change as the cursor is moved from surface to surface.

Summary of surface types

ZEMAX models planes, spheres, and conics; all of these surface types are grouped under the category of Standard surface. By double clicking on the Surface Type column other surface models can be selected. A pop-up dialog will appear which lists all of the available surface types which can be selected. ZEMAX supports many different types of surfaces in addition to the Standard surface.

User defined surfaces

No matter how many surfaces are added to ZEMAX, there always seems to be the need to add another surface type to solve a particular design, modeling, or tolerancing problem. If the surface type needed for a problem is not already included with ZEMAX, it is fairly easy to add new surface types using the User Defined surface described on page "User Defined" on page 331. User defined surfaces are created by writing software that defines the properties about the surface, and then dynamically linking the software into ZEMAX.

If you have a need for a custom surface type, and you do not wish to write the software for it yourself, please contact ZEMAX Technical Support for a quote on developing a custom surface to meet your requirements. We

have considerable experience in developing ray tracing algorithms, and can generally write user defined surface code for a small fee, often on very short notice.

Surface types

The types of surfaces supported by ZEMAX are summarized in the following table. The most common surface type is the Standard surface described on page 325. For types not listed, see “User Defined” on page 331.

SUMMARY OF SURFACE TYPES

Surface Model	Description	Page
ABCD	Uses ABCD matrix to model a “black box”.	280
Alternate Even	Even Asphere surface with alternate solution selected.	281
Alternate Odd	Odd Asphere surface with alternate solution selected.	281
Atmospheric Refraction	Refraction caused by looking through Earth's atmosphere.	281
Biconic	A conic asphere in X and Y independently.	282
Biconic Zernike	A biconic surface with x, y, and Zernike polynomial terms added.	282
Binary Optic 1	Uses 230 term polynomial to define phase.	283
Binary Optic 2	Uses radial polynomial to define phase.	284
Binary Optic 3	Dual zone aspheric and diffractive surface.	286
Birefringent In/Out	For modeling uniaxial crystals; supports tracing of rays in ordinary or extraordinary mode. Birefringent gradient materials with variable crystal axis orientation are also supported.	287
Black Box Lens	A surface that models a range of ZEMAX surfaces, but the defining data is hidden from the user.	292
Conjugate	Defines surface with perfect imaging at two points.	293
Coordinate Break	Permits rotation and decentration.	293
Cubic Spline	Rotationally symmetric fit to 8 points.	294
Cylindrical Fresnel	Polynomial cylindrical Fresnel on a polynomial cylindrical surface.	295
Diffraction Grating	Ruled grating on Standard surface.	296
Elliptical Grating 1	Elliptical grating with aspheric terms and polynomial grooves.	296
Elliptical Grating 2	Elliptical grating with aspheric terms and grooves formed by tilted planes.	297
Even Asphere	Standard surface plus polynomial asphere terms.	298
Extended Asphere	Uses radial polynomial to define sag.	298
Extended Cubic Spline	Rotationally symmetric fit of up to 250 points.	299
Extended Fresnel	Polynomial Fresnel on a polynomial surface.	299
Extended Odd Asphere	Uses odd terms of radial powers.	301
Extended Polynomial	Uses 230 term polynomial expansion to define sag.	301
Extended Toroidal Grating	An aspheric toroidal grating with extended polynomial terms	302

Surface Model	Description	Page
Filter	A surface to filter or apodize the beam with an arbitrary function.	303
Fresnel	Plane surface with refractive power.	303
Generalized Fresnel	XY polynomial Fresnel on an aspheric substrate.	304
Gradient 1	Radial gradient index material surface.	305
Gradient 2	Radial gradient index material surface.	305
Gradient 3	Radial and axial gradient index material surface.	306
Gradient 4	X, Y, and Z gradient index material surface.	306
Gradient 5	Radial and axial gradient index material surface with dispersion model.	307
Gradient 6	Radial gradient index material surface with Gradient Lens Corp. dispersion model.	308
Gradient 7	Spherical gradient profile.	309
Gradium™	Axial gradient index material surface with dispersion model.	309
Gradient 9	Radial gradient index material surface with NSG SELFOC lens dispersion model.	311
Gradient 10	Y gradient index material surface with dispersion model.	312
Grid Phase	A phase surface described by a grid of points.	313
Grid Sag	A surface shape described by a grid of points.	313
Hologram 1	Two-point optically fabricated hologram.	316
Hologram 2	Two-point optically fabricated hologram.	317
Irregular	A Standard surface with decenter, tilt and other deformations.	317
Jones Matrix	General Jones Matrix for modifying polarization state.	318
Lenslet Array	An array of lenslets.	318
Non-Sequential Components	For tracing rays non-sequentially through a collection of 3D surfaces and objects.	318
Odd Asphere	Standard surface plus polynomial asphere terms.	318
Odd Cosine	Odd Asphere surface plus cosine polynomial terms.	318
Optically Fabricated Hologram	Optically fabricated hologram with arbitrary construction optics and elliptical substrate.	319
Paraxial	Thin lens surface, has ideal behavior.	321
Paraxial XY	Thin lens with separate specification in X,Y.	322
Periodic	Cosine shaped surface.	322
Polynomial	Polynomial expansion in x and y.	323
Radial Grating	A diffraction grating with radial phase profile	323
Radial NURBS	Uses NURBS curve to define a rotationally symmetric surface.	324

Surface Model	Description	Page
Retro Reflect	Retro reflects rays back along incident path.	325
Standard	Includes planes, spheres, and conics.	325
Superconic	Superconic asphere with fast convergence.	326
Tilted	Defines a tilted surface without changing coordinate systems.	327
Toroidal	Models conic and aspheric toroids and cylinders, with added Zernike terms.	328
Toroidal Grating	Ruled grating on a conic toroid.	329
Toroidal Hologram	Toroidal substrate with two point optically fabricated hologram.	329
Toroidal NURBS	Uses NURBS curve to define a toroidally symmetric surface.	330
User Defined	A general surface which uses an arbitrary user defined function to describe the refractive, reflective, diffractive, transmissive, or gradient properties of the surface.	331
Variable Line Space Grating	Variable line space grating surface.	339
Zernike Fringe Phase	Uses 37 Zernike Fringe polynomials to define phase.	340
Zernike Fringe Sag	Uses 37 Zernike Fringe polynomials to define sag.	341
Zernike Standard Phase	Uses up to 231 Zernike Standard polynomials to define surface phase.	341
Zernike Standard Sag	Uses up to 231 Zernike Standard polynomials to define surface sag.	343
Zone Plate	Fresnel Zone Plate model using annular rings of varying depth.	343

ABCD

The ABCD surface provides a powerful method for modeling "black box" optical systems. If you have a lens (or a complete optical system) which is only a subsection of what you wish to model, and you do not have prescription data for the individual components, you can still model the behavior to first order.

The ABCD surface accepts eight parameters: Ax, Bx, Cx, Dx, Ay, By, Cy, and Dy. These are used to form two-by-two matrices (one for the x-direction and one for the y) which are used to alter the ray as it crosses the surface. The exit ray is related to the incident ray by

$$\begin{bmatrix} x' \\ \omega'_x \end{bmatrix} = \begin{bmatrix} A_x & B_x \\ C_x & D_x \end{bmatrix} \begin{bmatrix} x \\ \omega \end{bmatrix},$$

with a similar expression for the y components. Because there is no reliable way to compute the phase through an ABCD surface, any calculation that requires OPD data, such as OPD fans, MTF, and Zernike coefficients, will not be supported if an ABCD surface is present in the lens description.

PARAMETER DEFINITIONS FOR ABCD SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
Ax	Bx	Cx	Dx	Ay	By	Cy	Dy

Alternate Even

There are two solutions to the ray-surface intercept equations used in tracing a ray to the next optical surface, if that surface is of the form of the standard or even aspheric surface types. ZEMAX selects the correct solution in the vast majority of cases. However, in certain systems, so-called "strange" rays are actually intended to intersect the next surface at the other, "alternate" solution. Strange rays are most likely to occur after a grazing incidence reflection in which the ray is still traveling in the same direction (the Z component of the ray vector does not change sign). The alternate even surface model is identical to the even asphere surface model, except the alternate solution is used. ZEMAX may not be able to correctly compute the optical path difference when the alternate even surface is being used.

Alternate Odd

There are two solutions to the ray-surface intercept equations used in tracing a ray to the next optical surface, if that surface is of the form of the standard or even aspheric surface types. ZEMAX selects the correct solution in the vast majority of cases. However, in certain systems, so-called "strange" rays are actually intended to intersect the next surface at the other, "alternate" solution. Strange rays are most likely to occur after a grazing incidence reflection in which the ray is still traveling in the same direction (the Z component of the ray vector does not change sign). The alternate odd surface model is identical to the odd asphere surface model, except the alternate solution is used. ZEMAX may not be able to correctly compute the optical path difference when the alternate odd surface is being used.

Atmospheric Refraction

This surface is used to simulate the effects of refraction through the Earth's atmosphere when viewing a star or point source. The atmosphere has a small but non-zero dispersion which introduces a tilt term to the incoming wavefront which depends upon wavelength. ZEMAX uses a model based upon those described in the following publications:

P. K. Seidelmann, Ed., "Refraction - Numerical Integration", Section 3.281, Explanatory Supplement to the Astronomical Almanac, pp. 141-143, University Science Books, Mill Valley, 1992.

C. Y. Hohenkerk and A. T. Sinclair, NAO Technical Note 63, Royal Greenwich Observatory Science and Engineering Research Council, 1985.

Six parameters are supplied to the model: The observed zenith angle of the source in degrees, the height of the observer above sea level in meters, the ambient temperature at the observer in Kelvin, the atmospheric pressure at the observer in millibars, the relative humidity (a number between 0.0 and 1.0), and the latitude of the observer in degrees.

ZEMAX computes the atmospheric refraction angle in radians for all defined wavelengths, then subtracts from all wavelengths the amount of refraction at the primary wavelength; thereby using the primary wavelength as a reference. To disable the referencing to the primary wavelength, set the "Absolute" flag to 1. Atmospheric refraction manifests itself as a small tilt in the OPD fan plot, or a slight chief ray offset similar to lateral color in the ray fan plot. All refraction is assumed to occur in the y direction only.

This surface models color separation. To model atmospheric turbulence, use the Grid Phase surface with an externally provided data file. For information on the Grid Phase surface, see "Grid Phase" on page 313.

PARAMETER DEFINITIONS FOR ATMOSPHERIC REFRACTION SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7
Zenith	Height	Temperature	Pressure	Humidity	Latitude	Absolute

Biconic

The biconic surface is similar to a toroidal surface, except the conic constant and base radius may be different in the X and Y directions. The biconic surface allows specification of Rx, Ry, Kx, and Ky directly. The sag of a biconic is given by:

$$z = \frac{c_x x^2 + c_y y^2}{1 + \sqrt{1 - (1 + k_x)c_x^2 x^2 - (1 + k_y)c_y^2 y^2}},$$

where

$$c_x = \frac{1}{R_x}, c_y = \frac{1}{R_y}.$$

The radius in the x direction is set in the parameter 1 column. If set to zero, the x radius is interpreted to be infinity.

PARAMETER DEFINITIONS FOR BICONIC SURFACES

Parameter 1	Parameter 2
R_x	k_x

Biconic Zernike

The biconic Zernike surface is similar to a biconic surface, with the added capability to add X, Y, and Zernike polynomial deformations. The sag of a biconic Zernike is given by:

$$z = \frac{c_x x^2 + c_y y^2}{1 + \sqrt{1 - (1 + k_x)c_x^2 x^2 - (1 + k_y)c_y^2 y^2}} + \sum_{i=1}^{16} \alpha_i x^i + \sum_{i=1}^{16} \beta_i y^i + \sum_{i=1}^N A_i Z_i(\rho, \phi),$$

where

$$c_x = \frac{1}{R_x}, c_y = \frac{1}{R_y},$$

and the Z_i terms are the Zernike standard terms as described in "Zernike Standard Sag" on page 343. Parameter 0 is the extrapolate flag. If the extrapolate flag is set to 0, the Zernike terms are ignored outside of the normalization radius. If the extrapolate flag is set to 1, then the Zernike terms are considered no matter where the ray lands on the surface; even if the ray lands beyond the normalization radius. The radius and conic in the x direction are set in the parameter 1 and 2 columns. If the radius is set to zero, the x radius is interpreted to be infinity. Parameter 3 is the maximum number of Zernike terms, this value may be between 0 and 210, inclusive. Parameter 4 is the normalization radius used in the Zernike expansion.

PARAMETER DEFINITIONS FOR BICONIC ZERNIKE SURFACES

Parameter 0	Parameter 1	Parameter 2	Parameter 3	Parameter 4
Extrapolate	R_x	k_x	N	Normalization Radius

EXTRA DATA DEFINITIONS FOR BICONIC ZERNIKE SURFACES

Extra Data Number	Description
1-16	Coefficients on powers of x
17-32	Coefficients on powers of y
33-242	Coefficient on the Zernike Standard polynomials

Binary Optic 1

Binary optics, also called kinoforms, are similar to holograms and diffraction gratings in that small grooves or lines across the optical surface impart a change in phase of the wavefront passing through the surface. ZEMAX does not model the wavelength-scale grooves directly. Instead, ZEMAX uses the phase advance or delay represented by the surface locally to change the direction of propagation of the ray. Other effects, such as efficiency or multiple order diffraction are ignored. Binary surfaces can have zero thickness, with no index change across the surface, or may have different media on either side of the surface. In this case, the refraction due to the material change will be accounted for as well as the diffraction effects. Full exact polychromatic ray and OPD tracing is supported.

The binary optic 1 surface is similar to the extended polynomial surface, except the polynomial terms represent the variation in phase (rather than surface height) across the optic surface. The coefficients therefore have units of radians rather than lens units. The shape of the binary optic 1 surface is identical to the even asphere surface type; planes, spheres, conics, and polynomial aspheres up to order 16 are supported. The sag of the surface is given by the following expression:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \alpha_1r^2 + \alpha_2r^4 + \alpha_3r^6 + \alpha_4r^8 + \alpha_5r^{10} + \alpha_6r^{12} + \alpha_7r^{14} + \alpha_8r^{16},$$

where the terms are identical to those in the even asphere surface model. See "Even Asphere" on page 298 for a complete discussion. The binary optic 1 surface adds phase to the ray according to the following polynomial expansion:

$$\Phi = M \sum_{i=1}^N A_i E_i(x, y)$$

where N is the number of polynomial coefficients in the series, A_i is the coefficient on the i^{th} extended polynomial term, and M is the diffraction order. The polynomials are a power series in x and y , as described in "Extended Polynomial" on page 301. The coefficients A_i all have units of radians (2π radians is one wave). If Parameter 9 is any value other than zero, then the absolute value of the x and y coordinates are used for computing the phase in the above expression. This "Absolute" option yields a different, generally discontinuous phase profile.

PARAMETER DEFINITIONS FOR BINARY OPTIC 1 SURFACES

Param 0	Param 1-8	Param 9
M, the diffraction order	$\alpha_1 - \alpha_8$	Absolute? (see discussion)

EXTRA DATA DEFINITIONS FOR BINARY OPTIC 1 SURFACES

Extra Data Number	Description
1	Maximum term number.
2	Normalization radius. All ray-intercept points are divided by this number to determine the x and y coordinates for polynomial evaluation.
3-232	Polynomial terms.
233 and above	Not used.

The "Maximum term number" is used to specify the maximum polynomial term to be used in calculating the surface phase. This number is provided to speed the ray tracing calculation, as terms beyond this number are ignored. For example, if the last term in the series you wish to use is xy^3 , which is term number 13, then specify "13" in the maximum term number column. Note that the term number is 13, because it is the 13th term in the polynomial expansion, and NOT 15, which is the "extra data number", which is the position in the extra data list of parameters. The "extra data number" is always two greater than the term number!



The "extra data number" is always two greater than the term number!

The normalization radius scales the X and Y intercept coordinates of the ray so that the polynomial terms are all dimensionless and the coefficients are all in radians. For modeling phase surfaces that are not well described by polynomials, see "Grid Phase" on page 313.

Binary optic coefficients sign conventions

It is important to understand the binary optic surface phase coefficients sign convention. Imagine a collimated beam being focused by an infinitesimal thin binary optic surface. If the collimated beam becomes a focused, converging beam on the other side of the binary optic, then the path length of a marginal ray is longer than the path length of the axis ray. Therefore, the binary optic must add a negative path length to the marginal ray. Using this convention, a positive power binary optic has a negative quadratic phase coefficient. Although the sign convention chosen by ZEMAX is arbitrary, it may be important to know what the convention is as far as fabrication decisions are concerned. It is always a good idea to test a few cases to verify the sign convention being used by the software before committing a design to fabrication.

Binary Optic 2

Binary optics are similar to holograms and diffraction gratings in that small grooves or lines across the optical surface impart a change in phase of the wavefront passing through the surface. ZEMAX does not model the wavelength-scale grooves directly. Instead, ZEMAX uses the phase advance or delay represented by the surface locally to change the direction of propagation of the ray. Other effects, such as scattering, efficiency, or multiple order diffraction are ignored. Binary surfaces can have zero thickness, with no index change across the surface, or may have different media on either side of the surface. In this case, the refraction due to the material change will be accounted for as well as the diffraction effects. Full exact polychromatic ray and OPD tracing is supported.

The binary optic 2 surface is similar to the extended asphere surface, except there are additional polynomial terms which represent the variation in phase (rather than surface height) across the optic surface. The coefficients therefore have units of radians rather than lens units. The shape of the binary optic 2 surface is identical to the even asphere surface type; planes, spheres, conics, and polynomial aspheres up to order 16 are supported. The sag of the surface is given by the following expression:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \alpha_1r^2 + \alpha_2r^4 + \alpha_3r^6 + \alpha_4r^8 + \alpha_5r^{10} + \alpha_6r^{12} + \alpha_7r^{14} + \alpha_8r^{16},$$

where the terms are identical to those in the Even Asphere Surface model. See “Even Asphere” on page 298 for a complete discussion. The Binary Optic 2 surface adds phase to the ray according to the following polynomial expansion:

$$\Phi = M \sum_{i=1}^N A_i \rho^{2i}$$

where N is the number of polynomial coefficients in the series, A_i is the coefficient on the $2i^{th}$ power of ρ , which is the normalized radial aperture coordinate, and M is the diffraction order. The first extra data term is N , the number of terms, which can be zero to exclude all binary effects, or any integer up to 240 (or the ρ^{480} term). Extra data value number 2 is the normalization radius. Extra data values 3 through 242 are the coefficients.

PARAMETER DEFINITIONS FOR BINARY OPTIC 2 SURFACES

Param 0	Param 1	Param 2	Param 3	Param 4	Param 5	Param 6	Param 7	Param 8
M	α_1	α_2	α_3	α_4	α_5	α_6	α_7	α_8

EXTRA DATA DEFINITIONS FOR BINARY OPTIC 2 SURFACES

Extra Data Number	Description
1	Maximum term number.
2	Normalization radius. All ray-intercept points are divided by this number to determine the x and y coordinates for polynomial evaluation.
3	Coefficient in radians on ρ^2
4	Coefficient in radians on ρ^4
n+2	Coefficient in radians on ρ^{2n}

The "Maximum term number" is used to specify the maximum polynomial term to be used in calculating the surface phase. This number is provided to speed the ray tracing calculation, as terms beyond this number are ignored. For example, if the last term in the series you wish to use is ρ^{14} , which is term number 7, then specify "7" in the maximum term number column. Note that the term number is 7, because it is the 7th term in the polynomial expansion, and NOT 9, which is the "extra data number", which is the position in the extra data list of parameters. The "extra data number" is always two greater than the term number!



The "extra data number" is always two greater than the term number!

The normalization radius scales the X and Y intercept coordinates of the ray so that the polynomial terms are all dimensionless and the coefficients are all in radians. For modeling phase surfaces that are not well described by polynomials, see “Grid Phase” on page 313.

Binary optic coefficients sign conventions

See “Binary Optic 1” on page 283 for a discussion of sign conventions.

Binary Optic 3

The Binary Optic 3 surface is very similar to the Binary Optic 2 surface. The key difference is that the Binary Optic 3 supports two concentric radial zones, with independent radius, conic, and polynomial aspheric deformation and diffractive phase data for each zone. The surface is divided into two zones by two radial coordinates, A_1 , and A_2 . The inner radial zone extends from the center of the surface to the radial coordinate A_1 . The outer radial zone extends from A_1 outward. The radial coordinate A_2 is used for normalizing the phase coefficients in the outer zone, even though the surface may extend past the coordinate A_2 . The outer zone is offset from the inner zone to make the surface sag continuous across the zone boundary, unless the optional "break" parameter is set to 1. ZEMAX requires that $0 < A_1 < A_2$.

The sag of the surface in the inner zone is given by the following expression:

$$z_1 = \frac{c_1 r^2}{1 + \sqrt{1 - (1 + k_1)c_1^2 r^2}} + \sum_{i=1}^N \alpha_{1i} r^{2i}, \text{ for } r \leq A_1,$$

where N is the number of aspheric terms which may be set by the user as described below. The value for the inner zone curvature, c_1 , is the reciprocal of the radius of curvature specified in the Lens Data Editor. The inner zone conic constant is also set in the Lens Data Editor, in the usual conic column. A similar expression with different coefficients and an offset value is used to define the sag of the surface in the outer zone:

$$z_2 = z_o + \frac{c_2 r^2}{1 + \sqrt{1 - (1 + k_2)c_2^2 r^2}} + \sum_{i=1}^N \alpha_{2i} r^{2i}, \text{ for } r > A_1,$$

where the term z_o is chosen to make the surface continuous across the boundary between the inner and outer zones at the radial coordinate A_1 , or $z_o = z_1(A_1) - z_2(A_1)$ (the value of z_o is temporarily set to zero while evaluating z_2 for this calculation). The outer zone radius of curvature (from which the value c_2 is computed) and the outer zone conic are set in the parameter data as described below. For a flat outer zone radius, use zero.

Both the inner and outer zones have a diffractive phase profile, with independent coefficients. The phase of the inner zone is given by:

$$\Phi_1 = M_1 \sum_{i=1}^N \beta_{1i} \rho_1^{2i}, \text{ and } \rho_1 = \frac{r}{A_1},$$

where N is the number of polynomial coefficients in the series, β_{1i} is the coefficient on the $2i^{\text{th}}$ power of ρ_1 , which is the normalized radial aperture coordinate, and M_1 is the diffraction order. A similar expression describes the phase for the outer zone:

$$\Phi_2 = \delta_o + M_2 \sum_{i=1}^N \beta_{2i} \rho_2^{2i}, \text{ where } \rho_2 = \frac{r}{A_2}.$$

The phase offset serves a similar purpose to the sag offset, and is defined by $\delta_o = \Phi_1(1) - \Phi_2(A_1/A_2)$ (the value of δ_o is temporarily set to zero while evaluating Φ_2 for this calculation). Note that the normalization radius A_2 is used only for defining the radial aperture to normalize the phase coefficients for the outer zone. The outer zone of the surface, and the associated phase profile, may extend beyond this value.

The dual zone nature of this surface creates a complication when computing the phase of the surface as the zone boundary is crossed. The phase offset value δ_o makes certain the phase is continuous across the zone boundary. This is desirable for design and analysis purposes, because phase jumps of hundreds of waves make interpretation and analysis difficult. However, the phase offset is artificial, and this must be accounted for in the actual design. For optimum imaging properties, the outer zone should be in phase with the inner zone. This condition is met when the inner and outer zones differ in phase by an integral number of wavelengths at the boundary, or more to the point, if $\delta_o = J2\pi$, where J is some arbitrary integer. This condition can usually be met by a small change in A_1 , as long as there is some difference in the slope of the phase on either side of the boundary. To make this boundary condition simple to meet, ZEMAX computes $\sin\delta_o$ and places this value in parameter 7. The merit function boundary operand PMVA can then be used to target this value to be zero. Note this value is computed from the phase data, and should not be user defined or made variable.

The difference in sag between the inner and outer zones can be accounted for two ways. If the "Break?" parameter is set to 0, then the outer one is shifted along the local z axis so that the surface is continuous at the zone boundary. If "Break?" is set to 1, then the outer sag is defined by the sag equation with no offset, and in general a discontinuity in the sag will result. This discontinuity will cause a large discontinuity in the OPD computations that ZEMAX cannot subtract out.

PARAMETER DEFINITIONS FOR BINARY OPTIC 3 SURFACES

Param 1	Param 2	Param 3	Param 4	Param 5	Param 6	Param 7	Param 8
R_2	k_2	A_1	A_2	M_1	M_2	$\sin\delta_o$	Break?

EXTRA DATA DEFINITIONS FOR BINARY OPTIC 3 SURFACES

Extra Data Number	Description
1	Maximum term number, N.
2, 3, 4, 5	The coefficients α_{11} , β_{11} , α_{21} , and β_{21} if $N \geq 1$.
6, 7, 8, 9	The coefficients α_{12} , β_{12} , α_{22} , and β_{22} if $N \geq 2$.
10, 11, 12, 13, etc.	The pattern continues in groups of 4 coefficients for N groups.

Binary optic coefficients sign conventions

See "Binary Optic 1" on page 283 for a discussion of sign conventions.

Birefringent In and Birefringent Out



This feature is only available in the EE edition of ZEMAX.

This pair of surfaces models uniaxial crystals, such as calcite. These types of crystals are described by a crystal axis; which defines the axis of symmetry for the material, and two dispersion curves; one defining the "ordinary" index and the other the "extraordinary" index. Such materials refract rays differently depending upon the polarization state of the ray and the angle the ray makes with respect to the crystal axis. These different types of refraction yield two different possible refraction angles for any specific ray, and so they are called birefringent

or double refraction materials. For details on birefringent materials, see Saleh and Teich, *Fundamentals of Photonics*, Wiley Interscience. Only a brief description of these materials and their ray tracing properties is given here. Note that when rays propagate close to parallel to the crystal axis, some crystal materials, such as quartz, may exhibit an effect called optical activity that ZEMAX does not model. Optical activity may be accounted for using the Jones Matrix surface, see "Jones Matrix" on page 318. Birefringent In/Out surfaces may be planes, spheres, conic aspheres, or toroidal in shape.

ZEMAX can also model uniaxial crystals whose crystal axis orientation varies within the material. Several different polynomial models for the form of the variation are supported. For more information see "Defining the crystal axis for inhomogeneous materials" on page 289.

Birefringent materials refract rays according to Snell's law, but the effective index of refraction in the media depends upon the input polarization state and the angle the refracted ray makes with respect to the crystal axis. Ordinary rays are refracted according to

$$n \sin \theta = n_o \sin \theta',$$

where the subscript "o" indicates the ordinary index. This is just Snell's law, of course. For extraordinary rays, the refraction law is

$$n \sin \theta = n(\theta_w) \sin \theta'.$$

This is also Snell's law, but the effective index of refraction in the birefringent material is a function of the angle θ_w , which is the angle between the crystal axis vector \hat{a} and the refracted wave vector \hat{k} . Additionally, the ray vector \hat{s} , which is the vector pointing in the direction of energy propagation (also called the Poynting vector), does not follow the wave vector \hat{k} , but makes a small angle with respect to \hat{k} . In isotropic media \hat{k} and \hat{s} are the same, so for most optical design we keep track of \hat{k} ; here we must consider the ray and the wave vector being different. The angle θ_w is defined by

$$\cos \theta_w = \hat{k} \cdot \hat{a}.$$

The effective index of refraction is defined by

$$\left(\frac{1}{n(\theta_w)} \right)^2 = \left(\frac{\cos \theta_w}{n_o} \right)^2 + \left(\frac{\sin \theta_w}{n_e} \right)^2,$$

where n_o is the ordinary and n_e is the extraordinary index of refraction.

The angle α between \hat{k} and \hat{s} is defined by

$$\cos \alpha = \hat{k} \cdot \hat{s}, \text{ where}$$

$$\tan \alpha = \frac{(n_e^2 - n_o^2) \tan \theta_w}{n_e^2 + (n_o \tan \theta_w)^2},$$

and the vectors \hat{k} and \hat{s} are both coplanar with the crystal axis vector \hat{a} . The wave vector \hat{k} points along the normal to the wavefront, while \hat{s} points along the direction of energy propagation. For ray tracing purposes, ZEMAX uses the components of \hat{s} as the ray direction cosines.

Because the birefringence effects both the refraction into the media, as well as the refraction out of the media, two surfaces are required to model a birefringent element: the "birefringent in" and the "birefringent out" surfaces. Each "in" must be followed by an "out" surface; if this is violated ZEMAX will issue an error message and not trace any rays. The only exception is that any number of coordinate break surfaces (or surface tilts and decenters) may be placed between the "in" and "out" surfaces, unless the Gradient Mode is not zero. The ray is propagated to the next surface using the \hat{s} vector components; however the refraction out of the birefringent media is governed again by the \hat{k} vector components.

Defining the substrate shape

Birefringent In/Out surfaces may be planes, spheres, conic aspheres, or toroids. The default substrate shape is the same as a Standard surface (see "Standard" on page 325) defined by the radius of curvature and a conic constant. This standard substrate is called shape 0. The other substrate shape currently supported is shape 1, which is a surface similar to the Toroidal surface (see "Toroidal" on page 328) except the aspheric coefficients are limited to 10th order and no Zernike terms are supported. The shape number is defined by parameter 6.

Defining the index

To define the ordinary index, ZEMAX uses the glass catalog and glass name in the usual way. For example, to use Calcite, define the surface to be a birefringent "in" type, and enter Calcite for the glass name. ZEMAX uses this name to compute the ordinary index at any defined wavelength.

To compute the extraordinary index, ZEMAX appends "-E" to the glass name, and then looks for this material in the glass catalog. For the case of Calcite, ZEMAX searches for "Calcite-E" in the glass catalog, and if found, uses that for the extraordinary index. If not found, an error message is issued. This technique allows definition of any material with any ordinary and any extraordinary index to be defined and used.

Most ZEMAX calculations, such as EFL, EPD, F/#, etc. use the ordinary index of refraction. The extraordinary index is only used to trace extraordinary rays. The Prescription report will list both the ordinary and the extraordinary index, and this data should be carefully checked before using this feature for any critical analysis.

Defining the crystal axis for homogeneous materials

For homogeneous materials, the crystal axis orientation is constant throughout the media. The crystal orientation is defined using parameter columns 2, 3, and 4 for the x-, y-, and z- direction cosines of the axis. For example, to define the crystal axis along the x axis, the vector values are (1, 0, 0). For the axis aligned to the surface vertex normal, which is along the z axis, the vector is (0, 0, 1). Note the coordinates are local to the surface and ZEMAX internally converts to unit vector components.

Defining the crystal axis for inhomogeneous materials

Birefringent materials whose crystal axis orientation varies as a function of position are supported. For a discussion of the method used to model these types of materials, see C. Jenkins, R. Bingham, K. Moore, and G. Love, "Ray equation for a spatially variable uniaxial crystal and its use in the optical design of liquid-crystal lenses", J. Opt. Soc. Am. A, Vol. 24, No. 7, pp2089-2096 (2007). The variation in the crystal axis orientation is defined using one of several polynomials. The selection of the polynomial used is called the Gradient Mode, and the Gradient Mode is defined by Extra Data Value 1. The supported Gradient Modes are defined in the following table.

GRADIENT MODE POLYNOMIALS

Gradient Mode	Definition of \hat{a}
0	Material is homogeneous, \hat{a} is defined by parameters 2-4 on the Lens Data Editor.

1	$a_x = \sum_{i=0}^6 c_{xi}r^i, a_y = \sum_{i=0}^6 c_{yi}r^i, a_z = \sqrt{1 - a_x^2 - a_y^2}.$
2	$a_z = 1 + \sum_{i=1}^8 c_{zi}r^i, a_x = \frac{x}{r}\sqrt{1 - a_z^2}, a_y = \frac{y}{r}\sqrt{1 - a_z^2}.$
3	$\theta = \sum_{i=0}^8 p_i r^i, a_z = \sin\theta, a_x = \frac{x}{r}\cos\theta, a_y = \frac{y}{r}\cos\theta. \text{ If } \theta = 0 \text{ then } \hat{a} = \{0, 1, 0\}.$ The angle θ is measured in radians.
4	$a_x = 0, a_y = \sum_{i=0}^8 a_i y^i, a_z = \sqrt{1 - a_y^2}.$
5	$\theta = \sum_{i=0}^8 p_i r^i, a_z = \cos\theta, a_x = \frac{x}{r}\sin\theta, a_y = \frac{y}{r}\sin\theta.$ The angle θ is measured in radians.
6	$\sin\theta = \frac{\sqrt{O^2(E^2 - (E - A\phi)^2)}}{B(E - A\phi)^2},$ where ϕ is the polar angle and $a_x = 0, a_y = \sin\theta, a_z = \cos\theta.$ The angle ϕ is measured in radians. This formula models a radial vortex phase profile, where O and E are nominally the ordinary and extraordinary indices and A and B are used to define the index gradient around the vortex.
7	$a_x = \sum_{i=0}^{16} c_{xi}r^i, a_z, a_y = \sqrt{1 - a_x^2 - a_z^2}.$
-	Other equations may be added in the future based upon user suggestions.

After computing the component values of the \hat{a} vector using the equations in the table, the vector is normalized to be a unit vector. Care should be taken to ensure the polynomial coefficients define valid, real values for the crystal axis vector, otherwise the behavior of the crystal is undefined. All polynomial coefficients are defined on the Extra Data Editor in the columns following the Gradient Mode value.

When using a non-zero Gradient Mode, rays will generally take curved paths within the crystal. The reason is the continuously varying angle between the \hat{k} and \hat{a} vectors creates a varying effective index of refraction. Ray tracing must be done in a piece-wise fashion, and the accuracy of the ray trace depends upon the maximum size of the step allowed. The maximum step size allowed is called "Delta T" and is defined by Extra Data Value 2. The situation is very similar to that of gradient index glass. For more information see "Discussion on maximum step size for GRIN surfaces" on page 305.

Determining which ray is traced

ZEMAX will trace the ordinary ray if the mode is set to 0 or 2. If the mode is set to 1 or 3, the extraordinary ray will be traced. The mode is parameter 1 in the lens data editor. ZEMAX cannot trace both rays at once, however it is easy to create a multi-configuration lens with the mode set to 0 in configuration 1, and to 1 in configuration 2; this allows inspection of both possible paths as well as simultaneous optimization and layouts of the traced rays. The difference between modes 0 and 2 for ordinary rays and modes 1 and 3 for extraordinary rays is described in the next paragraph.

Accounting for phase rotation

To account for phase rotation when performing polarization ray tracing, there are two limiting cases to consider. Conceptually, the beam splits into an ordinary beam and an extraordinary beam with a small angle between the two beams. If the beams then propagate a large distance, the ordinary and extraordinary beams will separate and become two distinct beams. If however the propagation distance is short, the ordinary and extraordinary beams coherently interfere and recombine, resulting in a single beam with a (usually) rotated polarization vector. Conceptually this is like tracing both the ordinary and extraordinary rays, then coherently recombining them after propagation through the birefringent media. Which model is used to determine the transmission and polarization properties may be selected by the mode.

If the mode is 0 or 1, then only the ordinary or extraordinary part, respectively, is traced. The fraction of ray energy in the part of the ray that is not traced is discarded. No polarization rotation will be modeled using mode 0 or 1. For transmission computations, both rays need to be traced separately and the total energy computed. It is difficult to compute the effects of phase rotation using mode 0 or 1.

If the mode flag is 2, then the ordinary ray is traced, however the phase rotation due to the extraordinary ray is accounted for. If the mode is 3, the extraordinary ray is traced, and the phase rotation due to the ordinary ray is accounted for. For modes 2 and 3, no energy is discarded, and the polarization of the ray is properly rotated by the birefringent media.

Transmission and other properties of birefringent surfaces

When light propagates through birefringent media, the index of refraction of the glass is different for the S and P polarizations. The ordinary index is what is seen by the perpendicular, or S-polarized light, while the effective index is seen by the parallel or P-polarized light. Note the S and P polarization directions used in this context are not in general the same as those used by the coating and Fresnel surface effects computation. Here S and P refer to the perpendicular and parallel orientations relative to the crystal axis rather than the surface normal vector. The plane that contains the refracted ray and the crystal axis vector is the parallel plane; and the P vector lies in this plane normal to the ray vector. The S vector is perpendicular to both P and the ray vector.


For polarization analysis of birefringent materials, a few assumptions are needed. If the mode is 0, the ordinary ray is traced, which only has an S component, so the P component transmission is set to zero. If the mode is 1, the extraordinary ray is traced, and the S component is therefore set to zero. This technique yields the correct transmission results for each possible path separately. However, to get the total transmission requires analysis of each possible combination of modes for every pair of birefringent surfaces. If there are 2 pairs of birefringent surfaces in the system, 4 separate ray traces are required; and if there are 3 pairs of birefringent surfaces, 8 traces required, etc.

If the mode is 2, then the ray follows the ordinary path, but the ordinary index is used to phase rotate the S component of the field, and the effective index (properly referenced to the ordinary ray direction) is used to rotate the P component of the electric field. Only 1 trace is needed to properly model both transmission and phase rotation. If the mode is 3, then the ray follows the extraordinary path, with a similar phase rotation as in the case of mode 2.

When modeling bulk transmission through birefringent media, only the internal transmittance as defined by the ordinary index data is considered. No accounting for dichroism is supported.

Sample ZEMAX files

For some samples of the birefringent surface type, see the "<data>\Samples\Sequential\Birefringent prisms" folder. Shown below is the "Wollaston prism" sample file showing two configurations superimposed. The two halves of the prism are both birefringent, with the crystal axis oriented at 90 degrees to each other.



PARAMETER DEFINITIONS FOR BIREFRINGENT "IN" SURFACES

Parameter 0	Parameter 1	Parameters 2,3,4	Parameter 5	Parameter 6
Length to draw crystal axis, 0 to not draw.	Mode, 0, 1, 2, or 3.	X, Y, and Z-cosines	Parax Ignore. If 1, paraxial rays ignore the extraordinary index.	Shape. If the shape is not zero, parameters 7 and up are used to define the surface shape

PARAMETER DEFINITIONS FOR BIREFRINGENT "OUT" SURFACES

Parameter 0-5	Parameter 6
Unused.	Shape. If the shape is not zero, parameters 7 and up are used to define the surface shape

Black Box Lens

The Black Box Lens surface is used to "play back" a ZEMAX Black Box (BB) file previously created using the export Black Box data tool, described in "Export ZEMAX Black Box Data" on page 260.

To load a BB file, set the surface type to "Black Box Lens" and then enter the name of the file, with the extension ZBB, in the comment column of the Lens Data Editor. The file must be placed in the <data>\BlackBoxes folder. ZEMAX will load the file, decrypt the prescription data contained within the file, and then use the data to trace rays through the hidden prescription data whenever a ray is incident upon that surface.

If the first surface of a lens file is a BB surface, and the object is at infinite conjugates, insert a dummy STANDARD surface in front of the BB surface. This is done to simplify how ZEMAX internally handles systems at infinite conjugates. Neither the surface before nor after the BB should have any glass defined. The BB surface itself cannot have a glass type. The thickness, aperture type, and aperture size of the BB are automatically set by the data defined within the ZBB file, and cannot be modified by the user. The object and image surfaces may not be BB surfaces.

Features not supported when using Black Box surfaces

When using BB surfaces, not all analysis features are supported to such an extent as to be able to reproduce the analysis data for the original system. These include features where the individual surface data must be modified to complete the analysis, such as for the thermal analysis features. Also, analysis features that require significant surface-by-surface computations cannot yield the same analysis data as the original file, and are therefore not supported. These include Physical Optics Propagation and Skew Gaussian Beam analysis. It is highly recommended that the creator of the ZBB data verify that the computations of interest to the end user of

the ZBB work as they did for the original file, and provide guidance to the end user as to what aspects of the lens performance the BB is intended to illustrate. All common features that rely directly on ray tracing, such as ray fans, OPD fans, spot diagrams, MTF, PSF, image analysis, and most other analysis features work well.

Conjugate

The conjugate surface is defined by two user-specified points. ZEMAX always uses the surface vertex as the reference point; the two points required to define the conjugate surface are specified relative to this vertex. The conjugate surface will always perfectly image one point to the other point, assuming the surface is a mirror. Although the conjugate surface can have any material type, it is useful to think of it as being defined by its reflective properties.

If the z-coordinates of the two points are either both positive or both negative, then the image formed from one of the points to the other is real. In this case, the distance from one of the points to an arbitrary point on the surface, plus the distance from the arbitrary point on the surface to the second point, is constant for all points on the surface. One additional constraint is needed to make the surface unique: the surface must pass through the vertex of the local coordinate system. If the surface is reflective, then one point is the conjugate of the other, hence the name.

The surface generated by these two points satisfies the following expression if both z1 and z2 have the same sign:

$$\sqrt{(x-x_1)^2 + (y-y_1)^2 + (z-z_1)^2} + \sqrt{(x-x_2)^2 + (y-y_2)^2 + (z-z_2)^2} = \sqrt{x_1^2 + y_1^2 + z_1^2} + \sqrt{x_2^2 + y_2^2 + z_2^2}.$$

Note that the surface must intersect the point (0,0,0). Several types of surfaces can be formed with this model. For example, a sphere can be formed by setting the x and y values to zero, and the two z values each to the radius of the sphere. An elliptical surface of arbitrary orientation can be formed by specifying non-zero values for either the x or y values.

If z1 and z2 have opposite signs, then the image formed from one of the points to the other is virtual. In this case, the distance from one of the points to an arbitrary point on the surface, minus the distance from the arbitrary point on the surface to the second point, is constant for all points on the surface. Like the real imaging case, the surface must pass through the vertex of the local coordinate system.

The surface generated by these two points satisfies the following expression if z1 and z2 have opposite signs:

$$\sqrt{(x-x_1)^2 + (y-y_1)^2 + (z-z_1)^2} - \sqrt{(x-x_2)^2 + (y-y_2)^2 + (z-z_2)^2} = \sqrt{x_1^2 + y_1^2 + z_1^2} - \sqrt{x_2^2 + y_2^2 + z_2^2}.$$

Note that the surface must intersect the point (0,0,0). Several types of surfaces can be formed with this model. For example, a hyperbola can be formed by setting the x and y values to zero, and the two z values to opposite values. If the z values are equal but opposite, then a plane will be generated.

The coordinates of the two construction points are specified in the parameter columns, as shown in the following table. Neither the z1 nor the z2 value can be zero.

PARAMETER DEFINITIONS FOR CONJUGATE SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6
x_1	y_1	z_1	x_2	y_2	z_2

Coordinate Break

The coordinate break surface is used to define a new coordinate system in terms of the current system. It is always considered a "dummy" surface for ray tracing purposes. There are six parameters used to describe the new coordinate system: x-decenter, y-decenter, tilt about x, tilt about y, tilt about z, and a flag to indicate the order of tilting and decentration. Coordinate break decenters are always specified in lens units. Coordinate tilts are

specified in degrees, and are right-handed with respect to the positive axes. Coordinate breaks are always relative to the coordinate system of the previous surface.

An alternate way of implementing coordinate breaks is to use surface tilts and decenters, see “Surface tilt/decenter tab” on page 84 for details.

The order flag



The order of the decenters and tilts matters!

If the "order" flag is set to zero, ZEMAX first decenters in x and y (decenters are orthogonal so the order does not matter). Then ZEMAX tilts about the local x axis (which rotates the y and z axes to new orientations), then tilts about the new y axis (which rotates the x and z axes), then finally tilts about the new z axis.

If the "order" flag is any other value (such as unity), then the tilts are done first, in the order around local z, then around the new y, and then around the new x, and then finally the decenters are done. This "order" flag is extremely useful because a single coordinate break can undo an earlier coordinate break, even for compound tilts and decenters.

The coordinate break acts like a plane surface oriented in the coordinate system after the decenters and tilts have been applied. However, the surface is never drawn, and cannot be used to define the boundary between two media. The glass type will always be the same as the prior surface, and ZEMAX will display "-" for the glass name, which is meant to indicate that a glass type cannot be entered there. Coordinate breaks themselves can never be mirrors, nor can the object surface be a coordinate break.

PARAMETER DEFINITIONS FOR COORDINATE BREAK SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6
Decenter X	Decenter Y	Tilt About X	Tilt About Y	Tilt About Z	Order

Cubic Spline

The cubic spline surface is described by sag values which are the distances between the vertex tangent plane and the surface. Spline surfaces are used to describe unusual correctors, headlamps, and other non-standard optical surfaces, but rarely for imaging applications because of the fundamental properties of splines. See "Comments about spline surfaces" below for more discussion.

The Cubic Spline surface uses eight values to represent the sag at one-eighth, two-eighths, and so on to eight-eighths of the semi-diameter of that surface. Cubic spline surfaces are rotationally symmetric. All eight points must be defined. A subset cannot be used, although the semi-diameter may be defined to exceed the useful aperture of the surface. This is often required because of the steep curvatures occasionally introduced by spline fitting. If eight points provides an overly coarse sampling, see “Extended Cubic Spline” on page 299. For a more general non-rotationally symmetric surface, see “Grid Sag” on page 313.

Comments about spline surfaces

Cubic splines are formed by a piece-wise concatenation of curved segments. Within the bounds of each segment, the curve is defined by a third order polynomial. The polynomial coefficients describing each segment are determined from the sag values of the defined segment boundaries. The determination of the coefficients is driven by the boundary requirements that the curve goes through the defined points, and both the first and second derivatives be continuous across segment boundaries. For a third order spline, it is not possible to require higher order derivatives, such as third, to be continuous across segment boundaries. For this reason, splines are of limited accuracy and usefulness in high precision optical design.

A common characteristic of tracing rays through spline surfaces is rough or noisy looking ray data, with discontinuities in some results. These ray trace discontinuities are a fundamental limitation of splines, and they are not due to a flaw in ZEMAX, or a lack of numerical precision.

Higher order splines of course exist, and one way to eliminate the discontinuities is to use a higher order spline and fewer segments. In the limit, this is essentially the same as using a single high order polynomial for the whole

surface, see for example the “Even Asphere” on page 298. This is why high order polynomials, and not splines or NURBS, dominate in precision optical design; they are continuously smooth and differentiable to all orders.

For an excellent discussion of spline theory, properties and algorithms, see Numerical Recipes in C, by Press et al., Cambridge University Press.

PARAMETER DEFINITIONS FOR CUBIC SPLINE SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
Sag at 1/8	Sag at 2/8	Sag at 3/8	Sag at 4/8	Sag at 5/8	Sag at 6/8	Sag at 7/8	Sag at 8/8

Cylinder Fresnel

The Cylinder Fresnel surface has a polynomial aspheric cylindrical substrate with an independent polynomial aspheric cylindrical function defining the Fresnel lens properties. The surface substrate sag is given by:

$$z_s = \frac{c_s y^2}{1 + \sqrt{1 - (1 + k_s) c_s^2 y^2}} + \alpha_1 y^2 + \alpha_2 y^4 + \alpha_3 y^6 + \alpha_4 y^8 + \alpha_5 y^{10} + \alpha_6 y^{12} + \alpha_7 y^{14} + \alpha_8 y^{16}.$$

The previous expression is used to compute the ray-surface intercept. Once the intercept is found, the refraction (or reflection) of the surface is determined by the local slope of the Fresnel facets. The Fresnel facet shape is described by an identical expression, with independent coefficients:

$$z_f = \frac{c_f y^2}{1 + \sqrt{1 - (1 + k_f) c_f^2 y^2}} + \beta_1 y^2 + \beta_2 y^4 + \beta_3 y^6 + \beta_4 y^8 + \beta_5 y^{10} + \beta_6 y^{12} + \beta_7 y^{14} + \beta_8 y^{16}.$$

Note that the curvature (symbol c), the conic constant (symbol k), and all of the polynomial coefficients are independent for the substrate sag and Fresnel portions of the surface. The refraction at the surface accounts for only the Fresnel sag, while the ray-surface intercept depends only upon the substrate sag.

The substrate sag radius, conic, and polynomial terms are all specified in the Lens Data Editor. The Fresnel sag terms are specified in the Extra Data Editor. However, the Extra Data Editor uses curvature (the reciprocal of radius) rather than radius for the Fresnel sag. The Extra Data Values used are summarized in the following table.

EXTRA DATA DEFINITIONS FOR CYLINDER FRESNEL SURFACES

Extra Data Number	Description
1	Maximum aspheric polynomial term number, n. The maximum is 8.
2	The Y direction curvature (NOT RADIUS) of the Fresnel surface. This parameter will effect the refraction of the surface, but not the shape of the substrate.
3	The Y direction conic constant of the Fresnel surface. This parameter will effect the refraction of the surface, but not the shape of the substrate.
4	Coefficient on y^2
n+3	Coefficient on y^{2n}

The "Maximum aspheric polynomial term number" is used to specify the maximum polynomial term to be used in calculating the Fresnel sag. This number is provided to speed the ray tracing calculation, as terms beyond this

number are ignored. As with any complex surface model, extreme care should be taken to evaluate the accuracy and appropriateness of this model, especially where fabrication decisions are concerned.

Because there is no reliable way to compute the phase through a Fresnel surface which is not a plane, any calculation that requires OPD data, such as OPD fans, MTF, and Zernike coefficients, will not be supported if a non-plane substrate Fresnel surface is present in the lens description.

Diffraction Grating

The diffraction grating surface can be used to model straight-line ruled gratings. The grating lines are parallel to the local x-axis. Other orientations can be simulated by using a coordinate break surface before and after the grating surface. For a plane grating, rays traced to the grating are refracted according to the equation

$$n_2 \sin \theta_2 - n_1 \sin \theta_1 = \frac{M\lambda}{d} = M\lambda T,$$

where d is the grating spacing (always in micrometers), θ_2 is the refracted angle, θ_1 is the incident angle, M is the diffraction order, λ is the wavelength (always in micrometers), n_1 and n_2 are the indices of refraction before and after the grating, and T is the grating frequency in lines per micrometer. Note that the sign convention for M is completely arbitrary. ZEMAX uses the definition for T (lines per micrometer) rather than d (micrometers per line). The grating surface can be plane, spherical, or conical, and the medium before the grating, as well as the grating itself, can be air, glass, "MIRROR" or any other valid glass type. The grating is described by the y-spacing of the grating lines measured in lines per micrometer (independent of the system units) and the diffraction order. ZEMAX only models gratings to the extent of deviating ray paths. Other properties, such as efficiency, and relative transmission are not supported. If the grating spacing is too small (or if T is too large) to satisfy the grating relation, then a "Ray missed surface" error will be reported.

PARAMETER DEFINITIONS FOR DIFFRACTION GRATING SURFACES

Parameter 1	Parameter 2
Grating lines per micrometer	Diffraction order

Elliptical Grating 1

The elliptical grating 1 surface is a combination of an elliptical surface with polynomial aspheric terms and a grating with grooves that are straight when projected on to the vertex tangent plane (although the surface itself is usually not planar). The grating lines are parallel to the local x-axis, and may be equally spaced or have a variable spacing defined by a 4 term polynomial (up to the cubic term in y). Other orientations can be simulated by using a coordinate break surface before and after the grating surface. For a plane grating, rays traced to the grating are refracted according to the equation

$$n_2 \sin \theta_2 - n_1 \sin \theta_1 = \frac{M\lambda}{d} = M\lambda T,$$

where d is the effective grating spacing (always in micrometers), θ_2 is the refracted angle, θ_1 is the incident angle, M is the diffraction order, λ is the wavelength (always in micrometers), n_1 and n_2 are the indices of refraction before and after the grating, and T is the effective grating frequency in lines per micrometer. ZEMAX uses the definition for T (lines per micrometer) rather than d (micrometers per line). The effective grating frequency T is defined by:

$$\frac{1}{T} = d_{eff} = \frac{1}{T_0} + \alpha y + \beta y^2 + \Gamma y^3 + \Delta y^4 + \epsilon y^5.$$

Note the constant spacing term is defined in reciprocal units to be compatible with other grating surfaces in ZEMAX; while the higher order terms are defined in a polynomial in the y coordinate. If T evaluates to zero, the grating component is ignored, which is useful for modeling elliptical aspheres which are not gratings.

The medium behind the grating surface may be air, glass, or a mirror. ZEMAX only models gratings to the extent of deviating ray paths. Other properties, such as efficiency and relative transmission are not supported. If the grating spacing is too small (or if T is too large) to satisfy the grating relation, then a "Ray missed surface" error will be reported.

Elliptical Grating surface shape

The surface shape is an ellipse with added polynomial aspheric terms. The sag of the surface is given by the following expression:

$$z = \frac{cu^2}{1 + \sqrt{1 - u^2}} + \sum_{i=1}^N A_i E_i(x, y),$$

where

$$u^2 = a^2 x^2 + b^2 y^2.$$

In these expressions, x, y, and z are the ray-surface intercept coordinates, a, b, and c are the coefficients which define the shape of the ellipse, and the extended polynomials represented by E are defined in the section "Extended Polynomial" on page 301.

Note the symbols a, b, and c do not have the same meaning in this sag expression as these same symbol names do in the description of the Standard surface. To make a spherical elliptical surface, the values a and b should be identical and equal to one over the radius of curvature of the surface, and the value c should be equal to the radius of curvature:

$$a = b = \frac{1}{R} = \frac{1}{c}.$$

PARAMETER DEFINITIONS FOR ELLIPTICAL GRATING 1 SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameters 6-10
Lines per micrometer, T_0	Diffraction order	a	b	c	$\alpha, \beta, \Gamma, \Delta, \epsilon$

Elliptical Grating 2

The elliptical grating 2 surface is very similar to the elliptical grating 1 surface. The distinction is that the grating lines are formed by the intersection of tilted, parallel, evenly spaced planes with a substrate defined by an elliptical surface with polynomial aspheric terms. The intersection of the surface and tilted planes yield a grating with grooves that are straight when projected on to a plane tilted by an angle θ_y , measured in the local y direction. When projected onto a surface which is not plane, the grooves are curved, not straight.

The grating is defined by the parameter T, which is the lines per micrometer (the grating frequency) when the groove lines are projected onto the tilted XY plane. This is the reciprocal of the spacing between the planes which form the grooves. The diffraction order is M, and the elliptical substrate is defined by the parameters a, b, c, and extra data terms, exactly as for the elliptical grating.

The medium behind the grating surface may be air, glass, or a mirror. ZEMAX only models gratings to the extent of deviating ray paths. Other properties, such as efficiency and relative transmission are not supported. If the grating spacing is too small (or if T is too large) to satisfy the grating relation, then a "Ray missed surface"

error will be reported. The surface shape is an ellipse with added polynomial aspheric terms, see the “Elliptical Grating surface shape” on page 297.

PARAMETER DEFINITIONS FOR ELLIPTICAL GRATING 2 SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6
Lines per micrometer, T	Diffraction order	a	b	c	θ_y

Even Asphere

Rotationally symmetric polynomial aspheric surfaces are described by a polynomial expansion of the deviation from a spherical (or aspheric described by a conic) surface. The even asphere surface model uses only the even powers of the radial coordinate to describe the asphericity. The model uses the base radius of curvature and the conic constant. The surface sag is given by

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \alpha_1 r^2 + \alpha_2 r^4 + \alpha_3 r^6 + \alpha_4 r^8 + \alpha_5 r^{10} + \alpha_6 r^{12} + \alpha_7 r^{14} + \alpha_8 r^{16}.$$

Note that the coefficients have units. The coefficients are entered in the corresponding parameter columns, as shown in the following table.

PARAMETER DEFINITIONS FOR EVEN ASPHERE SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
α_1	α_2	α_3	α_4	α_5	α_6	α_7	α_8

Extended Asphere

The extended asphere surface is similar to the Even Asphere surface; see “Even Asphere” on page 298. However, the extended asphere can support aspheric coefficients up to order 480, whereas the even asphere is limited to 16. Also, a slightly different method is used for computing the coefficients of the polynomial terms. The sag of the surface is given by the following expression:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=1}^N \alpha_i \rho^{2i},$$

where the first expression is identical to the Standard surface, and the second term is a power series sum over a normalized radial coordinate. The normalized radial coordinate ρ is used because then the coefficients α_i all have units of lens units. The extra data terms are defined in the following table.

EXTRA DATA DEFINITIONS FOR EXTENDED ASPHERE SURFACES

Extra Data Number	Description
1	Maximum term number. The maximum is 240.
2	Normalization radius. All ray-intercept points are divided by this number to determine the x and y coordinates for polynomial evaluation.
3	Coefficient in lens units on ρ^2

Extra Data Number	Description
4	Coefficient in lens units on ρ^4
n+2	Coefficient in lens units on ρ^{2n}

The "Maximum term number" is used to specify the maximum polynomial term to be used in calculating the surface sag. This number is provided to speed the ray tracing calculation, as terms beyond this number are ignored. For example, if the last term in the series you wish to use is ρ^{14} , which is term number 7, then specify "7" in the maximum term number column. Note that the term number is 7, because it is the 7th term in the polynomial expansion, and NOT 9, which is the "extra data number", which is the position in the extra data list of parameters. The "extra data number" is always two greater than the term number!



The "extra data number" is always two greater than the term number!

The normalization radius scales the X and Y intercept coordinates of the ray so that the polynomial terms are all dimensionless and the coefficients are all in lens units. For modeling surfaces that are not well described by polynomials, see "Grid Sag" on page 313 or "User Defined" on page 331.

Extended Cubic Spline

The Extended Cubic Spline surface is similar to the Cubic Spline surface (see "Cubic Spline" on page 294), with the key difference being more terms are allowed. The Extended Cubic spline supports between 4 and 240 terms. The surface sag is defined by a list of sag points evenly spaced along the radius of a rotationally symmetric shape. The sag must always be zero at zero radius, and so this point is not entered into the spreadsheet. For important information about the limitations of splines, see "Comments about spline surfaces" on page 294. For a non-rotationally symmetric spline surface, see "Grid Sag" on page 313.

If you have a list of sag points that describes your optical surface (as might be generated from a computer program other than ZEMAX) and you wish to enter them into ZEMAX, use the load capability of the extra data editor. This feature is described in "Extra Data" on page 93.

EXTRA DATA DEFINITIONS FOR EXT. CUBIC SPLINE SURFACES

Extra Data Number	Description
1	Maximum number of spline points.
2	Step size in lens units between points along the radius.
3	The sag at a distance of one step size away from the vertex.
4...n	The sag at a distance of 2... (n-2) step sizes away from the vertex.

The "Maximum number of spline points" is used to specify how many points define the surface. This number may be between 4 and 240. At least 4 points are required to define the curved surface.

Spline surfaces can cause rough or ragged ray tracing results. A more general and much smoother solution is to use the grid sag surface, which is not restricted to rotational symmetry. See "Grid Sag" on page 313.

Extended Fresnel

The Fresnel surface described in "Fresnel" on page 303 models a plane surface which has the refractive (or reflective) power of a curved surface. The Extended Fresnel surface increases the flexibility of this model by supporting a plane, spherical, conic, or polynomial aspheric substrate on which a plane, spherical, conic, or polynomial Fresnel pattern is etched.

The surface sag is identical to the even aspheric surface:

$$z_s = \frac{c_s r^2}{1 + \sqrt{1 - (1 + k_s)c_s^2 r^2}} + \alpha_1 r^2 + \alpha_2 r^4 + \alpha_3 r^6 + \alpha_4 r^8 + \alpha_5 r^{10} + \alpha_6 r^{12} + \alpha_7 r^{14} + \alpha_8 r^{16}.$$

See the section “Even Asphere” on page 298 for details. The previous expression is used to compute the ray-surface intercept. Once the intercept is found, the refraction (or reflection) of the surface is determined by the local slope of the Fresnel facets, which depends upon both the Fresnel facet shape expression for Zf (below) and the substrate shape expression for Zs (above). The Fresnel facet shape is described by an expression virtually identical to the even asphere sag expression:

$$z_F = \frac{c_F r^2}{1 + \sqrt{1 - (1 + k_F)c_F^2 r^2}} + \sum_{i=0}^n \alpha_i r^{2i}.$$

The only difference is that the latter expression need not use all 8 terms if n is less than 8, the maximum accepted value. Note that the curvature (symbol c), the conic constant (symbol k), and all of the polynomial coefficients (symbol α) are independent for the substrate sag and Fresnel portions of the surface. The refraction at the surface accounts for both the substrate sag and the Fresnel sag, while the ray-surface intercept depends only upon the substrate sag. The intention was to model a Fresnel lens molded on a plane which is curved or warped into a new substrate shape after it is fabricated.

The substrate sag radius, conic, and polynomial terms are all specified in the Lens Data Editor, just like the even asphere surface. The Fresnel sag terms are specified in the Extra Data Editor. However, the Extra Data Editor uses curvature (the reciprocal of radius) rather than radius for the Fresnel sag. The Extra Data Values used are summarized in the following table.

EXTRA DATA DEFINITIONS FOR EXTENDED FRESNEL SURFACES

Extra Data Number	Description
1	Maximum aspheric polynomial term number, n. The maximum is 8.
2	The curvature (NOT RADIUS) of the Fresnel surface. This parameter will effect the refraction of the surface, but not the shape of the substrate.
3	The conic constant of the Fresnel surface. This parameter will effect the refraction of the surface, but not the shape of the substrate.
4	Coefficient on r^2
n+3	Coefficient on r^{2n}

The "Maximum aspheric polynomial term number" is used to specify the maximum polynomial term to be used in calculating the surface sag. This number is provided to speed the ray tracing calculation, as terms beyond this number are ignored. As with any complex surface model, extreme care should be taken to evaluate the accuracy and appropriateness of this model, especially where fabrication decisions are concerned.

Because there is no reliable way to compute the phase through a Fresnel surface which is not a plane, any calculation that requires OPD data, such as OPD fans, MTF, and Zernike coefficients, will not be supported if a non-plane substrate Fresnel surface is present in the lens description.

Extended Odd Asphere

The Extended Odd Asphere surface is similar to the Odd Asphere surface described in "Odd Asphere" on page 318. However, the Extended Odd Asphere can support aspheric coefficients up to order 240, whereas the Odd Asphere is limited to 8. Also, a slightly different method is used for computing the coefficients of the polynomial terms. The sag of the surface is given by the following expression:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=1}^N \alpha_i \rho^i,$$

where the first expression is identical to the Standard surface, and the second term is a power series sum over a normalized radial coordinate. The normalized radial coordinate ρ is used because then the coefficients α_i all have units of lens units. The extra data terms are defined in the following table.

EXTRA DATA DEFINITIONS FOR EXTENDED ODD ASPHERE SURFACES

Extra Data Number	Description
1	Maximum term number. The maximum is 240.
2	Normalization radius. All ray-intercept points are divided by this number to determine the x and y coordinates for polynomial evaluation.
3	Coefficient in lens units on ρ^1
4	Coefficient in lens units on ρ^2
n+2	Coefficient in lens units on ρ^n

The "Maximum term number" is used to specify the maximum term to be used in calculating the surface sag. This number is provided to speed the ray tracing calculation, as terms beyond this number are ignored.

Extended Polynomial

The extended polynomial surface is similar to the polynomial surface, except more terms are allowed. The surface also supports a base conic asphere surface upon which the polynomial aspheric terms are added. The surface sag is of the form:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=1}^N A_i E_i(x, y).$$

where N is the number of polynomial coefficients in the series, and A_i is the coefficient on the i^{th} extended polynomial term. The polynomials are a power series in x and y . The first term is x , then y , then $x*x$, $x*y$, $y*y$, etc. There are 2 terms of order 1, 3 terms of order 2, 4 terms of order 3, etc. The maximum order is 20, which makes a maximum of 230 polynomial aspheric coefficients. The position values x and y are divided by a normalization radius so the polynomial coefficients are dimensionless.

For example, the 12th term in the polynomial expansion, which is extra data term number 14, is the coefficient on the term in x^2y^2 . The coefficients A_i all have units which are the same as the lens units, such as millimeters or inches.

EXTRA DATA DEFINITIONS FOR EXT. POLYNOMIAL SURFACES

Extra Data Number	Description
1	Maximum term number.
2	Normalization radius. All ray-intercept points are divided by this number to determine the x and y coordinates for polynomial evaluation.
3-232	Polynomial terms.
233 and above	Not used.

The "Maximum term number" is used to specify the maximum polynomial term to be used in calculating the surface sag. This number is provided to speed the ray tracing calculation, as terms beyond this number are ignored. For example, if the last term in the series you wish to use is xy^3 , which is term number 13, then specify "13" in the maximum term number column. Note that the term number is 13, because it is the 13th term in the polynomial expansion, and NOT 15, which is the "extra data number", which is the position in the extra data list of parameters. The "extra data number" is always two greater than the term number! The maximum term number is the same as "N" in the surface sag equation.



The "extra data number" is always two greater than the term number!

The normalization radius scales the X and Y intercept coordinates of the ray so that the polynomial terms are all dimensionless and the coefficients are all in lens units.

Extended Toroidal Grating

Extended toroidal grating surfaces are a combination of the toroidal (page 328), toroidal grating (page 295), and extended polynomial (page 301) surfaces. Toroidal surfaces are described by defining a curve in the Y-Z plane, and then rotating this curve about an axis parallel to the Y axis and intersecting the Z axis. The resulting shape is then deformed by the addition of extended polynomials. Finally, a linear diffraction grating is placed upon the surface. The curve in the Y-Z plane is defined by:

$$z = \frac{cy^2}{1 + \sqrt{1 - (1+k)c^2y^2}} + \alpha_1y^2 + \alpha_2y^4 + \alpha_3y^6 + \alpha_4y^8 + \alpha_5y^{10} + \alpha_6y^{12} + \alpha_7y^{14} + \alpha_8y^{16}.$$

This curve is similar to the even asphere (page 298) surface sag formula, except the coordinate argument is y, not r. This curve is then rotated about an axis a distance R from the vertex. This distance R is referred to as the radius of rotation, and may be positive or negative. The Y-Z radius of curvature is specified in the same column on the spreadsheet editor as the radius for Standard surfaces. The radius of rotation is set on parameter column number 1. To model a cylinder lens which is flat in the X direction use zero, which ZEMAX interprets as infinite radius.

The resulting surface is then modified by the addition of the extended polynomial terms:

$$z = z_{toroid} + \sum_{i=1}^N A_i E_i(x, y).$$

where N is the number of polynomial coefficients in the series, and A_i is the coefficient on the i^{th} extended polynomial term. The polynomials are a power series in x and y. The first term is x, then y, then x^2 , x^2y , y^2 , etc. There are 2 terms of order 1, 3 terms of order 2, 4 terms of order 3, etc. The maximum order is 20, which makes

a maximum of 230 polynomial aspheric coefficients. The position values x and y are divided by a normalization radius so the polynomial coefficients are dimensionless.

For example, the 12th term in the polynomial expansion, which is extra data term number 14, is the coefficient on the term in x^2y^2 . The coefficients A_i all have units which are the same as the lens units, such as millimeters or inches.

The diffraction grating is defined in terms of the number of lines per micrometer and the diffraction order. These values are specified in parameter columns 2 and 3, respectively. The grating lines are parallel to the local x axis, and are evenly spaced when projected onto a plane.

PARAMETER DEFINITIONS FOR EXTENDED TOROIDAL GRATING SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameters 4-11
Radius of Rotation	Grating Lines per micrometer	Diffraction order	$\alpha_1 - \alpha_8$

EXTRA DATA DEFINITIONS FOR EXTENDED TOROIDAL GRATING SURFACES

Extra Data Number	Description
1	Maximum term number.
2	Normalization radius. All ray-intercept points are divided by this number to determine the x and y coordinates for polynomial evaluation.
3-232	Polynomial terms.
233 and above	Not used.

The "Maximum term number" is used to specify the maximum polynomial term to be used in calculating the surface sag. This number is provided to speed the ray tracing calculation, as terms beyond this number are ignored. For example, if the last term in the series you wish to use is xy^3 , which is term number 13, then specify "13" in the maximum term number column. Note that the term number is 13, because it is the 13th term in the polynomial expansion, and NOT 15, which is the "extra data number", which is the position in the extra data list of parameters. The "extra data number" is always two greater than the term number! The maximum term number is the same as "N" in the surface sag equation. The normalization radius scales the X and Y intercept coordinates of the ray so that the polynomial terms are all dimensionless and the coefficients are all in lens units.

Filter

Filter surfaces are modeled using the User Defined Surface described in "User defined surface apodization using DLLs" on page 333.

Fresnel

The Fresnel surface model is used to simulate flat surfaces which have been etched to have a spherical (or optionally aspheric) profile on a small scale (for Fresnel lenses which are not flat see the extended Fresnel surface type description). The surface intercept is determined by computing the intersection of the incoming ray with a plane. Once the plane intercept points are found, the surface is then treated as generally curved for the purposes of refraction into the next medium. This is only an approximation to a real Fresnel lens, however. The real Fresnel lens has grooves which may alter the exact intercept point. The model used here is adequate for Fresnel lens which have fine grooves (the groove depth is very shallow compared to the aperture). Extreme Fresnel lenses, such as those used in lighthouses, are not modeled well. The radius of curvature and conic constant, if any, are specified in the same manner as a Standard surface. The parameter values are exactly the same as for the even asphere model; polynomial aspherics are supported to 16th order.

For other types of Fresnel surfaces, see “Extended Fresnel” on page 299 and “Generalized Fresnel” on page 304.

PARAMETER DEFINITIONS FOR FRESNEL SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
α_1	α_2	α_3	α_4	α_5	α_6	α_7	α_8

Generalized Fresnel

This surface represents a generalization of the Fresnel surface described in “Fresnel” on page 303. The basic concept of a Fresnel surface, as far as ray tracing is concerned, is that the overall shape of the surface substrate is independent of the local slope. The generalized Fresnel uses a polynomial aspheric substrate model, identical to the Even Aspheric surface described in “Even Asphere” on page 298. The sag of the surface is given by the following expression:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \alpha_1r^2 + \alpha_2r^4 + \alpha_3r^6 + \alpha_4r^8 + \alpha_5r^{10} + \alpha_6r^{12} + \alpha_7r^{14} + \alpha_8r^{16},$$

where the terms are identical to those in the even asphere surface model. See that section for a complete discussion.

Once the ray has intercepted the surface, the ray reflects or refracts as if the surface had a shape described by a polynomial expansion of the form:

$$z = \sum_{i=1}^N A_i E_i(x, y),$$

where N is the number of polynomial coefficients in the series, and A_i is the coefficient on the i^{th} extended polynomial term. The polynomial terms are identical to those described in “Extended Polynomial” on page 301. It is the gradient of the extended polynomial terms alone that determines the reflective or refractive properties of the surface, not the surface substrate shape.

One application for this surface type is modeling faceted surfaces. For example, a flat substrate may consist of a series of small faceted planes, which would reflect or refract the light as though the surface was tilted. This can be simulated using a flat substrate and a linear x or y tilt term in the polynomial coefficients.

Because there is no reliable way to compute the phase through a Fresnel surface which is not a plane, any calculation that requires OPD data, such as OPD fans, MTF, and Zernike coefficients, will not be supported if a non-plane substrate Fresnel surface is present in the lens description.

Note that the coefficients have units. The coefficients are entered in the corresponding parameter columns, as shown in the following table.

PARAMETER DEFINITIONS FOR GENERALIZED FRESNEL SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
α_1	α_2	α_3	α_4	α_5	α_6	α_7	α_8

EXTRA DATA DEFINITIONS FOR EXT. POLYNOMIAL SURFACES

Extra Data Number	Description
1	Maximum term number.
2	Normalization radius. All ray-intercept points are divided by this number to determine the x and y coordinates for polynomial evaluation.
3-232	Polynomial terms.
233 and above	Not used.

The "Maximum term number" is used to specify the maximum polynomial term to be used in calculating the surface sag. This number is provided to speed the ray tracing calculation, as terms beyond this number are ignored.

Gradient 1

This surface has the same shape as the Standard surface, with media whose index of refraction is described by

$$n = n_0 + n_{r2}r^2 + n_{r1}r,$$

where $r^2 = x^2 + y^2$. Four parameters are required: the maximum step size, Δt , the base index, n_0 , the radial quadratic index, n_{r2} , and the radial linear index, n_{r1} . Note that n_{r2} and n_{r1} have units.

Discussion on maximum step size for GRIN surfaces

The maximum step size Δt determines the trade-off between ray tracing speed and accuracy. The exact value required depends upon the numerical aperture of the system and the magnitude of the coefficients. To determine a suitable step size, start at a large value (on the order of the thickness of the surface), and then perform a spot diagram. Note the RMS spot radius. Now decrease the step size by a factor of two. If the RMS spot radius changes by less than a few percent, the new step size is probably small enough. Otherwise, decrease the step again. For the final design phase, you may want to decrease the step size again. Using too small a step size needlessly slows the ray tracing speed without enhancement of accuracy. OPD tracing generally converges more slowly than ray tracing, so you may want to repeat the above procedure while inspecting the OPD fan. Occasionally check to ensure that the step size is appropriate as the design progresses.

Restrictions on surfaces following GRIN surfaces

The ray trace through a gradient index media requires iteration to determine the point of intersection of the ray with the surface following the gradient index surface. Because of this, not all surface types are allowed to follow gradient index surface types. If the type of the surface after a gradient index surface is not supported, a message will be presented indicating the error. Support for additional surface types to follow a GRIN surface may be added upon request; contact ZEMAX Technical Support for details.

PARAMETER DEFINITIONS FOR GRADIENT 1 SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4
Δt	n_0	n_{r2}	n_{r1}

Gradient 2

This surface has the same shape as the Standard surface, with media whose index of refraction is described by

$$n^2 = n_0 + n_{r2}r^2 + n_{r4}r^4 + n_{r6}r^6 + n_{r8}r^8 + n_{r10}r^{10} + n_{r12}r^{12},$$

where $r^2 = x^2 + y^2$. Eight parameters are required: the maximum step size, Δt , the base index squared, n_0 , and the remaining six coefficients of the preceding equation. Note that some of the coefficients have units.

The maximum step size Δt determines the trade-off between ray tracing speed and accuracy. For details see “Discussion on maximum step size for GRIN surfaces” on page 305. Also see “Restrictions on surfaces following GRIN surfaces” on page 305.

PARAMETER DEFINITIONS FOR GRADIENT 2 SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
Δt	n_0	n_{r2}	n_{r4}	n_{r6}	n_{r8}	n_{r10}	n_{r12}

Gradient 3

This surface has the same shape as the Standard surface, with media whose index of refraction is described by

$$n = n_0 + n_{r2}r^2 + n_{r4}r^4 + n_{r6}r^6 + n_{z1}z + n_{z2}z^2 + n_{z3}z^3,$$

where $r^2 = x^2 + y^2$. Eight parameters are required: the maximum step size, Δt , the base index, n_0 , and the remaining six coefficients of the preceding equation. Note that some of the coefficients have units.

The maximum step size Δt determines the trade-off between ray tracing speed and accuracy. For details see “Discussion on maximum step size for GRIN surfaces” on page 305. Also see “Restrictions on surfaces following GRIN surfaces” on page 305.

PARAMETER DEFINITIONS FOR GRADIENT 3 SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
Δt	n_0	n_{r2}	n_{r4}	n_{r6}	n_{z1}	n_{z2}	n_{z3}

Gradient 4

This surface has the same shape as the Standard surface, with media whose index of refraction is described by

$$n = n_0 + n_{x1}x + n_{x2}x^2 + n_{y1}y + n_{y2}y^2 + n_{z1}z + n_{z2}z^2.$$

There are eight parameters required: the maximum step size, Δt , the base index, n_0 , and the remaining six coefficients of the preceding equation. Note that some of the coefficients have units. This particular GRIN model is useful for gradient index optics with a cylindrical power profile. It is also useful for modeling thermal gradients in optical elements, if you have sufficient data to compute the coefficients.

The linear transverse terms, n_{x1} and n_{y1} , are ignored when performing paraxial ray tracing.

The maximum step size Δt determines the trade-off between ray tracing speed and accuracy. For details see “Discussion on maximum step size for GRIN surfaces” on page 305. Also see “Restrictions on surfaces following GRIN surfaces” on page 305.

PARAMETER DEFINITIONS FOR GRADIENT 4 SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
Δt	n_0	n_{x1}	n_{x2}	n_{y1}	n_{y2}	n_{z1}	n_{z2}

Gradient 5

The shape of the Gradient 5 surface is the same as the Standard surface plus a "tilt" term in both x and y directions:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + x \tan(\alpha) + y \tan(\beta),$$

where c is the curvature (the reciprocal of the radius), r is the radial coordinate in lens units, k is the conic constant, and $\tan \alpha$ and $\tan \beta$ are the tangents of the tilt angles in x and y. Note this is not the same surface shape as a tilted Standard surface, but it is a close approximation as long as the curvature is small, or if the tilt angles are small. The Gradient 5 surface has the following gradient profile:

$$n_{ref} = n_0 + n_{r2}r^2 + n_{r4}r^4 + n_{z1}z + n_{z2}z^2 + n_{z3}z^3 + n_{z4}z^4,$$

where $r^2 = x^2 + y^2$. Ten parameters are required: the maximum step size, Δt , the base index, n_0 , the remaining six coefficients of the preceding equation, and the two tilt terms. Note that some of the coefficients have units.

The important feature of the Gradient 5 model is that the model allows specification of dispersion properties of the media. The dispersion data is user defined and is stored in an ASCII file called SGRIN.DAT. The format for SGRIN.DAT will be described shortly.

The material name is entered in the glass column of the Gradient 5 surface type. If the glass column is blank, then dispersion effects are ignored.

To perform ray tracing, ZEMAX first computes the index at a "reference" wavelength using the preceding formula for n_{ref} . The index at any other wavelength is then computed using the following method based upon a general expansion of the Sellmeier formula:

$$n(\lambda)^2 = n(\lambda_{ref})^2 + \sum_{i=1}^3 \frac{K_i(\lambda^2 - \lambda_{ref}^2)}{\lambda^2 - L_i}, \text{ where}$$

$$K_i = \sum_{j=1}^{K_MAX} K_{ij}[n_{ref}]^{j-1}, L_i = \sum_{j=1}^{L_MAX} L_{ij}[n_{ref}]^{j-1}, \text{ and } n_{ref} = n(\lambda_{ref}).$$

The coefficients K_{ij} and L_{ij} define the dispersion of the material, whereas the gradient coefficients specified by parameters 2-8 (see the following table) define the index gradient profile at the reference wavelength. This very general dispersion model permits modeling of nearly arbitrary gradient index dispersion over a wide wavelength band. The parameters K_MAX and L_MAX may be as little as 1 term, or up to 8 terms for more accurate modeling.

The dispersion data is stored in the ASCII file SGRIN.DAT which must reside in the <glass> folder (see "Folders" on page 68).

The SGRIN.DAT file contains blocks of 10 lines each. The file format has the following structure:

```
MATERIALNAME
MIN_WAVELENGTH MAX_WAVELENGTH
REF_WAVELENGTH
K_MAX L_MAX
```

K11 K12 K13 ... K1K_MAX
 K21 K22 K23 ... K2K_MAX
 K31 K32 K33 ... K3K_MAX
 L11 L12 L13 ... L1L_MAX
 L21 L22 L23 ... L2L_MAX
 L31 L32 L33 ... L3L_MAX

Multiple materials may be defined in the same file by placing additional blocks of 10 lines one after the other, with no blank lines between. The supplied file SGRIN.DAT contains coefficients to describe some gradient materials offered by LightPath Technologies.

The maximum step size Δt determines the trade-off between ray tracing speed and accuracy. For details see "Discussion on maximum step size for GRIN surfaces" on page 305. Also see "Restrictions on surfaces following GRIN surfaces" on page 305.

PARAMETER DEFINITIONS FOR GRADIENT 5 SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8	Parameter 9	Parameter 10
Δt	n_0	n_{r2}	n_{r4}	n_{z1}	n_{z2}	n_{z3}	n_{z4}	$\tan \alpha$	$\tan \beta$

Gradient 6

This surface has the same shape as the Standard surface, with media whose index of refraction is described by

$$n = n_0 + n_1 r^2 + n_2 r^4 + n_3 r^6 + n_4 r^8 .$$

The n_x values are computed from the following formula:

$$n_x = A_x + B_x \lambda^2 + \frac{C_x}{\lambda^2} + \frac{D_x}{\lambda^4} ,$$

with identical expressions (but different values for A, B, C, and D) for n_0, n_1, n_2, n_3 and n_4 . This expression uses nanometers rather than micrometers for λ .

The dispersion data is user defined and is stored in an ASCII file called GLC.DAT. This file is stored in the folder <glass> (see "Folders" on page 68). The first line in the file is a file version format specifier of the form:

VERSION N

where N is the version number of the file format, currently 1. ZEMAX will read older GLC data files that omitted the version number, however the current version number and format should be used for any new files created or data added. The remaining lines are blocks of 21 lines defining the name and data for each material in the form:

NAME
 A0
 B0
 C0
 D0
 A1
 B1
 etc... ending with D4

The first line in the block is the name of the material, which can be any name (without special characters such as spaces or quotes) less than 10 characters long. The next 20 lines are the values for A, B, C, and D for n_0 through n_4 , respectively. There are no blank lines or comments allowed. ZEMAX can read in data for up to 25 different materials in the GLC.DAT file.

The dispersion data in the supplied GLC.DAT file was provided by the Gradient Lens Corporation (GLC) in Rochester, NY, (716) 235-2620. Contact GLC for detailed information on material properties. Not all of the materials offered by GLC are included in the GLC.DAT file. The following materials are included: ARS10, ARS20, ARS27, and ARS31.

To use the Gradient 6 surface materials, change the surface type to Gradient 6, and then enter the appropriate material name in the glass column of the Lens Data Editor.

The maximum step size Δt determines the trade-off between ray tracing speed and accuracy. For details see “Discussion on maximum step size for GRIN surfaces” on page 305. Also see “Restrictions on surfaces following GRIN surfaces” on page 305.

PARAMETER DEFINITIONS FOR GRADIENT 6 SURFACES

Parameter 1
Δt

Gradient 7

This surface has the same shape as the Standard surface, with media whose index of refraction is described by

$$n = n_0 + \alpha(r - R) + \beta(r - R)^2, \text{ where}$$

$$r = \frac{R}{|R|} \sqrt{x^2 + y^2 + (R - z)^2}.$$

The coordinates x , y , and z are the usual coordinates measured with respect to the vertex tangent plane, and R is the radius of the isoindex contours measured at the vertex. The isoindex contours are spherical shells centered about the point $z = R$. The starting index, n_0 , is measured at the vertex of the surface, not at the center of the isoindex contours. There are five parameters required: the maximum step size, Δt , the base index, n_0 , R , α , and β . Note that α and β have units.

The isoindex contour radius R may be specified independently from the front or rear radius of curvature of the lens. However, if R is set to be zero, then ZEMAX assumes the isoindex radius and the front radius of curvature are equal.

The maximum step size Δt determines the trade-off between ray tracing speed and accuracy. For details see “Discussion on maximum step size for GRIN surfaces” on page 305. Also see “Restrictions on surfaces following GRIN surfaces” on page 305.

PARAMETER DEFINITIONS FOR GRADIENT 7 SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5
Δt	n_0	R	α	β

GRADIUM™

This surface type models lenses made from stock gradient index blanks available from LightPath Technologies, Inc. The surface shape is the same as the Standard surface. The blanks have an axial gradient index profile, which describes a reference index of refraction as a function of axial position within the glass blank. All that is required to define the lens is the starting position within the blank, the name of the stock blank profile, and of

course the radii and thickness. The GRADIUM surface has a gradient index profile which is described by a polynomial of the form:

$$n = \sum_{i=0}^{11} n_i \left(\frac{z + \Delta z}{z_{max}} \right)^i.$$

The z coordinate is the distance from the front vertex of the surface, zmax is the maximum z coordinate of the blank (also known as the boule thickness), and the value Δz is the "offset" distance along the profile. Unlike most other gradient glass models in ZEMAX, the GRADIUM surface uses only fixed, predefined axial profile coefficients. The only design parameter required is the offset value, Δz . The available profiles provided with ZEMAX are defined in the ASCII file PROFILE.GRD. Other profile files may be defined, if they end in the extension GRD and are placed in the <glass> folder (see "Folders" on page 68). To select a different profile file, see "GRADIUM Profile" on page 110. For a list of the available profiles, see "Gradium™ Profile" on page 212.

GRADIUM profile file format

The file format is a series of blocks of 13 lines of data defined as follows:

```
PROFILE_NAME GLASS_FAMILY MAX_Z DENSITY UNUSED
n0
n1
...
n11
```

Each block of data starts with the name of the profile, which can be any valid ASCII name less than 20 characters. On the same line, the glass family name follows, which must be a gradient index material as defined in the SGRIN.DAT file, which is described in the section on the Gradient 5 surface type. The glass family name defines the reference wavelength which the profile describes. The last entry on the line is the maximum z coordinate of the blank. The 12 polynomial coefficients follow, from n0 to n11. The maximum allowed number of profiles is 100.

When ZEMAX performs ray tracing, the local z coordinate on the surface (which may be negative) is computed. The offset value is then added to determine where in the profile the coordinate lies.

Normally, the resulting value must always be positive and less than or equal to the maximum z, otherwise an error is generated (See the discussion of "Capping" below). The reference index is then evaluated. Once the reference index is computed, the index at the wavelength being traced is computed using the technique described in the section on the Gradient 5 surface.

The reference wavelength index at the front vertex is displayed as parameter 3. This value may be changed in the Lens Data Editor. When a new value is entered, ZEMAX computes the appropriate Δz to yield the specified reference index. However, Δz is the value that matters. The reference index is just displayed for convenience, and should not be made variable, or a multi-configuration operand. Note the reference index is the vertex index at the reference wavelength, which is the wavelength defined in the glass family definition file SGRIN.DAT. It is not necessarily the primary wavelength.

The GRADIUM surface model also supports 4 additional parameter terms that are intended for use in tolerancing: Decenter X, Decenter Y, Tilt X, and Tilt Y. These four terms model an axial gradient that is not perfectly centered and not perfectly parallel to the local Z axis. The tolerance terms modify the axial profile by redefining the axial coordinate z as follows:

$$z' = t_x(x - d_x) + t_y(y - d_y) + t_z z, \text{ where}$$

$$t_z = [1.0 - t_x^2 - t_y^2]^{\frac{1}{2}},$$

and t_x , t_y , t_z are the coefficients of the unit vector which points along the axial gradient axis and d_x and d_y are the decenters in lens units of the start of the profile. If t_x and t_y are both zero, then the d_x and d_y values do not matter (since the gradient is only along the z axis) and the t_z value is unity. The terms t_x and t_y determine the slope of the profile axis in x and y , which is intended to model the tolerance of axial alignment between the gradient axis and the mechanical axis of the lens. This expression is a linear approximation which is only valid for asymptotically small t_x and t_y .

The tolerance terms, d_x , d_y , t_x , and t_y , are ignored when performing paraxial ray tracing.

Usually, only the defined range of the profile is used. However, in some cases the profile may be extended in one or both directions to add additional glass to the ends of the profile, which allows the use of GRADIUM in thicker lenses. This technique is called "capping". By default, ZEMAX turns capping off, so that any ray trace which requires glass beyond the profile limits is flagged as an error. This automatically enforces boundary constraints during optimization. To remove this restriction, the capping flag can be set to be 1, 2, or 3. The default value of zero indicates the blank is bounded to the profile length at both ends. If the capping flag is 1, then only the left edge is bounded (the right edge is allowed to pass beyond the profile limit). If the capping flag is 2, then only the right edge is bounded. And if the capping is 3, then neither the left nor right edges of the blank are bounded, and both the thickness and the offset may take on any values. The additional material added to the beginning and end of the profile is assumed to be a homogeneous glass with the same index and dispersion as the respective end points of the profile. This assumption may be inaccurate if there is some slope to the profile at the defined endpoints, which is usually the case. Contact LightPath Technologies for more detailed information on designing with GRADIUM capping.

All GRADIUM profiles are defined in pairs; there is a "positive" and "negative" profile. These are identical profiles that are reversed in the axial direction. When modeling lenses in double pass, it is necessary to use one profile for the forward pass, and the other (reversed) profile for the backward pass.

The maximum step size Δt determines the trade-off between ray tracing speed and accuracy. For details see "Discussion on maximum step size for GRIN surfaces" on page 305. Also see "Restrictions on surfaces following GRIN surfaces" on page 305.

PARAMETER DEFINITIONS FOR GRADIUM SURFACES

Parameter 0	Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
boile thickness	Δt	Δz	n_{ref}	d_x	d_y	t_x	t_y	Capping

Gradient 9

The Gradient 9 surface may be used to model SELFOC® materials available from NSG America, Inc., or any gradient glass with a similar index variation. The shape of the Gradient 9 surface is the same as the Standard surface plus a "tilt" term in both x and y directions:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + x \tan(\alpha) + y \tan(\beta),$$

where c is the curvature (the reciprocal of the radius), r is the radial coordinate in lens units, k is the conic constant, and $\tan \alpha$ and $\tan \beta$ are the tangents of the tilt angles in x and y . Note this is not the same surface shape as a tilted Standard surface, but it is a close approximation as long as the curvature is small, or if the tilt angles are small. The Gradient 9 surface has the following gradient profile:

$$n = n_0 \left[1.0 - \frac{A}{2} r^2 \right].$$

Both A and n_0 are functions of wavelength:

$$A(\lambda) = \left[K_0 + \frac{K_1}{\lambda^2} + \frac{K_2}{\lambda^4} \right]^2,$$

$$n_0 = B + \frac{C}{\lambda^2},$$

where the wavelength is measured in micrometers. The dispersion data is user defined and is stored in an ASCII file called GRADIENT_9.DAT. This file is stored in the folder <glass> (see “Folders” on page 68). The GRADIENT_9.DAT file contains blocks of 6 lines each. The first line in the file is the name of the material, which can be any name (without special characters such as spaces or quotes) less than 20 characters long. The next 5 lines are the values for B, C, K0, K1, and K2. There are no blank lines allowed between blocks. ZEMAX can read in data for up to 25 different materials in the GRADIENT_9.DAT file.

The dispersion data in the supplied GRADIENT_9.DAT file was provided by NSG America, Inc., in Somerset, NJ, (732) 469-9650, or by Corning, Inc., Corning NY, (607) 974-4755. Contact the respective manufacturers for detailed information on material properties. Not all of the materials offered by these manufacturers are necessarily included in the GRADIENT_9.DAT file. The included materials are: SLS-1.0, SLS-2.0, SLW-1.0, SLW-1.8, SLW-2.0, SLW-3.0, SLW-4.0, and SLH-1.8 for NSG and Corning-Grin for Corning.

To use the Gradient 9 surface materials, change the surface type to Gradient 9, and then enter the appropriate material name in the glass column of the Lens Data Editor.

The maximum step size Δt determines the trade-off between ray tracing speed and accuracy. For details see “Discussion on maximum step size for GRIN surfaces” on page 305. Also see “Restrictions on surfaces following GRIN surfaces” on page 305.

PARAMETER DEFINITIONS FOR GRADIENT 9 SURFACES

Parameter 1	Parameter 2	Parameter 3
Δt	$\tan \alpha$	$\tan \beta$

Gradient 10

This surface has the same shape as the Standard surface, with media whose index of refraction is described by

$$n = n_0 + n_{y1}y_a + n_{y2}y_a^2 + n_{y3}y_a^3 + n_{y4}y_a^4 + n_{y5}y_a^5 + n_{y6}y_a^6,$$

where $y_a = |y|$ and the || symbols indicate the absolute value. This form of gradient has a discontinuity at the plane $y = 0$, and the gradient is bisymmetric about the $y = 0$ plane. If the glass type is left blank, then there is no dispersion. If a glass type is entered in the glass column, then it must be one of the Gradient 5 materials defined in the section on the Gradient 5 surface type on page 307. The above formula then defines the index profile at the reference wavelength for the material, and the index at other wavelengths is computed according to the dispersion model defined for the Gradient 5 surface. The linear transverse term, n_{y1} , is ignored when performing paraxial ray tracing.

The maximum step size Δt determines the trade-off between ray tracing speed and accuracy. For details see “Discussion on maximum step size for GRIN surfaces” on page 305. Also see “Restrictions on surfaces following GRIN surfaces” on page 305.

PARAMETER DEFINITIONS FOR GRADIENT 10 SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
Δt	n_0	n_{y1}	n_{y2}	n_{y3}	n_{y4}	n_{y5}	n_{y6}

Grid Phase

This surface is nearly identical to the Grid Sag surface (see “Grid Sag” on page 313). The key differences are:

- The units of sag are radians of phase instead of units of length.
- The unitflag data is only used to scale the delx, dely, and derivative values.
- The surface shape is a plane.
- A diffraction order is supported. The diffraction order is a multiplier on the phase values. A value of zero will turn "off" the phase effects. Setting the order to -1 will reverse the sign of all defined phase values.
- A "shear distance" is supported. See "Using the shear distance" below.

The file format and general information provided in the grid sag description is otherwise valid.

Using the shear distance

One method of modeling atmospheric turbulence is to use a Grid Phase surface with data generated by an external atmospheric modeling program. Accurate modeling may require multiple surfaces with different phase data modeling atmospheric layers separated by long distances. Because of the separation between these layers, different field angles pass through different parts of the atmosphere, depending upon the distance between the layers and the ray angle. The separation causes a field dependent lateral shearing between the surfaces.

To model this shearing, one could in principle place multiple Grid Phase surfaces separated by large distances in front of the optical system. In practice, this works poorly, because the bending of the rays as they pass through a Grid Phase surface may cause the rays to miss the entrance pupil entirely. To avoid this problem, the Grid Phase surface may be placed at the entrance pupil of the system, and then the "shear distance" corresponding to the actual position of the phase layer may be defined. When ray tracing to the Grid Phase surface, for phase purposes only, ZEMAX will adjust the ray coordinates to correspond to the coordinates the incoming ray would have had if the Grid Phase surface had been placed at the shear distance. Normally, shear distances are negative as the distance is measured from the surface to the effective position of the phase layer. If the shear distance is zero, no lateral shifting occurs.

Phase coefficients sign conventions

See “Binary Optic 1” on page 283 for a discussion of sign conventions.

PARAMETER DEFINITIONS FOR GRID PHASE SURFACES

Parameter 0	Parameter 1	Parameter 2
Diffraction order	Shear Distance	Interpolation method. Use 0 for bicubic spline, 1 for linear. See “Bicubic spline vs. linear interpolation” on page 315.

Grid Sag

The grid sag surface has a shape defined by a base plane, sphere, conic asphere, polynomial asphere, or Zernike asphere plus additional sag terms defined by a rectangular array of sag values. The surface shape is

determined by either a linear or a bicubic spline interpolation of the sag values. The sag can be described at discrete points by

$$Z_{ij} = Z_{base} + Z(x_i - xdec, y_j - ydec), \text{ where}$$

$$z_{base} = \frac{cr^2}{1 + \sqrt{1 - (1 + k)c^2 r^2}} + \sum_{i=1}^8 \alpha_i r^{2i} + \sum_{i=1}^N A_i Z_i(\rho, \varphi),$$

$$1 \leq i \leq nx, \quad 1 \leq j \leq ny, \text{ and}$$

and the Z_i terms are the Zernike standard terms as described in "Zernike Standard Sag" on page 343. The terms delx and dely are the spacings between grid points in the x and y directions, xdec and ydec are optional decentration coordinates, and nx and ny are the number of points in the grid in each direction. The nx and ny values must be no smaller than 5. The terms defining Z_{base} are identical to those defined for the even asphere surface. Note that the sag defined by the grid points may be decentered relative to the sag defined by the base aspheric shape. There are also Zernike decenter terms in parameters 9 and 10, which are optionally used to decenter just the Zernike expression relative to the surface vertex. ZEMAX computes the effective paraxial curvature using the vertex radius of curvature, asphere coefficient data, and grid data points when computing paraxial power for this surface.

Importing grid data

All of this data must be computed and tabulated outside of ZEMAX, arranged in the proper file format, and then read using the Extra Data Editor "Import" feature (see "Extra Data" on page 93). The proper file format is as follows:

```
nx ny delx dely unitflag xdec ydec
z dz/dx dz/dy d2z/dxdy nodata
.
.
.
```

The first line in the file contains seven values, which define the (integer) number of points in the x and y directions, the (floating point) increment in the x and y directions, an (integer) flag indicating the units of the data: 0 for mm, 1 for cm, 2 for in, and 3 for meters, and the (floating point) decenter of the grid points relative to the base surface in x and y. Any required scaling to the current lens units will be performed. Note that sag and cross derivative values have dimensions, and are therefore scaled, but first derivative values are dimensionless, and are not scaled.

The remaining nx*ny lines of the file contain four (floating point) numbers and (optionally) one integer each. The four floating point values are the sag, the x derivative of the sag, the y derivative of the sag, and the cross derivative d/dxdy. The optional fifth data entry is an integer flag that indicates if the data is invalid. Valid measured data should either have a zero or blank space for the nodata flag. Points for which the data is not valid should have a nodata value of unity. Currently, ZEMAX reads in the nodata flag value but does not use it; therefore, the sag data and derivatives should be set to zero or some other suitable value for points with no valid data.

The first data line in the file corresponds to the upper left corner of the surface, that is, the corner defined by -x and +y limits. Each point that follows is read across the face of the surface from left to right. After nx points, the nx+1 point is read in as the first value in row 2, and so on until nx*ny points have been read in. The file must be in ASCII and end in the extension .DAT (*for files used by NSC objects, the extension should be .GRD*).

The derivative values are required for smooth bicubic interpolation of the sag between the data points. The derivative values are not used by the linear interpolation algorithm. If all of the derivative values (dz/dx, dz/dy, and d2z/dxdy) are zero for every point in the file, then ZEMAX will automatically estimate the derivatives using a finite difference method.

The grid sag file format also supports comment lines. Any line starting with the "!" character is ignored.

Bicubic spline vs. linear interpolation

The sag of the surface at any arbitrary point is computed using either of two interpolation algorithms. The bicubic spline algorithm is used by default, and provides an interpolation that is smooth in both sag and the first derivative of the sag. The bicubic spline is the preferred algorithm for reasonably smooth data. For noisy pseudo-random like data, or grid data with sharp discontinuities, bicubic spline can cause "overshoot", which means the interpolation of the surface data can yield sag points which are far from the grid points. For these cases, a linear interpolation algorithm may provide more useful results. The choice of bicubic spline or linear interpolation is user selectable as one of the parameters.

Suggestions for using the Grid Sag surface

For important information about the limitations of splines, see "Comments about spline surfaces" on page 294.

To use the grid sag surface, first change the surface type to grid sag. Then from the Extra Data Editor, choose Tools, Import, and specify the correct surface number and DAT file name. The data will be loaded into memory, but will not be displayed on the Extra Data Editor. The reason for not displaying the data is the potentially large number of points. Once the grid file is loaded, the values can be checked using the Surface Data Report. The data is thereafter stored with the lens file, and the original DAT file is ignored. If the DAT file must be changed, the file needs to be reloaded using the import feature in the Extra Data Editor.

The grid size is only limited by available memory. Each point in the grid requires 4 8-byte double precision values and 1 1-byte value, or 33 bytes. A 255 by 255 grid file would require approximately 2 Megabytes of memory. Surfaces using Grid Sag surfaces with many points can be slow to edit and read and write from files. One way to speed up editing is to turn off the "disk multi step" undo feature in ZEMAX. See "Preferences" on page 68.

The bicubic interpolation algorithm is smooth to third order, with exact results at the grid points, and large grid files are generally not required for reasonably smooth surface shapes. However, like all low order spline models of surfaces, the Grid Sag surface can never closely model a high order asphere surface. The reason is that no reasonable number of piece-wise third order polynomials can accurately follow a higher order shape. The Grid Sag surface is intended to model arbitrary low order shapes without higher order undulations.

The Grid Sag surface is not defined outside of the bounds of the grid. Rays traced outside the grid area are treated as a ray miss error. It is a good idea to make the valid data portion of the grid slightly larger than the maximum area illuminated by rays; specifically, do not define the grid to be exactly the same size as the beam print. Rays nearly exactly at the edge of the grid may not trace as intended.

For Grid Sag surfaces that model steeply curved optics, such as fast mirrors, it is better to use the base radius and conic to define the parent shape, and then use the Grid Sag data to define the deviations from the base shape; rather than to imbed the entire sag definition for the surface in the grid data. The reason is ZEMAX can use the base radius and conic to find a good "first guess" to the ray-surface intercept, then use the additional grid sag data to iterate to the exact solution. It is difficult and slow for ZEMAX to find the correct intercept between a steeply incident ray and a steeply curved Grid Sag surface if the base spherical surface is flat.

An example C language program, GRIDSAMP.C, can be found in the folder <objects>\Grid Files (see "Folders" on page 68). This sample program shows the correct method for creating a grid DAT file; the example used is a spherical surface. The GRIDSAMP program also shows how to use finite differences to compute the derivatives when only the sag formula of a surface is known analytically.

No optimization or tolerancing of grid data is possible.

PARAMETER DEFINITIONS FOR GRID SAG SURFACES

Parameter 0	Parameter 1-8	Parameter 9	Parameter 10
Interpolation method. Use 0 for bicubic spline, 1 for linear. See "Bicubic spline vs. linear interpolation" on page 315.	$\alpha_1 - \alpha_8$, see "Even Asphere" on page 298	Zernike decenter X in lens units	Zernike decenter Y in lens units

EXTRA DATA DEFINITIONS FOR GRID SAG SURFACES

Extra Data Number	Description
1	Number of terms.
2	Normalization radius. Coordinates are normalized by this value.
3 - 233	Coefficients on Zernike polynomials 1 - 231, respectively, in lens units.

Hologram 1

The hologram 1 surface can be used to model optically constructed holograms. The hologram surface can be plane, spherical, or conical, and the medium behind the hologram can be air or glass. The glass can also be "MIRROR" which indicates the hologram is constructed and used in reflection. The hologram itself is described by the x, y, and z coordinates of two different construction points, a construction wavelength, and the diffraction order. The hologram deviates ray paths according to the equation

$$\hat{n} \times (r'_o - r'_r) = \left(\frac{m\lambda_p}{\lambda_c} \right) \hat{n} \times (r_o - r_r),$$

where \hat{n} is the unit vector normal to the surface of the hologram at ray intersection point, r_o is the unit vector along the first construction beam, r_r is the unit vector along the second construction beam, r'_r is the unit vector along the incident readout beam, r'_o is the refracted ray, λ_c and λ_p are the construction and playback wavelengths, and m is the diffraction order. A value of $m = 0$ means the ray is undeviated, while other integral values of m refer to higher diffracted orders. The notation used here is from the book "Aberrations of Optical Systems" by Welford, Adam Hilger (1986). Modeling holograms requires an understanding of their behavior which is beyond the scope of this manual, and the user is advised to see the discussion in Welford, or some other reference, before using this feature.

Most holograms are constructed and used in transmission or reflection. There are occasions where the hologram is constructed in transmission, and then the substrate is aluminized and used in reflection. This special case can be simulated with the hologram surface by specifying a negative construction wavelength. Although ray tracing will be correct for this special case, OPD tracing will not work.

ZEMAX only models holograms to the extent of deviating ray paths. Other properties, such as efficiency and relative transmission are not supported.

The two construction beams are defined in terms of their source points. The x, y, and z coordinates of the source points are measured relative to the hologram vertex coordinate are defined in current lens units. ZEMAX computes the unit vector at the ray-surface intersection point using the local coordinate data and the construction point data for the two construction beams. The construction wavelength always has units of micrometers. If the distance to a source point is greater than 1.0E+08 lens units, the point is assumed to be infinitely distant and the construction beam is a perfect plane wave for that point. This assumption yields more accurate OPD computations than simply entering large coordinate values.

The hologram is defined by the interference between the two defined construction beams, with no aberrations assumed on the construction beams. Optically fabricated holograms with aberrated construction beams may be modeled in a very general way, see "Optically Fabricated Hologram" on page 319.



Optically fabricated holograms with aberrated construction beams may be modeled using the "Optically Fabricated Hologram" surface.

The hologram 1 surface assumes that both construction beams diverge from the specified construction points. Because of the reciprocity of the construction beams, this is identical to the case where both construction beams

converge toward the construction points. Some hologram fabrication methods require one beam to be converging while the other beam is diverging. See "The hologram 2 surface" for information on this latter type of hologram.

PARAMETER DEFINITIONS FOR HOLOGRAM SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
X_1	Y_1	Z_1	X_2	Y_2	Z_2	λ_c	m

Hologram 2

The hologram 2 surface is very similar to the hologram 1 surface. The key difference is that the hologram 1 surface assumes both construction beams either diverge from or converge to the construction points, whereas the hologram 2 surface assumes one construction beam converges to one construction point, and the other construction beam diverges from the other construction point. Which beam is first or second does not matter due to reciprocity. The parameter data is the same for both hologram 1 and hologram 2 surfaces.

Irregular

The irregular surface is a Standard surface shape (plane, spherical, or conic) that has additional aspheric deviations in terms of decenter, tilt, spherical aberration, astigmatism, and coma. This surface type is primarily used by the tolerancing algorithm to model irregularities in a standard shape surface. The surface sag is given by:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + Z_s\rho^4 + Z_a\rho_y'^2 + Z_c\rho^2\rho_y',$$

where

$$\rho_x = \frac{x}{r_{max}}, \rho_y = \frac{y}{r_{max}}, \rho = \sqrt{\rho_x^2 + \rho_y^2}, \rho_y' = \rho_y \cos\theta - \rho_x \sin\theta,$$

and rmax is the maximum radial aperture of the lens, defined by the semi-diameter value for the surface. The coefficients Zs, Za, and Zc represent the amount of spherical aberration, astigmatism, and coma, respectively, in lens units at the maximum radial aperture. The astigmatism and coma are oriented along a line that makes an angle θ in degrees with respect to the y axis.

The x and y coordinates of the previous equations are in a decentered and tilted coordinate system defined by the decenter x, decenter y, tilt about x, and tilt about y values. The decenters are in lens units, and the tilt is in degrees. The tilt and decenter values work exactly like the coordinate break surface defined in this chapter, however, the tilts and decenters are undone after the ray is traced to the surface. Ray tracing is done according to this algorithm:

The surface is decentered, tilted about x, then about y.

The ray is traced to the surface.

The surface is untilted about y, untilted about x, then undecentered.

The irregular surface uses the first seven parameters to define the decenter, tilt, and Z coefficients, and the eighth parameter to define the angle. All the coefficients are measured in lens units, except the tilt angles which are in degrees.

PARAMETER DEFINITIONS FOR IRREGULAR SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
Decenter X	Decenter Y	Tilt About X	Tilt About Y	Zs	Za	Zc	θ

Jones Matrix



This feature is only available in the EE edition of ZEMAX.

This surface is used to define an arbitrary polarizing component. The surface shape is always plane. The Jones matrix modifies a Jones vector (which describes the electric field) according to

$$\begin{bmatrix} E'_x \\ E'_y \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} E_x \\ E_y \end{bmatrix},$$

where A, B, C, D, E_x , and E_y are all complex numbers. See “Defining polarizing components” on page 623 for a complete discussion of the Jones matrix. Only polarization analysis features consider the effects of this surface type. This surface uses eight parameter values, and none of the extra data values.

PARAMETER DEFINITIONS FOR JONES MATRIX SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
Ar	Ai	Br	Bi	Cr	Ci	Dr	Di

Lenslet Array

Lenslet arrays are modeled using the User Defined Surface described in “Lenslet arrays using DLLs” on page 333.

Non-Sequential Components



This feature is only available in the EE edition of ZEMAX.

The sag of the Non-Sequential Components surface is the same as the Standard surface.

This surface enables ray tracing through one or more non-sequential objects. See the chapter “Non-Sequential Components” for details.

Odd Asphere

The odd asphere model deviation is similar to the even asphere, except both even and odd powers of r are used. The sag is given by

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \beta_1 r^1 + \beta_2 r^2 + \beta_3 r^3 + \beta_4 r^4 + \beta_5 r^5 + \beta_6 r^6 + \beta_7 r^7 + \beta_8 r^8.$$

Note that the coefficients have units. The coefficients are entered in the corresponding parameter columns, as shown in the following table.

PARAMETER DEFINITIONS FOR ODD ASPHERE SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
β_1	β_2	β_3	β_4	β_5	β_6	β_7	β_8

Odd Cosine



This feature is only available in the EE edition of ZEMAX.

The odd cosine surface is an extension of the Odd Asphere surface, with 16 radial terms plus up to 6 additional cosine terms. The sag is given by

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=1}^{16} \beta_i r^i + \sum_{i=1}^m A_i s^{P_i} \cos(B_i \theta + C_i).$$

The first term is the sag for a Standard surface (plane, sphere, or conic). The second term is similar to the Odd Asphere surface, but the number of coefficients is fixed at 16. The third term supports m cosine terms, where the integer m must be between 0 and 6, inclusive. The coordinate s is the normalized radial coordinate given by $s = r / R$, where R is the user defined normalization radius. If R is zero or negative, the cosine terms are ignored. The surface may not be continuous in the first derivative if any B value is not an integer. The β_i coefficients have units which depend upon the index i, and the A_i coefficients have units of length in lens units. The angle θ is measured in radians and is related to the x and y coordinates on the surface by

$$\tan \theta = \frac{y}{x}.$$

The coefficient B is dimensionless and C is in units of radians. The coefficients are entered in the corresponding extra data columns, as shown in the following table.

EXTRA DATA DEFINITIONS FOR ODD COSINE SURFACES

Parameters 1-16	Parameter 17	Parameter 18	Parameters 19-42
The odd aspheric coefficients β_i .	The number of cosine terms m, must be between 0 and 6.	The normalization radius R for the coordinate s in the cosine portion of the sag formula.	The A, P, B, and C terms in the cosine portion of the sag formula.

Optically Fabricated Hologram



This feature is only available in the EE edition of ZEMAX.

The Hologram 1, Hologram 2, and Toroidal Hologram surfaces described elsewhere in this chapter are used to define optically fabricated holograms assuming the construction optics are "perfect" and no aberrations are imposed on the construction beams.

The Optically Fabricated Hologram is far more general. The hologram is defined using 3 ZEMAX lens files:

- 1) The playback file, in which the Optically Fabricated Hologram surface is placed.
- 2) The construction file for beam #1
- 3) The construction file for beam #2

When rays are traced in the playback system, ZEMAX automatically calls the two construction systems to determine the vectors defining the interference of the two construction beams at the interception point in the playback system. There are significant advantages to this method of defining the hologram:

- 1) The construction optics may be completely arbitrary. Each construction file may consist of multiple lenses, mirrors, or even other holograms; anything ZEMAX can model. The aberrations introduced in the construction beams are therefore fully considered.
- 2) The two construction optics files may include completely different optics.
- 3) Any variables set in the construction system automatically become variable in the playback system; allowing simultaneous closed loop optimization of the construction and playback systems.

Defining the substrate shape

The shape of the Optically Fabricated Hologram surface substrate can be any of three possible shapes, defined by the "Shape" parameter zero:

Shape 0: Same as the "Standard" on page 325, using c and k.

Shape 1: Same as the "Elliptical Grating surface shape" on page 297, using a, b, c, and the polynomial terms.

Shape 2: Same as the "Toroidal Grating" on page 329, using c, k, and R.

Defining the construction optics

To define an Optically Fabricated Hologram, first create two lens files that each define the optics used to illuminate and record the hologram. There are several important rules which MUST be followed when defining the construction optics files:

- 1) The files must have the same name, with a "_1" and "_2" appended to the end. For example, the file names "FAB_1" and "FAB_2" are valid file names for a pair construction optics files.
- 2) The construction files must reside in the same folder as the playback lens file.
- 3) The surface at which the two construction beams will interfere must be the stop surface in each file. Only the ray intercept vectors at the respective stop surfaces will determine the hologram properties. The stop surface may be a different surface number in each file.
- 4) The stop surface in each construction file must be either standard, elliptical grating, or toroidal grating surface types.
- 5) Only 1 configuration is allowed in each construction file. The playback system may use multiple configurations.
- 6) Only 1 field point is allowed in each construction file, although the field definition is arbitrary.
- 7) Only 1 wavelength is allowed in each construction file, and the wavelengths in each construction file must be identical. The playback wavelength is arbitrary, and need not be the same as in the construction optics.
- 8) Ray aiming must be turned on in each construction file.
- 9) The system aperture type must be "float by stop size" in each construction file.
- 10) No virtual propagations are allowed in the construction files; the rays must all take physically significant paths.

With the construction optics files defined, the name of the files is given to ZEMAX via the comments column. In the playback system, the hologram surface is set to "Optically Fabricated Hologram" and the comment for that surface is set to the name that defines the construction optics file names. For example, if "FAB_1" and "FAB_2" are used for the construction optics file names, enter "FAB" in the comment column. ZEMAX will append the "_1" and "_2" and automatically read in the construction optics files. If no name is given, the surface will ignore the effects of the hologram.

Optimizing the construction and playback optics

As ZEMAX traces rays in the playback system, any variables which affect the shape of the playback hologram (the semi-diameter, a, b, and c parameters, or extra data parameters) are automatically copied into the construction optics stop surfaces. This is why the stop surfaces in the construction and playback optics must be of the same type, so the construction beams overlap on a surface with the same shape as the playback surface.

Any variables set in the construction optics automatically become variable in the playback system, so the construction optics may be optimized simultaneously with any and all variables in the playback system.

Whenever the playback lens file is saved; the construction optics files are automatically saved as well.

To constrain design parameters in the construction files while optimizing both the playback and construction files simultaneously, use the CMFV operand in the playback system's merit function to call the merit functions defined in each of the construction beams. See the Chapter "Optimization" for information about the CMFV operand.

Because this surface uses external ZEMAX files to define the construction beams, the Global Search algorithm should not be used. The other optimization algorithms are acceptable to use.

Selecting the hologram type

There are two ways of interfering the hologram construction beams:

- 1) By assuming that both construction beams either diverge or converge from their respective sources, or
- 2) By assuming that one source beam converges and one diverges from their respective sources.

Assumption #1 is analogous to the model for the Hologram 1 surface, while assumption #2 is analogous to the Hologram 2 surface. The Optically Fabricated Hologram surface supports both modes of interference. To choose the hologram 1 type, set the "Holo Type" in parameter 1 to be 1, otherwise, set the "Holo Type" in parameter 1 to be 2. The type must be either 1 or 2.

In the majority of cases, ZEMAX can correctly compute the OPD through the hologram using the default "OPD Mode" of zero. However, there are some special cases where the OPD is not computed correctly. For these cases, the user must manually determine the proper OPD algorithm by setting the OPD Mode equal to 1, 2, 3, or 4. The correct OPD Mode is a function of the construction and playback geometry and the total number of mirrors in the three files. There is no reliable algorithm for determining the proper OPD Mode automatically in all cases. If the OPD is obviously wrong, try each of the OPD mode values 0 - 4 until the OPD values are computed correctly.

PARAMETER DEFINITIONS FOR OPTICALLY FABRICATED HOLOGRAM SURFACES

Parameter 0	Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6
Shape	Holo Type	Diffraction order	c, a, or c	k, b, or k	-, c, or R	OPD Mode

Paraxial

The paraxial surface acts as an ideal thin lens. Two parameters are required to model the paraxial surface: the focal length and the OPD mode. The focal length is that which would be measured in air (unity index) although the paraxial model will support imaging into a non-unity index medium. If the focal length is zero the paraxial lens has no optical power. The OPD mode indicates how ZEMAX should calculate the optical path difference for rays refracted by the paraxial lens. Although the ray tracing through paraxial lenses is well defined (see below), computing the OPD is more difficult, especially in the presence of significant aberrations, as discussed below.

Setting the OPD Mode

Mode = 0 is very fast and accurate in most cases if the aberrations are not severe. If the Mode = 0, then the OPD computation is based upon the conjugate positions computed by tracing a real paraxial ray. This is preferred for axial optical systems with modest aberrations.

Mode = 1 integrates the actual phase introduced by the surface for each ray traced. Mode = 1 is substantially slower than the other modes, but is generally the most accurate. If the aberrations in the optical spaces on either side of the paraxial lens are large, the paraxial lens OPD computation cannot assume that the lens is working at fixed conjugates for all incoming rays. For these systems, the OPD mode should be set to 1. Mode = 1 can also be used to check if one of the other modes is accurate. If the same OPD results are produced by Mode = 1 as another mode, then the other mode may be used as it will generally compute the OPD much faster.

Mode = 2 assumes that the lens is used at infinite conjugates regardless of the actual conjugates. This option will return incorrect results if the incoming beam is not reasonably collimated.

Mode = 3 is similar to Mode = 0, except paraxial rather than real rays are used. Mode = 3 is only valid for systems with modest aberrations that are well described by first order optical theory.

For maximum OPD computation accuracy, the working F/# of the paraxial lens should be no faster than about F/4, regardless of the OPD mode selected.



To model systems with collimated output, consider using Afocal Image Space rather than a paraxial lens; see "Afocal Image Space" on page 103.

The paraxial surface refracts rays using the following equations:

$$n'u_x' = nu_x - x\phi$$

$$n'u_y' = nu_y - y\phi$$

where ϕ is the surface power, n is the index of refraction, primes indicate values on the image side of the surface, and the angles are slopes which are computed from the ray direction cosines:

$$u_x = \frac{l}{n}$$

$$u_y = \frac{m}{n}$$

The paraxial surface shape is plane.

PARAMETER DEFINITIONS FOR PARAXIAL SURFACES

Parameter 1	Parameter 2
Focal length	OPD Mode

Paraxial XY

The paraxial XY surface type is similar to the paraxial surface except the optical power can be specified in the X and Y directions separately. This surface can therefore be used as a paraxial cylindrical or toroidal lens. Two parameters are required to define the paraxial XY surface: the X power, and the Y power. OPD through paraxial XY surfaces is computed using mode = 1 described in the paraxial surface description. Optical power is measured in inverse length lens units, such as inverse millimeters if the lens units are millimeters. See “Units” on page 104.

The paraxial XY surface shape is plane.

PARAMETER DEFINITIONS FOR PARAXIAL XY SURFACES

Parameter 1	Parameter 2
X Power	Y Power

Periodic

This surface shape is described by the following expression:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} - A \left[\frac{1}{4} [1 + \cos(2\pi\alpha x)][1 + \cos(2\pi\beta y)] - 1 \right],$$

where A is the peak to valley height in lens units of the modulation, and the values α and β are the spatial frequencies of the oscillations in the x and y directions, respectively. Note this expression is the sag of a sphere plus the cosine modulation. The sag of the surface is exactly zero at the vertex, and the amplitude is the peak to valley amplitude. The frequencies are measured in inverse lens units.

This surface uses three parameter values, and none of the extra data values.

PARAMETER DEFINITIONS FOR PERIODIC SURFACES

Parameter 1	Parameter 2	Parameter 3
A	α	β

Polynomial

The base radius of curvature and the conic constant are not used by this surface model. The sag of the polynomial surface is given by

$$z = \gamma_1 x^2 + \gamma_2 x^4 + \gamma_3 x^6 + \gamma_4 x^8 + \gamma_5 y^2 + \gamma_6 y^4 + \gamma_7 y^6 + \gamma_8 y^8.$$

A more general polynomial surface is also available, see “Extended Polynomial” on page 301.

PARAMETER DEFINITIONS FOR POLYNOMIAL SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
γ_1	γ_2	γ_3	γ_4	γ_5	γ_6	γ_7	γ_8

Q-Type Asphere

The Q-Type asphere is a radially symmetric asphere similar to the Extended Asphere surface (see “Extended Asphere” on page 298) with an alternate polynomial formulation. This surface supports two different variations of the Q-Type asphere, commonly called Qbfs and Qcon. The surfaces are described in G. W. Forbes, "Robust, efficient computational methods for axially symmetric optical aspheres", Opt. Express 18, 19700-19712 (2010). The mathematics of these surfaces are fairly involved compared to traditional aspheric polynomials and the details are not reproduced here. If the surface type is Qbfs, the conic term is always set to zero and no variables or solves are allowed on the conic term.

The extra data terms this surface uses are defined in the following table.

EXTRA DATA DEFINITIONS FOR Q-TYPE ASPHERE SURFACES

Extra Data Number	Description
1	Type. Use 0 for Qbfs and 1 for Qcon.
2	Maximum term number. The maximum is 100.
3	Normalization radius in lens units. This is used by the algorithm to normalize the radial coordinates, analogous to the Extended Asphere surface.
4	Coefficient A0
n+4	Coefficient An, where n has a maximum given by the maximum term number.

Radial Grating

The radial grating surface is similar to the diffraction grating surface (see “Diffraction Grating” on page 296), except the grating lines have radial symmetry, the grating line spacing is variable over the surface, and the substrate is an even aspheric surface shape (see “Even Asphere” on page 298). For a plane grating, rays traced to the grating are refracted according to the equation

$$n_2 \sin \theta_2 - n_1 \sin \theta_1 = \frac{M\lambda}{d},$$

where d is the grating spacing (always in micrometers), θ_2 is the refracted angle, θ_1 is the incident angle, M is the diffraction order, λ is the wavelength (always in micrometers), and n_1 and n_2 are the indices of refraction before and after the grating. The radial grating surface allows d to vary over the surface according to the equation:

$$d(p) = A_0 + A_1p^1 + A_2p^{-1} + A_3p^2 + A_4p^{-2} + \dots,$$

where the A_i coefficients all have units of micrometers, and p is the normalized radial coordinate defined by

$$p = \frac{r}{R},$$

where r is the radial coordinate on the surface in lens units and R is the user defined normalization radius. The grating spacing d can be interpreted in two different ways. The normal ZEMAX convention is to measure d along the projection of the grating on the XY plane, ignoring any sag or curvature of the underlying surface. The Radial Grating surface supports an additional "Grating Mode", where d is interpreted as being measured along the local surface tangent. The Grating Mode can be set to 0 or 1 in parameter 9 in the Lens Data Editor.

Note that the sign convention for M is arbitrary. The grating surface can be plane, spherical, conical, or even aspheric, and the medium before the grating, as well as the grating itself, can be air, glass, "MIRROR" or any other valid glass type. ZEMAX only models gratings to the extent of deviating ray paths. Other properties, such as efficiency, and relative transmission are not supported. If the grating spacing is too small to satisfy the grating relation, then a "Ray missed surface" error will be reported.

PARAMETER DEFINITIONS FOR RADIAL GRATING SURFACES

Parameter 0	Parameter 1-8	Parameter 9
Diffraction Order	Even aspheric coefficients $\alpha_1 - \alpha_8$	Grating Mode

EXTRA DATA DEFINITIONS FOR RADIAL GRATING SURFACES

Extra Data Number	Description
1	Maximum number of terms
2	Normalization Radius
3	Coefficient on p^0
4	Coefficient on p^1
n	Coefficient on $p^{(n-3)}$

Radial NURBS



This feature is only available in the EE edition of ZEMAX.

The acronym NURBS stands for Non-Uniform Rational B-Spline. NURBS are a very general class of curves and surfaces. For a complete discussion of NURBS, which is well beyond the scope of this discussion, see The NURBS Book, Second Edition, by Les Piegl and Wayne Tiller, Springer-Verlag, ISBN 3-540-61545-8.

The radial NURBS surface is defined by a series of weighted control points. The control points define a curve which starts at the origin and lies in the YZ plane along the +Y direction. Once this curve is defined, a figure of

revolution is formed by rotating the curve a full 360 degrees around the Z axis. Unlike a spline surface, a NURBS curve does not actually go through the control points, except for the first and last control points.

Each control point has a positive "y" coordinate, a "z" coordinate (which may be positive or negative), and a weight "w". The first point is always at y = 0 and z = 0, so the curve begins at the vertex origin of the surface. Subsequent points, numbered from 1 to the maximum number of points desired, extend monotonically outward in the +Y direction. The y values must be spaced at least 1.0E-3 lens units apart for numerical stability in the spline fit.

The weights should all be set to 1.0 initially. The higher the weight, the closer the curve will be to the actual control point. Lower weights mean the curve is less constrained to be near the control point. See the above reference for details on the influence of weights on a NURBS curve.

The y value of the last point defined determines the maximum radial clear aperture of the surface. This value should generally be fixed, and not be a variable. Any rays which do not intercept the surface within the defined areas of the curve will be terminated with a "ray miss" error.

The advantage to the NURBS description is that any shape may be defined and traced reliably. Unusual aspheric correctors which cannot be described by polynomials may be modeled as NURBS.

The disadvantages to the NURBS description include very slow ray tracing, and sometimes great difficulty in finding suitable starting values for the control curve points and weights.

This surface does not use any of the parameter columns.

EXTRA DATA DEFINITIONS FOR RADIAL NURBS SURFACES

Extra Data Number	Description
1	Number of control points. At least 4 points are required, and no more than 60 allowed.
2	y coordinate for control point 1.
3	z coordinate (sag) for control point 1.
4	w (weight) value for control point 1.
3n-1, 3n, 3n+1	y, z, and w for control point n.

Retro Reflect

The Retro Reflect surface is a planar shape surface. If the surface is made of glass or air, rays refract according to Snell's law in the usual way. If the surface is made of "MIRROR", rays do not reflect in the usual way; instead the rays are undeviated. This means under subsequent propagation in the opposite direction the rays will follow the incident path; that is, the rays retro-reflect.

ZEMAX cannot accurately compute the optical path difference of rays retro-reflected from this surface if the aberrations are large (hundreds of times the diffraction limit) and the beam incident upon the retro-reflect surface is very fast (about F/2 or faster). For this reason, the retro reflect surface should only be used with reasonably slow beams and modest aberrations.

This surface uses no parameter or extra data values.

Standard

The most commonly used optical surface is a spherical surface. The sphere is centered on the current optical axis, with the vertex located at the current axis position. ZEMAX treats planes as a special case of the sphere (a sphere with infinite radius of curvature) and conics as a special case as well. The "sag" or z-coordinate of the Standard surface is given by

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}}$$

where c is the curvature (the reciprocal of the radius), r is the radial coordinate in lens units and k is the conic constant. The conic constant is less than -1 for hyperbolas, -1 for parabolas, between -1 and 0 for ellipses, 0 for spheres, and greater than 0 for oblate ellipsoids. For more information on conic constants, see "REFERENCES ON LENS DESIGN" on page 37. The Standard surface does not use any of the parameter values.

Modeling an ellipse with the Standard surface

There are a few handy formulas for converting the semi major and semi minor axis lengths of an elliptical surface to a radius and conic description. If "a" is the semi major axis length, and "b" is the semi minor axis length, then

$$\frac{1}{c} = R = \pm \frac{b^2}{a},$$

$$k = -\varepsilon^2 = -\left[\frac{a^2 - b^2}{a^2}\right].$$

Modeling an axicon with the Standard surface

The Standard surface can be used to make an almost perfect axicon. If $(1+k)c^2r^2 \gg 1$, the Standard surface reduces to

$$z = \frac{r}{\sqrt{-(1+k)}} \text{ or } z = r \tan \alpha, \text{ where } \tan \alpha = \frac{1}{\sqrt{-(1+k)}},$$

and α is the axicon angle, measured from the XY plane to the axicon surface. To create an axicon, calculate the conic constant value (k) from the angle (α) desired, and use any small value for the radius of curvature. The resulting value of k must be negative. The exact value of the radius or curvature is not important, as long as it is roughly three or more orders of magnitude *smaller* than the radial aperture of the axicon. The axicon is not perfect in the sense that there is no cusp at the origin; the region around the surface vertex will be rounded off over a size given approximately by the radius value. This is actually a desirable property for ray tracing, as the surface is everywhere smooth. However, paraxial data, such as effective focal length, magnification, and other common first-order optical properties are not generally meaningful for axicon surfaces. Even though the surface shape is still a conic asphere, the shape of the axicon implies the paraxial properties of the system are not representative of the optical surface as a whole. This is the case for any optical surface that is not well described by the radius of curvature alone.

Superconic

The most common form of a polynomial aspheric surface uses a power series expansion in the radial coordinate r to define the surface sag, where r is defined by

$$r^2 = x^2 + y^2.$$

For example, the Even Aspheric surface described in "Even Asphere" on page 298 uses such an expansion. Since r does not depend upon z , the expansion term is the distance from the vertex to the point on the surface as projected on to the tangent plane. Generally, the departure of the asphere from the tangent plane increases with radial aperture. As the departure increases, the power series expansion parameter r corresponds to a point on the tangent plane which is farther from the point on the surface. This causes the expansion to have poor convergence.

A novel solution proposed by Alan Greynolds of Breault Research Organization is to instead expand in powers of the distance from the vertex to the point on the surface. The expansion is then in terms of

$$s^2 = x^2 + y^2 + z^2.$$

Starting with the conic equation for a surface

$$kz^2 - 2Rz + x^2 + y^2 + z^2 = 0,$$

where k is the conic constant and R is the radius of curvature, a general power series expansion can be made of the form

$$Az^2 - 2Bz + C = 0.$$

The constants are defined as

$$A = \frac{k}{R},$$

$$B = 1 + U_1s^2 + U_2s^4 + \dots, \text{ and}$$

$$C = \frac{s^2}{R} + V_1s^4 + V_2s^6 + \dots,$$

where U and V are coefficients which define the aspheric shape. Note that if all the U and V terms are zero, a standard conic results. If A is also zero, then the superconic becomes a sphere. The coefficients A, U_1 , and V_1 together form a Cartesian oval. These properties make the superconic stable when optimizing for the coefficients. The superconic can be used to model surfaces which otherwise would require aspheric terms of very high order. ZEMAX models superconics with up to 240 terms, in practice designs rarely use more than 5 terms.

EXTRA DATA DEFINITIONS FOR SUPERCONIC SURFACES

Extra Data Number	Description
1	Maximum term number. The maximum is 250, but 4-10 is typical.
2	U_1
3	V_1
even n	$U_{n/2}$
odd n	$V_{(n-1)/2}$

Tilted

The tilted surface is a plane that makes an angle with respect to the x and y axes. The surface is easily defined in terms of the tangent angle between the plane and the X and Y axes:

$$z = x \tan \theta_x + y \tan \theta_y.$$

The tilted surface uses the first two parameters to define the tangents of the x and y angles. This surface is very useful for implementing tilted object and image surfaces, as well as tilted faces on prisms. The tilted surface does not change the local coordinate system, and so should not be used for implementing fold mirrors; use the coordinate break surface instead.

PARAMETER DEFINITIONS FOR TILTED SURFACES

Parameter 1	Parameter 2
$\tan \theta_x$	$\tan \theta_y$

Toroidal

Toroidal surfaces are formed by defining a curve in the Y-Z plane, and then rotating this curve about an axis parallel to the Y axis and intersecting the Z axis. Toroids are defined using a base radius of curvature in the Y-Z plane, as well as a conic constant and polynomial aspheric coefficients. The curve in the Y-Z plane is defined by:

$$z = \frac{cy^2}{1 + \sqrt{1 - (1+k)c^2y^2}} + \alpha_1y^2 + \alpha_2y^4 + \alpha_3y^6 + \alpha_4y^8 + \alpha_5y^{10} + \alpha_6y^{12} + \alpha_7y^{14}.$$

This curve is similar to the even aspheric surface sag formula, except the sixteenth order term has been omitted, and the coordinate argument is y, not r. This curve is then rotated about an axis a distance R from the vertex. This distance R is referred to as the radius of rotation, and may be positive or negative. The Y-Z radius of curvature is specified in the same column on the spreadsheet editor as the radius for Standard surfaces. The radius of rotation is set on parameter column number 1. To model a cylinder lens which is flat in the X direction use zero, which ZEMAX interprets as infinite radius.

The surface may optionally be modified by the addition of the Zernike Standard sag terms:

$$z = z_t + \sum_{i=1}^N A_i Z_i(\rho, \varphi),$$

where z_t is the base toroidal sag and the Zernike Standard terms are defined in the section "Zernike Standard Sag" on page 343. The Zernike Standard terms are defined on the Extra Data Editor.

Note if the Y-Z radius is set to infinity, a surface with power in x but not in y can be described, therefore, the cylinder may be oriented in either direction. The other parameter columns are used for the optional aspheric coefficients, as specified in the following table. If aspheric coefficients are required in the X direction, then rotate the toroid with a pair of coordinate break surfaces and rotate about Z. If different aspheric surfaces are required in both the X and Y directions, see the "biconic", "polynomial" and "extended polynomial" surfaces described elsewhere in this chapter.

PARAMETER DEFINITIONS FOR TOROIDAL SURFACES

Parameter 0	Parameter 1	Parameter 2-8
Extrapolate	Radius of Rotation	$\alpha_1 - \alpha_7$

EXTRA DATA DEFINITIONS FOR TOROIDAL SURFACES

Extra Data Number	Description
1	Number of terms.
2	Normalization radius. Coordinates are normalized by this value.
3 - 233	Coefficients on Zernike polynomials 1 - 231, respectively, in lens units.

The "Number of terms" is used to specify the maximum Zernike polynomial term to be used in calculating the surface sag. This number is provided to speed the ray tracing calculation; terms beyond this number are ignored.

Zernike polynomials are orthogonal over the unit circle, and so the normalization radius should be set to the radius over which the coefficient data was normalized. Zernike polynomials tend to diverge quite rapidly beyond the normalization radius, and so care should be taken that rays do not strike the surface beyond this radius. Although the ray tracing algorithm may work, the data may be inaccurate. The extrapolate flag may be set to zero to ignore the Zernike terms for rays that land outside the normalization radius.

Toroidal Grating

Toroidal grating surfaces are similar to regular toroidal surfaces, except no aspheric sag terms are supported, and a diffraction grating may be placed on the toroidal surface. Toroidal gratings are described by defining a curve in the Y-Z plane, and then rotating this curve about an axis parallel to the Y axis and intersecting the Z axis. Toroidal gratings are defined using a base radius of curvature in the Y-Z plane, as well as a conic constant. The curve in the Y-Z plane is defined by:

$$z = \frac{cy^2}{1 + \sqrt{1 - (1+k)c^2y^2}}$$

This curve is similar to the Standard surface sag formula, except the coordinate argument is y , not r . This curve is then rotated about an axis a distance R from the vertex. This distance R is referred to as the radius of rotation, and may be positive or negative. The Y-Z radius of curvature is specified in the same column on the spreadsheet editor as the radius for Standard surfaces. The radius of rotation is set on parameter column number 1. To model a cylinder lens which is flat in the X direction use zero, which ZEMAX interprets as infinite radius.

Note if the Y-Z radius is set to infinity, a surface with power in x but not in y can be described, therefore, the cylinder may be oriented in either direction.

The diffraction grating is defined in terms of the number of lines per micrometer and the diffraction order. These values are specified in parameter columns 2 and 3, respectively. The grating lines are parallel to the local x axis, and are evenly spaced when projected onto a plane.

For a toroidal grating with aspheric deformation terms see "Extended Toroidal Grating" on page 302.

PARAMETER DEFINITIONS FOR TOROIDAL GRATING SURFACES

Parameter 1	Parameter 2	Parameter 3
Radius of Rotation	Grating Lines per micrometer	Diffraction order

Toroidal Hologram

The Toroidal Hologram surface shape is identical to the Toroidal surface described in "Toroidal" on page 328. The curve in the Y-Z plane is defined by:

$$z = \frac{cy^2}{1 + \sqrt{1 - (1+k)c^2y^2}} + \alpha_1y^2 + \alpha_2y^4 + \alpha_3y^6 + \alpha_4y^8 + \alpha_5y^{10} + \alpha_6y^{12} + \alpha_7y^{14}$$

The radius and conic values, along with the parameter values, are used to define the surface shape. The extra data values are used to define the holographic properties of the surface. The terms used to define the holographic construction beams are identical to those described in the earlier sections on the Hologram 1 and Hologram 2 surface types. Extra data values 1-6 are the x, y, and z construction points for sources 1 and 2, value 7 is the construction wavelength, and value 8 is the playback order M.

The hologram is defined by the interference between the two defined construction beams, with no aberrations assumed on the construction beams. Optically fabricated holograms with aberrated construction beams may be modeled in a very general way using the "Optically Fabricated Hologram" on page 319.



Optically fabricated holograms with aberrated construction beams may be modeled using the "Optically Fabricated Hologram" surface.

There is also a ninth extra data value which is used as a "flag" to indicate if the two construction beams are converging or diverging. If both beams are converging or diverging to or from the construction points, then the flag value should be +1. If one beam is converging and one is diverging, then the flag should be -1. Any value less than zero is the same as -1, and any value greater than zero is the same as +1. This flag value serves to distinguish these two cases, which is the only difference between the Hologram 1 and 2 surface types.

PARAMETER DEFINITIONS FOR TOROIDAL HOLOGRAM SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6	Parameter 7	Parameter 8
Radius of Rotation	α_1	α_2	α_3	α_4	α_5	α_6	α_7

EXTRA DATA DEFINITIONS FOR TOROIDAL HOLOGRAM SURFACES

Extra Data Number	Description
1	Construction coordinate X1 for first construction beam.
2	Construction coordinate Y1 for first construction beam.
3	Construction coordinate Z1 for first construction beam.
4	Construction coordinate X2 for second construction beam.
5	Construction coordinate Y2 for second construction beam.
6	Construction coordinate Z2 for second construction beam.
7	Construction wavelength.
8	Playback order.
9	Converge/Diverge flag.

Toroidal NURBS



This feature is only available in the EE edition of ZEMAX.

The toroidal NURBS surface is similar to the radial NURBS surface (page 324). The key difference is that instead of rotating the curve about the Z axis to form a surface; the curve is instead mirror-imaged to the -Y axis, then rotated about an offset Y axis to form a toric shape.

The offset axis may be in the +Z or -Z direction, depending upon the sign of the radius of rotation. If the radius of rotation is zero, then the radius is assumed to be infinite and the surface has cylindrical symmetry. The X width

of the surface is determined by either the range of rotation angle or by direct specification of the -x and +x limits; the latter being used if the radius of rotation is infinite (which is defined by a zero value).

PARAMETER DEFINITIONS FOR TOROIDAL NURBS SURFACES

Parameter 1	Parameter 2	Parameter 3
Radius of Rotation	Min X or angle	Max X or angle

EXTRA DATA DEFINITIONS FOR TOROIDAL NURBS SURFACES

Extra Data Number	Description
1	Number of control points. At least 4 points are required, and no more than 60 allowed.
2	y coordinate for control point 1.
3	z coordinate (sag) for control point 1.
4	w (weight) value for control point 1.
3n-1, 3n, 3n+1	y, z, and w for control point n.

User Defined



This feature is only available in the EE edition of ZEMAX.

The user defined surface (UDS) is a powerful, flexible, and fast way of implementing surfaces not already built into ZEMAX. A UDS may be any shape, have any refractive, reflective, or diffractive properties, may impart any arbitrary phase to the beam, and may be followed by either homogeneous media or by a gradient index media of arbitrary form. A UDS may also apodize or attenuate the beam arbitrarily, or be used to define any electric field or coating data for polarization analysis. This latter capability allows the beam to be partially transmitted according to an arbitrary formula or table as defined by the user at any surface in the optical system.

The secret to this great flexibility is that all the properties of the surface are defined by the user in a separate C or C++ program compiled and linked into ZEMAX dynamically using the Windows Dynamic Link Library (DLL) capability. The DLL must contain functions which return to ZEMAX all the data required to draw the surface, trace rays, compute refraction angles, and for gradient index, to determine the index as a function of position within the media following the surface.

Because the DLL is compiled code, and because ZEMAX passes to the DLL a pointer to the data that needs to be computed, UDS are very fast; nearly as fast as native ZEMAX code (there is a small amount of overhead due to the function call).

The power of the UDS comes at a price, although it is a reasonable one. Use of the UDS does require that the user have a suitable compiler or development tool that can generate 32 bit Windows compatible DLLs. It is also assumed that the user can write the required code, and most importantly, ensure that the code is reliable and bug-free. To maximize speed, ZEMAX performs very little error checking on data returned by the DLL, and so buggy UDS DLLs are quite capable of bringing ZEMAX to a crash.

For this reason, technical support on the implementation of UDS is strictly limited to illustrating that the provided sample files work correctly. If you need a UDS DLL, and do not possess the desire or ability to write them yourself, please feel free to contact ZEMAX Technical Support for a quote on developing a custom DLL to meet your requirements. We have considerable experience in developing ray tracing algorithms, and can generally write DLLs at very competitive rates on very short notice.

The UDS DLL

The best way to learn how to create a UDS DLL is to start by copying the source code file (the one that ends in .C) of a sample DLL to a new file name, edit the sample to suit your requirements, and then recompile to make a new DLL. A good portion of the DLL source code is "boiler plate" which is common to all DLLs.

There are two files that are required to define a UDS DLL: a C (or C++) source code file such as MY_SURF.C, and a header file called USERSURF.H. Only the C file needs to be modified. The C file defines 2 functions; the first is "DLLMain" which is used to (optionally) initialize data used by the DLL. The other function is "UserDefinedSurface3", and this is the function that is modified for each surface type. Earlier versions of the ZEMAX User Defined surface interface used the function names UserDefinedSurface and UserDefinedSurface2; these names are still supported but are obsolete. New DLL files should use UserDefinedSurface3.

ZEMAX passes two structures to the function, and these structures are defined in the USERSURF.H file. The USERSURF.H file defines a data structure used by UserDefinedSurface3 called "FIXED_DATA3". Inside the UserDefinedSurface3 function is a C "switch-case" construct. ZEMAX passes a "type" parameter, which is a number between 0 and 10, inclusive, that indicates what data ZEMAX needs computed and returned to ZEMAX. The type codes are described in the following table:

UDS DLL Type Codes

TYPE	What ZEMAX wants the DLL to compute
0	The surface name, radial symmetry status, and grin status
1	The name of the parameter columns
2	The name of the extra data columns
3	The sag and alternate sag of the surface
4	A paraxial ray trace to and refraction through the surface
5	A real ray trace to and refraction through the surface, including transmittance and polarization data if any
6	The index and first derivatives for gradient index propagation
7	The default data when the user first selects the surface type
8	The initialization of the DLL, if required. This includes allocating static memory the DLL will require, loading data files, or other one-time calculations or initializations. This type code is called once each time the DLL is loaded. However, ZEMAX may load the same DLL multiple times. For example, each analysis window in ZEMAX gets it's own copy of the lens data, and each copy will load and initialize the DLL every time the analysis window is updated. ZEMAX may update multiple analysis windows at the same time (parallel execution) and so multiple DLLs may be running simultaneously. For this reason, using fixed file names to store temporary data or to make log files is not recommended; the same file will be created or written to multiple times as ZEMAX loads the DLL multiple times. The DLL only needs to handle this type code if the DLL requires one-time initialization before the DLL can trace rays.
9	The termination of the DLL, if required. This includes releasing allocated memory or closing data files. This type value is called when the DLL is removed from the surface or the lens is closed. See the comments for type code 8 above for details on how ZEMAX may load, initialize, and terminate multiple DLLs at once. The DLL only needs to handle this type code if the DLL must release memory or perform some other "clean-up" before terminating.
10	The scaling of parameter and extra data values used by the DLL. The scale factor is stored in the USER_DATA structure in the path argument.

Extensive comments are provided in the sample DLL source code files. The DLL files can be found in the folder <program>\DLL\Surfaces, and any new DLL files must be placed there as well.

Refractive and reflective UDS DLLs

For conventional homogeneous surfaces which are reflective or refractive, but not gradient or diffractive, start with the file US_STAND.C. That file is the source code for US_STAND.DLL, which is a clone of the built-in ZEMAX Standard surface. It is not exactly the same code, and is somewhat slower than the form of equations used by

ZEMAX, however, it is functionally equivalent. US_STAND.C includes a generic Snell's' law refraction function which also works for reflection. The DLL, like all UDS DLLs, provides the sag, ray intercept equations, surface normals, optical path, and paraxial refraction functionality.

Gradient index UDS DLLs

Gradient surfaces should use US_GRIN1.C as a starting point. The code is very similar to US_STAND.C, except a flag is returned indicating that the surface is a gradient type, and data type "6" is implemented to return the index of refraction given X, Y, Z, and all the parameter and extra data. The first derivative of the index in the X, Y, and Z directions must also be provided.

For gradient index UDS DLLs, the index of refraction used for paraxial ray tracing is the index at the front vertex of the surface, at the coordinates (0, 0, 0), as long as the glass column is left blank. If a glass type is specified for the DLL (in the usual glass column of the LDE) then the paraxial reference index is computed from the glass catalog data for that material. It is not required, or even advisable, to specify a glass name when using gradient index DLLs. The only time a glass name should be provided is when the DLL computes an index offset based upon the catalog glass properties.

Diffractive UDS DLLs

Diffractive optics are very similar to standard optics, except the rays are further deviated by the derivative of the phase as a function of X and Y, and the optical path length needs to be modified to include the phase change.

Generally, "refraction" for diffractive optics is defined by:

$$l' = l + \frac{\lambda}{2\pi} \frac{\partial \varphi}{\partial x}$$

$$m' = m + \frac{\lambda}{2\pi} \frac{\partial \varphi}{\partial y}$$

where l and m are the direction cosines, and the z direction cosine, n, is computed to make the magnitude of the direction vector unity, and φ is the phase in radians.

A sample DLL, US_GRATE.DLL, illustrates the diffractive computation by cloning the ZEMAX grating surface.

Lenslet arrays using DLLs

Lenslet arrays are easily modeled using the user defined surface. Basically, the ray trace determines which segment of the array is struck, then uses the local lens curvature to determine the refraction. The sample source code and DLL are provided as US_ARRAY.C and US_ARRAY.DLL, respectively.

User defined surface apodization using DLLs

The real ray trace portion of the DLL, type code 5, allows definition of the surface transmittance. The surface transmittance must be a number between 0.0 and 1.0, and indicates the relative fraction of intensity that the ray transmits through the surface. In this context, transmit means "continues on", so for reflective surfaces, the transmitted portion is that which normally reflects to the next surface.

The surface transmittance can be used to define arbitrary surface apodizations. The transmittance function can be any formula based upon the ray coordinates, direction cosines, surface parameters, or other data; or may be derived from a look up table, or any other method that can be implemented within the DLL.

The surface transmittance need not be defined. If it is not defined by the DLL, ZEMAX assumes it is 1.0. Whatever the surface transmittance is, the ray intensity will still be modified by ZEMAX to account for the Fresnel surface, thin film, and bulk absorption affects that are normally accounted for when doing polarization ray tracing. If ZEMAX is not doing polarization ray tracing, then the DLL defined transmittance is the only attenuating effect.

Unlike pupil apodization, surface apodizations defined using UDS may be placed on any surface anywhere in the optical system. Note that for apodizations which are a function of ray position (and most are) different fields will then see a different effective apodization. Note that this technique can be used to model arbitrary neutral density filters, or filters whose transmission are wavelength dependent.

Polarization and coating data using DLLs

The real ray trace portion of the DLL, type code 5, allows definition of either the transmitted or reflected electric field directly, or the definition of the s and p orientation complex reflection and transmission coefficients. The field or coating data can be based upon any data available within the DLL, including ray cosines, normal vectors, index, or other user defined data, or any other method that can be implemented within the DLL. The field and coating data need not be defined. If it is not defined by the DLL, ZEMAX uses the default algorithm for the surface. For a source code example of how to define the electric field or coating data see the sample DLL US_POLARIZATION.

Error handling and UDS

The convention ZEMAX uses internally is to return a zero value if the DLL computed a meaningful result, and no error occurred. Otherwise, the DLL should return -1. The exception is when ray tracing, either paraxial or real. If the ray misses the surface, it should return the surface number. If the ray total internally reflects (TIR's) then the return value should be the negative of the surface number. ZEMAX uses these error codes to help provide meaningful diagnostics to both the user and various ZEMAX features.

Sample DLLs

Numerous sample UDS DLLs have been written and are provided as both a ready to use DLL and as C source code for study and modification. The easiest way to write a DLL is to find one most similar to the one you need, and copy and edit the C source code file as required. The following table lists the DLLs available and a brief description of each. Note that although the sample DLLs have been tested and are generally considered reliable, they are provided "as is".

SAMPLE UDS DLL'S

DLL name	Description
QBFS, QCON	These two DLLs are an implementation of the surface types described in "Shape specification for axially symmetric optical surfaces", G. W. Forbes, Optics Express Vol. 15(8) pp 5218-5226 (April 16, 2007).
US_ANAMR	Anamorphic asphere. This aspheric surface is defined by the following expression: $Z = \frac{CxX^2 + CyY^2}{1 + \sqrt{1 - (1 + Kx)(Cx^2X^2) - (1 + Ky)(Cy^2Y^2)}} + AR[(1 - AP)X^2 + (1 + AP)Y^2]^2 + BR[(1 - BP)X^2 + (1 + BP)Y^2]^3 + CR[(1 - CP)X^2 + (1 + CP)Y^2]^4 + DR[(1 - DP)X^2 + (1 + DP)Y^2]^5$
US_ANASP	Annular aspheric surface. This surface is described in the paper "Annular surfaces in annular field systems" by Jose Sasian, Opt. Eng. 36(12) 3401-3403 (December 1997). This surface is similar to an odd aspheric polynomial, with a radial offset in the polynomial expansion to better model surfaces used in annular field systems.
US_APGXY	This DLL is a clone of the ZEMAX "Standard" surface type, with the added capability to modify the surface transmission. This surface defines an apodization function of the form $T(x, y) = \text{Exp}\left(-2G_x\left(\frac{x}{R}\right)^2 - 2G_y\left(\frac{y}{R}\right)^2\right),$ <p>where R is the semi-diameter of the surface and x and y are the coordinates at which the ray intercepts the surface.</p>

DLL name	Description
US_ARRAY	Models an n x m array of lenses, each of dimension H x W; all 4 parameters may be specified. The lenses are rectangular in outline but are "standard" in shape; that is, plane, spherical, or conic asphere. The number of lenses in each direction must be odd.
US_ARRAYEVEN	Models an n x m array of lenses, each of dimension H x W; all 4 parameters may be specified. The lenses are rectangular in outline but are "even aspheric" in shape; that is, plane, spherical, conic asphere, and radial polynomial asphere. The number of lenses in each direction must be odd.
US_CYLAR	Models a vertical array of n cylindrical lenses, each H high. The lenses are cylindrical or conic aspheres in the YZ plane.
US_DGCYL	Models a diffraction grating on a cylindrical surface, where the grating lines are equally spaced along the arc of the surface rather than equally spaced along the y coordinate of the tangent plane. The grating lines are parallel to the axis of the cylinder.
US_EAPERIODIC	<p>This surface is similar to the Even Asphere, with an additional term allowing for a sinusoidal variation of the sag. Such variation may be useful for modeling surface irregularities. The full sag expression for the surface is given by:</p> $z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=1}^8 \alpha_i \rho^{2i} - \frac{A}{2} [\cos(2\pi\omega_0 r + \phi_0) - \cos(\phi_0)]$ <p>where the parameters c, k, and α_i are the same as for the Even Asphere surface. The parameter A represents the magnitude of the sinusoidal sag perturbation, and should be specified in lens units. The parameter ω_0 represents the frequency of the sinusoidal sag perturbation, and should be specified in inverse lens units. The parameter ϕ_0 is the initial phase of the sinusoidal sag perturbation, and should be specified in degrees (the value is converted to radians inside of the DLL for calculation of the sag).</p>
US_FILT1	This surface defines an apodization function of the form $T = 1.0 - Exp(-G\rho^2)$ where p is the radial coordinate normalized to the semi-diameter of the surface.
US_FILT2	This surface defines an apodization function where the optical density is Dmax from zero to a normalized radial coordinate p1, then linearly goes to D=0 at p2, and D= 0 thereafter. The density is $D = -Log_{10}(T)$ where T is the transmission.
US_FILT3	This surface defines an apodization function of the form $D = 0.5D_{max}(1.0 + \cos\pi\rho)$ where p is the radial coordinate normalized to the semi-diameter of the surface.

DLL name	Description
US_FILT4	<p>This surface has an apodization function defined by:</p> $T(r) = 1 \text{ if } r < R - \Delta R,$ $T(r) = 0 \text{ if } r > R + \Delta R,$ $T(r) = \frac{1 + \cos\left(\frac{\pi(r - (R - \Delta R))}{2\Delta R}\right)}{2} \text{ if } R - \Delta R < r < R + \Delta R,$ <p>where r is the radial coordinate, R is the aperture radius, and Δ is a dimensionless fraction of the aperture radius. The transmission varies smoothly from 1 to 0 over twice the distance defined by ΔR. The value for Δ is defined as 0.001 if the provided value for Δ is zero. This function makes a convenient "soft-edged" aperture which is useful for optimization of transmission through a circular aperture.</p>
US_FILT5	<p>This surface has an apodization function defined by:</p> $T(na) = 1.0 \text{ if } na < NA - \Delta,$ $T(na) = 0.0 \text{ if } na > NA + \Delta, \text{ otherwise}$ $T(na) = 0.5 \left(1 + \cos\left(\frac{(na - (NA - \Delta NA))\pi}{(2\Delta NA)}\right) \right),$ <p>where na is the numerical aperture of the incoming ray, NA is a user defined parameter, and Δ is fixed at 0.001 (although this value may be edited in the DLL source code). This type of filter demonstrates apodization of a ray based upon its incident angle, rather than position. Rays whose na exceeds the $NA - \Delta$ value will be attenuated smoothly to zero. This permits easier, more continuous optimization.</p>
US_FILT6	<p>This surface has an apodization function defined by:</p> $T(x, y) = (1/4)(1 + \sin(2\pi\alpha x))(1 + \sin(2\pi\beta y)),$ <p>where α and β are the spatial frequencies of the sinusoidal transmission in x and y, respectively.</p>
US_FILT7	<p>This surface apodizes the surface transmission via linear interpolation of values specified on the extra data editor as a function of either the radial coordinate on the surface or as a function of the angle the ray makes with respect to the local Z axis after refraction from the surface.</p>
US_FILT8	<p>This surface models a soft edged rectangular aperture. The MaxX and MaxY parameters are the X and Y half widths of the rectangular aperture. The transmission varies smoothly from 1 to 0 over twice the distance defined by the DeIX and DeIY distances from the aperture boundaries. This function makes a convenient "soft-edged" aperture which is useful for optimization of transmission through a rectangular aperture.</p>

DLL name	Description
US_GCYL	Models a cylindrical surface with a gradient index medium. The GRIN medium has concentric cylindrical shells of constant index, and a polynomial GRIN profile centered on the center of curvature. This DLL was intended to model light transversing a fiber perpendicular to the axis of the fiber (through the side).
US_GRATE	Models a Standard surface with grating lines parallel to the x axis, equally spaced along y in the tangent plane. This is a clone of the ZEMAX "Diffraction Grating" surface. The purpose is to show how diffractive surfaces are modeled in ZEMAX UDS DLLs.
US_GRIN1	Models a quadratic GRIN medium. This is a clone of the ZEMAX "Gradient 1" surface. The purpose is to show how GRIN surfaces are modeled in ZEMAX UDS DLLs.
US_IGRIN	An example of a GRIN with dispersion.
US_ITERA	This DLL is a clone of the ZEMAX "Standard" surface, except it uses "dumb" iteration to find the intercept point rather than a closed form expression. The point of this example is to show how to find the surface intercept point if only the sag expression is known. Most polynomial aspheres require this type of iteration because the ray-surface intercept formulas cannot be determined in closed form.
US_LUNE	Models a perfect Luneberg lens GRIN surface.

DLL name	Description
US_MEMS	<p>Models a Micro-Electromechanical System (MEMS), such as a Digital Mirror Device (DMD). The MEMS consists of a 2D array of small rectangular mirrors. The mirrors may tip at any of three angles, each rotated about an axis to point the mirror in any direction. The mirrors may be set to any of the three states by addressing rows, columns, or by individual mirrors if desired to model any state the MEMS can be in. Although the model is strictly geometric, it effectively models where rays are reflected by such a device for an arbitrary setting of the individual mirrors. The parameters on the model are:</p> <p>Nx: Number of X direction mirrors Ny: Number of Y direction mirrors Wx: Total width in lens units in X direction Wy: Total width in lens units in Y direction A0, A1, A2: Angle of mirror in state 0, 1, or 2 Rotation Angle: Angle of rotation about the Z axis to tip the mirror Extra Data Value 1: If 0, mirrors are addressed by rows, if 1, by columns, if 2, by individual mirrors. Extra Data Value 2: State of rows/columns/mirrors 1 - 15 Extra Data Value 3: State of rows/columns/mirrors 16 - 30 Extra Data Value n: State of rows/columns/mirrors 1+15*(n-2) - 15*(n-1)...</p> <p>The rotation angle effectively rotates the plane of tip of the mirror around the local Z axis; with the initial tip plane being around the local x axis. The rotation angle is measured clockwise from the +y axis. The mirror angles are then tipped about the rotated tip direction.</p> <p>The Extra Data Values (EDV) are used to define the state of the rows/columns/mirrors using a base 3 integer value. To determine the values for any logic state of the MEMS, construct a table similar to the one below, which shows the values for 3 rows or columns or mirrors:</p> <pre>r/c/m: 3 2 1 EDV 0 0 0 0 0 0 1 1 0 0 2 2 0 1 0 3 0 1 1 4 0 1 2 5 0 2 0 6 etc..</pre> <p>Note the EDVA is given by: $EDV = M1*(3^0) + M2*(3^1) + M3*(3^2) + \dots$ where M1 is the logic state (0, 1, or 2) of the first row/column/mirror and M2 is the logic state of the second row/column/mirror, etc. Up to 15 row/column/mirror values are defined by each EDV.</p>
US_MULTI_ZONE_AS PHÈRE	<p>This surface models a series of annular even aspheric zones. Each zone has a radius, conic, even asphere coefficients to order 12, and a maximum radial aperture for the zone. Up to 20 concentric zones are supported. Each zone is shifted so the surface is continuous across the zone boundaries.</p>
US_OFFST	<p>The offset surface. This surface simulates an additional propagation of distance that is hidden within the surface. The propagation distance is independent for each wavelength. This allows modeling of systems that have independent focusing of each color channel, such as a digital projector.</p>

DLL name	Description
US_OGIVE	<p>Ogive surface shape. An ogive is identical to the Standard surface, except the axis of rotation of the surface is offset by an amount r_o. The surface sag is given by</p> $z = \frac{cr_g^2}{1 + \sqrt{1 - (1+k)c^2r_g^2}}, \text{ where}$ $r_g = r_o + \sqrt{x^2 + y^2}.$
US_POLARIZATION	This DLL shows how to define the electric field or coating data.
US_STAND	This DLL is a clone of the ZEMAX "Standard" surface type. This is the simplest DLL surface and is a good place to start the study of UDS DLLs.
US_STAND2	This DLL is a clone of the ZEMAX "Standard" surface type using the UserDefinedSurface2 function and structure definition.

Variable Line Space Grating

The variable line space grating surface models a special grating with straight, but unevenly spaced lines. The grating is described in M. Hettrick and S. Bowyer, "Variable line-space gratings: new designs for use in grazing incidence spectrometers", Applied Optics, Vol. 22, No. 24, p3921 (1983). The grating is designed to form images upon a concave spherical surface. A converging beam is used to illuminate the grating in reflection. The grating lines are parallel to the local x axis, and lie in the x-z plane. Only plane gratings are supported, and the feature has only been tested when the grating is used in reflection (the glass type must be "MIRROR").

Although the hologram surface type could be used to model the ideal grating, in practice holographic construction is not always possible. A compromise is made in the fabrication, where the grating lines are straight and parallel (as opposed to curved for the hologram case). This introduces aberrations into the beam, which this surface model accounts for.

Five parameters are used to describe the surface: M , the diffraction order upon playback; L , the (positive) radius of the concave focal surface, $\cos(\alpha_0)$, the cosine of the angle the incident beam makes with the plane of the grating, $\cos(\beta_0)$, the cosine of the angle the reflected beam makes with the plane of the grating, and λ_0 , the construction wavelength. This is the wavelength at which the grating will reflect the axis ray along the angle β_0 .

The grating frequency in lines per micrometer is given by

$$T = \frac{1}{\lambda_0} \left[\frac{L \cos \beta_0 - y}{\sqrt{y^2 + L^2 - 2Ly \cos \beta_0}} - \frac{L \cos \alpha_0 - y}{\sqrt{y^2 + L^2 - 2Ly \cos \alpha_0}} \right],$$

where y varies across the grating surface. The variable line space grating surface uses five of the parameter surfaces, but none of the extra data surfaces.

PARAMETER DEFINITIONS FOR VARIABLE LINE SPACE GRATING SURFACES

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5
M	L	$\cos \alpha_0$	$\cos \beta_0$	λ_0

Zernike Fringe Phase

The Zernike Fringe Phase surface has a substrate shape identical to the Standard surface (which supports planes, spheres, and conics) plus additional phase terms defined by the Zernike Fringe coefficients. The surface sag is identical to the Standard surface formula. The additional phase terms deviate and add optical path to the rays as they cross the surface. This surface model is well suited to modeling system aberrations for which measured interferometer data is available. The Zernike Fringe phase surface can also be used to model some holograms and binary optics surfaces. The phase of the surface is given by

$$\Phi = M \sum_{i=1}^N 2\pi A_i Z_i(\rho, \varphi)$$

where N is the number of Zernike coefficients in the series, A_i is the coefficient on the i^{th} Zernike Fringe polynomial, ρ is the normalized radial ray coordinate, φ is the angular ray coordinate, and M is the diffraction order. The Zernike Fringe polynomials are defined in the table given in "Zernike Fringe Coefficients" on page 196. ZEMAX supports up to 37 terms. The coefficients A_i all have units of waves. One wave is 2π radians. Note that the phase of the surface is independent of the wavelength; ZEMAX accounts for the deviation in the ray path based upon the wavelength being traced. If the "Extrapolate" flag is set to 0, the Zernike terms are ignored outside of the normalization radius. If the "Extrapolate" flag is set to 1, then the Zernike terms are considered no matter where the ray lands on the surface; even if the ray lands beyond the normalization radius.

Note that the Zernike Fringe Phase surface describes phase variations (or wavefront error), not surface deformations directly. If you have Zernike coefficient data in terms of surface deformations, as might have been measured with a profilometer, see "Zernike Fringe Sag" on page 341. See also "Zernike Standard Phase" on page 341.

PARAMETER DEFINITIONS FOR ZERNIKE FRINGE PHASE SURFACES

Parameter 0	Parameter 1
Diffraction order	Extrapolate

EXTRA DATA DEFINITIONS FOR ZERNIKE FRINGE PHASE SURFACES

Extra Data Number	Description
1	Number of terms.
2	Normalization radius. Coordinates are normalized by this value.
3 - 39	Coefficients on Zernike polynomials 1 - 37, respectively, in units of waves.
40 and above	Not used.

The "Number of terms" is used to specify the maximum Zernike polynomial term to be used in calculating the phase. This number is provided to speed the ray tracing calculation; terms beyond this number are ignored.

Zernike polynomials are orthogonal over the unit circle, and so the normalization radius should be set to the radius over which the coefficient data was normalized. Zernike polynomials tend to diverge quite rapidly beyond the normalization radius, and so care should be taken that rays do not strike the surface beyond this radius. Although the ray tracing algorithm may work, the data may be inaccurate. The extrapolate flag may be set to zero to ignore the Zernike terms for rays that land outside the normalization radius.

Zernike Fringe Phase coefficients sign conventions

See "Binary Optic 1" on page 283 for a discussion of sign conventions.

Zernike Fringe Sag

The Zernike Fringe Sag surface is defined by the same polynomial as the Even Aspheric surface (which supports planes, spheres, conics, and polynomial aspheres) plus additional aspheric terms defined by the Zernike Fringe coefficients. The surface sag is of the form:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=1}^8 \alpha_i r^{2i} + \sum_{i=1}^N A_i Z_i(\rho, \varphi)$$

where N is the number of Zernike coefficients in the series, A_i is the coefficient on the i^{th} Zernike Fringe polynomial, r is the radial ray coordinate in lens units, ρ is the normalized radial ray coordinate, and φ is the angular ray coordinate. The Zernike terms may be decentered from the conic and aspherical terms using parameters 9 and 10. The Zernike Fringe polynomials are defined in the table "Zernike Fringe Coefficients" on page 196. ZEMAX supports up to 37 Zernike Fringe terms. The coefficients A_i all have units which are the same as the lens units, such as millimeters or inches. The coefficients α_i have units, and are as defined in the section "Even Asphere" on page 298. If the "Extrapolate" flag is set to 0, the Zernike terms are ignored outside of the normalization radius. If the "Extrapolate" flag is set to 1, then the Zernike terms are considered no matter where the ray lands on the surface; even if the ray lands beyond the normalization radius.

Note that the Zernike Fringe Sag surface describes surface deformations, not wavefront error directly. If you have Zernike coefficient data in terms of waves of OPD, as might have been measured with an interferometer, see "Zernike Fringe Phase" on page 340. Also see "Zernike Standard Sag" on page 343.

PARAMETER DEFINITIONS FOR ZERNIKE FRINGE SAG SURFACES

Parameter 0	Parameters 1-8	Parameter 9	Parameter 10
Extrapolate	$\alpha_1 - \alpha_8$	Zernike decenter X in lens units	Zernike decenter Y in lens units

EXTRA DATA DEFINITIONS FOR ZERNIKE FRINGE SAG SURFACES

Extra Data Number	Description
1	Number of terms.
2	Normalization radius. Coordinates are normalized by this value.
3 - 39	Coefficients on Zernike polynomials 1 - 37, respectively, in lens units.
40 and above	Not used.

The "Number of terms" is used to specify the maximum Zernike polynomial term to be used in calculating the surface sag. This number is provided to speed the ray tracing calculation; terms beyond this number are ignored.

Zernike polynomials are orthogonal over the unit circle, and so the normalization radius should be set to the radius over which the coefficient data was normalized. Zernike polynomials tend to diverge quite rapidly beyond the normalization radius, and so care should be taken that rays do not strike the surface beyond this radius. Although the ray tracing algorithm may work, the data may be inaccurate. The extrapolate flag may be set to zero to ignore the Zernike terms for rays that land outside the normalization radius.

Zernike Standard Phase

The Zernike Standard phase surface has a substrate shape identical to the Standard surface (which supports planes, spheres, and conics) plus additional phase terms defined by the Zernike Standard coefficients. The surface sag is identical to the Standard surface formula. The additional phase terms deviate and add optical path

to the rays as they cross the surface. This surface model is well suited to modeling system aberrations for which measured interferometer data is available. The Zernike Standard Phase surface can also be used to model some holograms and binary optics surfaces. The phase of the surface is given by

$$\Phi = M \sum_{i=1}^N 2\pi A_i Z_i(\rho, \varphi)$$

where N is the number of Zernike coefficients in the series, A_i is the coefficient on the i^{th} Zernike Standard polynomial, ρ is the normalized radial ray coordinate, φ is the angular ray coordinate, and M is the diffraction order. The Zernike Standard polynomials are defined in the table given in "Zernike Standard Coefficients" on page 199. ZEMAX supports up to 231 terms. The coefficients A_i all have units of waves. One wave is 2π radians. Note that the phase of the surface is independent of the wavelength; ZEMAX accounts for the deviation in the ray path based upon the wavelength being traced. If the "Extrapolate" flag is set to 0, the Zernike terms are ignored outside of the normalization radius. If the "Extrapolate" flag is set to 1, then the Zernike terms are considered no matter where the ray lands on the surface; even if the ray lands beyond the normalization radius.

Note that the Zernike Standard Phase surface describes phase variations (or wavefront error), not surface deformations directly. If you have Zernike coefficient data in terms of surface deformations, as might have been measured with a profilometer, see "Zernike Standard Sag" on page 343.

PARAMETER DEFINITIONS FOR ZERNIKE STANDARD PHASE SURFACES

Parameter 0	Parameter 1
Diffraction order	Extrapolate

EXTRA DATA DEFINITIONS FOR ZERNIKE STANDARD PHASE SURFACES

Extra Data Number	Description
1	Number of terms.
2	Normalization radius. Coordinates are normalized by this value.
3 - 233	Coefficients on Zernike polynomials 1 - 231, respectively, in units of waves.
234 and above	Not used.

The "Number of terms" is used to specify the maximum Zernike polynomial term to be used in calculating the phase. This number is provided to speed the ray tracing calculation; terms beyond this number are ignored.

Zernike polynomials are orthogonal over the unit circle, and so the normalization radius should be set to the radius over which the coefficient data was normalized. Zernike polynomials tend to diverge quite rapidly beyond the normalization radius, and so care should be taken that rays do not strike the surface beyond this radius. Although the ray tracing algorithm may work, the data may be inaccurate. The extrapolate flag may be set to zero to ignore the Zernike terms for rays that land outside the normalization radius.

Zernike Standard Phase coefficients sign conventions

See "Binary Optic 1" on page 283 for a discussion of sign conventions.

Zernike Standard Sag

The Zernike Standard Sag surface is defined by the same polynomial as the Even Aspheric surface (which supports planes, spheres, conics, and polynomial aspheres) plus additional aspheric terms defined by the Zernike Standard coefficients. The surface sag is of the form:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=1}^8 \alpha_i r^{2i} + \sum_{i=1}^N A_i Z_i(\rho, \varphi)$$

where N is the number of Zernike coefficients in the series, A_i is the coefficient on the i^{th} Zernike Standard polynomial, r is the radial ray coordinate in lens units, ρ is the normalized radial ray coordinate, and φ is the angular ray coordinate. The Zernike terms may be decentered from the conic and aspherical terms using parameters 9 and 10. The Zernike Standard polynomials are defined in the table given in "Zernike Standard Coefficients" on page 199. ZEMAX supports the first 231 Zernike Standard terms. The coefficients A_i all have units which are the same as the lens units, such as millimeters or inches. The coefficients α_i have units, and are as defined in the section on the "Even Asphere" surface type. If the "Extrapolate" flag is set to 0, the Zernike terms are ignored outside of the normalization radius. If the "Extrapolate" flag is set to 1, then the Zernike terms are considered no matter where the ray lands on the surface; even if the ray lands beyond the normalization radius.

Note that the Zernike Standard Sag surface describes surface deformations, not wavefront error directly. If you have Zernike coefficient data in terms of waves of OPD, as might have been measured with an interferometer, use the Zernike Standard Phase surface instead.

PARAMETER DEFINITIONS FOR ZERNIKE STANDARD SAG SURFACES

Parameter 0	Parameters 1-8	Parameter 9	Parameter 10
Extrapolate	$\alpha_1 - \alpha_8$	Zernike decenter X in lens units	Zernike decenter Y in lens units

EXTRA DATA DEFINITIONS FOR ZERNIKE STANDARD SAG SURFACES

Extra Data Number	Description
1	Number of terms.
2	Normalization radius. Coordinates are normalized by this value.
3 - 233	Coefficients on Zernike polynomials 1 - 231, respectively, in lens units.

The "Number of terms" is used to specify the maximum Zernike polynomial term to be used in calculating the surface sag. This number is provided to speed the ray tracing calculation; terms beyond this number are ignored.

Zernike polynomials are orthogonal over the unit circle, and so the normalization radius should be set to the radius over which the coefficient data was normalized. Zernike polynomials tend to diverge quite rapidly beyond the normalization radius, and so care should be taken that rays do not strike the surface beyond this radius. Although the ray tracing algorithm may work, the data may be inaccurate. The extrapolate flag may be set to zero to ignore the Zernike terms for rays that land outside the normalization radius.

Zone Plate

This surface is used to model a Fresnel Zone Plate (FZP). A FZP consists of a refractive or reflective surface upon which annular zones of varying depth are cut or etched. Typically, the spacing between grooves is large compared to the wavelength, and therefore FZP's are entirely refractive, and not diffractive, devices. The varying

depth of the annular zones imparts a slight phase change to different zones of the beam, due to the change in path length through the material.

The FZP is modeled by altering the local depth of the surface. Edge effects, such as rays striking the side of the zone, and diffraction from the edges, are ignored. The FZP is defined by a normalized optical thickness of each zone. Each zone is spaced a fixed radius apart, and up to 240 zones may be defined. Three parameter values are required: the mode, the Δr , and the reference wavelength, λ_0 .

If the mode is zero, the steps are not interpolated, and each zone has a constant optical thickness. This results in discontinuous phase steps, like a staircase. If the mode is unity, then the optical thickness is interpolated between adjacent zones. Although the phase will be continuous, the first derivative of the phase will generally be discontinuous at the zone boundaries. The " Δr " is the step in radius between adjacent zones in lens units. The "reference wavelength" is the wavelength in micrometers used to normalize the optical thickness. For example, if the FZP defines the boundary between a material with index 1.55 and air, with index 1.00, and the reference wavelength is 0.45 micrometers, then an optical thickness of 1 wave corresponds to a groove depth of

$$T = \frac{\lambda_0}{n_2 - n_1} = \frac{0.45\mu}{0.55} = 0.818182\mu.$$

This thickness is the actual groove depth if the optical thickness is given to be 1.0. ZEMAX uses this normalized optical thickness for convenience. To specify a zone that introduces one-half wave of phase at the reference wavelength, enter an optical thickness of 0.5. At other wavelengths, the phase will be greater or less depending upon the dispersion of the two adjacent media.

The optical thickness is defined at n points, separated by Δr : $r_1, r_2, r_3, \dots, r_n$. The number of points, " n ", is specified in the first column of the Extra Data Editor. The remaining terms are defined in columns 2 through $n+1$ in the Extra Data Editor.

PARAMETER DEFINITIONS FOR ZONE PLATE SURFACES

Parameter 1	Parameter 2	Parameter 3
Mode	Δr	Ref. Wave

EXTRA DATA DEFINITIONS FOR ZONE PLATE SURFACES

Extra Data Number	Description
1	Number of points
2	Optical thickness at Δr
3	Optical thickness at $2 \Delta r$
4	Optical thickness at $3 \Delta r$
$n+1$	Optical thickness at $n \Delta r$

Introduction



This feature is only available in the EE edition of ZEMAX.

Most imaging systems are well described by optical surfaces which are sequential, which means that rays always trace from the object surface to surface 1, then 2, then 3, etc... in a strict sequence. Each ray "hits" each surface once and only once in this predetermined sequence. Rays in sequential mode never trace from surface 3 to 9 then back to 1, for example. The sequential model is simple, numerically fast, and extremely useful and complete for many important cases.

However, there are times when a non-sequential trace is required. Non-sequential means the rays trace in the actual physical order they hit various objects or surfaces, and not necessarily in the order the objects are listed in the software user interface. Note rays in a non-sequential trace may hit the same object repeatedly, and entirely miss other objects. Generally, the order in which objects are hit by rays depends upon the object geometry and the angle and position of the input ray. Objects which require or at least benefit from non-sequential ray tracing include prisms, light pipes, lens arrays, reflectors, and Fresnel lenses. Certain types of analysis, such as stray or scattered light effects, are only practical in a completely non-sequential environment.

Traditionally, lens design programs that supported surfaces (rather than 3D objects) for sequential ray tracing would implement non-sequential ray tracing using the same surface model; the rays would simply intersect surfaces in a possibly out of sequence order.

The disadvantage of using surfaces in a more general non-sequential approach is that surfaces do not adequately describe most optical components. For example, lenses not only have a front and back surface, they also have edges and perhaps flattened outer faces for mounting. Light may intercept, and then refract or reflect from these additional surfaces normally ignored by sequential surface ray tracing codes. Complex prisms, such as a dove or roof prism, contain many faces, and the rays may intersect these faces in a complex order depending upon the input angle and position of the ray.

To support these types of components in a very general and accurate way requires the use of full 3D solid object models instead of just 2D surfaces. ZEMAX calls this type of ray tracing non-sequential components, or NSC, which is different from non-sequential surfaces or NSS. NSC ray tracing in ZEMAX supports all of the following features:

- Definition and placement of multiple sources, objects and detectors.

- Real radiometric and photometric units; including watts, lumens, lux, phot, footcandles, and others.

- Automatic determination of ray-object intersection order.

- Automatic detection of reflection, refraction, and total internal reflection (TIR).

- Support for a very wide range of 3D objects, including diffractive optics.

- Polarization ray tracing and arbitrary thin film coatings.

- Statistical models for scattering, including Lambertian, Gaussian, and ABg.

- Automatic ray splitting for efficient analysis.

This chapter provides information on setting up a NSC group, defining objects, and tracing rays through NSC.

Paraxial data and ray tracing with NSC

There is no paraxial ray tracing in NSC. When paraxial rays are traced to a non-sequential surface, the equivalent real rays are traced instead. Nearly all paraxial data, such as focal length and F/#, are meaningless for non-sequential systems. Some NSC systems are well behaved in the sense that real rays close to the axis behave much as idealized paraxial rays would. However, NSC systems with obscurations or objects only partially in the path of the beam will not be able to correctly predict any paraxial data, and paraxial data should be considered meaningless for these systems.

Optimization with NSC

See “Optimizing objects in a non-sequential group with sequential rays” on page 531 and “Optimizing with sources and detectors in non-sequential mode” on page 531.

The two methods of using NSC ray tracing

ZEMAX supports two distinct methods of using NSC ray tracing:

Tracing rays through an NSC group that is part of an otherwise sequential system (NSC with ports).

Tracing rays through an NSC group that contains all the objects of interest (NSC without ports).

Although the method of defining and placing objects within a NSC group is the same for both methods, the details of how rays are launched, what analysis may be performed, how energy distributions are determined, and what types of systems are best modeled by each method, are considerably different. NSC with and NSC without ports are described below.

NSC with ports

When non-sequential object or objects are part of an otherwise sequential system, NSC with ports is generally the best approach to use. An example of this system would be a point or extended surface object represented by rays which follow a sequential path through one or more conventional lenses, then follow a non-sequential path through a prism or light pipe before illuminating the image surface.

This method requires the use of ports for rays to enter and leave the NSC group. Ports are described in detail in the section "Overview of NSC ray tracing with ports" below. When using ports, rays are launched from defined field positions on the object surface, and these rays ignore source and detector objects within the NSC group. Rays must leave the NSC group via the exit port; then continue through the remainder of the sequential system.

When using analysis features such as ray fans, spot diagrams, and MTF, only rays which enter and leave the NSC group through ports are considered. All the usual sequential ZEMAX system data, such as field positions and pupil sizes, determine the properties of rays entering the NSC group. All normal ZEMAX analysis, such as ray fans and spot diagrams, are still available (although the data may be meaningless depending upon the properties of the NSC system).

NSC without ports

Some systems have no sequential paths or portions at all; such as headlamp reflectors, complex light pipes, or general illumination systems. It is also possible to analyze ghost, stray, and scattered light properties of nominally sequential systems (such as a camera lens or telescope) by placing the entire system in a non-sequential group and performing non-sequential ray tracing on the entire model. These types of systems and analysis are better suited to NSC without ports.

This method requires the use of NSC sources to launch rays. When not using ports, rays are launched only from defined NSC sources, and these rays consider detector objects within the NSC group.

Since rays launched from sources within the NSC group cannot leave through the exit port; the only available analysis is the ray distribution and energy as determined using detector objects.

Combining NSC with and without ports

The 3D Layout, Wireframe, Solid Model, and Shaded Model analysis features are all capable of simultaneously displaying BOTH rays coming from the sequential entry port as well as any sources defined for use in NSC without ports. Rays entering from the entry port do not interact with any NSC detectors; and rays launched from NSC sources do not interact with the entry port, the exit port, or any optics defined outside the NSC group.

Therefore it is advised, though not required, that when using NSC without ports (i.e. you have defined sources among your NSC objects) that you set the number of rays in these layout features to zero. This will make for a less confusing display. There are times when showing the sequential rays coming from the entry port and non-sequential rays on the same plot is useful, especially when placing baffles in a stray light analysis.

Alternatively, there is a 3D Layout feature which is specific to NSC, which shows only objects and rays defined within a single NSC group.

Overview of NSC ray tracing with ports

Ray tracing through a group of NSC objects using ports is accomplished with the following basic steps:

- 1) A Non-Sequential Components surface is inserted in the Lens Data Editor. This surface becomes the entry port for the non-sequential group.
- 2) The Non-Sequential Components surface parameters are used to define the location of the exit port for the non-sequential group.
- 3) Objects are defined in a list associated with the Non-Sequential Components surface.
- 4) ZEMAX traces a ray sequentially to the entry port, then non-sequentially within the NSC group, until the ray strikes the exit port.
- 5) Rays entering the NSC group through the entry port cannot split.

The Non-Sequential Components surface has parameters which determine where the rays will exit the NSC group as described below.

The entry port

The Non-Sequential Components surface acts like a plane, sphere, or conic aspheric surface whose location is determined by the previous surfaces in the Lens Data Editor, in the usual way. The surface shape may be hyperhemispheric to allow acceptance of rays over a full 4π steradians. The Non-Sequential Components surface is the entry port into a group of objects which will be traced non-sequentially. The entry port is how rays get into the NSC group. No objects should be placed as to touch or surround the entrance port.

The exit port

There are 9 parameters used in the definition of the Non-Sequential Components surface:

Draw Ports?: If 0, no ports are drawn, if 1, draw entry, 2, draw exit, 3 draw both.

Exit Location X: The x position of the exit port relative to the entry port.

Exit Location Y: The y position of the exit port relative to the entry port.

Exit Location Z: The z position of the exit port relative to the entry port.

Exit Tilt About X: The rotation about the local X axis of the exit port.

Exit Tilt About Y: The rotation about the local Y axis of the exit port.

Exit Tilt About Z: The rotation about the local Z axis of the exit port.

Order: If the order flag is zero, then the above locations and tilts are done in the following order: decenter x, decenter y, decenter z, rotate around global z, rotate around global y, rotate around global x. If the flag is any value other than zero, then the order is reversed. This follows the same convention as the sequential coordinate break surface when using order flag = 0 (see page 293).

Reverse Rays: If this flag is 0, then ZEMAX assumes the non-sequential group acts like a refractive lens. If this flag is 1, then ZEMAX assumes the non-sequential group acts like a mirror. For example, if rays entering the non-sequential group travel in a positive direction with respect to the local z axis, and leave the exit port still traveling in the positive local z direction, the flag should be zero. If the rays reverse direction relative to the incoming direction, then the reverse rays flag should be 1.

The diameter of the circular exit port in lens units. This value is defined by the semi-diameter of the surface following the Non-Sequential Components surface. Note any additional aperture may be placed on the exit surface if another aperture shape is required.

These parameters define the location and size of the exit port relative to the entry port. If the exit port is located exactly at the entry port, then rays will immediately exit the non-sequential group without striking any objects. This generally means the exit location z parameter must not be zero.

The glass column of the Non-Sequential Components surface is also used to define the "background" material and index of refraction of the media in which NSC objects are placed. The surface after the Non-Sequential Components surface acts like a Standard plane surface oriented in the coordinate system after the decenters and tilts have been applied.

Note the exit port position is the same as the surface following the Non-Sequential Components surface, and its location in 3D space is determined by the parameters of the Non-Sequential Components surface. The thickness of the Non-Sequential Components surface is not used; only the location and tilt parameter values. The exit port should not be placed at the same location as, or within the glue distance ("Glue Distance In Lens Units" on page 114) of the entrance port; otherwise rays entering the NSC group will immediately strike the exit port and not intersect any of the objects within. No objects should be placed as to touch or surround the exit port.

Getting rays in

A ray leaves the object surface, and traces through the lens in the usual sequential fashion until it reaches the Non-Sequential Components surface. The ray is then sent into the group of components associated with that surface, and the non-sequential tracing begins.

Tracing rays within the NSC

Once inside the NSC group, 3 things can happen to a ray:

- 1) It can hit the exit port.
- 2) It can hit nothing at all.
- 3) It can hit one of the objects within the group.

If the ray hits the exit port, the ray coordinates and the direction cosines are computed on the exit port, and the ray then traces sequentially again through the remaining surfaces of the lens.

If the ray hits nothing at all, then the ray tracing is terminated, and the ray trace function returns a "ray missed" error on the following surface (since the ray never struck the exit port, which is always the next surface in the sequential portion of the ray trace).

If the ray strikes an object in the NSC group, then the ray will either reflect, refract, total internal reflect (TIR), or be absorbed, depending upon the properties of the object struck. Rays entering the NSC group through the entry port cannot split. If the ray is absorbed, the ray trace is terminated and a ray miss error is returned, otherwise the new ray coordinates and direction cosines are computed, and the process repeats until one of the following conditions is met:

- 1) The ray hits the exit port.
- 2) The ray hits no object.
- 3) The ray is absorbed.
- 4) The ray has intercepted more than the maximum allowed number of objects (see "Maximum Intersections Per Ray" on page 113).

Cases 1, 2, and 3 are handled exactly as described above. In case 4, even though the ray technically still can be traced, it is terminated to prevent infinite loops from occurring. In this case, the ray trace returns a ray miss error.

Getting rays out

When a ray strikes the exit port, the coordinates and direction cosines of the ray in the coordinate system of the exit port are computed, and then the ray traces sequentially through any remaining surfaces. If one of the following surfaces is another Non-Sequential Components surface, then the process begins again for the components defined for that group. Note rays within one NSC group cannot "see" objects defined in another group, even if they physically share the same location in space; nor can the rays "see" surfaces outside the present NSC group.

Overview of NSC ray tracing without ports

Ray tracing through a group of NSC objects without using ports is accomplished with the following basic steps:

- 1) Choose File, Non-Sequential Mode from the main menu.
- 2) Insert sources, objects, and detectors in the Non-Sequential Components editor.

There are no entry or exit ports to consider; these are ignored when using NSC without ports. The only data that needs to be defined outside the NSC Editor are:

The wavelengths to be used for ray tracing (on the wavelength data editor)

The glass catalogs to be used (on the general system data dialog box)

The coating definitions (in the coating file associated with the current lens)

Getting rays in

To get rays into the NSC group, define one or more sources. ZEMAX supports point, rectangular, elliptical, user defined, and other source models. Each source object has the following parameters defined (and perhaps other parameters as well):

Layout Rays: Defines how many random rays to launch from the source when creating layout plots.

Analysis Rays: Defines how many random rays to launch from the source when performing analysis.

Power (units): Power is the total power over the defined range of the source. The power units are specified by the system source units. See "Source Units" on page 105 for details.

Wavenumber: The wavenumber to use when tracing random rays. Zero means polychromatic; which chooses ray wavelengths randomly with the weighting defined on the wavelength data dialog box.

Note multiple sources may be superimposed with different powers and wavenumbers to create correct polychromatic sources. Sources may be placed anywhere without restriction (even inside objects).

Once the ray is launched, non-sequential tracing begins.

Tracing rays within the NSC

Once inside the NSC group, 2 things can happen to a ray:

1) It can hit nothing at all.

2) It can hit one of the objects within the group.

If the ray hits nothing at all, then the ray tracing for this ray is terminated.

If the ray strikes an object, then the ray will either reflect, refract, total internal reflect (TIR), scatter, split, diffract, or be absorbed; or a combination of these, depending upon the properties of the object struck.

Detecting rays

If the ray strikes a detector object; the pixel which the ray struck is determined, and the total pixel energy is incremented by the ray energy. Detectors may be absorbing, reflecting, transmissive, or refractive.

The process repeats until one of the following conditions is met:

1) The ray hits no object.

2) The ray is absorbed.

3) The ray has intercepted more than the maximum allowed number of objects.

4) The total number of ray segments exceeds the maximum allowed number.

5) The relative or absolute energy of the ray falls below the minimum threshold.

The limitations on number of objects, ray segments, and ray energy are defined on the System dialog box, see "Non-Sequential" on page 113. Cases 1 and 2 are handled exactly as described above. In case 3, 4, and 5, even though the ray technically still can be traced, it is terminated to prevent infinite loops from occurring.

Launching rays for analysis

Rays may be launched, and detectors reset, from the Detectors, Ray Trace/Detector Control.

Object Placement

The conventions and restrictions on placing objects in the NSC group is critically important. Objects may be placed anywhere in 3D space; and objects may be placed with respect to any other object. Objects may also be placed entirely inside of other objects, or may be placed adjacent to other objects.

The object coordinate system

Each object's position is defined by 6 parameters: the x, y, and z coordinates, and the rotation about the x, y, and z axes. ZEMAX first decenters in x, y, and z (decenters are orthogonal so the order does not matter). Then ZEMAX tilts about the local x axis (which rotates the y and z axes to new orientations), then tilts about the new y axis (which rotates the x and z axes), then finally tilts about the new z axis. This is the same convention as for the coordinate break surface when using order flag = 0 (see page 293). The conversion from object local coordinates to global coordinates can be written in equation form as

$$\begin{bmatrix} x_g \\ y_g \\ z_g \end{bmatrix} = \begin{bmatrix} x_o \\ y_o \\ z_o \end{bmatrix} + \begin{bmatrix} R_{11} & R_{12} & R_{13} \\ R_{21} & R_{22} & R_{23} \\ R_{31} & R_{32} & R_{33} \end{bmatrix} \begin{bmatrix} x_l \\ y_l \\ z_l \end{bmatrix},$$

where the g subscript indicates the global coordinate, o is the offset, and l is the local object coordinate. The matrix R is the rotation matrix, which relates the orientation of local and global coordinates. These equations can be written more compactly as

$$G = O + RL,$$

Where G is the global coordinate vector, O is the offset vector, R is the rotation matrix, and L is the local coordinate vector. For more information on properties of the rotation matrix, see "Global Coordinate Reference Surface" on page 111. The order of rotations may also be reversed, see "Use Global XYZ Rotation Order" under "Type tab" on page 353.

Reference objects

It is often useful to reference an object's position and rotation relative to another object. This is useful especially when placing related objects in a group, and then decentering or tilting the entire group.

The object to which the coordinates are relative to is the "reference object". The default reference object is object 0, which is the vertex of the Non-Sequential Component surface. If a positive number greater than zero is specified, then the coordinates of the object are referenced to the location and rotation of the specified object. This is an "absolute" reference object. If the reference object number is negative, then the reference object is determined by adding the current object number to the negative reference object number. This is a "relative" reference object. For example, if the reference object is -3 on object 8, the reference object will be 5 because $8 - 3 = 5$. Relative reference objects are particularly useful when copying and pasting groups of objects; this is easiest if all the objects in the group use relative references to the first object in the group.

When a reference object is used, the rotation and offset matrixes then become:

$$G' = O' + R'G$$

$$G' = O' + R'[O + RL]$$

$$G' = [O' + R'O] + [R'R]L$$

Any number of coordinate reference nesting is supported; so that object 9 could be placed in the coordinate frame of object 5 which in turn is placed in the frame of object 3. The only restriction is that the reference object must precede in the object list the object whose coordinates are being referenced. To modify the reference object without the need to manually recompute the object coordinates and angles, see the "Modify Reference Objects" on page 98.

Placing objects inside of, adjacent to, or overlapping other objects

Placing sources within other objects, and the use of absolute and relative values for the "Inside Of" column are described in "Placing sources inside objects" on page 449.

Objects can be combined to make more complex objects, by placing one object inside of, adjacent to, or overlapping another object. What determines the ray tracing properties of such a compound object depends upon the position and type of the various objects and whether or not they touch or overlap. Here the word touch means that one or more points on the boundary surface of one object is in the same location in 3D space as a point on a boundary surface of another object. Mirror objects may be placed anywhere, even in contact with or partially or fully inside any other object without restriction. Rays will always reflect from mirror surfaces back into the medium the rays had been traveling through.

Glue distance

For a discussion of glue distance, see "Glue Distance In Lens Units" on page 114.

Nesting object limits

There is a user defined limit on the maximum number of nested objects. This defines an upper limit on how many objects can be placed inside each other. For example, if the maximum number of nested objects is 3, then object 3 may be placed inside of 2 which is placed inside of object 1. There may be any number of groups of objects each nested 3 deep in this case. The limit applies to the total nesting in any one collection of objects, however, there may be any number of such collections within the NSC group. The maximum number of nested objects is set on the Non-Sequential tab of the General dialog box. Setting the nesting limit no higher than required for the system being modeled conserves memory usage, although the minimum setting is 4.

Nesting volumes


Volumes of refracting material are more complicated, because ZEMAX must keep track of the index of refraction through which the ray is propagating. The rule to remember is: if a ray strikes more than one object at the exact same point in space; the last object listed in the NSC Editor determines the properties of the surface or volume at that point.



If a ray strikes more than one object at the exact same point in space; the last object listed in the NSC Editor determines the properties of the surface or volume at that point.

For example, if a diffraction grating lens is object #1, and a non-grating lens of the same thickness and radii made of air or glass is placed inside the first lens is object #2; then rays that strike in the zone including both objects #1 and #2 will act as though they just hit object #2.

This allows defining objects with "holes" and other compound objects. Objects may be touching at one or more faces, nested or not, or may partially overlap, to create a wide range of compound solid shapes.



Nesting surfaces

The nesting rules defined above apply to solid volumes of refractive or reflective material. Some special surfaces may also be "nested" in the sense that more than one surface may exist at the same point in space. The rules are similar to those that apply to volumes, but surfaces cannot be refracting.

The same rule applies to surfaces as does for solids: if a ray strikes more than one object at the exact same point in space; the last object listed in the NSC Editor determines the properties of the surface at that point.



If a ray strikes more than one object at the exact same point in space; the last object listed in the NSC Editor determines the properties of the surface at that point.

For surfaces, there are a few rules to consider when more than one surface exists at the ray intercept point:

- 1) The last surface listed will determine the properties of the surface.
- 2) If the last surface listed is a mirror, the ray will reflect.
- 3) If the last surface listed is an absorber, the ray will be absorbed.
- 4) If the last surface listed is neither a mirror nor an absorber, the ray will ignore the surface.
- 5) Surface objects may not share boundaries with volume objects unless the surface object is reflective or absorbing, or unless the volume object is listed after the surface object; in which case the volume defines the properties of the common boundary.
- 6) Only standard surface objects may share boundaries, support for nesting of other surface object types may be added in future versions of the software.

Refraction and reflection from NSC objects

All NSC objects allow specification of a material on the NSC Editor (gradient index media are defined separately, see “Defining GRIN media for non-sequential ray tracing” on page 373). The reflective, refractive, and absorptive properties of an object depend upon the object type, whether or not the object describes a surface or a volume, and the name of the material specified. The material name may be blank, "MIRROR", "ABSORB", a glass name, a table glass name, a MIL number designation, or a model glass.

No material defined

Blank materials are assumed to indicate an index of unity, which may or may not be the background index of the non-sequential space (which is set by the index of the NSC surface on the Lens Data Editor).

MIRROR or ABSORB

A material of MIRROR indicates the object reflects all incident rays. A material of ABSORB indicates the object absorbs all incident rays.

Material defined by a glass name

Homogeneous glasses are defined in the glass catalogs. For more information, see “Specifying which glass catalogs to use” on page 587.

Material defined by a table glass

To define a table glass, see “Using table glasses” on page 601.

Material defined by MIL number

A 6 digit MIL number, such as 517642, may be used to define an approximate material model. For more information on MIL number glasses, see “Using MIL number glasses” on page 600.

Material defined by a glass model

The glass model uses 3 independent terms to define a typical dispersion curve. The model glass is more flexible and accurate than the MIL glass, but not as good as actual measured data. For more information see “Using model glasses” on page 601.

Object properties by material

The properties of objects are described by the following table.

OBJECT PROPERTIES BY MATERIAL

Object Type	Material is blank or glass	Material is "MIRROR"	Material is "ABSORB"
Volume	All faces are refractive.	All faces are reflective.	All faces are absorbing.
Surface	Object is ignored.		
Polygon Object (volume)	If a face in the POB file is marked as reflective or absorbing, then the face remains reflective or absorbing. Refractive faces remain refractive.	If a face in the POB file is marked as absorbing, then the face remains absorbing. All other faces become reflective.	If a face in the POB file is marked as reflective, then the face remains reflective. All other faces become absorbing.
Polygon Object (surface)	Object is ignored.	If a face in the POB file is marked as absorbing, then the face remains absorbing. All other faces become reflective.	If a face in the POB file is marked as reflective, then the face remains reflective. All other faces become absorbing.

Reflective objects are usually surfaces rather than volumes. Hollow light pipes can also be modeled as a group of flat reflectors. Some reflectors do enclose a volume; the interior of which no ray will ever see. Examples include rectangular and spherical volumes with all surfaces set to be reflective. Any volume with a material name of "MIRROR" will be all reflective.

Refractors must always enclose a volume. Volume refractors include lenses, prisms, and solid light pipes. Absorbers may be surfaces or volumes with the material name set to "ABSORB". Polygon Objects are composed of faces, and each individual faces may be refractive, reflective, or absorptive.

Diffraction from NSC objects

Some NSC objects have one or more diffractive faces; such as the Diffraction Grating, Binary 1, and Binary 2 objects. These objects refract or reflect rays as well as diffract them, according to the grating period or phase and the diffraction order and wavelength. For any ray, if the diffraction order being traced does not satisfy the grating equation, then the energy of that ray will refract or reflect along the zero order path.

Ray splitting is supported on diffractive surfaces, but splitting is only allowed by order and not by reflected and transmitted Fresnel coefficients. For this reason, no "ghost" rays are generated from diffractive surfaces.

All diffractive NSC objects support a parameter to defined the "order" for diffraction from the object. This order is called the "primary" diffraction order. Sequential rays which enter the non-sequential group through the entry port will diffract only along the primary order.

Rays which originate from a non-sequential source will diffract only along the primary order if ray splitting is off. If ray splitting is on, then the diffraction may be optionally controlled by the settings on the "Diffraction" tab of the Object Properties dialog box. This tab includes optional settings which will split the ray by order; allowing more than one diffraction order to be simultaneously traced.

For more information on diffraction, see "Diffraction tab" on page 357.

The object properties dialog box

For all objects inserted into the NSC Editor, the object properties dialog box is used to define properties such as the object type, special aperture settings, scattering and coating properties, and gradient index properties. The previous/next object buttons allow rapid navigation through all objects. The object properties dialog box has multiple tabs as described below.

Type tab

The Type tab supports the following controls:

Type: Used to select the general object type, for example, sphere, ellipse, rectangle, or other type.

Data File: If the object type is defined by an external file, such as a polygon object, the file name may be selected here.

User Defined Aperture: If selected, a user defined aperture (UDA) file will be used to define the extent of the object. Not all objects support this feature. For more information see page 364.

File: If a user defined aperture is selected, the name of the aperture file may be selected from this control.

Edit User Aperture: This button will invoke a text editor to allow user editing of the selected UDA file. The UDA file needs to be saved and the object reloaded to make the changes effective.

UDA Scale: The UDA Scale is a dimensionless multiplier that scales the aperture defined in the UDA file.

Row Color: This control chooses the color of the row in the NSC Editor for the object. By default, objects made of glass, "MIRROR", or "ABSORB" material, as well as sources and detectors are color coded. Any object may use either no color, the default color, or a user defined color. The user defined colors are described in "Colors 1-12, Colors 13-24" on page 73. The coloring of rows may be disabled, see "Editors" on page 71.

Consider Objects: For a detailed discussion of this feature see "Defining paths using the Consider Objects list" which follows on page 363.

Ignore Objects: For a detailed discussion see "Defining an Ignore Objects list" on page 364.

Use Consider/Ignore Objects When Splitting: If checked, then when rays are split the child rays will only look for intersections with the objects listed on the Consider Objects list, or only ignore objects on the Ignore Objects list. If unchecked, all possible object intersections will be checked. See "Defining paths using the Consider Objects list" which follows on page 363.

Rays Ignore This Object: If an object is included for rendering or reference purposes only, ray tracing will be faster if this box is checked, as ZEMAX will not bother to check if the ray intercepts the object. Checking this option can cause incorrect ray trace results if any rays could intersect the object, as rays will travel through the object as if it were not there.

Use Global XYZ Rotation Order: If checked, the rotation convention for object tilts is to first rotate about the X axis, then the Y axis, then the Z axis. If unchecked, the rotations are done about the Z axis, then the Y axis, then the X axis. This latter ZYX convention is what a sequential coordinate break surface does if the order flag is zero. Note that rotating using the ZYX convention is exactly the same as rotating about the X axis, then rotating about the "new" Y axis, then finally the "new" Z axis. See "The object coordinate system" on page 350.

Object Is A Detector: If checked, then rays striking the object will increment a detector associated with the object.

Show As: This control selects the color scheme used to represent objects that are detectors on the shaded model plot. Color by flux will color according to the total power incident on each area of the object. Color by irradiance will color each area by the power per area. Only objects that are detectors use this feature.

Fast Ray Trace (Slow Update): There is a trade-off between ray trace speed and the time it takes to create the internal representation of certain objects. These objects include most non-radial symmetric aspheres, toroids, user defined, and other object types for which no radial symmetry and no exact ray trace solution exists. The default for this setting is checked, which means the object will ray trace faster but update more slowly than if this feature is not checked. For many aspheric lenses, rays will trace from 10 to 50 times faster with Fast Ray Trace on than off. The usual reason to turn off Fast Ray Trace is if the editing of the object parameters in the NSC Editor becomes tedious because of the slow update. This switch has no effect at all on objects not described above. If the object is used as a detector (as selected with the "Object Is A Detector" checkbox described above) and Fast Ray Trace is on, ZEMAX will alter the order of the object's detector triangles to speed up the ray trace. If Fast Ray Trace is unchecked, the order of the triangles is unmodified, which slows the ray tracing but preserves the original order of the triangles. Additionally, if Fast Ray Trace is on, the rendering of the object will not be smoothly shaded on Shaded Model plots, because this smooth shading depends upon the original order of the triangles being maintained.

Convert Imported Files To ZOF: If checked (the recommended setting), any imported file which requires more than 2.0 seconds to read and convert will be saved in a new file that ZEMAX can read quickly. For a detailed description, see "Imported objects and ZOF files" on page 424. Only imported objects use this setting.

Use Pixel Interpolation: If checked (the recommended setting), the energy from rays striking a pixel on the Detector Rectangle will be apportioned among the neighboring pixels. The amount of energy allocated to each pixel depends upon the position the ray strikes within the pixel. Rays striking the exact center of the pixel will have 100% of their energy allocated to that pixel. Rays striking very near the edge of the pixel will have 50%

of their energy to each of the pixels adjacent to the edge. Pixel Interpolation eliminates the abrupt edge response of pixels and generally yields less noisy plots and smoother optimization. Pixel interpolation is not used if the detector is using PSF mode.

Normalize Coherent Power: See "Comments on coherent data computations" on page 467.

Coat/Scatter tab

The Coat/Scatter tab supports the following controls:

Face: The face to which the properties will be applied. For imported objects, see "Face tab" on page 355 to define the faces. See "Object faces" on page 364.

Profile: A profile is a collection of settings related to the thin film coatings and scattering model data applied to an object face. Suppose that many objects in an optical system are all composed of the same material and have the same coating and scattering properties. Rather than type this identical data in for each object and face; the data may be typed in once, then saved to a profile. Once saved, any other object face may use the same profile. Profiles are stored in the file whose name is specified on the Files tab of the System, General dialog box; see "Scatter Profile" on page 109 for details. If unique settings for this object are desired, then choose "Use definitions below" to edit the scattering parameters. If a profile is selected; then the coating and scattering data boxes will be disabled, as this data is defined by the selected profile. However, the settings will be displayed.

Save: Saves the current scattering settings as a new profile. Pressing this button will prompt for the name of the new profile.

Delete: Removes the currently selected profile from the Scatter Profile file.

Face Is: This setting controls whether the face is refractive, reflective, or absorbing. "Object Default" means the face is either refractive, reflective, or absorbing, depending upon the material type defined in the NSC Editor. "Reflective" means the face is always reflective, and "Absorbing" means the face is always absorbing.

Coating: The name of the coating to apply to the surfaces with the selected face number. Coatings are defined in the system coating file, see "Defining coatings in ZEMAX" on page 613 for details.

For all other scattering controls, see the discussion "Scattering" on page 364.

Scatter To tab

Scatter To Method: Choose either "Scatter To List" or "Importance Sampling". For a detailed discussion of these features see "How to model scattering efficiently" on page 369.

Face tab

The controls on this tab are used to assign face numbers to the various surfaces defined in an imported CAD format file. Only imported objects (see "Imported" on page 423) use this tab. The Face number assignments are stored in a file with the same name as the imported CAD file, with the extension ZEN appended. The ZEN file is stored in the same folder as the CAD file. For example, if the imported object is MyObject.STP, the face number assignments will be stored in the file MyObject.STP.ZEN. If the object definition or Face Mode changes, the ZEN file is automatically regenerated.

Imported Object Surface #, Face #: This control selects the desired surface number. Changing the selection in this control will update the NSC Object Viewer if the viewer is opened from the "View Object" button at the bottom of the tab. Note more than one surface may be selected at one time. This allows groups of surfaces to be selected and renumbered together using the "Change To->" button. An alternate method to select surfaces is to click on the surface in the NSC Object Viewer. See "Viewer Highlights" below.

Change To-> Face #: Pressing "Change" will change the highlighted surface numbers to the selected Face number.

Reset All: Automatically renumbers all surfaces to the default. The default is to assign each surface a unique Face number starting at 0 through the maximum number of Faces supported. All remaining surfaces are assigned Face number 0.

Select All: Selects all surfaces.

Clear All: Unselects all surfaces.

Invert All: Reverses the selected and unselected surfaces.

Viewer Highlights: Sets the NSC Object Viewer to highlight in red all selected surfaces or all selected faces. Note if multiple surfaces are assigned the same face number, then highlighting by face will show all surfaces with the same face number.

To change how ZEMAX initially determines the number of unique surfaces and sets the default face numbering, see "Imported" on page 423.

Bulk Scatter tab

The Bulk Scattering tab is used to define bulk (volumetric) scattering of solids. To set whether a ray from a given source can bulk scatter once or multiple times, see "Sources tab" on page 358.

The tab supports the following controls:

Model: Select either no bulk scattering, angle scattering, or DLL defined scattering. For a detailed description of bulk scattering, see "Bulk scattering" on page 370.

Mean path: The mean path between scatter events.

Angle: The angular cone in which the rays scatter.

DLL: The name of the DLL which defines the scattering function.

Parameter values: These are used to define the parameters to be passed to the DLL. To make these parameters variable or under multi-configuration control, use the NPRO operand; for details see "NPRO" on page 582.

Wavelength Shift: The wavelength shift control allows definition of wavelength transitions during bulk scatter events. The syntax is "in, out, prob" where "in" is the input wavelength number, "out" is the output wavelength number, "prob" is the relative probability that this shift will occur when tracing the in wavelength. Multiple transitions may be defined using a semi colon separator.

Example 1. A single input wavelength (#1) will shift to another single wavelength (#2): "1,2,1.0".

Example 2. A single input wavelength (#3) will shift to either wavelength #4 or #5 with a relative probability of 6.0 to 1.5: "3, 4, 6.0; 3, 5, 1.5".

Example 3. A single input wavelength (#1) will shift to wavelength #2 50% of the time, wavelength #3 40% of the time, and the remaining 10% of the time will remain at the input wavelength: "1, 2, 50.0; 1, 3, 40.0; 1, 1, 10.0".

Example 4: Wavelength #2 and #3 both shift to wavelength #5 with 100% probability: "2, 5, 1.0; 3, 5, 1.0".

Example 5: Wavelength #1 shifts to wavelength #2, and if bulk scatter occurs again, the wavelength will shift to #4 with 50% probability and #5 with 50% probability: "1, 2, 1.0; 2, 4, 50.0; 2, 5, 50.0".

The probability values are only relative to other wavelength shifts defined for the same input wavelength.

This feature is primarily used for modeling fluorescence. To control how many times a ray can bulk scatter, see "Bulk Scatter" under "Sources tab" on page 358.

Wavelength Shifting is not supported when using Source Color models other than the System Wavelengths.

GRIN tab

The GRIN tab is used to define the properties of solids made of a gradient index material. The tab supports the following controls:

Use DLL defined GRIN media: If checked, then the object will use an externally provided DLL to define the properties of the gradient index media.

DLL: The name of the DLL to use.

Maximum Step Size: The maximum step size to use during piecewise ray tracing.

Remaining controls: The remaining controls are used to define the parameters to be passed to the DLL. To make these parameters variable or under multi-configuration control, use the NPRO operand; for details see "NPRO" on page 582.

For more information on defining GRIN media, see page 373. The media may not be both gradient index and birefringent at the same time.

Birefringence tab

This tab is used to define the birefringent properties of solids made from uniaxial anisotropic materials, such as calcite. The tab supports the following controls:

Birefringent media: if checked, the media is anisotropic, otherwise, the media is isotropic.

Mode: This control allows selection of which rays are actually traced. The distribution of energy does not change with this setting, and energy assigned to rays that are not traced are placed in the "lost energy (thresholds)" sum. Either ordinary or extraordinary rays, or both, may be traced. Waveplate Mode traces the energy along the path of the ordinary ray, however, the extraordinary electric field is also propagated with the ray, and the total electric field will be phase rotated by the differential index in the ordinary and extraordinary directions.

Reflections: This control determines whether refracted or reflected rays, or both, will be traced. Note rays that TIR at a surface are considered reflected rays.

Ax/Ay/Az: The dimensionless components of the uniaxial crystal orientation vector in object local coordinates. These coordinates define the orientation of the axis of symmetry for the media. Internally, the coordinates are normalized to a magnitude of unity. If all three terms are zero, the orientation used is (0, 0, 1).

Axis Length: The length in lens units of the crystal axis drawn on layout plots. Use a value of zero to not draw the crystal axis.

Birefringence is ignored unless both polarization and ray splitting are on. If birefringence is ignored, the media acts like an isotropic media whose index is defined by the material associated with the volume. If birefringence is considered, the material associated with the object defines the "ordinary" refractive index. The "extraordinary" refractive index is defined by the material whose name is determined by appending "-E" to the ordinary material name. For example, if the object material name is CALCITE, then the extraordinary material must be named CALCITE-E. If both material names are not found in the current glass catalogs, an error is issued. If the material is defined as either a table or model glass, the extraordinary and ordinary indices are set to the same value and the media is effectively not birefringent.

When a ray strikes a boundary where either the incident or substrate media (or both) is anisotropic, as many as four separate rays are computed - two for ghost reflection and two for refraction. In each direction, there may be a separate ordinary and extraordinary ray. The ordinary ray carries the portion of the electric field that lies orthogonal to the plane that contains the direction of propagation and the crystal axis vector, while the extraordinary ray carries the electric field that lies in that same plane. If the media the rays are traveling in is isotropic (such as the ghost reflection off of a birefringent object) then the ordinary and extraordinary rays are combined into a single ray for convenience, as the energy propagates along the same path in this case.

Although there are important differences between the non-sequential and the sequential model for birefringent media in ZEMAX, most of the discussion about the birefringent media itself is common to both models. This discussion is found in "Birefringent In and Birefringent Out" on page 287.

Birefringent media may not also be gradient index or use bulk scattering, and may not be diffractive or scatter rays at the media boundaries. Simple ray splitting is not supported for birefringent media. Sources may not be placed inside birefringent media.

Diffraction tab

For important information that applies to all diffractive objects, see "Diffraction from NSC objects" on page 353.

The Diffraction tab is used to define the properties of diffractive surfaces. The tab supports the following controls:

Split: Selects how rays are split off from diffractive surfaces. The options are:

Don't split by order: The rays will not split at the surface. Only the order defined by the object parameters will be traced, and all the transmitted energy goes into this one order.

Split by table below: A user defined number of rays will be traced over a range of integral orders. The fraction of energy given to each order is defined by the user in a table.

Split by DLL function: An external DLL program is used to define which orders are traced, how much energy each order is given, and optionally, what the output ray properties are (for user defined diffraction). For more information on defining Diffraction DLLs, see page 375.

DLL: The name of the DLL to use. This DLL must be placed in the <program>\DLL\Diffraction folder. See "Folders" on page 68.

Start/Stop order: The beginning and ending order number. These numbers determine how many cells in the table are active; and how many times the DLL will be called to compute the output ray properties.

Remaining controls: The remaining controls are used to define the parameters to be passed to the DLL. To make these parameters variable or under multi-configuration control, use the NPRO operand; for details see "NPRO" on page 582.

Sources tab

The Sources tab is used to define properties of source objects, including polarization state, coherence length, and initial phase, position, and direction of rays of light emanating from NSC sources. For important information see "Defining the initial polarization" on page 622. The tab supports the following controls:

Random Polarization: If checked, the source will emit randomly polarized light. If unchecked, the polarization state may be defined using other controls on this tab.

Reverse Rays: Checking this option will reverse the direction cosines of every ray. This is useful for reversing the initial direction of rays from the source. Reverse rays is done after the Pre-Propagation distance is considered, if any.

Jx, Jy: The magnitude of the electric field in the local x and y directions, respectively.

X-Phase, Y-Phase: The phase in degrees of the electric field in the local x and y directions, respectively.

Initial Phase: The initial phase of the ray in degrees, with 360 degrees being equal to one wave of optical path. This setting only affects coherent ray computations which depend upon the phase of the ray.

Coherence Length: The length of ray propagation in lens units over which the phase is known. For details on Coherence Length effects see "Coherence length modeling" on page 373.

Pre-Propagation: The distance in lens units the ray is propagated before beginning the actual ray trace through NSC objects. Pre propagation moves the starting point of rays forward or backward along the ray direction cosines. This feature is useful for defining rays at one position but beginning the ray trace at a different position along the ray path; such as prior to or after an object near the source. The Pre-Propagation distance will alter the initial phase and electric field of the ray to account for the propagation length. Pre-Propagation occurs before reversing the rays if the reverse rays option is selected.

Bulk Scatter: Normally, if a ray travels through an object with a bulk scattering media, the rays may scatter multiple times within the media. This is the default "Many" option. If "Once" is selected, each branch of a ray can only bulk scatter once. Note that if a ray splits before scattering, each of the child rays may scatter, since each child's branch is scattering for the first time. If "Never" is selected, then no bulk scattering will occur for rays from this source. This control is useful for modeling fluorescence. See also "Bulk Scatter tab" on page 356.

Sampling Method: The available options are random and Sobol. See the following discussion of random vs. Sobol sampling for details on sampling method.

Array Type: This feature creates an array of identical sources, all with the same properties as the "parent" source. For more information see "Modeling arrays of sources" on page 360.

Source Color, Spectrum, and Wavelengths From/To: These settings choose the method for generating the spectral model for the source. for more information see "Defining the color and spectral content of sources" on page 360.

Since rays in general occupy a 3D space, the electric field in the Z direction is determined by the above data assuming the Jx and Jy values are the projections down on to the local YZ plane of the source. The resulting electric field will always be perpendicular to the ray propagation vector.

The polarization properties of the source should only be set non-random if the source is reasonably well collimated. Partially polarized sources may be simulated by superimposing sources with different polarization properties at the same location.

Random vs. Sobol sampling

Once the parameters of a source model are defined, rays are generated randomly to model the light leaving the source. This process involves generating random numbers which are scaled to represent values of the starting coordinates and direction cosines of the rays. Note the starting coordinates describe the near-field source properties, and the direction cosines represent the far-field properties.

Truly random values are not always desirable. The reason is that random numbers tend to *not* uniformly sample parameter space if the number of samples is small. Random numbers may group together, leaving relatively large gaps in sampling space. In practice, this means that it takes a great many rays to get sufficient sampling to produce smoothly varying results.

A solution is to use a sampling method which looks qualitatively random, but is in fact a carefully selected distribution that optimally "fills in" previously unsampled space. A widely used algorithm is a pseudo-random sequence called Sobol sampling. For a good discussion of this algorithm see Press, Flannery, Teukolsky, and Vetterling, *Numerical Recipes*, Cambridge Press.

The following graphs show the near-field ray patterns from two rectangular sources. The first ray pattern shown here on the right uses random sampling. Note the distribution has no obvious periodic artifacts.

Some of the rays are very close to one another, while other relatively large areas contain no rays at all. In practice, this will yield a "noisy" irradiance distribution on a detector, since some pixels will contain one or two rays, while other pixels will contain no rays, even though the source is uniform for this sample case.



For random rays, the Signal-to-Noise Ratio (SNR) is approximately given by $N / \sqrt{N} = \sqrt{N}$ where N is the number of rays *per pixel*. For large detectors with many pixels, a very large number of rays may be required to obtain a high SNR.

For example, a detector with 10,000 pixels would require roughly 100 million rays to achieve a SNR of 100, assuming the best-case scenario of a uniform source. For more practical and realistic optical systems, lower light levels would have much lower SNR.

The graph at the right shows the same number of rays traced to an identical detector using the Sobol pseudo-sampling grid. Note the rays are much more uniformly distributed, with relatively fewer gaps and fewer rays exceedingly close to one another. As additional rays are traced, the rays are always selected to lie between rays already traced, creating an ever finer and more uniform distribution. For reasons which are beyond the scope of this discussion, the SNR using the Sobol sampling goes as approximately N , rather than \sqrt{N} . For the previous example of 10,000 pixels, only about 1 million rays would need to be traced to achieve a SNR of 100, a 100X increase in efficiency.

There are two disadvantages of the Sobol algorithm. First, there are sometimes subtle artifacts in the ray distribution perceptible to the human eye.

For some optical systems, if the sampling is too low, these artifacts can show up in the detected irradiance as false pseudo-periodic structures. Second, the rays generated with the Sobol algorithm are always the same. Tracing M rays, and then tracing M more rays, will yield the exact same results as the first set of M rays; whereas



tracing 2M rays would yield a different result. This is why sources using Sobol sampling as drawn on layout plots will always trace the exact same rays. The only way to get different rays is to trace at one time a larger number of rays. Nonetheless, using Sobol sampling will generally produce faster convergence with fewer rays than random sampling.

Not all source models require or support Sobol sampling. If a source model does not support Sobol sampling, this control will be greyed out. Some source models only support Sobol sampling for part of the ray generation, such as the ray starting coordinates, and not for other parts, such as the ray starting direction cosines.

Modeling arrays of sources

The Source Array Type feature can be used to create an array of identical sources. Depending upon the Array Type, additional parameters become available on the Source tab to define the number of array elements and the size of the array. The array types supported are:

Rectangular: This array can be used to make a 1D line or 2D array of sources with uniform spacing in the local X and Y directions. The options available include the number of sources in X and Y, and the source-to-source spacing in lens units along each direction. The minimum number of sources in either direction is 1, the maximum is 2000. The sources are numbered starting from 1 in the first (x) column and first (y) row at the location of the parent object. Each subsequent source is numbered across the columns along the first row, until the number of x direction sources is reached. The next source will be placed at the next row in the first column, and then numbering continues across the columns again, until all sources are placed.

Circular: This array is a single circle of sources centered on the parent coordinate. The sources are equally spaced in angle at the specified radial coordinate in lens units. The first source is at 0 degrees on the XY plane, and the sources continue counter-clockwise (looking down the -Z axis) around the circle until all sources are placed.

Hexapolar. This array consists of equally spaced rings of sources. The first "ring" is at the location of the parent, and the first source is placed there. The second ring contains 6 sources, equally spaced in angle, starting at source number 2 at +90 degrees on the XY plane, and the sources continue counter-clockwise (looking down the -Z axis) around the circle until all 6 sources (source numbers 2-7) are placed. The third ring contains 12 sources, the fourth ring 18, and the pattern repeats until the last ring is reached. A maximum of 20 rings is supported. The spacing parameter is the radial spacing between adjacent rings.

Hexagonal. This type forms an array with hexagonal symmetry. The first "ring" contains one source at the location of the parent source. The second "ring" contains 6 sources, placed around the center source in hexagonal fashion. The third "ring" consists of 12 sources placed outside of the second ring, and the pattern repeats until the last ring is reached. A maximum of 20 rings is supported. The spacing parameter is the full height of the hexagonal region, which is the vertical spacing between sources in the same column. The numbering convention for this array starts at 1 at the bottom (-y coordinate) of the leftmost (-x coordinate) column, and then goes up the leftmost column, then starts at the bottom of the next column to the right, and proceeds up that column. The pattern repeats until the top source in the right most column is reached.

For all arrays of sources, if the total number of sources exceeds 10,000, the source objects will not be drawn. The parameter settings for the source, such as number of layout rays, number of analysis rays, and power apply to *each* source in the array. For example, a 3 x 3 array of 1 watt sources will produce 9 watts.

Defining the color and spectral content of sources

There are a number of ways to model the spectral content of a source. Sources may be either monochromatic, or may cover some region of the visible spectrum to represent a composite color, such as orange or white. The available source color models are described below.

System Wavelength: If this option is selected, then the monochromatic wavelengths defined on the system wavelength dialog box are used. The system wavelengths are described in "Wavelengths" on page 116. Once the system wavelengths are defined, the wavelength used for ray tracing is defined by the wavenumber used by the source. The wavenumber is one of the NSC Editor parameters used by all sources. See "Parameters common to all source objects" on page 450 for a full description. This is the default setting for Source Color and is the only setting that uses the system wavelengths.

CIE 1931 Tristimulus XYZ: This source color model defines the color of the source using the three CIE

tristimulus values commonly given the names X, Y, and Z. See “The spectrum fitting algorithm” on page 361 for important information on the spectrum generated to represent the desired color. Note that although the Y value normally represents the overall brightness of the source in lumens, ZEMAX does not use the Y value for this purpose when defining the source color. The power or brightness of the source is set independently; see “Parameters common to all source objects” on page 450 for a full description. Internally, ZEMAX will normalize the XYZ values to 1 lumen for purposes of computing the spectrum.

CIE 1931 Chromaticity xy: This source color model is essentially identical to the XYZ model above, except the normalized chromaticity coordinates are used instead. ZEMAX converts the xy values to XYZ and then follows the procedure described in “The spectrum fitting algorithm” on page 361.

CIE 1931 RGB (Saturated): RGB values are converted to CIE XYZ coordinates, and the method for computing the spectrum is then the same as for CIE 1931 Tristimulus values. Note the RGB color is the fully saturated RGB color, which means colors such as grey on a 8 bit scale (128, 128, 128) are saturated to the maximum brightness (256, 256, 256) and will appear white. There is no difference in the color spectrum of grey and white, or any dark and light colors, as long as the relative color values are maintained.

Uniform Power Spectrum: This creates a uniform power spectrum over the specified wavelength range at the number of discrete wavelengths defined.

D65 White: Defines $X = 0.9505$, $Y = 1.0000$, $Z = 1.0890$, which is the D65 white color on computer monitors.

Color Temperature: Based upon the temperature in Kelvin, the XYZ tristimulus values are computed to yield the same color as a black body of the specified temperature, and the spectrum is computed as described for the CIE 1931 Tristimulus values. Note this is not a true black body spectrum, but a spectrum fit to yield the same color as the blackbody of that temperature.

Black Body Spectrum: Based upon the temperature in Kelvin, this yields a true blackbody spectrum over the specified wavelength range. This color model is not limited to the visible spectrum.

User Defined Spectrum: This color model reads the wavelength values and weights from a file. For details on the file format, see “Defining a spectrum file” on page 362.

CIE 1976 Chromaticity $u' v'$: This source color model is essentially identical to the XYZ model above, except the normalized u' and v' chromaticity coordinates are used instead. ZEMAX converts the u' and v' values to XYZ and then follows the procedure described in “The spectrum fitting algorithm” on page 361.

Note that for all Source Color settings other than System Wavelengths, Uniform Power Spectrum, Black Body Spectrum, and User Defined Spectrum, a spectrum must be fit to represent the selected color. The spectrum fitting algorithm is described below.

The spectrum fitting algorithm

ZEMAX cannot trace a non-monochromatic "ray". Every ray traced must have a unique wavelength, so that the index of refraction, reflection, transmission, bulk absorption, diffraction, and other effects can be computed accurately. For this reason, sources that are not monochromatic must be represented by a spectrum of discrete wavelengths.

It is important to grasp that for a given composite color of light, such as light blue, the underlying monochromatic spectrum is not unique. There are an infinite number of weighted discrete spectrums that all have the same tristimulus values. When the color of a source is defined using tristimulus values or their equivalent, ZEMAX must compute an appropriate spectrum that represents that color. The resulting discrete wavelengths, appropriately weighted, are then traced.

There are 3 inputs that allow user control over the spectrum generation: the minimum wavelength, the maximum wavelength, and the number of wavelengths. ZEMAX will create a spectrum of evenly spaced monochromatic wavelengths that span the defined range. For example, if the minimum wavelength is 0.500, the maximum is 0.700, and 5 wavelengths are used, then the wavelengths chosen will be 0.500, 0.550, 0.600, 0.650, and 0.700 (all wavelengths in ZEMAX are in micrometers). ZEMAX will compute the XYZ response at each of these wavelengths, and determine the optimal weighting at each wavelength to yield the total desired XYZ color. The resulting "best fit" tristimulus XYZ values are displayed underneath the spectrum input. Additionally, the RGB equivalent of the best fit color is displayed for reference. When using the fitting to produce a desired

XYZ color, the minimum and maximum wavelengths must be between 0.380 and 0.830 micrometers, which is the practical extent of the visible range of light and the tristimulus Y response is non zero.

Note that the best fit XYZ values may differ from the desired value. This can happen if there are two few wavelengths in the spectrum, or if the wavelength range is not sufficiently broad, or if the desired XYZ values cannot be represented. In the latter case, this is usually because the XYZ values are non-physical and do not lie within the visible region on the CIE chromaticity chart - for example, the tristimulus values (1, 1, 18) do not correspond to any visible color.

Once the spectrum is determined, ZEMAX randomly chooses wavelengths to trace from among the spectrum values, weighted according to the relative weights determined by the fitting process. To see the wavelengths and weights actually used in the fit, use the prescription report described in "Prescription Data" on page 271.

Defining a spectrum file

Spectrum files are in text format, end in the extension *.SPCD, and are placed in the <objects>\Sources\Spectrum Files folder (see "Folders" on page 68). The files are used to describe a user-defined spectrum for source color modeling. The file format is pairs of numbers in the following format:

```
# comment <optional>
wavelength1 weight1
wavelength2 weight2
etc...
```

The wavelength values are in micrometers. The weight values are in dimensionless relative power units. Between 3 and 100 pairs of data points may be defined.

How rays are selected from the spectrum

Once a spectrum is defined, ZEMAX will randomly choose wavelengths from the spectrum based upon the relative weight of each wavelength. A wavelength with twice the relative weight as another wavelength will trace twice as often, and therefore twice as many rays of that wavelength will be traced. If a wavelength has a low relative weight, fewer rays of that wavelength will be traced, and this may lead to undersampling (more noise) in data associated with that wavelength.

If the source units are in Watts or Joules, every ray has equal power (or energy), and the total power of all the rays traced will exactly equal the total power of the source. However, if the source units are Lumens, this will not be precisely true. To correctly compute the true color of the detected light, ZEMAX must scale the ray units from Lumens to Watts. The conversion back to Lumens is done when the ray is detected. For this reason, different rays will have different power values in Watts. Since the number of rays is fixed, and the ray wavelengths are randomly chosen, the total power in Lumens will not be exactly the same value as that specified for the source. Generally speaking, if a reasonable number of rays is traced, the total Lumens will be very close to the defined source Lumens value.

Using ABg scattering with source color models

For important information on ABg coefficient scaling, see "Wavelength scaling of ABg data" on page 251.

Draw tab

The draw tab is used to set options related to drawing each NSC object. The tab supports the following controls:

Do Not Draw Object: If checked, then the object will not be drawn on layout plots. Rays will still act as though the object were there. This control is useful to remove from the layout an object that encloses other objects, so the inner objects may be seen more easily.

Draw Local Axis: If checked, the local x, y, and z axis will be drawn on the 3d Layout plot. The length of the x and y axis drawn is half that of the "missed ray draw distance", see "Missed Ray Draw Distance in Lens Units" on page 114. The z axis is twice this length. The vertex of the coordinate indicator is at the local vertex.

Drawing Resolution: Choose Standard, Medium, High, Presentation, or Custom. The options increase the resolution of the drawing at the expense of computation time. Different objects use the drawing resolution for different purposes depending upon the symmetry of the object. The Custom setting allows a user defined resolution to be specified in the edit controls below the Drawing Resolution control.

Increase Resolution On Shaded Model Plots: If checked, the drawing resolution will be temporarily increased

when generating Shaded Model plots. Although this produces much better looking graphics, it does take more time and memory. To save time and memory, disable this option.

Make All Objects Draw Like This One: Pressing this button will copy the drawing settings to all objects.

Opacity: If the opacity is set at 100%, then the object will be rendered on the shaded model plot as a solid color, and the object may fully obscure other objects from view. If the opacity is less than 100%, then the object is partially transparent, which allows other objects to be visible through the partially opaque surface.

Object Color: This control is used to select the color the object will be drawn with on the shaded model display.

Export As Triangles: If checked, the object will be exported to CAD format files as a tessellation rather than a smooth object. This feature is only intended as a work around if the CAD format export feature fails to work properly. The recommended setting is not checked.

Comments about drawing resolution and bounding regions

ZEMAX uses the drawing resolution settings to create a "boundary region" around each object. This region is used to quickly determine if a ray comes close enough to an object to possibly intercept the object. If the ray does not cross the boundary region, ZEMAX assumes the ray will not hit the object. For this reason, it is possible for rays to miss an object if the drawing resolution is set too low. For these rare cases, increase the drawing resolution to ensure the ray strikes the object.



Do not set the drawing resolution needlessly high as this will slow down ray tracing and rendering without increasing accuracy!

It is important to note that the drawing resolution settings are only used to render the object on the layout plots, and for some objects, to provide a "first guess" as to the location of the ray-object intercept point. For accurate ray tracing, iteration is used to find the exact intercept point. The actual surface shape of all ZEMAX objects are modeled exactly for ray tracing purposes; the accuracy of the ray trace is not limited by the drawing resolution, as long as the drawing resolution is sufficient to provide an adequate first guess.

Defining paths using the Consider Objects list

It is often the case that rays propagating within an NSC group follow a reasonably well defined path. For example, when several lenses are mounted in a tube, rays striking lens #5 will either hit lens #4 or lens #6 next; or possibly the tube itself. If the list of possible objects to intersect is known, and the number of objects on this list is small compared to the total number of objects; then significant ray tracing speed gains may be made by telling ZEMAX to only consider the subset of objects the ray may hit.

The Object Type tab has a user-definable list labeled "Consider Objects" for doing this. If left blank (the default), then rays leaving this object may strike any object, and ZEMAX will use it's own internal algorithms for determining which object is actually hit.

To specify the subset of objects the ray may hit when leaving an object, list the object numbers separated by a space. For example, to specify that rays leaving an object may hit objects 4, 6, or 23, enter the string "4 6 23" on the Consider Objects data field. Note the object's own object number should normally be included in the list. ZEMAX will automatically update this table as new objects are inserted or deleted. The maximum number of objects that may be listed is 10. If a ray may hit more than 10 different objects; leave this field blank. Use zero for the object number of the exit port.

If you are unsure as to what objects a ray leaving any object will hit, leave this field blank. If any object is listed; only the listed objects will be tested for ray intercepts! This means incorrect ray tracing results will occur if the list does not explicitly include all possible objects that may be hit by rays leaving the current object. If the ray can leave the object, and strike the same object again, the consider objects list for that object must explicitly list itself as one of the possible objects. For example, if rays striking object 5 can refract or reflect and hit object 5 again, object number 5 must be listed in the consider objects list for object 5 if a consider objects list is defined.

Note that it makes no sense to define both a "Consider Objects" and an "Ignore Objects" list. The Consider Objects list takes precedence over the Ignore Objects list if both lists are defined.

Defining an Ignore Objects list

The Ignore Objects list allows definition of objects that should be ignored when a ray leaves any object. A common application of this feature is for Source Objects to ignore the object from which they were defined, so that the ray does not immediately interact with object located where the ray starts.

The Object Type tab has a user-definable list labeled "Ignore Objects" for doing this. If left blank (the default), then rays leaving this object may strike any object, and ZEMAX will use it's own internal algorithms for determining which object is actually hit.

To specify the objects the ray may NOT hit when leaving an object, list the object numbers separated by a space. For example, to specify that rays leaving an object should ignore objects 5 and 7, enter the string "5 7" on the Ignore Objects data field. Note the object's own object number should normally not be included in the list. ZEMAX will automatically update this table as new objects are inserted or deleted. The maximum number of objects that may be listed is 10. Use zero for the object number of the exit port.

Note that it makes no sense to define both a "Consider Objects" and an "Ignore Objects" list. The Consider Objects list takes precedence over the Ignore Objects list if both lists are defined.

User defined apertures

User defined apertures (UDA) may be placed on some NSC surface objects. UDA's are defined exactly as described for sequential surfaces, see "User defined apertures and obscurations" on page 80 for a complete description. To place an aperture on an object surface, open the Object properties dialog box, choose the "Type" tab, select the "User Defined Aperture" checkbox, and select the UDA data file name from the menu. The UDA file must be placed in the <objects>\Apertures folder (see "Folders" on page 68).

Object faces

A face is a collection of one or more surfaces on an object upon which the same optical properties are applied. For example, for a singlet lens, there are 3 faces: the front face, the back face, and all remaining faces (which include the edges and squared faces around the edges). In this case, 3 different coatings may be applied; one to each face group, if all faces of the lens are coated. Note that more than one surface may be grouped together in a single face number, and all surfaces in the same face group have the same optical properties. The valid face numbers are 0 through 50, although most objects have less than 4 faces.

Polarization and thin film coatings

Rays traced through non-sequential components may be done while accounting for polarization effects, or polarization may be ignored. The initial polarization state for a ray is determined by the source properties, see "Sources tab" on page 358. If polarization ray tracing is being used, transmission, reflection, and absorption of optical energy is accounted for at all surfaces. Bulk absorption is also accounted for. Thin film coatings significantly affect transmission and reflection properties of optical surfaces. Surfaces are initially uncoated, but coatings may be applied to surfaces or group of surfaces.

Coatings on surfaces in contact

If two surfaces are in contact, such as two 45-45-90 prisms placed so that one face of each prism is in contact with the other, then a coating may be applied "between" the surfaces in contact.

This is accomplished using the same convention described above in the section "Object Placement". The rule is: the LAST object listed in the NSC Editor determines the properties of the interface between two objects.

For example, to place a thin metal coating on the interface between two prisms arranged to form a beam splitter, the first prism object listed should have the contact surface "uncoated", while the second object listed should have the contact surface coated with the appropriate thin film coating. Rays striking this interface from either side will see the correct coating, and the ray transmission and/or reflection will be correctly computed. Note prisms modeled as POB objects can have different coatings applied to different faces, so some faces may be anti-reflection coated while others are coated with a reflective coating.

Scattering

Scattering may occur at any ray-surface intersection. Scattering properties are defined for each face on an object. See the section "Object faces" on page 364 for information on faces and face numbers. The default scatter

model is "No scattering", which means no scattering will occur. The resulting ray is called the unscattered or the "specular" ray (even if the surface is not actually reflective).

Fraction to scatter and number of scatter rays


If a scattering model other than "No scattering" or "ABg" is selected, the "Fraction to Scatter" must be defined. This fraction must be between 0.0 (no rays will be scattered) and 1.0 (every ray will be scattered). For the ABg model scattering, the fraction to scatter is determined by the ABg parameters, see "Defining ABg data" on page 368.

If ray splitting is off, the decision to scatter or not to scatter is made by the generation of a single random number between 0.0 and 1.0. If this random number is larger than the fraction to scatter, the ray will not scatter, otherwise, the ray will scatter. For example, if the fraction to scatter is 1.0, the ray always scatters. If the fraction to scatter is 0.0, the ray will never scatter. If the fraction to scatter is 0.25, then on average one out of four rays will scatter. All of the energy of the ray follows the randomly generated scatter path. The number of scatter rays has no affect if ray splitting is off.

If ray splitting is on, then ZEMAX will split the specular ray into one or more scattered rays, while still possibly tracing the specular ray. The specular ray will receive a fraction of the original energy equal to $(1.0 - f)$ where f is the fraction to scatter. The remaining energy will be divided equally among the one or more scattered rays. The number of scatter rays determines how many scatter rays will be generated. For example, if the fraction to scatter is 1.0, then the specular ray will receive zero energy and will no longer be traced; and all the energy will be divided equally among the scattered rays. If the fraction to scatter is 0.0, no scattered rays will be traced, and the specular ray retains all the original energy. If the fraction to scatter is 0.25 and the number of scatter rays is 5, then the specular ray will receive a relative energy of 0.75, and each of the 5 scattered rays will have a relative energy of 0.05. If the number of scatter rays is set to zero, then the fraction to scatter is ignored and no scattering occurs. The maximum number of scatter rays is 300, although much smaller values are typically used.

Scatter models

Scatter models are defined in terms of a probability distribution function. When ZEMAX scatters a ray, a new direction of propagation is chosen. The direction is chosen using a probability function and one or more random numbers. The net effect is that if many rays are traced, the resulting scattered ray distribution would approach the probability distribution function. The following figure defines the vectors used to describe the scattering models.



The normal vector \hat{N} defines the orientation of the surface at the ray-surface interception point. The incident ray vector is \hat{I} , the specular ray vector is \hat{R} , and the scattered ray vector is \hat{S} . The specular vector may be the reflected or the refracted ray vector; the figure shows this ray as the reflected ray vector for simplicity. Note \hat{N} , \hat{I} , \hat{R} , and \hat{S} are all unit vectors. The projection of the specular and scattered ray vectors down to the surface are denoted by $\hat{\beta}_0$ and $\hat{\beta}$, respectively. The projections are not unit vectors; the vector $\hat{\beta}_0$ has a magnitude equal

to $\sin\theta_r$, while $\vec{\beta}$ has a magnitude of $\sin\theta_s$, where θ_r and θ_s denote the angle between the normal vector and the specular and scattered ray angles, respectively. The vector $\vec{\beta} - \vec{\beta}_0$ is denoted \vec{x} , and if $|\vec{x}|$ tends to zero, the scattered and specular vector become the same. How \vec{x} is determined depends upon the scatter model chosen.

The maximum magnitude for the vectors $\vec{\beta}_0$ is $\vec{\beta}$ 1.0. However, because these vectors do not generally point in the same direction the maximum possible magnitude for the vector \vec{x} is 2.0. The vector \vec{x} must also result in a vector $\vec{\beta}$ that lies within the unit circle of the projected plane.

Bi-Directional Scatter Distribution Function

The Bi-Directional Scatter Distribution Function (BSDF) is defined as the scattered radiance per unit incident irradiance, or

$$BSDF(\theta_i, \phi_i, \theta_s, \phi_s) = \frac{dL_s(\theta_s, \phi_s)}{dE_i(\theta_i, \phi_i)}$$

where θ is measured from the normal, and ϕ is the azimuthal angle, and the subscripts i and s refer to incident and scattered directions, respectively. Note BSDF has units of inverse steradians. The BSDF may also be defined as a function of the vector \vec{x} rather than the polar coordinates θ and ϕ . In general, the BSDF is a function of the incident angle and wavelength. The term BSDF can refer to two separate functions, the BRDF and BTDF, for reflective and transmitted distributions, respectively.

The integral of the BSDF over all possible scatter angles (a hemisphere) is called the Total Integrated Scatter (TIS) and is defined by

$$TIS(\theta_i, \phi_i) = \iint BSDF(\theta_i, \phi_i, \theta, \phi) \cos\theta \sin\theta d\theta d\phi,$$

where the i indicates the incident direction. For all scatter models other than ABg, the BSDF is normalized to yield a TIS equal to the "fraction to scatter" parameter. For the ABg model, the TIS must be less than 1.0, and the TIS indicates the total fraction of energy that scatters. All remaining energy is assumed to be specular.

Available scatter models

There are five scattering models available: none, Lambertian, Gaussian, ABg, and user defined. Each available scatter model is defined in the following sections.

No scattering

No modification to the specular ray is made, the BSDF is zero and the magnitude of the vector \vec{x} is zero.

Lambertian scattering

For Lambertian scatter, the scattered ray projection vector $\vec{\beta}$ has equal probability anywhere in the unit circle, and the BSDF is just $1/\pi$. The scattered intensity is proportional to $\cos\theta_s$. Note Lambertian scattering is independent of the ray incident angle. Most diffuse surfaces are nearly Lambertian.

Gaussian scattering

Gaussian scattering is of the form:

$$BSDF(\vec{x}) = A e^{-\frac{|\vec{x}|^2}{\sigma^2}}$$

where A is a normalizing constant. The resulting distribution is rotationally symmetric in direction cosine space, no matter what angle the specular ray makes with respect to the surface normal. The dimensionless value σ determines the width of the Gaussian distribution on the projected plane. Values of σ greater than about 5.0 yield a BSDF that is nearly Lambertian. For this reason, the maximum allowed value of σ is 5.0.

ABg model scattering

The ABg scattering model is a widely used method for defining the BSDF. This scattering model is generally a good model to use when the scattering is mainly due to random isotropic surface roughness, and the scale of the roughness is small compared to wavelength of light being scattered. These assumptions are generally valid for polished optical surfaces. The ABg BSDF is given by


$$BSDF(\hat{x}) = \frac{A}{B + |\hat{x}|^g}$$

The following restrictions are placed on the coefficients of this model: A must be greater than or equal to 0.0, and B must be greater than 1E-12 unless g = 0.0. If g is 0.0, B may be 0.0. If A is 0.0, no scattering will occur. If g is 0.0 (note g may be any value, positive or negative, but it typically is between 0.0 and 3.0), then the BSDF is constant:

$$BSDF(\hat{x}) = \frac{A}{B + 1}$$

and the resulting scattering is effectively Lambertian. If A, B, and g are set to provide a relatively flat BSDF curve, it is far faster numerically to use the Lambertian model instead.

A typical ABg BSDF curve is shown below for the parameters A = 0.002, B = 0.0005, g = 2.0.



The ABg BSDF model has several interesting properties (when g is not zero):

When $|\hat{x}|^g \ll B$, the BSDF curve flattens out to a value of A/B as $|\hat{x}|$ goes to zero.

When $|\hat{x}|^g \gg B$, the BSDF becomes a straight line on the log-log plot, with a slope of -g.

The transition point between the flat and sloped parts of the curve occurs at $|\lambda| = \beta_t$, where $\beta_t = 10^{\left(\frac{\log B}{g}\right)}$.

Defining ABg data

It is frequently the case where the same ABg data will be applied to many objects in the optical system. To avoid the need to type in the same data redundantly, and to simplify the task of editing the ABg data, ZEMAX provides an ABg data catalog. The catalog may be edited directly in ZEMAX; see "ABg Scatter Data Catalogs" on page 250. Note the ABg data defined must integrate over all possible scatter angles to a total value of less than 1.0 to conserve energy.

ZEMAX allows definition of separate ABg data for reflection and refraction. If the specular ray reflects or refracts, the ray is subsequently scattered using the appropriate coefficients. The ABg data names which appear in the dialog box control are those defined by the ABg scatter data file. The ABg data file to be used is selected from the "Files" tab of the General dialog box, see page 109 for details.

ZEMAX first decides whether or not to scatter the ray, as described in the section "Scattering" above. If the ray scatters, a scattered ray is randomly generated. The ray direction is generated so that if a large number of rays were scattered, the appropriate BSDF function would result.

For important information on using scatter models with a source defined by a Source Color spectrum, see "Wavelength scaling of ABg data" on page 251.

User defined scattering

Completely general surface scattering may be defined via an external program called a Dynamic Link Library (DLL). Sample DLLs are provided with ZEMAX with source code. New DLLs may be easily created with a suitable compiler. See also "Comments about DLLs" on page 392.

Defining an object to use DLL defined surface scattering

To make an object use a DLL defined scatter function, choose the "User Defined" scattering type from the object properties dialog box, then select from the available DLL scatter functions listed in the "DLL Name" box.

DLL parameters

Each DLL may use up to 6 user defined data values as parameters in the computation of the scattering properties. These values are defined by the DLL and are only used by the DLL.

Creating a new DLL

The DLL must include two functions:

UserScatterDefinition

UserParamNames

When a ray is scattered, ZEMAX passes to the UserScatterDefinition function the local x, y, and z coordinates on the surface, the local normal vector, the local specular ray vector, polarization data, and other parameter data as defined by the user. UserScatterDefinition then is required to determine the following values:

Whether or not the ray actually scatters

The direction cosines of the scattered ray

The attenuation of the ray, if any

Optionally the new electric field vector (if this is not provided ZEMAX will make a reasonable guess)

Optionally the BSDF and TIS if Importance Sampling is used (see "Importance Sampling" on page 370).

These values are returned to ZEMAX and are used to continue the trace. The function UserParamNames is used to define the names of all used parameters. These names appear in the Coating/Scattering tab of the object properties dialog box. Surface scattering DLLs must be placed in the <program>\DLL\SurfaceScatter folder. See "Folders" on page 68.

Note that the "Fraction to scatter" control is active when User Defined scatter is selected. Only the fraction of energy which is allowed to scatter is sent to the DLL for possible scattering. However, the DLL may choose to not scatter the ray (see for example the sample DLL in the next section).

More details can be found in the source code of the sample surface scattering DLL files provided with ZEMAX.

Sample surface scattering: TwoGaussian.DLL

The surface scattering sample TwoGaussian.DLL supports a scatter model which is the superposition of 4 scattering models: specular, Lambertian, and two individual Gaussian models. The DLL uses the parameter data to define the fraction of energy which scatters according to the Gaussian Sigma 1, the fraction which scatters according to Sigma 2, and the fraction that is Lambertian. The remaining fraction, if any, is specular. To see how this is accomplished in the DLL, review the source code file Two Gaussian.c provided with ZEMAX.

Modeling measured BSDF data with the RI_BSDF scattering DLL

RI_BSDF.DLL is a DLL surface scattering model which supports the use of tabular BSDF data for defining the scattering properties of a surface. Data are provided to the DLL via a text file. This file must follow the BSDF Data Interchange file format developed by Radiant Imaging, Inc. A full description of that format is provided in the Knowledge Base article "BSDF Data Interchange File Format Specification" available on the ZEMAX web site www.zemax.com.

To use the DLL, choose "RI_BSDF.DLL" for the "DLL Name:". Any input files used by the DLL will appear in a list under the "File Name:" option. For a file to appear in the list, it must have a .BSDF extension (as indicated in the article describing the file format), and be located in the <data>\Miscellaneous folder (see "Folders" on page 68). The maximum number of characters that the file name can contain (including the .BSDF extension) is 60. A full description of the model and its use is provided in the Knowledge Base article "How to Use Tabular Data to Define the Surface Scattering Distribution" also available on the ZEMAX web site.

Sample surface scattering: K-correlation.DLL

The surface scattering sample K-correlation.DLL supports a scatter model which follows the K-correlation distribution, as described in the paper "K-correlation power spectral density & surface scatter model" by Michael G. Dittman (Proc. of SPIE Vol. 6291, 62910R, 2006). The inputs to the DLL include the surface roughness (in micrometers), the characteristic surface wavelength (multiplied by 2π), and the logarithmic slope of the BSDF. More details are provided in the article entitled "How to Model Surface Scattering via the K-correlation Distribution" in the ZEMAX Knowledge Base. The source code K-correlation.c is provided with ZEMAX as a sample file.

How to model scattering efficiently

A very large number of rays may need to be traced to find a relatively small number of scattered rays that strike an object of interest, such as a detector. There are two ways to improve the efficiency of the scattering analysis.

The first method is to scatter a ray according to the scatter distribution, but only trace the ray if the ray propagates towards an object of interest. This method may be implemented by defining a "Scatter To" list of objects. The Scatter To method works well for wide angle scatter (such as Lambertian scatter) and when the object of interest subtends a relatively large angle as seen from the scattering surface.

The second method is to always scatter the ray towards the object of interest, and then to normalize the energy the ray carries to account for the probability the ray would have actually scattered in that direction. This method is called "Importance Sampling". Importance Sampling is generally superior to the Scatter To method if the scatter is narrow angle or the object of interest subtends a relatively small angle as seen from the scattering surface.

The Scatter To list

The Scatter To feature speeds up scattering analysis by ignoring scattered rays which do not propagate *directly* towards an object of interest. The Scatter To list is very similar to the Consider Objects list (see "Defining paths using the Consider Objects list" on page 363). The Scatter To list is a string of integer object numbers separated by spaces. A scattered ray will only be traced if the ray intersects one of the objects listed or the object that scattered the ray. For example, suppose object 2 defines the Scatter To list as "3 4". If a ray scatters off object 2, it will only be traced if the ray will intersect objects 2, 3, or 4. Note that any other objects are ignored, so scattered rays will miss any object not explicitly listed (or the scattering object itself), even if that object lies directly in the path between the scatter point and the listed scatter to object.

Choosing an object on the Scatter To list does not guarantee that a scattered ray will be traced to that object. The scattered rays are generated based upon the scattering properties of the object the incident ray strikes (see “Scattering” on page 364). ZEMAX will generate the scattered rays, then ignore those rays that do not intersect any of the listed objects. This method does not ensure that a ray will scatter toward a listed object. However, the number of scattered rays (see “Fraction to scatter and number of scatter rays” on page 365) may be made large so that it is more probable some of them will head in the desired directions.

To turn off the Scatter To feature, leave the list blank. In this case, all scattered rays will be traced. The key difference between the Scatter To list and the Consider Object list is that the Scatter To list only applies to scattered rays, while the Consider Object list applies to both specular and scattered rays.

Importance Sampling

Importance Sampling speeds up scattering analysis by *always* choosing scattered rays that will propagate in a desired direction. To determine the direction for the scattered ray, a target sphere is defined. The target sphere is defined by position, size, and limit. The target sphere position is defined by the selection of an object number. The target sphere will then be placed at the origin of the selected object. Any object type may be used, including Null objects. The size parameter defines the radius in lens units of the target sphere. The limit parameter is used to define the maximum solid angle in Steradians of the target sphere as seen from the scatter point. If the solid angle of the target sphere as seen from the scatter point exceeds the limit value, Importance Sampling will not be used for that scattered ray. The purpose of this feature will be discussed below.

When the ray scatters from the surface, the direction and subtended solid angle of the target sphere as seen from the scatter point are computed. A ray direction is randomly selected from within this solid angle cone with uniform probability. This will be the scattered ray direction. Note that the direction is not dependent upon the BSDF of the surface.

Because the ray direction is determined without consideration of the BSDF, the energy the ray carries must be modified to get correct ray tracing results. The amount of energy the scattered ray will carry depends upon the BSDF and TIS of the surface as follows:

$$E' = E \frac{BSDF}{TIS} \Omega \cos \theta ,$$

where E is the energy allocated to the scattered ray before importance sampling is applied, E' is the amount of energy in the ray traced towards the target sphere, Ω is the solid angle of the target sphere as seen from the scatter point, θ is the angle between the local normal and the scattered ray, and BSDF and TIS are the usual scattering function and total integrated scatter (see “Bi-Directional Scatter Distribution Function” on page 366). Note this energy renormalization is only accurate if the BSDF is reasonably uniform over the solid angle of the target sphere. If Ω becomes large compared to the variation of the BSDF, the energy normalization will not be accurate. For this reason, a limit value may be defined. If the target sphere is close enough to the scatter point so that Ω exceeds the limit, then Importance Sampling is not used for that scattered ray.

Up to 6 independent target spheres may be defined for each object. If one scatter ray is selected, the first target sphere is used. If two scatter rays are selected, then the first two target spheres are used, and so on up to 6 scattered rays. If more than 6 scattered rays are selected, the seventh and subsequent rays scatter according to the BSDF and Importance Sampling is not used. Energy normalization is complex and somewhat prone to statistical errors when the number of scattered rays is not equal to the number of target spheres. For the fastest, most accurate results, use a number of scattered rays that is exactly equal to the number of target spheres.

Bulk scattering

Bulk (volumetric) scattering models random scattering of rays while propagating through a solid object. ZEMAX supports 3 modes for the bulk scattering of a solid object:

- No Bulk Scattering
- Angle Scattering
- DLL Defined Scattering

Each model is described in a subsequent section.

No Bulk Scattering

If no bulk scattering is selected, rays will propagate through the solid without scattering.

Angle Scattering

Angle scattering uses a simple model for scattering within a solid. Rays traveling a distance x within the media have an integrated probability of having been scattered given by

$$p(x) = 1.0 - e^{-\mu x}, \text{ where}$$

$$\mu = \frac{1}{M},$$

and the symbol M is the mean free path in lens units. Note that as x increases, the probability that the ray has scattered asymptotically approaches 1.0. Setting this expression to a randomly selected value between 0.0 and 1.0, then solving for x , yields a randomly generated path length with the correct statistics. If this path length is greater than the distance the ray propagates to the next ray-object intersection, then no scattering occurs; otherwise, the ray scatters at the specified position along the propagation direction.

Once the position is determined, the scattering is modeled by a new ray direction being chosen. For angle scattering, a random angle is chosen so that the new ray direction lies uniformly distributed within a cone that makes an angle with respect to the current ray direction. The semi-angle of the cone is one of the parameters to the model, and this parameter should be set to between 0.0 and 180 degrees. The latter value will scatter the ray randomly in any direction. Once the ray scatters, ZEMAX automatically will adjust the polarization vector and randomize the phase of the scattered ray.

DLL Defined Scattering

If the angle scattering model is not sufficient, more complex bulk scattering functions may be defined via an external program called a Dynamic Link Library (DLL). Sample DLLs are provided with ZEMAX with source code. New DLLs may be easily created with a suitable compiler. See also "Comments about DLLs" on page 392.

Sample bulk scattering: Poly bulk scat.DLL

The sample bulk scattering file `poly_bulk_scat` defines the scatter probability as follows:

$$P(\theta) = \sum_{i=0} c_i \theta^i,$$

where θ is the polar angle of the scattered ray with respect to the unscattered ray direction. The source code file `poly_bulk_scat.c` is provided with ZEMAX as a sample file.

Sample bulk scattering: Henyey-Greenstein-bulk.DLL

The Henyey-Greenstein bulk scatter model defines the scatter probability as follows:

$$P(\theta) = \frac{1 - g^2}{4\pi(1 + g^2 - 2g \cos\theta)^{3/2}},$$

where θ is the polar angle of the scattered ray with respect to the unscattered ray direction. The parameter g can be any value from -1 to 0.9999. If the input value for g is less than -1, g is set to -1. If the input value for g is greater than 0.9999, then g is set to 0.9999. If the input value for $|g|$ is less than 1.0e-4, then g is set to 1.0e-4. The source code file `Henyey-Greenstein-bulk.c` is provided with ZEMAX as a sample file.

Sample bulk scattering: Rayleigh.DLL

The Rayleigh bulk scattering model defines the scatter probability according to Rayleigh theory for an unpolarized input beam:

$$P(\theta,\lambda) = (3(1 + (\cos\theta)^2))/(8\lambda^4),$$

where θ is the polar angle of the scattered ray with respect to the unscattered ray direction and λ is the wavelength. The inputs to the DLL are the Reference Wavelength (i.e. the wavelength associated with the input value for the Mean Path) and the Transmission (which allows for energy to be lost during each scattering event). More details are provided in the article entitled "Bulk Scattering with the Rayleigh Model" in the ZEMAX Knowledge Base. The source code file Rayleigh.c is provided with ZEMAX as a sample file.

Sample bulk scattering: Mie.DLL

The Mie bulk scattering model defines the scatter probability according to Mie theory. The algorithm used in ZEMAX was taken from the book "Absorption and Scattering of Light by Small Particles" by Craig F. Bohren and Donald R. Huffman (John Wiley & Sons, 1983). The inputs to the DLL include the particle index, the particle size (in micrometers), and the particle density (in cm⁻³). More details are provided in the article entitled "How to Simulate Atmospheric Scattering using a Mie model" in the ZEMAX Knowledge Base. The source code file is not distributed with ZEMAX.

Sample bulk scattering: Phosphor.DLL

The Phosphor bulk scattering model is meant to represent a phosphor material in which both fluoresce and scattering can occur. Light in a particular wavelength range (defined as "blue") enters the medium in which the DLL is applied, and that light can either be converted to another wavelength (i.e. undergo fluorescence) or simply bulk scatter. When the input light undergoes fluorescence, the resultant light is emitted into a 180 degree cone angle with equal probability (i.e. the input light "scatters" according to the Angle Scattering model). When the input light scatters with fluorescing, it does so according to the Mie distribution. The fluoresced light (or any other input light not in the "blue" wavelength range) may also scatter according to the Mie distribution. The inputs to the DLL include the mean-free path for fluorescence, the wavelength range defining the "blue" wavelength range, and the particle parameters associated with the Mie distribution. More details are provided in the article entitled "Modeling a White Light Source using a Phosphor" in the ZEMAX Knowledge Base. The source code file is not distributed with ZEMAX.

Defining an object to use DLL defined bulk scattering

To make an object use a DLL defined bulk scatter function, choose the Bulk Scattering tab from the object properties dialog box, then select Model: "DLL Defined Scattering". The available DLL functions are then listed in the "DLL:" control.

DLL parameters

Each DLL may use between zero and 16 user defined data values as parameters in the computation of the media properties. These values are defined by the DLL and are only used by the DLL.

Creating a new DLL

The DLL must include two functions:

UserBulkDefinition

UserParamNames

While ray tracing through a solid using DLL defined bulk scattering, ZEMAX passes to the UserBulkDefinition function the current propagation length along the ray, various ray data values, and other parameter data as defined by the user. UserBulkDefinition then is required to determine if the ray will scatter at some point along the propagation length. If the ray will scatter, then UserBulkDefinition must determine the following values:

The position at which the ray will scatter

The new direction cosines of the ray

The attenuation of the ray, if any

Optionally the new electric field vector (if this is not provided ZEMAX will make a reasonable guess)

These values are returned to ZEMAX and are used to continue the trace. The function UserParamNames is used to define the names of all used parameters. These names appear in the Bulk Scatter tab of the object properties dialog box. Bulk scattering DLLs must be placed in the <program>\DLL\BulkScatter folder. See "Folders" on page 68.

Coherence length modeling

By default, non-sequential sources in ZEMAX are either purely monochromatic, or simulated to be polychromatic by tracing more than one monochromatic wavelength. Physically, no source is perfectly monochromatic. A more realistic model for a nominally monochromatic source is to account for the range of wavelengths emitted by the source. A convenient parameter is the coherence length; defined as

$$\Delta x = \frac{c}{\Delta f},$$

where Δx is the coherence length, c is the speed of light in vacuum and Δf is the range of frequencies emitted by the source, also called the bandwidth. A more convenient representation is the range of wavelengths emitted defined in terms of the coherence length:

$$\Delta \lambda = \frac{\lambda^2}{\Delta x}.$$

Using this expression, it is easy to compute the variation of optical path length (phase) as the ray propagates. This is modeled in ZEMAX by randomly choosing a wavelength within $\pm(\Delta \lambda)/2$ for purposes of computing the optical path and thus the phase. ZEMAX chooses the random wavelength with a uniform distribution. The greater the path difference between two rays from the same source, the greater the difference in phase between the rays will be; which means coherent effects such as fringe visibility will be reduced.

ZEMAX actually traces the ray using the central wavelength. The randomized wavelength is only used to compute the effects of coherence length on the phase. Therefore, only coherent computations will exhibit the phase degradation. To turn off coherence degradation, either set the coherence length to a very large number or to zero. To make sources completely incoherent, set the coherence length to a small but non-zero number.

Defining GRIN media for non-sequential ray tracing

Any solid enclosed volume may be made of a homogeneous material (the default assumption) or the media may be of a gradient index. All gradient index materials used by non-sequential objects are defined in a separate program called a Dynamic Link Library (DLL). Numerous DLLs are provided with ZEMAX with source code. New DLLs may be easily created with a suitable compiler. See also "Comments about DLLs" on page 392.

Defining an object to be of gradient index material

To make an object media by of a gradient index material, choose the GRIN tab from the object properties dialog box, then select "Use DLL defined GRIN media". The available gradient index materials are then listed in the "DLL:" control.

Discussion on maximum step size for GRIN objects

The maximum step size determines the tradeoff between ray tracing speed and accuracy. The exact value required depends upon the rate at which the index of refraction changes and the desired accuracy of the computations. If the maximum step size is too large, the ray trace results will have large errors, and some rays will miss objects they would otherwise hit, or vice-a-versa.

GRIN DLL parameters

Each DLL may use between zero and 12 user defined data values as parameters in the computation of the media properties. These values are defined by the DLL and are only used by the DLL.

Creating a new GRIN DLL

The DLL must include two functions:

UserGrinDefinition

UserParamNames

While ray tracing through a GRIN media, ZEMAX passes to the UserGrinDefinition function the current x, y, and z location of the ray (in the objects local coordinate system), the wavelength, and other parameter data as defined by the user. UserGrinDefinition then is required to compute the following four values:

n (the index of refraction)

n * dn/dx (the index times the derivative with respect to the x direction)

n * dn/dy (the index times the derivative with respect to the y direction)

n * dn/dz (the index times the derivative with respect to the z direction)

These four values are returned to ZEMAX and are used to trace the ray through the grin media.

The function UserParamNames is used to define the names of all used parameters. These names appear in the GRIN tab of the object properties dialog box. GRIN DLLs must be placed in the <program>\DLL\GradientIndex folder. See "Folders" on page 68.

Sample GRIN DLLs

The following sample GRIN DLLs are provided with ZEMAX.

SAMPLE GRIN DLLS

DLL name	Description
GRIN1	<p>Defines a GRIN media of the form:</p> $n = N_0 + N_{r1}r + N_{r2}r^2,$ <p>similar to the sequential Gradient1 surface type. This sample DLL also supports array dx and array dy values to specify a grin variation which is periodic over the specified intervals. This option is used for arrays of gradient index lenslets. If both the array dx and array dy values are zero, the gradient is not periodic.</p>
GRIN2	<p>Defines a GRIN media of the form:</p> $n^2 = n_0 + n_{r2}r^2 + n_{r4}r^4 + n_{r6}r^6 + n_{r8}r^8 + n_{r10}r^{10} + n_{r12}r^{12},$ <p>similar to the sequential Gradient2 surface type.</p>
GRIN3	<p>Defines a GRIN media of the form:</p> $n = n_0 + n_{r2}r^2 + n_{r4}r^4 + n_{r6}r^6 + n_{z1}z + n_{z2}z^2 + n_{z3}z^3,$ <p>similar to the sequential Gradient3 surface type.</p>
GRIN4	<p>Defines a GRIN media of the form:</p> $n = n_0 + n_{x1}x + n_{x2}x^2 + n_{y1}y + n_{y2}y^2 + n_{z1}z + n_{z2}z^2,$ <p>similar to the sequential Gradient4 surface type.</p>
GRIN5	<p>Defines a GRIN media of the form:</p> $n = n_0 + n_{r2}r^2 + n_{r4}r^4 + n_{z1}z + n_{z2}z^2 + n_{z3}z^3 + n_{z4}z^4,$ <p>similar, <i>but not identical to</i>, the sequential Gradient5 surface type. Note there is no dispersion in this grin model (no dependence on wavelength).</p>

DLL name	Description
GRIN6	<p>Defines a GRIN media of the form:</p> $n = n_0 + n_1 r^2 + n_2 r^4, \text{ where}$ $n_x = A_x + B_x \lambda^2 + \frac{C_x}{\lambda^2} + \frac{D_x}{\lambda^4}, \text{ for all three values of } n_x,$ <p>which is similar to the sequential Gradient6 surface type.</p>
GRIN9	<p>Defines a GRIN media of the form:</p> $n = n_0 \left[1.0 - \frac{A}{2} r^g \right], \text{ where}$ <p>where both A and n_0 are functions of wavelength:</p> $A(\lambda) = \left[K_0 + \frac{K_1}{\lambda^2} + \frac{K_2}{\lambda^4} \right]^2, \text{ and}$ $n_0 = B + \frac{C}{\lambda^2}.$ <p>This is similar to the sequential Gradient9 surface type, with the additional parameter "g" which is fixed at 2.0 for the sequential surface Gradient9.</p>
GRIN10	<p>Defines a GRIN media of the form:</p> $n = n_0 + n_{y1} y_a + n_{y2} y_a^2 + n_{y3} y_a^3 + n_{y4} y_a^4 + n_{y5} y_a^5 + n_{y6} y_a^6,$ <p>where $y_a = y$ and the symbols indicate the absolute value. This form of gradient has a discontinuity at the plane $y = 0$, and the gradient is bisymmetric about the $y = 0$ plane. This is similar to the sequential Gradient10 surface type, but without the dispersion model.</p>
TORUS	<p>Defines a GRIN media of the form:</p> $n = N_0 + N_{r2} t^2,$ <p>where t is the distance from an axis of rotation defining the toroidal symmetry of the gradient index. The torus radius is one of the defining parameters. The torus radius is measured from the coordinate $z = 0.0$ on the object to the location of the axis of rotation defining the torus.</p>

Defining DLLs for ray splitting at diffractive surfaces

Any diffractive surface may split rays that started at a non-sequential source, by order, to allow for simultaneous tracing of multiple diffracted orders. How many rays split, and how much energy is in each order, may be defined externally in a separate program called a Dynamic Link Library (DLL). Numerous DLLs are provided with ZEMAX with source code. New DLLs may be easily created with a suitable compiler. See also "Comments about DLLs" on page 392.

Defining an object to use the diffraction DLL

To make a diffractive object use the DLL, choose the Diffraction tab from the object properties dialog box, then select "Split by DLL function". The available DLLs are then listed in the "DLL:" control.

Diffraction DLL parameters

Each DLL may use between zero and 12 user defined data values as parameters in the computation of the diffractive surface properties. These values are defined by the DLL and are only used by the DLL.

Creating a new Diffraction DLL

The DLL must include two functions:

UserDiffraction

UserParamNames

While ray tracing through a diffractive surface, ZEMAX passes to the UserDiffraction function the current x, y, and z location of the ray (in the object's local coordinate system), the wavelength, the present order being traced, and other ray specific data and parameter data as defined by the user. UserDiffraction then is required to compute the relative energy in the specified order, and return this value to ZEMAX. ZEMAX will refract or reflect the ray into the specified diffraction order and give the resulting ray the specified fraction of the input ray energy computed by the DLL. If the new child ray has sufficient energy, the ray is traced.

The Diffraction DLL format also supports as an option for the DLL to compute ALL output properties of the ray; including direction cosines, electric field polarization vector, and transmitted energy. If the DLL also computes these properties, the DLL must set a flag indicating the values were passed back to ZEMAX for further ray tracing. See the sample DLLs for documentation on this feature.

The function UserParamNames is used to define the names of all used parameters. These names appear in the Diffract tab of the object properties dialog box. See the sample DLLs for the use of this feature. Diffraction DLLs must be placed in the <program>\DLL\Diffractive folder. See "Folders" on page 68.

Ray splitting

Generally, when a ray strikes a surface, part of the energy will be reflected, part will be transmitted, and depending upon the surface properties, part may be absorbed. Ray splitting refers to the ability of ZEMAX to compute both the reflected and refracted paths, and then continue to trace both rays. ZEMAX also supports an option to randomly choose either the reflected or the refracted path rather than both, see "Simple Ray Splitting" on page 114. Rays reflected from otherwise refractive interfaces are commonly called ghost reflections.

Once a ray splits, each of the "child" rays will in general strike another object, and the rays may split again and again. After many ray-object intersections, the total number of rays can become extremely large, so controls must be placed on the ray tracing to ensure that the computation will eventually end.

There are several ways to restrict the amount of ray splitting that occurs:

Maximum number of ray-object intersections: This control defines how many times a ray along any path, from the original source parent ray to the final ray-object intersection, may intersect another object.

Maximum number of ray segments: A ray segment is that portion of a ray path from one intersection to the next. When a ray is launched from the source, it travels to the first object. That is 1 segment. If the ray then splits into 2 rays, each of those are another segment (for a total of 3). If each of those rays split again, there will be 7 segments. Generally, the number of ray segments grows far faster, and needs to be set much larger, than the number of ray-object intersections does.

Minimum relative ray intensity: As each ray splits, the energy decreases. The relative ray intensity is a lower limit on how much energy the ray can carry and still be traced. This parameter is a fraction, such as 0.001, relative to the starting ray intensity from the source. Once a child ray falls below this relative energy, the ray is terminated.

Minimum absolute ray intensity: This parameter is very similar to the minimum relative ray intensity, except it is absolute in system source units rather than relative to the starting intensity. If this is zero, the absolute ray intensity threshold is ignored. The initial intensity of each ray is always given by the source intensity divided by the total number of analysis rays for that source. The number of layout rays is not used to determine the initial intensity of rays, even for those rays drawn on layout plots. Source units are defined on the Units tab of the System, General dialog box. See "Source Units" on page 105 for details.


All of these settings are defined on the Non-Sequential tab of the System, General dialog box. See "Non-Sequential" on page 113 for details.

Ray splitting and polarization

Because accurate reflection and transmission computation requires polarization information, ray splitting is only allowed when performing polarization ray tracing.

Ray splitting can be turned off, and in this case the transmitted path is always taken at a refractive interface unless the ray TIR's; the reflected path is of course always taken if the surface is a mirror.

The figure at right shows some of the ray paths possible when rays are split. There is only 1 input ray drawn.



When using ray splitting, the number of rays being traced can become enormous if the minimum relative ray energy is set too low. For example, for rays traced inside a cube of glass with 50% coatings on all surfaces, a relative ray intensity of 0.01 will yield about 256 ray segments. If the relative ray intensity is set to 1E-8, about 270 million ray segments will be generated for EACH ray traced into the cube! Systems with relatively low reflectivity (such as AR coated optics) will not generate as many ray segments because the intensity of the reflected paths falls so much more quickly. Still, it is advised to set the relative ray transmission reasonably high, around 0.001, until the model is working well and more detailed results are needed.

Because the total number of rays gets so large, the 3D Layout plots can become very cluttered. One way to reduce the amount of rays drawn is to use the "filter" string to suppress drawing of some of the rays, see "The filter string" on page 386.

Putting it all together

The real benefit of the non-sequential ray tracing, polarization, scattering, and ray splitting features occurs when they are all used together to determine ray distributions and detected energies. The order in which these effects are all accounted for is as follows:

- The object interface the ray strikes is computed first.
- If the surface is reflective or the ray TIR's, the reflected path is chosen.
- Otherwise,
- If ray splitting is on, the reflected and the refracted path are both determined.
- If ray splitting is off, the refracted path is chosen.
- The ray(s) are scattered if scattering is enabled.
- The ray(s) intensity and polarization vectors are adjusted.
- The ray(s) continue on to the next object interface.

On 3D Layouts and other analysis features which draw NSC objects and rays, the effects of sources, ray splitting, and scattering may be seen from the randomly generated rays on these plots. Both ray splitting and ray scattering may be selected on or off on these plots to isolate their effects. The layout plots are extremely useful for qualitatively seeing the effects of sources, objects, scattering, and ghost or other reflections.

To accurately compute quantitative energy distributions on the placed detectors generally requires a larger number of rays than the qualitative analysis of the layout plots. This is why the number of layout and number of analysis rays are defined separately in the source parameter listing.

The Ray Trace Control

The Ray Trace Control is used to launch and trace rays for analysis using defined sources, objects, and detectors. The controls on the dialog box are described below.

RAY TRACE CONTROL OPTIONS

Item	Description
Clear Detectors	Select either "All" or any one specific detector to clear. This resets all of the detector's pixels to zero energy.

Trace	Begins tracing rays from all defined sources. Once the trace begins, only the Terminate and Auto Update controls are available. See the discussion below on other effects of initiating a ray trace.
Terminate	Stops ray tracing. This action may require a brief time.
Auto Update	If checked, all detectors will be redrawn periodically so progress may be monitored.
# CPU's	Selects the number of CPU's over which to spread the ray tracing task. More than 1 may be selected, even on a single CPU computer, in which case the single CPU will time share the multiple simultaneous tasks. The default is the number of processors detected by the operating system. Using more CPU's may speed the computation but requires significantly more memory. If ZEMAX reports insufficient memory errors or begins using the hard disk as a memory cache, try reducing the number of CPU's.
Use Polarization	If checked, polarization will be used to determine ray energy, reflection, transmission, absorption, and thin film effects. If Split Rays is checked, Use Polarization will be checked automatically, as it is required for accurate ray splitting.
Split Rays	If checked, rays will split according to the interface properties as described in "Ray splitting" on page 376.
Scatter Rays	If checked, rays will scatter according to the scatter properties of the interface.
Ignore Errors	If checked, errors generated by ray tracing are ignored, and any errors that occur will not terminate the ray trace. This should be turned on with caution, since errors may indicate a serious flaw in the system being traced. However, some perfectly good systems occasionally have a few rays fail because of finite computer precision or other minor problems. If the "lost energy (errors)" given at the end of the ray trace is small relative to the total power of all sources defined, then the few rays which trigger errors can be safely ignored. See "Lost energy" on page 379.
Save Rays	If checked, ray data will be saved to the file name specified. The file name only should be provided with the proper extension and no folder name. See also "Filter" below.
Filter	This option is only available if "Save Rays" is checked on. If the filter is not left blank, then the defined filter will be applied to each ray tree traced before it is saved to a file. Only if the ray tree passes the filter will the ray tree be saved. This may substantially reduce the amount of data saved, and therefore may reduce the resulting ZRD or DAT file size. For a description of the filter syntax, see "The filter string" on page 386.
Exit	Closes the dialog box.

There are some important considerations to understand when performing a Monte Carlo ray trace. When the detectors are cleared, all the energy counts in each pixel are reset to zero. Once the Trace button is selected, ZEMAX will trace exactly the number of specified analysis rays for each source listed in the NSC Editor. Each source has an associated power in watts or other units ("Analysis Units" on page 105). The initial intensity of each ray from a given source is the total power divided by the number of rays to be traced. For example, for a 1 watt source tracing 1,000,000 rays; each ray initially has 1.0E-06 watts of power associated with it.

If the ray tracing is terminated before all the rays are traced; the total power and peak power as listed on the detector objects will be inaccurate; although the ray distribution will be accurate for the number of rays traced. The analysis cannot be paused and resumed; so once a ray trace is terminated, selecting Trace again will start a new trace.

If the detectors are cleared, then Trace is pressed, the ray trace completes, and Trace is pressed again; all detectors will show the total power from both traces. ZEMAX has no reliable way of knowing how many times rays have been traced between detector clearings; since sources may be added, deleted, or modified at any time.

For accurate power computations, the detectors should be cleared before tracing; then a single trace should be performed without early termination.

Saving ray data to a file

If the Save Rays option is checked, some or all of the traced ray data may be saved to a file. The nature of the data saved depends upon the file name provided.

If the file extension is ZRD, the complete (optionally filtered) ray data will be saved. The folder the file will be placed in is the same as the lens file being traced. The ZRD format is used by ZEMAX in the ray data base viewer (see "The Ray Database Viewer" on page 382). For details on the ZRD format, see "Ray database (ZRD) files" on page 380.

If the extension is DAT or SDF, then only rays (optionally filtered) that strike a particular object are saved to a source file. This format is used by Source File objects (see "Source File" on page 454). To use the source file format, the file name must be of the form n-filename.DAT or n-filename.SDF where n is an integer corresponding to the object number desired. In this case, the filename.DAT or filename.SDF (with the n- part removed) will be created listing rays that hit object n. For example, the file name "11-ReferenceObject.DAT" will save all rays that hit object 11 into the file "ReferenceObject.DAT". The file will be placed in the folder <objects>\Sources\Source Files (see "Folders" on page 68). If the file extension DAT is used, the ray data is stored in the older "flux only" format and the wavelength data is lost. This results in smaller files and is recommended for monochromatic sources only. If the file extension SDF is used, the ray data is stored in the newer "spectral color" format, and wavelength data is stored with every ray. For more information on source file formats see "Binary Source File format" on page 456.

Saving rays slows down the calculation significantly, as writing data to a file is much slower than the ray tracing. For this reason the Save Rays feature should only be used when the data is needed for subsequent analysis.

Lost energy

When tracing some NSC systems, not every ray can be traced successfully. When a ray trace is terminated, the energy associated with the ray is called "lost energy". The lost energy is measured in source units (see "Source Units" on page 105). There are different causes for terminating the trace of a particular ray:

- If a ray falls below the minimum energy threshold for tracing, the ray will be terminated. The minimum relative and absolute energy thresholds are described in "Non-Sequential" on page 113. The ray trace control reports this type of lost energy as the energy lost due to threshold settings.

- Sometimes geometry errors or round-off errors lead to conditions where the ray can no longer be traced. Geometry errors are usually caused by poorly formed solids or improper nesting of objects. For some rays, correct intercept or refraction data cannot be computed. This may happen when iteration to a surface fails, or the coating data cannot be computed. Rays that hit the maximum number of object intercepts, nesting levels, or segment levels, are also terminated with an error condition. The ray trace control reports these types of lost energy as the energy lost due to errors.

- Rays that strike the edge of an object may be terminated if the local refraction and/or reflection cannot be computed accurately. A ray is considered to strike the edge of an object if the intercept point lies within a threshold distance to the edge, or if two or more objects are struck at the same point and the objects have non-parallel normals at the intercept point. For most objects, the threshold distance is *approximately* the glue distance in lens units. Along the edge of a faceted object, the threshold distance is *approximately* the glue distance (interpreted here as a dimensionless number) multiplied by the dimensions of the facet. Note that rays that strike very close to the edge of a part have no useful physical interpretation anyway; an actual object would diffract (or scatter) all energy incident within a few wavelengths of the edge. Rays that miss all objects, or are absorbed by either coatings, bulk transmission, or detectors, are also terminated. ZEMAX absorbs all these rays without error or report, as the energy is not considered lost because the energy is known to be absorbed or propagated out of the system and is thus accounted for.

If the fraction of lost energy is relatively small, lost energy can be ignored. Even well designed and accurate optical models will occasionally have spurious ray errors. Rays that strike within edges or seams between parts and are thus absorbed will usually represent an insignificant fraction of the total energy in the system. If the lost energy is unacceptably high, then the settings (see for example "Minimum Relative Ray Intensity" on page 114) or the object geometry needs to be checked. For assistance in studying geometry errors, see "Create Source Ray From Last Geometry Error Tool" on page 98.

Ray database (ZRD) files

Ray data may be saved for subsequent analysis by checking on "Save Rays" and providing a file name with a ZRD extension in the edit box provided. As the rays are traced, the ray data will be saved to the file specified. Saving the ray data may slow down the ray tracing significantly, because of the (relatively) slow speed of the hard disk access. However, it is often convenient to trace the rays once; then do multiple analysis on the resulting data. Rays contained in a ZRD database may be drawn on layout plots rather than tracing new rays; see "NSC 3D Layout" on page 125 for details.

ZEMAX supports three different ZRD file formats, each of which offers various combinations of file size, compression, and data content. The ZRD file formats are binary to keep the file as compact as possible.

 ***All ZRD formats use the same file extension "ZRD". The three letter format names are only used to distinguish the formats in the documentation, and are not used as file extensions.***

These are the three different ZRD file formats available:

Uncompressed Full Data (UFD) - All data from the ray trace in an uncompressed format. This is the format to use if you have an external program that reads or writes ZRD files. UFD format makes the largest files, and is also the slowest to read and write. For a detailed description of the UFD format see "The ZRD Uncompressed Full Data Format (UFD)" on page 380.

Compressed Basic Data (CBD) - This format stores a subset of the full data, including only status, level, hit object, parent, xybin, lmbin, x, y, z, and intensity. The ray coordinates and intensity are stored only to floating point (32 bit) rather than the usual double (64 bit) precision to save disk space. Although direction cosine data is not explicitly stored in the file, the cosines are automatically recreated from the x, y, and z coordinate data during the file read. CBD files are dynamically compressed in a way that makes the file unsuitable for read/write operations outside of ZEMAX. CBD results in the smallest file size and fastest writing. Generally a file in CBD format will be about 10% as large as the UFD format. Note that data missing from this format is not available for subsequent analysis when reading the ZRD file.

Compressed Full Data (CFD) - All data from the ray trace in a compressed format. All double precision numbers are converted to floating point precision to save space, but in all other respects this format contains all the data from the ray trace, including polarization, index, phase and other data. This format is about 20% as large as the UFD format if polarization is not used, and about 35% as large if polarization is used.

ZRD Format And Version Numbers

ZRD files contain a version number and format number. ZEMAX will check these numbers when reading the file. Older version files are automatically converted by ZEMAX into the current version upon reading. The current version number is 2000. The format number is 0 for UFD format, 10000 for CBD format, and 20000 for CFD format. The version and format numbers are added together and stored in the first integer record in the file, thus the first integer will be 2000 for UFD format, 12000 for CBD format, and 22000 for CFD format ZRD files.

The ZRD Uncompressed Full Data Format (UFD)

The first record in the file is a single 32-bit integer indicating the version and format number of the ZRD file.

The second record in the file is a single 32-bit integer indicating the maximum possible number of segments in the remaining records. The actual number of segments will always be less than or equal to this number.

The remainder of the file consists of a 32-bit integer indicating the number of segments that follows, then that many segments, then another 32-bit integer indicating the number of segments in the next record, and so on until the end of the file. A single ray segment is described by a "C" structure with the following format:

```
typedef struct
{
    unsigned int status;
    int level;
    int hit_object;
    int hit_face;
    int in_object;
    int parent;
```

```

int storage;
int xybin, lmbin;
double index, starting_phase;
double x, y, z;
double l, m, n;
double nx, ny, nz;
double path_to, intensity;
double phase_of, phase_at;
double exr, exi, eyr, eyi, ezr, ezi;
} RAYPATH;

```

Note that like typical "C" structures in Windows, the structure is packed on 8 byte boundaries to yield a size of 208 bytes. This packing means that 2-byte members are aligned on a memory address that is an integer multiple of 2, 4-byte members are aligned on multiples of 4, and 8-byte members are aligned on multiples of 8. This packing speeds memory access but creates some small gaps in the structure. That is why the structure size is slightly larger than the sum of the byte sizes of all the members. The gap in this case is 4 bytes and appears between the last int member (lmbin) and the first double member (index).

Each data item is described below.

status: bitwise flags indicating the status of the ray. The flags (in integer notation) are: terminated: 1, reflected: 2, refracted: 4, scattered: 8, diffracted: 16, ghost: 32, diffracted from previous object: 64, scattered from previous object: 128, ray segment had fatal error: 256, bulk scattered: 512.

level: The number of ray segments between the ray segment and the original source. This is the number of ray-object intercepts the ray has accumulated.

hit_object: The object number the ray intercepted. If zero, the ray did not hit anything. For the zero segment, this value will be the source segment object number.

hit_face: The face number the ray intercepted. Valid only if hit_object is not zero.

in_object: The object number the ray is propagating inside of. This generally determines the index of refraction the ray is traveling through. A value of zero means inside the "background media".

parent: The (prior) ray segment from which the ray originated. If zero, the ray came from the original source point. Note more than 1 child ray may have the same parent; but each ray has only one parent. The parent segment number will always be less than the child rays segment number.

storage: A temporary buffer used by ZEMAX for indexing. For segment 0 only, this integer is the ray wavelength number. See nx below for the actual wavelength value. Other segments may contain meaningless data.

xybin, lmbin: The pixel number on a detector object which the ray struck. The xybin is for spatial data, the lmbin is for angular (intensity) data.

index: the index of refraction of the media. Not an adequate description for gradient index media.

starting_phase: the initial optical path length the ray starts with. This is usually the contribution from diffraction gratings when rays are split off.

x, y, z: the global coordinates of the ray intercept.

l, m, n: the global direction cosines of the ray. Not an adequate description for gradient index media.

nx, ny, nz: the global normal vector of the object at the intercept point. For segment 0 only, the nx value stores the wavelength of the ray being launched.

path_to: the physical (not optical) path length of the ray segment.

intensity: the intensity of the ray segment.

phase_of: the phase of the object. This is the optical phase path length equivalent in lens units added by gratings, holograms, and other phase modifying surfaces.

phase_at: the accumulated total phase of the ray. This value is only computed at detector surfaces.

exr, exi, eyr, eyi, ezr, ezi: the electric field in the global x, y, and z coordinates, real and imaginary parts. These values will all be zero if the ray trace did not use polarization.

The ZRD UFD format is subject to change, however, ZEMAX will automatically convert older formats to the current format. Note the structure size is 208 bytes for each ray segment. If 100,000 rays are traced with an average of 50 segments each, the resulting file will be about 100,000 x 50 x 208 bytes = 1Gb in size. The file sizes can quickly become enormous, so great care must be taken to limit the number of segments before deciding to save the ray data to a file. Use of a filter string to limit the saved data to the rays of interest may substantially reduce the amount of data written to the file.

The ZRD file consists of alternating blocks of data. For example, if three rays are traced, with 5, 7, and 9 segments each, then the ZRD file will be formatted as follows:

```
4 bytes: version
4 bytes: maximum number of segments possible, will be 9 or larger in this example
4 bytes: 5 (the number of segments in ray 1)
208 bytes x 5 (the segment data for ray 1)
4 bytes: 7 (the number of segments in ray 2)
208 bytes x 7 (the segment data for ray 2)
4 bytes: 9 (the number of segments in ray 3)
208 bytes x 9 (the segment data for ray 3)
```

For information on what ZEMAX can do with ZRD files, see the next section.

The Ray Database Viewer

The Ray Database Viewer reads in a previously saved ZRD file, and displays the data in a text format. For information on creating a ZRD file, see the previous section.

RAY DATABASE VIEWER OPTIONS

Item	Description
File	The name of the ZRD file to use. ZEMAX will look in the folder where the current lens is saved, and list all files with the ZRD extension found there.
Show Unprocessed Data	If checked, the complete data tree for each ray traced will be displayed.
Expand Branches	Into If checked, each branch of the tree will be expanded to show a single complete path from source to ray termination.
XYZ	If checked, the global X, Y, and Z coordinate data of the ray segment-object intercept point will be displayed.
LMN	If checked, the global L, M, and N direction cosine data of the segment will be displayed.
Normal	If checked, the global Nx, Ny, and Nz object normal vector components at the ray segment-object intercept point will be displayed.
Path	If checked, the path length in lens units from the parent segment is displayed. This is the physical, not the optical path length. ZEMAX will also show the phase in radians, but only for segments striking a detector surface.
Exyz	If checked, the global electric field in the X, Y, and Z directions will be displayed.
First/Last Ray	The first and last number of the rays to display. These controls are very useful for limiting the length of the displayed data. If the subset data is to be saved in a new ZRD file, only rays within this range will be included. The value -1 may be used for the Last Ray, which indicates the last ray defined in the file. If -1 is used for the Last Ray, text output will be suppressed.
Apply Filter	If checked, the filter string entered will be used to reduce the amount of data displayed. The filter string is also used to determine the subset of the segments which are saved to a new ZRD file. For information on the filter string syntax, see "The filter string" on page 386.

Save Subset Data As	If checked, the data displayed which passes the filter string will be saved in ZRD format in the file name specified. The file name should end in the extension ZRD, with no path provided. The path used will be the same as the current lens file. To prevent accidental overwrite of the subset data file, this control is automatically "unchecked" after every save.
ZRD Format	This control selects the format to be used for the ZRD file created if "Save Subset Data As" is checked. For information on ZRD formats, see "Ray database (ZRD) files" on page 380.
Save Rays on Object "n" As	If checked, ray segments which hit the specified object number and which pass the defined filter (if any) will be saved as source rays in a binary source file format of the specified name. Note the ray position data is saved in global coordinates. This file may subsequently be used as a source file. For information on the ray source file format, see "Source File" on page 454. The file name should end in the extension DAT (for flux only monochromatic data) or SDF (for polychromatic data), with no path provided. The folder used will be <objects>\sources\Source Files. To prevent accidental overwrite of a previously created data file, this control is automatically "unchecked" after every save. Only rays within the range of the first/last ray setting are saved. To save all rays that hit the specified object from the ZRD file, choose 1 for the first ray and any integer greater than the number of rays in the ZRD file for the last ray. If the total number of rays exceeds 1000, then printing of the individual ray data in the ray database report will be suppressed to speed up the process. To see the text listing, choose "Update" again from the text window menu bar. The second update will show the text data because the "Save Rays On Object" box is automatically unchecked after every save.

The text display for the ray database viewer can be quite long, so it is best to limit the amount of data displayed by using the first/last controls and the filter string.

A typical unprocessed ray listing looks like this, with some columns removed for brevity:

```
Ray 1, Wavelength 1:
Seg# Prnt Lev1 In Hit XRTS DGEF BZ X Y Z
0 0 0 0 0 ---- ---- -- 0.0000E+000 5.0000E+000 0.0000E+000
1 0 1 0 2 -*- ---- -- 0.0000E+000 3.0870E+000 9.0000E+000
2 1 2 0 0 *--- *--- -- 0.0000E+000 2.7044E+000 7.2002E+000
3 1 2 2 2 -*- ---- -- 0.0000E+000 2.5468E+000 1.2838E+001
4 3 3 2 2 -*- *--- -- 0.0000E+000 9.5506E-001 9.0000E+000
5 4 4 2 2 -*- *--- -- 0.0000E+000 -6.9886E-001 1.2988E+001
6 5 5 2 2 -*- *--- -- 0.0000E+000 -2.0350E+000 9.0000E+000
7 6 6 2 2 -*- *--- -- 0.0000E+000 -3.2848E+000 1.2730E+001
8 7 7 0 3 *--- ---- -- 0.0000E+000 -5.7054E+000 1.8400E+001
9 6 6 0 0 *--- ---- -- 0.0000E+000 -2.9069E+000 7.3797E+000
10 5 5 0 3 *--- ---- -- 0.0000E+000 -4.2858E+000 1.8400E+001
11 4 4 0 0 *--- ---- -- 0.0000E+000 -9.6429E-002 7.4900E+000
12 3 3 0 3 *--- ---- -- 0.0000E+000 9.6510E-001 1.8400E+001
```

"Seg#" is the segment #, and is sequential from 0 to the total number of segments the ray generated.

"Prnt" is the parent number, which indicates which prior segment generated that particular segment. For example, the parent of segment 11 is segment 4, whose parent is 3, whose parent is 1, whose parent is 0. Segment 0 is the source segment, and the starting coordinates indicate the location of the source point. Any segment can be traced back to the source in this manner.

"Levl" is how many segments away from the source ray the segment is. For segment 11, there are 4 segments back to the source (4, 3, 1, and 0).

"In" indicates what object number the segment is traveling through.

"Hit" indicates the object number struck at the end of the segment. The source segment has a hit number of 0, since the XYZ coordinates listed are the starting and not the terminating coordinates of the segment. For all other segments, the XYZ coordinates and other data are at the segment end on the "hit" object. If the hit object is 0, the ray missed all objects and is terminated.

"XRTS DGEF BZ" is a quick indicator of what happened to the segment. X means terminated, R reflected, T transmitted, S scattered, D diffracted, and B bulk scattered. All these values indicate what happened at the end

of the segment, when the segment struck the object "hit". The values G, E, and F indicate the segment was a ghost, diffracted, or scattered, respectively, from the parent segment's "hit" object number, so these events occurred at the start of the segment. The symbol Z indicates a ray error occurred and the ray was terminated. Note more than one event may occur for any ray segment.

The other data displayed includes the global XYZ coordinates of the end of the segment, the global normal vector coordinates of the object struck at the end of the segment, the ray intensity, the object comments, and possible other data as this feature is expanded.

The Detector Viewer

The Detector Viewer displays data from detectors in either a graphic or text format. For information on tracing rays to record data on detectors, see "The Ray Trace Control" on page 377.

DETECTOR VIEWER OPTIONS

Item	Description
Surface	Selects the NSC surface number. This is always surface 1 when in NSC mode.
Detector	Selects the detector object. If the window is a text window, there is also an option to summarize all detectors.
Show As	Choose grey scale or false color display, or cross section plots of rows or columns. Detector color and detector polar objects support a "true color" option which shows a representation of the detected visible color light. Detector polar objects support either a full or half polar directivity plot, which displays normalized detected energy as a function of angle. Cross sections are not available in all text displays.
Row/Column	Selects the row or column to display for cross section plots. For Detector Surface objects, see "Comments regarding Detector Surface objects" on page 386.
Angle	Detector polar objects support cross sections displayed as linear in angle, or in full or half polar directivity plots. This angle control defines the angle of the plane defining the cross section through the data, relative to the positive local x axis of the detector. When displaying directivity plots only, the Angle control name becomes Angles(s). The control then supports multiple angles separated by commas. For example, the Angle(s) list may be 0, 45, 90, 135. All the angles are then shown on the same plot. A maximum of 10 angles may be listed.
Z-Plane	This control is only used by detector volume objects. The Z-Plane selects which cross section of a volume detector to display. The number of Z-Planes in a detector volume is set by the number of voxels in the Z direction. The text version of the Detector Viewer also supports an "All" option to display all cross sections as well as a table of all voxel center coordinates and data. Pressing the left and right arrow keys while viewing detector volume data will cycle through all the Z-Planes.
Smoothing	Smoothing reduces pixel noise due to low sampling by averaging the data in adjacent pixels. The data in every pixel is replaced by the average of the data in the pixel and the neighboring pixels. This averaging algorithm is applied the number of times specified by the smoothing parameter.
Show Data	See discussion below for important information.
Ray Database	If none is selected, the ray data currently stored on the detectors is used. Otherwise, the previously saved ray database data is used.
Filter	A filter string may be applied to select only rays with particular properties. This option is only available when using a ray database. For information on the filter string syntax, see "The filter string" on page 386.
Scale	Choose linear or logarithmic scales. See "Max/Min Plot Scale" below.

Max/Min Plot Scale	If non-zero values are provided, the plot will be scaled to the specified min and max data range. If these values are both zero, the default scaling will be applied. If the Scale (see above) is Logarithmic, the Min Plot Scale setting is ignored. If the Max Plot Scale setting is greater than zero, then the next greater power of 10 is used as the maximum log scale.
Contrast Enhancement	By default, no scaling of the data is done. If Contrast Enhancement, a non-linear weighting function is applied to the pixel data to exaggerate the less intense pixels. This can be useful to bring out fine detail in the image, however, the visual appearance may be misleading since dim regions will be brighter in the image than they would really be. This option is only supported for detector color and rectangle objects, and if log scaling is not used.
Output Image File	If a file name is provided, the detected image will be saved to the file. The file name should include the full path, or the default path will be the same as for IMA files. The file extension should be JPG or BMP. If the file extension is not recognized JPG will be used. This option is only supported for detector color and rectangle objects. The exported image is at the full resolution of the detector, not the screen, and will use the same number of pixels in the vertical and horizontal directions as the detector does. However, the pixels in the exported image are assumed to be square, so if the detector pixels are not square some distortion of the aspect ratio of the image will be introduced.

Discussion:

Once ray data has been recorded on the detectors, or from rays stored in a ray database, the Show Data option determines what aspect of the data is displayed as follows:

Incoherent Irradiance: This is incoherent power per area as a function of spatial position on the detector. The power of each ray that strikes the same pixel is summed without regard to the phase of the ray. This option becomes Incoherent Illuminance when using photometric rather than radiometric units.

Coherent Irradiance: This is coherent power per area as a function of spatial position on the detector. The complex amplitude of each ray that strikes the same pixel is summed, keeping track of the real and imaginary parts separately considering the phase of each ray. The final amplitude is then squared to yield the coherent power. This option becomes Coherent Illuminance when using photometric rather than radiometric units.

Coherent Phase: The phase angle of the complex amplitude sum used in Coherent Irradiance.

Radiant Intensity: The power per solid angle in steradians as a function of incident angle upon the detector. This option becomes Luminous Intensity when using photometric rather than radiometric units.

Radiance (Position Space): The power per area per solid angle in steradians as a function of spatial position on the detector. Since it is not possible to display the variation in angle space and position space at the same time, this plot is effectively identical to the Incoherent Irradiance display, with the data values divided by two pi. Note this scaling assumes the detector is collecting data over a hemisphere, no matter what angular range the detector actually collects rays over. This option becomes Luminance when using photometric rather than radiometric units.

Radiance (Angle Space): The power per area per solid angle in steradians as a function of incident angle upon the detector. Since it is not possible to display the variation in angle space and position space at the same time, this plot is similar to the Radiant Intensity display, with the data values divided by the product of the area of the detector and the cosine of the angle at the center of the pixel. This option becomes Luminance when using photometric rather than radiometric units.

Incident Flux: For detector volumes, the incident flux is the flux entering each voxel, measured in source units (watts, lumens, or joules).

Absorbed Flux: For detector volumes, the absorbed flux is the flux absorbed within each voxel, measured in source units (watts, lumens, or joules).

Absorbed Flux/Volume: For detector volumes, the absorbed flux/volume is the flux absorbed within each voxel, measured in source units (watts, lumens, or joules) per unit volume measured in cubic lens units.

Comments regarding Detector Surface objects

Detector Surface objects use a collection of triangles to approximate a curved surface that can detect flux and irradiance. The display of the detected data requires conversion of the triangles into square pixels. This conversion is done automatically when displaying the detector data as a false color, grey scale, or cross section plot. The number of square pixels selected is adjusted to provide a reasonable display resolution, and will vary depending upon the number of triangles. For text listings of Detector Surface objects, the X, Y, Z, R, and Angle data will be listed along with Flux and Irradiance values for each pixel. The X, Y, and Z values are the coordinates of the center of each detector triangle. The R value is the distance from the center of each detector triangle to the object vertex. The Angle values are the angle in degrees to the center of each detector triangle as projected to the XY plane. For more information on defining Detector Surface objects see "Detector Surface object" on page 468.

The filter string

It is frequently convenient to display only NSC rays or data which have certain properties. For example, when scattering and splitting are turned on, the layout diagram will become very confusing if many rays are traced. The filter string allows definition of a "test" rays must pass before they are drawn or displayed. The filter string syntax consists of logical operations between flags that indicate if one segment within a ray hit, miss, reflected, refracted, scattered, diffracted, or ghost reflected from an object within the NSC group. There are also "Extended" flags which support numerical arguments. The supported flags are described in the table below.

FILTER STRING FLAGS

Flag	Description
Bn	Ray bulk scattered inside of object n. If the n value is 0, then bulk scattered segments from any object will return true for this test.
Dn	Ray diffracted after striking object n. See En.
En	Diffracted from parent segment's object n. This flag only gets set for ray segments split from diffractive elements, for order numbers other than zero, when ray splitting is on.
Fn	Scattered from parent segment's object n. This flag only gets set for ray segments split from scattering surfaces when ray splitting is on. The specular segment does not get this flag, only scattered segments. If the n value is 0, then scattered segments from any object will return true for this test.
Gn	Ghost reflected from parent segment's object n. This flag only gets set for ray segments reflected from refractive objects when ray splitting is on. If the n value is 0, then ghost segments from any object will return true for this test.
Hn	Ray hit object n. To test whether a ray hit an object, the flag is of the form Hn. For example, to test if a ray hit object 5, the flag would be H5. See Ln.
Jn	Similar to Gn, except that all segments prior to the ghost reflection point are set to have zero intensity. This allows Detector Viewers to look only at ghost energy, not direct incident energy, even if the ray later ghosted off another object. The zero intensity values will only affect the Detector Viewer, not the ray database viewer or layouts.
Ln	Ray hit object n last. To test whether the last segment of a ray branch hit an object, the flag is of the form Ln. For example, to test if the last segment of a ray branch hit object 5, the flag would be L5. See Hn.
Mn	Ray missed object n. To test whether a ray missed an object, the flag is of the form Mn. For example, to test if a ray missed object 15, the flag would be M15.
On	Ray originated at source number n. O0 (that is "O" as in Origin and "0" as in zero) will select all sources.

Rn	Ray reflected after striking object n. The flag R7 would test if the ray reflected after striking object 7. See Gn.
Sn	Ray scattered after striking object n. This tests the "S" flag as listed in the ZRD file, which refers to scattering at the point a ray strikes an object. See also Fn and X_SCATTER.
Tn	Ray transmitted (refracted) in to or out of object n. The flag T4 would test if the ray refracted in or out of object 4 after striking the object.
Wn	Ray uses wavelength n. If the n value is 0, then rays with any wavelength will return true for this test. Note this filter only tests the <u>initial</u> wavelength for the ray as it leaves the source. If wavelength shifting is used (see "Bulk Scatter tab" on page 356) the wavelength may change during propagation.
X_AXYG(n,v)	Ray has incident angle (in degrees) on object n in the local x-y plane greater than v. The angle is measured with respect to the +y direction without regard to the direction of propagation. If the ray never strikes object n, this flag is false.
X_AXYL(n,v)	Ray has incident angle (in degrees) on object n in the local x-y plane less than v. The angle is measured with respect to the +y direction without regard to the direction of propagation. If the ray never strikes object n, this flag is false.
X_AXZG(n,v)	Ray has incident angle (in degrees) on object n in the local x-z plane greater than v. The angle is measured with respect to the +z direction without regard to the direction of propagation. If the ray never strikes object n, this flag is false.
X_AXZL(n,v)	Ray has incident angle (in degrees) on object n in the local x-z plane less than v. The angle is measured with respect to the +z direction without regard to the direction of propagation. If the ray never strikes object n, this flag is false.
X_AYZG(n,v)	Ray has incident angle (in degrees) on object n in the local y-z plane greater than v. The angle is measured with respect to the +z direction without regard to the direction of propagation. If the ray never strikes object n, this flag is false.
X_AYZL(n,v)	Ray has incident angle (in degrees) on object n in the local y-z plane less than v. The angle is measured with respect to the +z direction without regard to the direction of propagation. If the ray never strikes object n, this flag is false.
X_EXT(n,b)	Ray segment is an extraordinary ray generated from a birefringent interface after the parent ray has hit object n exactly b times. To apply this filter, a search is made for the parent segment that hit object n exactly b times, and only the children of that particular parent segment are considered. If no parent segment hit object n exactly b times, the filter returns false. See also X_ORD.
X_GHOST(n,b)	Ray segment has ghosted exactly b times, and has hit object n at least once. If n is zero, any ray segment that has ghosted b times will pass the test. For example, to consider only all second generation ghosts (ghost rays from ghost parents), use X_GHOST(0, 2). X_GHOST does not consider ghost ray segments that end in a TIR condition; although rays that TIR are considered ghosts. For example, if a third generation ghost ray leaves one surface, strikes another surface, and then TIR's from this second surface, X_GHOST(0, 3) will not include this segment because the segment ended in a TIR and not a ray termination (the ray reflected and continued). This same segment will however be included in the filter X_GHOST(0, 4) because the ray ghosted a fourth time (at the TIR point). This is an artifact of how ZEMAX defines segments and counts ghost rays. In all cases, all ghost rays can be found if sufficiently high values of b are tested. Note rays which TIR from refractive surfaces are considered ghosts, but rays reflected from mirror surfaces are not. See also Gn.
X_HIT(n,b)	Ray segment has hit object n exactly b times. See also Hn, X_HITS, X_HITFACE, and X_HITFACE2.

X_HITS(n1,n2)	If n2 is zero: Ray has n1 or more hits on any object(s). If n2 is not zero, then Ray has between n1 and n2 hits, inclusive.
X_HITFACE(n,f)	Ray segment has hit object n on face f. See also Hn, X_HIT, and X_HITFACE2.
X_HITFACE2(n,f,b)	Ray segment has hit object n on face f exactly b times. See also Hn and X_HIT.
X_IAGT(n,v)	Ray has absolute intensity greater than value v on object n. If the ray never strikes object n, this flag is false.
X_IALT(n,v)	Ray has absolute intensity less than value v on object n. If the ray never strikes object n, this flag is false.
X_IRGT(n,v)	Ray has intensity relative to initial intensity greater than value v on object n. If the ray never strikes object n, this flag is false.
X_IRLT(n,v)	Ray has intensity relative to initial intensity less than value v on object n. If the ray never strikes object n, this flag is false.
X_LGT(n,v)	Ray has local incident x ray direction cosine greater than value v at point on object n. If the ray never strikes object n, this flag is false.
X_LLT(n,v)	Ray has local incident x ray direction cosine less than value v at point on object n. If the ray never strikes object n, this flag is false.
X_MGT(n,v)	Ray has local incident y ray direction cosine greater than value v at point on object n. If the ray never strikes object n, this flag is false.
X_MLT(n,v)	Ray has local incident y ray direction cosine less than value v at point on object n. If the ray never strikes object n, this flag is false.
X_NGT(n,v)	Ray has local incident z ray direction cosine greater than value v at point on object n. If the ray never strikes object n, this flag is false.
X_NLT(n,v)	Ray has local incident z ray direction cosine less than value v at point on object n. If the ray never strikes object n, this flag is false.
X_ORD(n,b)	Ray segment is an ordinary ray generated from a birefringent interface after the parent ray has hit object n exactly b times. To apply this filter, a search is made for the parent segment that hit object n exactly b times, and only the children of that particular parent segment are considered. If no parent segment hit object n exactly b times, the filter returns false. See also X_EXT.
X_SCATTER(n,b)	Ray segment has scattered from parent exactly b times, and has hit object n at least once. If n is zero, any child ray segment split off from the parent ray that has scattered b times will pass the test. For example, to consider only first generation scatter rays, use X_SCATTER(0, 1). This filter tests only the scatter from parent or "F" flag as listed in the ZRD. See also Sn and X_SCATTERF.
X_SCATTERF(n,b)	Ray segment has scattered from object n after the parent of the segment hit object n exactly b times. To apply this filter, a search is made for the parent segment that hit object n exactly b times, and only that particular parent segment is considered. If no parent segment hit object n exactly b times, the filter returns false. For example, to consider only scattered rays that branch off from the parent ray after the third hit on object 5 (that is, the ray leaving the source has twice before hit this same object), use X_SCATTERF(5, 3). See also Fn and X_SCATTER.
X_SEGMENTS(n1,n2)	If n2 is zero: Ray has n1 or more segments. If n2 is not zero, then Ray has between n1 and n2 segments, inclusive.
X_WAVERANGE(n, a, b)	Ray has hit object n and has a wavelength between a and b micrometers, inclusive.

X_WAVESHIFT(i,j)	Ray has wave shifted during a bulk scatter event from wavelength i to wavelength j.
X_XGT(n,v)	Ray has local x coordinate greater than value v at point on object n. If the ray never strikes object n, this flag is false.
X_XLT(n,v)	Ray has local x coordinate less than value v at point on object n. If the ray never strikes object n, this flag is false.
X_YGT(n,v)	Ray has local y coordinate greater than value v at point on object n. If the ray never strikes object n, this flag is false.
X_YLT(n,v)	Ray has local y coordinate less than value v at point on object n. If the ray never strikes object n, this flag is false.
X_ZGT(n,v)	Ray has local z coordinate greater than value v at point on object n. If the ray never strikes object n, this flag is false.
X_ZLT(n,v)	Ray has local z coordinate less than value v at point on object n. If the ray never strikes object n, this flag is false.
Z	Ray has fatal error.

All flags except the O, G, E, and F flag indicate what happened at the end of the ray segment, at the point where the segment intercepted the object indicated by the "n" value. The G, E, and F flags indicate that the ray segment's parent struck object "n" and then launched a child ray that was a ghost, diffracted, or scattered segment, respectively. The O flag indicates the ray originated from source object "n". The W flag indicates the ray is of wavelength number "n". The flags of the form X_***(n,v) require two numerical arguments within the parentheses separated by a comma, or a syntax error will be issued.

Each flag is evaluated for each ray traced, and the flag is assigned a status of TRUE or FALSE. The flags may be used alone, or may be combined using logical operands. Logical operands generally act on two other logical flags (the exception is the NOT operand which acts only on the flag to the right of the operand). The supported logical operands are:

- &: Logical AND. Both flags on either side of the & symbol must be TRUE for the AND operation to return TRUE.
 - |: Logical OR. If either of the flags are TRUE, OR returns TRUE.
 - ^ Exclusive OR (XOR). If either of the flags are TRUE, but not both, XOR returns TRUE.
 - ! Logical NOT. Returns TRUE if the flag to the right was FALSE and vice-a -versa.
- The parentheses symbols may also be used to set operator precedence.

Filter string examples

To select only rays that hit object 7, the filter string would be:

H7

To select only rays that hit both objects 7 and 9, the filter string would be:

H7 & H9

To select rays that must have either a) hit object 7 and object 9, but did not reflect off object 6, or b) missed object 15, the filter string would be:

(H7 & H9 & !R6) | M15

If only ghost rays generated by striking object 3, 4, or 7 are desired, the filter string would be:

G3 | G4 | G7

To select only rays incident at an angle of less than 10 degrees on a flat object number 5 whose normal vector points along local +z, the filter string would be (note $\cos(10) = 0.984808$):

X_NGT(5, 0.984808)

Note only rays which hit the specified object in the X_NGT flag can possibly make the flag true.

The filter string is checked for basic syntax errors, such as mismatched parentheses, but not all possible syntax errors are checked and reported. The number of rays which meet the filter string test may be very small, and perhaps even zero. ZEMAX still traces the number of rays defined for the source (by the # of layout rays parameter), but only that fraction of these rays which pass the filter will actually be saved, displayed, or drawn. The filter string may also be applied to ray data base files as the ray data is being saved, see "The Ray Trace Control" on page 377. Because filter strings are used with both dynamic and static windows, ZEMAX does not automatically renumber the filter string arguments when objects are inserted or deleted.

Saving and Loading Detector Data

After a ray trace is performed, detector objects hold data about the energy incident upon the detector. This data may be recreated by "playing back" a ZRD file (see "Ray database (ZRD) files" on page 380) on the detector using the Detector Viewer (see "The Detector Viewer" on page 384). An alternate way of recreating the detector data is to store the data to a file after a ray trace, then loading the detector data back onto the detector object. The Save Detector and Load Detector features are described in "Save Detector Data" on page 99 and "Load Detector Data" on page 99.

Detector Data File Formats

Detector data files store energy data for Detector Rectangle, Color, Polar, and Volume objects. The expected file extension for these different detector types is DDR, DDC, DDP, and DDV, respectively. The binary file format for all file types consists of a single header structure, followed by one data block for each pixel in the detector.

Header Block Format

The header block uses the following format (in C syntax):

```
int version, type, lens_units, source_units, source_multiplier;
int i_data[50];
double d_data[50];
```

version: Currently 1002. Future versions of ZEMAX may use different version numbers as the format changes, but ZEMAX will recognize older versions and convert automatically when loading data. Any undefined data values in the format are reserved for future expansion of the format and may be ignored.

type: 1, 2, 3, or 4 for Detector Rectangle, Color, Polar, and Volume objects, respectively.

lens_units: 0, 1, 2, or 3 for mm, cm, in, or M respectively.

source_units: 0, 1, or 2 for Watts, Lumens, or Joules, respectively.

source_multiplier: 0 through 9 for source unit multipliers femto through Tera, respectively.

i_data: integer data that is detector type specific (see below).

d_data: double precision that is detector type specific (see below).

Data for Detector Rectangle And Detector Color Objects:

i_data[0], i_data[1]: number of x pixels, number of y pixels.

i_data[2], i_data[3]: number of rays striking the spatial detector, number of rays striking the angular detector.

d_data[0], d_data[1]: x half width of the detector, the y half width of the detector.

d_data[2] - d_data[5]: angular x minimum, x maximum, y minimum, y maximum.

Data for Detector Polar Objects:

i_data[0], i_data[1]: number of polar pixels, number of angular pixels.

i_data[3]: number of rays striking the detector.

d_data[0], d_data[1]: maximum angle of the detector, radial size of the detector.

Data for Detector Volume Objects:

i_data[0] - i_data[2]: number of x, y, and z direction pixels.

d_data[0] - d_data[2]: half width of the detector in the x, y, and z directions, respectively.

Data Block Format

After the header data is defined and written to the file, the data structures are written. One data block is written for each pixel on the detector. The file formats for the data structures depends upon the detector type as follows.

Data Structure For Detector Rectangle Objects:

```
double inc_int_pos, inc_int_ang, coh_real, coh_imag;
```

These are the incoherent intensity in position space, the incoherent intensity in angle space, and the coherent real and imaginary amplitudes in position space, respectively.

Data Structure For Detector Color Objects:

```
double pos_P, pos_X, pos_Y, pos_Z, ang_P, ang_X, ang_Y, ang_Z;
```

The first four are the position space power, tristimulus X, Y, and Z values, and the same 4 values for angle space.

Data Structure For Detector Polar Objects:

```
double unused1, unused2, unused3, unused4, ang_P, ang_X, ang_Y, ang_Z;
```

The first four are unused values, followed by the angular space power, tristimulus X, Y, and Z values.

Data Structure For Detector Volume Objects:

```
double inc_int_pos, inc_absorbed_flux, unused1, unused2;
```

These are the incoherent intensity, the absorbed flux, and two unused values, respectively.

Pixel Ordering

For Detector Rectangle and Color objects, pixel #1 is in the (-x, -y) lower left corner and the subsequent pixels move first across the columns in the +x direction, then up through the rows.

For Detector Polar objects, pixel #1 starts at a polar angle of zero, and subsequent pixels go along the radial arm at zero degrees to the +x axis to the maximum polar angle. The next pixel starts again at polar angle of zero and continues along the radial arm at the first azimuthal angle increment. The pattern repeats for each angular arm along the azimuthal direction.

For Detector Volume objects, pixel #1 is in the (-x, -y) lower left corner of the first plane on the -Z side of the detector, and subsequent pixels move first across the columns in the +x direction, then up through the rows to (+x, +y) of the first plane. The next pixel is at the (-x, -y) corner of the next Z plane, and the pattern continues through all the Z planes.

Special considerations for faceted objects

Ray tracing through faceted objects is generally straightforward in a mathematical sense, with one major exception: how to handle the special case of a ray intercepting "exactly" on the common edge of two adjacent facets. In this context, "exactly" means within the numerical precision limit of the computer. ZEMAX uses double precision 64 bit numbers for all ray tracing algorithms, and this yields about 12 decimal places of accuracy. However, because of the frequent use of the square root and other math operations, the practical accuracy of many computations is somewhat lower.

The most common difficulty is tracing rays through a roof prism, corner cube, or faceted Fresnel lens where one of the common edges is placed exactly along the x or y axis of the incoming beam. Since the rays hit exactly on the edge of two facets, each with a different normal vector, the reflection or refraction properties become ambiguous.

Physically, infinitely sharp edges cannot exist. Therefore, the solution ZEMAX uses is to absorb any ray which strikes within the glue distance (see page 351) of an edge of a facet that borders another facet with a different normal vector. An absorbed ray is terminated, and its energy is discarded.

This problem may be avoided by either of these techniques:

For sequential ray tracing, offset the object axis a very small amount in the X or Y directions; so the rays will "see" one facet at a time. An offset of about $1e-6$ lens units (which is only 1 nanometer if the lens units are mm) is generally sufficient.

For non-sequential ray tracing, either offset the object as for sequential ray tracing, or use sources with random ray distributions. Occasionally a ray will strike within the glue distance of the facet edge, and the ray energy will be absorbed; but the total amount of energy should be very small.

Generally, using randomly generated ray sets, such as those in the image analysis or dithered spot diagram, will eliminate this problem.

Note that any real roof prism or corner will physically not "correctly" reflect or refract rays within very small distances of the edge due to manufacturing limitations and physical optics interference.

Comments about DLLs

ZEMAX supports several different types of DLLs for extending the capabilities of the non-sequential features, including DLLs for sources, gradient index media, diffraction, and scattering.

To create a user defined DLL source, a suitable compiler or development tool that can generate 32 bit Windows compatible DLL's is required. It is also assumed that the user can write the required code, and most importantly, ensure that the code is reliable and bug-free. To maximize speed, ZEMAX performs very little error checking on data returned by the DLL, and so buggy DLL's are quite capable of bringing ZEMAX to a crash.

For this reason, technical support on the implementation of DLL's is strictly limited to illustrating that the provided sample files work correctly. If you need a DLL, and do not possess the desire or ability to write them yourself, please feel free to contact Radiant ZEMAX LLC for a quote on developing a custom DLL to meet your requirements. Radiant ZEMAX LLC has considerable experience in developing algorithms, and can generally write DLLs at competitive rates on short notice.

The best way to learn the use of the various DLL types is to study an existing DLL. The sample DLLs provided with ZEMAX include extensive documentation and comments on the data format; see any of the sample source code files for examples.

NSC Objects

ZEMAX NSC object types include ellipses, triangles, rectangles, spheres, cylinders, and other basic shapes. Complex objects such as arbitrary prisms, aspheric lenses, torics, toruses, and other optical components are also available. The reflective, refractive, and absorptive properties of these objects are determined by the material assigned to the objects. For details on reflective, refractive, and absorptive properties, see a following section.

Each NSC object type is described in the following summary table, and in greater detail in the following sections. Note these basic objects may be combined to form more complex objects. See “Object Placement” on page 423 for information on placing objects inside or adjacent to one another.

If an object type is required that is not listed, please contact technical support to suggest the new object type be added to ZEMAX.

SUMMARY OF NSC OBJECTS

Object Name	Description	Page
Annular Aspheric Lens	An annular volume with an aspheric surface on one face.	396
Annular Axial Lens	An annular lens formed by sweeping an aspheric cross section about a decentered Z axis.	397
Annular Volume	A volume formed by two tilted, elliptical annular planar faces.	398
Annulus	A planar ellipse with an elliptical hole in the middle.	398
Array	This object creates a 3D parallelogram array of objects based upon a master object.	399
Array Ring	This object creates a ring array of objects based upon a master object.	401
Aspheric Surface	A surface with a conic substrate plus even and odd radial polynomial power aspheric terms.	402
Aspheric Surface 2	Identical to the Aspheric Surface, with support for more complex aperture types.	402
Axicon Surface	An axicon formed by rotating a spherical arc around the z axis.	403
Biconic Lens	A circular or rectangular lens with separate curvature and conic in the X and Y directions on both front and back faces.	403
Biconic Surface	A surface with separate curvature and conic in the X and Y directions. May be partially hyperhemispheric under special circumstances.	404
Biconic Zernike Lens	A lens with a standard front face and a biconic Zernike rear face.	405
Biconic Zernike Surface	A surface with a biconic Zernike shape. The surface supports elliptical, rectangular, and user apertures, and the biconic and Zernike terms may be separately decentered.	406
Binary 1	A standard lens with a binary 1 phase profile on the front face. The binary 1 phase profile is an X-Y polynomial.	407
Binary 2	A standard lens with a binary 2 phase profile on the front face. The binary 2 phase profile is a radial polynomial in even powers of r.	408

Object Name	Description	Page
Binary 2A	An even aspheric lens (front and back faces) with a binary 2 phase profile on the front face. The binary 2 phase profile is a radial polynomial in even powers of r.	408
Boolean	A very general object defined by a Boolean combination of other objects.	409
CPC	A compound parabolic concentrator (CPC). The CPC may be hollow or solid.	411
CPC - Rectangular	A CPC with rectangular symmetry and separate X and Y aperture and acceptance angles.	412
Cone	A section of a cone defined by two points, r and z, which forms a line segment, which is then rotated around the local z axis.	412
Cylinder Pipe	A tapered cylindrical surface shape.	413
Cylinder Volume	A tapered cylindrical volume, with end caps.	413
Cylinder 2 Pipe	A cylindrical surface shape with tilted ends.	413
Cylinder 2 Volume	A cylindrical volume, with tilted end caps.	414
Diffraction Grating	A standard lens with a diffraction grating with constant or variable line spacing on the front face.	414
Dual BEF Surface	An idealized dual brightness enhancement film surface	414
Ellipse	A planar ellipse.	415
Elliptical Volume	An elliptical volume (or shell). The shape is a tapered cylinder with an elliptical cross section.	415
Even Asphere Lens	A rotationally symmetric lens with up to 16th order aspheres on both the front and back faces.	415
Extended Polynomial Lens	A lens with extended polynomial surfaces on front and back faces.	416
Extended Polynomial Surface	A surface with an extended polynomial shape.	417
Extruded	A solid volume formed by extruding a User Defined Aperture.	417
Freeform Z	A solid lens or shell defined by rotating a spline curve around the Z axis. The spline curve is defined by points along the Z axis.	417
Fresnel 1	A true Fresnel lens, either radial or cylindrical, defined in terms of a number of construction parameters.	418
Fresnel 2	An idealized Fresnel lens object having a circular or rectangular shape, with one face an ideal Fresnel lens with either radial or cylindrical polynomial aspheric terms over a base conic asphere.	419
Grid Sag Lens	A lens with a standard front face and grid sag rear face.	420
Grid Sag Surface	A surface defined by a table of sag points.	421
Hexagonal Lenslet Array	A lenslet array with hexagonal symmetry.	421
Hologram Lens	A circular or rectangular solid with an optically fabricated hologram on the front face.	421
Hologram Surface	A circular or user defined plane, conic, spherical or aspherical surface with an optically fabricated hologram on the face.	422

Object Name	Description	Page
Imported	An object or objects imported from a CAD program in either IGES or STEP format.	423
Jones Matrix	An elliptical flat object which supports 8 ABCD parameters which can be use to model neutral density filters, polarizers, rotators, and other arbitrary Jones Matrix type objects.	425
Lenslet Array 1	A rectangular shaped lens array object with a curved back face. This object may be diffractive.	426
Lenslet Array 2	A rectangular shaped lens array object with curved front and back faces.	427
MicroElectroMechanical System (MEMS)	An array of flat faces (intended to be made of material "MIRROR") where the individual faces may be tipped about any of 3 angles. The faces may be addressed either by rows, by columns, or by individual pixels.	427
Null Object	A non-existent object. This is the default when new objects are inserted into the NSC Editor, and is handy to use as a place holder.	428
Odd Asphere Lens	A rotationally symmetric lens with up to 12th order odd and even power aspheres on both the front and back faces.	428
Paraxial Lens	An idealized planar surface that acts like a thin lens with separable powers.	
Polygon Object	A user defined general polygon object which may be open or closed, and may have refractive, reflective, or absorptive faces, or a mixture of these types of faces.	429
Ray Rotator	An imaginary component that rotates incident rays in two directions.	432
Rectangular Corner	A corner composed of 3 rectangular surfaces.	432
Rectangle	A rectangle defined by a width and height.	433
Rectangular Pipe	A four sided box.	433
Rectangular Pipe Grating	A four sided box with a grating on the side faces.	433
Rectangular Roof	A roof composed of two rectangular surfaces meeting at an angle.	434
Rectangular Torus Surface	A section of a rectangular torus which defines a surface; open at the ends.	434
Rectangular Torus Volume	A section of a rectangular torus with closed ends to define a volume.	434
Rectangular Volume	A six sided box enclosing a volume.	435
Rectangular Volume Grating	A six sided box with a diffraction grating on side faces.	435
Slide	A color RGB file, in either BMP or JPG format. The slide may be scaled to any size and aspect ratio, and may be used to filter a source to create colored image sources.	435
Sphere	A sphere.	436
Standard Lens	A 5 surface lens object with two Standard faces, two flat edges, and a cylindrical outer edge.	436

Object Name	Description	Page
Standard Surface	A surface with the same shape as the ZEMAX Standard surface (a plane, sphere, conic asphere, or hyperhemisphere).	437
STL Object	A general polygon object defined in STL format, typically exported from a CAD program.	437
Swept Object	A surface or volume formed by sweeping the cross section of another object around an arbitrary axis through an angle.	438
Tabulated Faceted Radial	A faceted object generated from a table of coordinates rotated about the local z axis. This object type can be used to define surfaces or volumes.	439
Tabulated Faceted Toroid	A faceted object generated from a table of coordinates rotated about an axis parallel to the local y axis. This object type can be used to define surfaces which are cylindrical or toroidal.	440
Tabulated Fresnel Radial	A smooth Fresnel lens generated from a table of coordinates. This is very similar to the Tabulated Faceted Radial object except the surfaces are smooth, not faceted. See also the Fresnel object.	441
Toroidal Hologram	A circular, elliptical, or rectangular lens with aspheric toroidal surfaces on the front and back sides, and a hologram on the front face.	441
Toroidal Lens	A circular, elliptical, or rectangular lens with aspheric toroidal surfaces on the front and back sides.	442
Toroidal Surface	A rectangular surface with an aspheric toroidal shape.	442
Torus Surface	A section of a torus which defines a surface; open at the ends.	443
Torus Volume	A section of a volume torus; like a surface but with end caps to enclose a volume.	443
Triangular Corner	A corner composed of 3 triangles.	444
Triangle	A triangle defined by 3 points in a plane.	444
User Defined Object	An object defined by a user provided DLL.	444
Zernike Surface	A surface defined by radial aspheric and standard Zernike polynomial sag terms.	446


Most objects require specification of various parameters. The parameters required by each object type and detailed descriptions of each object follow.

Annular Aspheric Lens

The annular aspheric lens is an annular solid with a conic asphere with even aspheric polynomial terms to 16th order, identical to the Even Asphere Lens object, on both front and back faces. See “Even Asphere Lens” on page 415 for a description of the even aspheric polynomial.

The annular shape is defined by a min and max radial aperture which may be different on the front and back faces. The aspheric rear face is offset from the front by a thickness measured at a specified radial aperture, which may be zero.

For backward compatibility with the original definition of this object, the back face parameters are listed before the front surface parameters.



The following restrictions apply: all apertures must be positive, maximum apertures on each face must be larger than the minimum aperture on that face, and if one face has zero for a minimum aperture the other face must also have a zero minimum aperture.

The object is defined by the following parameters:

1-2: The Minimum and Maximum Radial Apertures on the front face.

3-4: The Minimum and Maximum Radial Apertures on the back face.

5: The Thickness Aperture, which is the radial aperture used to measure the object thickness. This value may be zero.

6: Back face radius of curvature.

7: Back face conic.

8: Thickness (measured at the Thickness Aperture).

9-16: The Even Aspheric coefficients α_1 through α_8 for the back face.

17: Front face radius of curvature.

18: Front face conic.

19-26: The Even Aspheric coefficients α_1 through α_8 for the front face.

Face Numbers: Outside face connecting the "maximum" apertures Face 0, front face Face 1, back face Face 2, inside face connecting the "minimum" apertures on the front and back Face 3.

Annular Axial Lens

The annular axial lens is defined by sweeping the cross section of a bi-aspheric lens around a decentered axis parallel to the axis of symmetry of the cross section. The cross section consists of two aspheric faces, each of which is a polynomial asphere. The cross section has the same shape as the Even Asphere Lens, see "Even Asphere Lens" on page 415. The cross section may however be decentered along the radial direction. The defining parameters are:

1: The radius of curvature of the front face, use zero for flat.

2: The conic constant of the front face.

3: The radius of curvature of the rear face, use zero for flat.

4: The conic constant of the rear face.

5: The thickness of the cross section. The thickness must be large enough so that the edge thickness of the part is at least 0.001 times the center thickness.

6: The radial aperture of the cross section.

7: The radius of rotation from the local Z axis to the axis of symmetry of the cross section. This value must be at least 1% larger than the radial aperture of the cross section defined by parameter 6.

8: The decenter of the cross section within the lens aperture. Positive decenters shift the axis of the cross section radially outward, negative decenters shift the axis radially inward.


9: Unused.

10-17: The front face coefficients $\alpha_1 - \alpha_8$.

18-19: Unused.

20-27: The rear face coefficients $\alpha_1 - \alpha_8$.


Face Numbers: Outside face connecting the "maximum" apertures Face 0, front face Face 1, back face Face 2, inside face connecting the "minimum" apertures on the front and back Face 3.



Annular Volume

The annular volume consists of tilted annular elliptical faces front and rear, separated by an axial distance. The front and rear face minimum and maximum coordinates are measured on the XY plane. These coordinates form an ellipse on the front and rear reference planes. The inner and outer elliptical cones formed by connecting the front and rear elliptical shapes are the inner and outer surfaces of the annular volume. The front and rear tilted faces are then tilted by the specified angles in x and y. The boundary shape of each face is defined by the intersection of the tilted plane and the annular elliptical cone. The defining parameters are:

- 1-2: Unused.
- 3: The front face minimum x coordinate.
- 4: The front face minimum y coordinate.
- 5: The rear face minimum x coordinate.
- 6: The rear face minimum y coordinate.
- 7: The front face maximum x coordinate.
- 8: The front face maximum y coordinate.
- 9: The rear face maximum x coordinate.
- 10: The rear face maximum y coordinate.
- 11: The length along the local Z axis.
- 12: The front face x tilt angle in degrees.
- 13: The front face y tilt angle in degrees.
- 14: The rear face x tilt angle in degrees.
- 15: The rear face y tilt angle in degrees.



Face Numbers: Outside face connecting the "maximum" apertures Face 0, front face Face 1, back face Face 2, inside face connecting the "minimum" apertures on the front and back Face 3.


Annulus

The annulus is a plane elliptical surface shape defined by 4 parameters:

- 1: The X Maximum Half Width.
- 2: The Y Maximum Half Width.
- 3: The X Minimum Half Width.
- 4: The Y Minimum Half Width.


The annulus resides entirely within the local XY plane. If the minimum half widths are set to zero, then no "hole" exists and rays will hit anywhere within the maximum half widths. In this case, the annulus is identical to the ellipse.

The reference coordinate is the center of the annulus. Face Numbers: All faces Face 0.



Array

This object is a collection of identical objects, based upon a single parent object, arranged in a 3D array. The individual array elements are all functionally identical to the parent object, just as if a copy of the parent object had been placed in each array location. The 3 Axes of the array need not be orthogonal, and are user definable; see the discussion below. The Array object is generally faster to trace and more memory efficient than the equivalent system created using the replicate tool (see "Replicate Object Tool" on page 468) or using an array of similar objects as supported by the Lenslet 1 and Lenslet 2 objects. The Array object is defined by the following parameters:



1: The parent object #. The parent object must precede the array object in the NSC Editor.

2-4: The number of array elements in the X', Y', and Z' directions.

5-7: The element spacing in the X', Y', and Z' directions in lens units. See the discussion for non-uniform spacing.

8-10: The direction cosines of the X' axis of the array.

11-13: The direction cosines of the Y' axis of the array.

14-16: The direction cosines of the Z' axis of the array.

17-19: The element tilt in degrees about the X, Y, and Z directions. Note these tilts are about the "normal" local X, Y, and Z axes, and not the user defined X', Y', and Z' axes.

20: Draw Limit. If the product of the number of elements in each direction ($N_x \times N_y \times N_z$) exceeds the Draw Limit, the individual array elements will not be drawn. Instead, a simple box that approximates the size of the array is drawn. This feature is intended to speed up the rendering of the array object when the number of elements is large.

21: Draw Boundary: If this is set to 1, the box described in parameter 20 above will be drawn in all cases. This feature is useful to compare the box dimensions and orientation with the actual array elements; which may extend beyond the box boundaries.

22-30: The non-linear spacing coefficients, see the discussion below.

31-33: The maximum dimensions of the array in the X', Y', and Z' directions. These values are not set by the user. The values are computed based upon the spacing coefficients and the number of elements defined. The optimization operands NPGT, NPLT, and NPVA may be used to constrain the overall size of the array while optimizing the non-linear spacing coefficients. The maximum dimensions reported in these parameter values is the center-to-center spacing from the first to the last array element in each direction, and does not consider the size of the individual array elements.

The parent object may be any object type, except for a source, detector, or another array object. If the parent object has the "Object Is A Detector" box checked, this property will be ignored in the array elements. To create an array of sources, see "Modeling arrays of sources" on page 433.

The number of array elements in the 3 directions may in principle be any number, however, very large arrays may require excessive computation time or computer memory to ray trace or render (see parameters 20 and 21 for settings to eliminate the rendering of large arrays).

The array elements are arranged on a 3D grid where the grid axes are not necessarily the usual orthogonal XYZ axes. The X', Y', and Z' axes default to the normal orthogonal local axes, however, any other orientation for

the axes may be defined by specification of the direction cosines for the primed axes. The array positions are computed for the element with index ijk (i for X' , j for Y' , k for Z') according to the following expression:

$$\vec{r} = \begin{bmatrix} i\delta_x \\ j\delta_y \\ k\delta_z \end{bmatrix}, \text{ where}$$

$$\begin{bmatrix} \delta_x \\ \delta_y \\ \delta_z \end{bmatrix} = \begin{bmatrix} X_x & Y_x & Z_x \\ X_y & Y_y & Z_y \\ X_z & Y_z & Z_z \end{bmatrix} \begin{bmatrix} \Delta_{1x} \\ \Delta_{1y} \\ \Delta_{1z} \end{bmatrix},$$

and the X_x, X_y, X_z values are cosines defining the X' axis (parameters 8, 9, and 10), and similarly for Y' and Z' , and the Δ_1 values are the element spacings along the axes (parameters 5, 6, and 7). Note that the direction cosine vectors are always normalized automatically. Once the array element is located, the local rotation about the $X, Y,$ and Z axis is performed, in a manner identical to the usual coordinate rotations (see "Object Placement" on page 423).

The Δ_1 values in the three directions are used to define a uniform spacing of the array elements. However, a non-uniform spacing may also be defined using parameters 22-30. The position of the array elements along the X' direction is in general defined as

$$x_i = \Delta_1 i + \Delta_2 i^2 + \Delta_3 i^3 + \Delta_4 i^4,$$


where Δ_1 is the uniform spacing (defined by parameters 5, 6, and 7 for $X', Y',$ and Z' , respectively), the index " i " goes from 0 to n_x-1 (n_x is the number of elements in the X' direction), and the Δ_k coefficients are the k^{th} order terms measured in lens units. If any of the Δ_k coefficients are non-zero, the array spacing will be non-uniform. A similar expression and set of coefficients is available for the Y' and Z' directions. The non-linear coefficients must be defined in such a way that the spacing between the elements cannot change sign so the array folds back on itself.

Most properties of the Array object are defined by the parent object. The exceptions, which may be distinct for the array object and the parent object, are the "Rays Ignore This Object" and the "Do Not Draw Object" settings. All other properties, including face definition, material, coatings, scattering, diffraction, or gradient index properties, are as defined for the parent object, and the array object settings will be ignored. If the parent object is not needed, use the "Rays Ignore This Object" and "Do Not Draw Object" settings on the parent object (See "Type tab" on page 426 and "Draw tab" on page 362). The array will still be created and ray traced even if the parent object is not drawn or ray traced. Note that rays striking an array object will report hitting the parent object number, not the array object number. For this reason, when using ZRD files, optimization operands, or filter strings, the object number used to refer to ray data should be the parent object number.

When nesting array objects inside or outside other objects, the usual nesting rules apply. For more information see "Nesting volumes" on page 424. However, it is the placement of the *parent* object in the NSC Editor, and not the *array* object, that determines the nesting rules. Specifically, if an array object is to be nested inside of another, outermost object, the *parent* object must follow the outermost object in the editor.

Array Ring

This object is a collection of identical objects, based upon a single parent object, arranged in a ring or arc. The individual array elements are all functionally identical to the parent object, just as if a copy of the parent object had been placed in each array location. The Array Ring object is generally faster to trace and more memory efficient than the equivalent system created by copying the parent object to each element location. The Array Ring object is defined by the following parameters:



1: The parent object #. The parent object must precede the array object in the NSC Editor.

2: The mode number.

3: The number of array elements, N .

4: The Radius of the ring or the reference sphere, R , in lens units.

5-9: The parameters α , β , γ , δ , and ε . Whenever these parameters are angles, the units are in degrees.

The parent object may be any object type, except for a source or detector. If the parent object has the "Object Is A Detector" box checked, this property will be ignored in the array elements. To create an array of sources, see "Modeling arrays of sources" on page 433.

The number of array elements may in principle be any number, however, very large arrays may require excessive computation time or computer memory to ray trace or render.

The array elements are arranged on a ring (or arc). There are two ways to define the positions of the array elements, defined by the mode flag. Mode 0 results in a ring of objects in a plane. Mode 1 creates a ring of objects on a reference sphere. The method for placing each array elements is as follows for each mode:

Mode 0: The parent object is first placed at the vertex of the Array Ring. The object is then rotated around the array local +Z axis by the angle α . The object is then shifted along the array local +X axis by the distance R . The object is then rotated around the array local Z axis by the angle γ . This is the position of the first element in the array. Subsequent array elements are further rotated around the array local Z axis by an angle δ/N . The angular position of the array element "k" is given by $\theta = \gamma + k*\delta/N$, where k varies from 0 to $N-1$. For a complete ring, δ should be 360. Mode 0 may also be used to form a Spiral of Archimedes, where R is replaced by $R + k*\beta*\delta/N$ and β is interpreted to be in lens units per degree. The value ε in lens units may also be used to define a step change in R every time the angle given by $k*\delta/N$ passes 360 degrees. Ray tracing will generally be much slower when either β or ε are not zero, since the symmetry of the array is broken.

Mode 1: The parent object is first placed at the vertex of the Array Ring. The object is then rotated around the array local +Z axis by the angle α . The object is then rotated an angle β about an axis parallel to the array local Y axis that is shifted by R in the array local +Z direction, so the array element's local Z points toward the center of a sphere with radius R (or away from the center if R is negative). The sign convention for the angle β used here is that β is always positive, and the object is always rotated toward the local +X axis, whether or not R is positive or negative. The object is then rotated around the array local Z axis by the angle γ . This is the position of the first element in the array ($k = 0$). Subsequent array elements are further rotated around the Z axis by an angle δ/N .

For all modes, to form a complete ring, δ should be 360. Smaller values of δ will form an arc of objects. Larger values of δ may cause elements to overlap, which may create an invalid geometry. The angle γ "clocks" the starting point of the ring or arc around Z.

Most properties of the Array Ring object are defined by the parent object. The exceptions, which may be distinct for the Array ring object and the parent object, are the "Rays Ignore This Object" and the "Do Not Draw Object" settings. All other properties, including face definition, material, coatings, scattering, diffraction, or gradient index properties, are as defined for the parent object, and the array object settings will be ignored. If the parent object is not needed, use the "Rays Ignore This Object" and "Do Not Draw Object" settings on the parent object (See "Type tab" on page 426 and "Draw tab" on page 435). The array will still be created and ray traced even if the parent object is not drawn or ray traced. Note that rays striking an array object will report hitting the

parent object number, not the array object number. For this reason, when using ZRD files, optimization operands, or filter strings, the object number used to refer to ray data should be the parent object number.

When nesting array objects inside or outside other objects, the usual nesting rules apply. For more information see “Nesting volumes” on page 424. However, it is the placement of the *parent* object in the NSC Editor, and not the *array* object, that determines the nesting rules. Specifically, if an array object is to be nested inside of another, outermost object, the *parent* object must follow the outermost object in the editor.

Aspheric Surface


An aspheric surface is defined by the following sag equation:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1 + k)c^2r^2}} + \sum_{i=1}^N \alpha_i r^i,$$

where c is the curvature of the surface, k is the conic constant, r is the radial coordinate, and the α terms are aspheric coefficients. The surface supports specification of both a minimum and a maximum radial aperture; so annular surfaces may be defined. Note that both even and odd terms are defined; up to approximately 240 coefficients may be used.

The following parameters are used to define the aspheric surface:

- 1: The radius of curvature. If this value is zero, then the curvature is assumed to be zero.
- 2: The conic constant k .
- 3: The maximum radial aperture in lens units.
- 4: The minimum radial aperture in lens units. This value may be zero.
- 5-6: Unused.
- 7: The number of terms to use in the aspheric expansion. Ray tracing will be faster if this term is no larger than the highest order non-zero coefficient.




8-249: The α coefficients on the polynomial expansion.

Face Numbers: All faces Face 0. This object supports user defined apertures, see “User defined apertures” on page 437. For other aperture shape options, see “Aspheric Surface 2” on page 402.

Aspheric Surface 2

This surface is similar to the Aspheric Surface above, with added support for elliptical or rectangular decentered annular apertures. The following parameters are used to define the aspheric surface:

- 1: The radius of curvature. If this value is zero, then the curvature is assumed to be zero.
- 2: The conic constant k .
- 3: The maximum X direction aperture in lens units.
- 4: The maximum Y direction aperture in lens units.
- 5: The minimum X direction aperture in lens units. This value may be zero.
- 6: The minimum Y direction aperture in lens units. This value may be zero.
- 7: The X direction aperture decenter in lens units.



8: The Y direction aperture decenter in lens units.

9: The "Is Rectangle" flag. Use 0 for elliptical symmetry, 1 for rectangular symmetry.

10-11: Unused.

12: The number of terms to use in the aspheric expansion. Ray tracing will be faster if this term is no larger than the highest order non-zero coefficient.


13-249: The α coefficients on the polynomial expansion.

Face Numbers: All faces Face 0.

Axicon Surface

There are many types of axicons. Common polynomial types can be modeled using the Aspheric Surface (see "Aspheric Surface" on page 402).

The Axicon Surface object uses a different method for defining the object shape. The Axicon surface is defined by a section of a circle (an arc) that lies in the YZ plane. The arc is defined by three points. The start of the arc is at (z=0, y=0). The center of the arc is at the user defined coordinates (cz, cy). The end of the arc is at the point where z = L, where L is a user defined parameter corresponding to the length of the axicon along the Z axis. The surface is formed by revolving the arc around the Z axis. If cz is positive, the negative of the absolute value of cy is taken, and the portion of the circle lying in the first quadrant is used to define the surface. If cz is negative, the absolute value of cy is taken, and the portion of the circle lying in the first quadrant is used to define the surface.



The following parameters are used to define the axicon surface:

1: The length of the axicon along the local z axis.

2: The z coordinate of the center of the arc measured in the YZ plane.

3: The y coordinate of the center of the arc measured in the YZ plane.

Face Numbers: All faces Face 0.

Biconic Lens

A biconic lens is similar to a toroidal lens, except the conic constant and base radius may be different in the X and Y directions. The biconic lens allows specification of Rx, Ry, Kx, and Ky for the front and back surfaces independently. The sag of a biconic surface is given by:

$$z = \frac{c_x x^2 + c_y y^2}{1 + \sqrt{1 - (1 + k_x) c_x^2 x^2 - (1 + k_y) c_y^2 y^2}},$$

where

$$c_x = \frac{1}{R_x}, c_y = \frac{1}{R_y}.$$

The biconic lens is defined by these parameters:

1: The radial height of the lens object in lens units. This value is used for the y direction half height if the lens is rectangular.

2: The x half width of the lens object in lens units. If zero, the lens is circular; otherwise, the lens is rectangular.

3: The center thickness of the lens in lens units.

4-5: Unused.

6: The base radius of curvature in the XZ plane for the front surface. If this value is zero, then the XZ curvature is assumed to be zero.

7: The base radius of curvature in the YZ plane for the front surface. If this value is zero, then the YZ curvature is assumed to be zero.

8: The X direction conic for the front surface.

9: The Y direction conic for the front surface.


10: The base radius of curvature in the XZ plane for the back surface. If this value is zero, then the XZ curvature is assumed to be zero.

11: The base radius of curvature in the YZ plane for the back surface. If this value is zero, then the YZ curvature is assumed to be zero.

12: The X direction conic for the back surface.

13: The Y direction conic for the back surface.

Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.



Biconic Surface

A biconic surface is similar to a toroidal surface, except the conic constant and base radius may be different in the X and Y directions. The biconic surface allows specification of R_x , R_y , K_x , and K_y directly. The sag of a biconic is given by:

$$z = \frac{c_x x^2 + c_y y^2}{1 + \sqrt{1 - (1 + k_x)c_x^2 x^2 - (1 + k_y)c_y^2 y^2}},$$

where

$$c_x = \frac{1}{R_x}, c_y = \frac{1}{R_y}.$$

The biconic surface is defined by 13 parameters:

1: The base radius of curvature in the XZ plane. If this value is zero, then the XZ curvature is assumed to be zero.


2: The base radius of curvature in the YZ plane. If this value is zero, then the YZ curvature is assumed to be zero.

3: The X direction conic.

4: The Y direction conic.

5-6: The maximum X/Y aperture in lens units.

7-8: The minimum X/Y aperture in lens units. These parameters are ignored if the "Is Rectangle?" parameter is zero.



9-10: Unused.

11: The "Is Rectangle?" flag. If zero, the resulting surface shape will have an elliptical boundary. If non-zero, the surface will have a rectangular boundary.

12: The "Is Top Hyper?" flag. If zero, the maximum Y aperture will lie on the non-hyperhemispheric portion of the surface. If non-zero, the maximum Y aperture of the surface will be hyperhemispheric. This parameter is ignored if the "Is Rectangle?" parameter is zero, if the Y direction conic is less than or equal to -1.0, if the maximum Y aperture is less than zero, or if the XZ base radius is non-zero.

13: The "Is Bot Hyper?" flag. If zero, the minimum Y aperture will lie on the non-hyperhemispheric portion of the surface. If non-zero, the minimum Y aperture of the surface will be hyperhemispheric. This parameter is ignored if the surface is elliptical, if the Y direction conic is less than or equal to -1.0, if the minimum Y aperture is greater than zero, or if the XZ base radius is non-zero.

Making a hyperhemispheric surface

The biconic surface can also be used to make a hyperhemispheric cylinder surface, but only under these conditions:

- 1) The Rx and Kx values are set to zero.
- 2) The "Is Rectangle?" flag is set to 1 (that is, the surface boundary is a rectangle).
- 3) The Ky value is greater than -1.

For a surface that is hyperhemispheric on the top half, set the maximum Y aperture to a positive value and set the "Is Top Hyper?" flag to unity.

For a surface that is hyperhemispheric on the bottom half, set the minimum Y aperture to a negative value and set the "Is Bot Hyper?" flag to unity.

The surface may be hyperhemispheric in both the top and bottom portions at once, but currently only for the case of Rx = 0.0. If either the top or bottom is hyperhemispheric, the actual maximum Y aperture is given by

$$y^2 = \frac{(R_y)^2}{(1 + k_y)}.$$

In this special case, a hyperhemispherical cylinder will result.

Face Numbers: All faces Face 0.

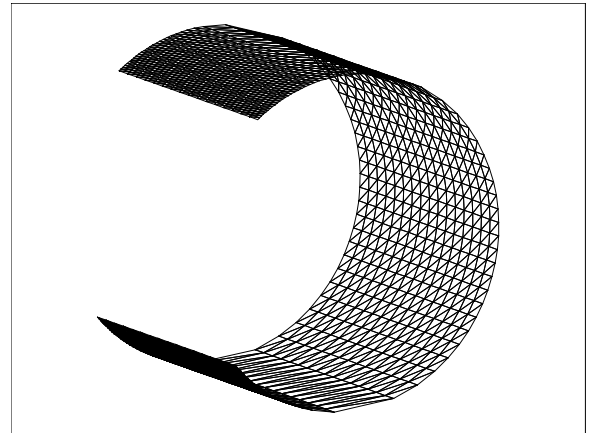
Biconic Zernike Lens

A Biconic Zernike Lens is either rectangular or elliptical in shape, with a standard surface on the front face and a biconic Zernike surface on the rear face. The rear surface is described by the following sag equation:


$$z = \frac{c_x x^2 + c_y y^2}{1 + \sqrt{1 - (1 + k_x)c_x^2 x^2 - (1 + k_y)c_y^2 y^2}} + \sum_{i=1}^{16} \alpha_i x^i + \sum_{i=1}^{16} \beta_i y^i + \sum_{i=1}^N A_i Z_i(\rho, \phi),$$

where

$$c_x = \frac{1}{R_x}, c_y = \frac{1}{R_y},$$



Rx and Ry are the radius of curvature values in the x and y directions, kx and ky are the conic constants in the x and y directions, the α terms are the x aspheric coefficients, the β terms are the y aspheric coefficients, N is the number of Zernike coefficients in the series, A_i is the coefficient on the i^{th} Zernike Standard polynomial, r is the radial coordinate in lens units, ρ is the normalized radial coordinate, and ϕ is the angular coordinate. The Zernike Standard polynomials are defined in the table given in “Zernike Standard Coefficients” on page 199.



The following parameters are used to define the Biconic Zernike Lens:

- 1: The Radial Height of the lens in the y direction in lens units.
 - 2: The X Half-Width in lens units. If this parameter is zero, then the outer boundary of the lens will be a circle with a radial size equal to the Radial Height. If this parameter is positive, the outer boundary of the lens will be rectangular. If this parameter is negative, the outer boundary of the lens will be elliptical.
 - 3: The center thickness of the lens.
 - 4, 5: Unused.
 - 6: The front face radius of curvature. Use zero for flat.
 - 7: The front conic constant.
 - 8/9: The rear face radius in the X/Y direction. Use zero for flat.
 - 10/11: The rear face conic in the X/Y direction.
 - 12: The number of Standard Zernike terms.
 - 13: The normalization radius for the Zernike terms.
 - 14-29: The x direction aspheric α coefficients.
 - 30-45: The y direction aspheric β coefficients.
 - 46 and up: The Zernike Standard coefficients.
- Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Biconic Zernike Surface


The Biconic Zernike Surface is either rectangular or elliptical in shape, or may have a shape defined by a user defined aperture. The surface is described by the following sag equation:

$$z = \frac{c_x x^2 + c_y y^2}{1 + \sqrt{1 - (1 + k_x)c_x^2 x^2 - (1 + k_y)c_y^2 y^2}} + \sum_{i=1}^{16} \alpha_i x^i + \sum_{i=1}^{16} \beta_i y^i + \sum_{i=1}^N A_i Z_i(\rho, \phi),$$

where

$$c_x = \frac{1}{R_x}, c_y = \frac{1}{R_y},$$

Rx and Ry are the radius of curvature values in the x and y directions, kx and ky are the conic constants in the x and y directions, the α terms are the x aspheric coefficients, the β terms are the y aspheric coefficients, N is the number of Zernike coefficients in the series, A_i is the coefficient on the i^{th} Zernike Standard polynomial, r is the radial coordinate in lens units, ρ is the normalized radial coordinate, and φ is the angular coordinate. The Zernike Standard polynomials are defined in the table given in “Zernike Standard Coefficients” on page 199.



This surface also supports optional X and Y decenters on the Biconic and Zernike terms separately. The biconic decenters are applied to the biconic sag and individual x and y aspheric coefficients, while the Zernike decenters are applied only to the Zernike terms.

The following parameters are used to define the Biconic Zernike Surface:

1/2: The X and Y Half-Width in lens units. If either value is less than zero, the aperture shape will be elliptical, otherwise, the aperture is rectangular, unless a user-defined aperture is applied to the surface.

3/4: The Rx and Ry radii of curvature. Use zero for flat.

5/6: The Kx and Ky conic constants.

7/8: The biconic X and Y direction decenters.

9/10: The Zernike X and Y direction decenters.

11: The number of Standard Zernike terms.

12: The normalization radius for the Zernike terms.

14-29: The x direction aspheric α coefficients.

30-45: The y direction aspheric β coefficients.

46 and up: The Zernike Standard coefficients.

Face Numbers: All faces Face 0.

Binary 1

The Binary 1 object is a standard lens with a diffractive optic phase profile similar to the Binary 1 surface type on the front face. The Binary 1 phase profile adds phase to the ray according to the following polynomial expansion:

$$\Phi = M \sum_{i=1}^N A_i E_i(x, y),$$

where N is the number of polynomial coefficients in the series, M is the diffraction order, and A_i is the coefficient on the i^{th} extended polynomial term. The polynomials are a power series in the normalized coordinates x and y , as described in “Extended Polynomial Lens” on page 416. The coefficients A_i all have units of radians (2π radians is one wave).

The definitions for the parameters are:

1-9: See the Standard Lens description for information on these parameters.

10: The diffraction order M .

11: The normalization radius. The x and y coordinates are normalized to this value. This keeps all coefficients in units of radians.


12: The maximum term number.

13-242: The values of the coefficients on the polynomial terms.

This object does not diffract rays correctly if the front surface is made hyperhemispheric. See also the Binary 2 object.

For important information on diffractive objects, see “Diffraction from NSC objects” on page 426.

Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.



Binary 2

The Binary 2 object is a standard lens with a diffractive optic phase profile similar to the Binary 2 surface type on the front face. The Binary 2 phase profile adds phase to the ray according to the following polynomial expansion:

$$\Phi = M \sum_{i=1}^N A_i \rho^{2i},$$

where N is the number of polynomial coefficients in the series, A_i is the coefficient on the $2i^{th}$ power of ρ , which is the normalized radial aperture coordinate. The coefficients A_i all have units of radians (2π radians is one wave).

The definitions for the parameters are:

1-9: See the Standard Lens description for information on these parameters.

10: The diffraction order M .

11: The normalization radius. The x and y coordinates are normalized to this value. This keeps all coefficients in units of radians.


12: The maximum term number.

13-242: The values of the coefficients on the polynomial terms.

This object does not diffract rays correctly if the front surface is made hyperhemispheric. See also the Binary 1 object.

For important information on diffractive objects, see “Diffraction from NSC objects” on page 426.

Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.



Binary 2A

The Binary 2A object has an even polynomial aspheric front and back face, with a diffractive optic phase profile similar to the Binary 2 surface type on the front face. The even aspheric sag formula is:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2 r^2}} + \sum_{i=1}^{12} \alpha_i r^{2i}.$$

The Binary 2A object consists of two of these faces, separated by a thickness.


The Binary 2A phase profile is placed on the front face, and adds phase to a ray according to the following polynomial expansion:

$$\Phi = M \sum_{i=1}^N A_i \rho^{2i},$$

where N is the number of polynomial coefficients in the series, A_i is the coefficient on the $2i^{th}$ power of ρ , which is the normalized radial aperture coordinate. The coefficients A_i all have units of radians (2π radians is one wave).

The definitions for the parameters are:

- 1: The Maximum Radial aperture.
- 2: The Thickness of the lens at the center.
- 3-4: Unused.
- 5: The front face radius of curvature.
- 6: The front face conic constant k .
- 7-18: The front face coefficients $\alpha_1 - \alpha_{12}$.
- 19: The back face radius of curvature.
- 20: The back face conic constant k .
- 21-32: The back face coefficients $\alpha_1 - \alpha_{12}$.



33: The diffraction order M .

34: The normalization radius. The x and y coordinates are normalized to this value. This keeps all coefficients in units of radians.

35: The maximum binary phase term number.


36-235: The values of the coefficients on the polynomial terms.

See also the Binary 1 and Binary 2 objects. For important information on diffractive objects, see “Diffraction from NSC objects” on page 426. The reference coordinate is the center of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Boolean

A Boolean object is defined by a series of Boolean operations performed on other objects. Boolean objects can be used to form very general shapes by adding, subtracting, and intersecting other standard shapes. For example, a lens with a hole can be created by defining a Standard Lens object (see “Standard Lens” on page 436) and then subtracting from the lens a Cylinder Volume (see “Cylinder Volume” on page 413). A wide range of complex objects may be created by performing successive Boolean operations on other objects.

The Boolean operations are performed by converting each component object into a NURBS based representation, and then a series of Boolean trimming and combining operations are performed to yield the resulting object. For this reason, Boolean objects have the same considerations that Imported objects do; see “Comments about imported objects” on page 424 for a discussion. Some loss of precision is possible when converting from ZEMAX native, high precision representations to NURBS representations; this is not a limitation of ZEMAX but is inherent in the NURBS



representation of arbitrary surfaces. Generally, precision can be increased as desired using the Spline and Tolerance parameters described below. When extreme optical precision is required, the user should verify that the Boolean object traces to a suitable accuracy.

The Boolean object is defined by 22 parameters and a string defining the desired Boolean operations:

1: Spline controls the number of points to use for fitting splines to aspheric surfaces during the conversion from ZEMAX format to NURBS format. The value for spline is an integer code, where 0 = 4 points, 1 = 8 points, 2 = 16 points, 3 = 32 points, 4 = 64 points, 5 = 128 points, and 6 = 256 points. Higher spline settings increase accuracy for some objects at the expense of slower object creation, rendering, and ray tracing.

2: Tolerance sets the maximum surface deviation between the actual object and the NURBS representation during creation of the object. Tolerance is an integer with 0, 1, 2 or 3 with the maximum deviation being 1.0E-4, 1.0E-5, 1.0E-6, and 1.0E-7 lens units, respectively. Setting the tolerance integer too high (and the corresponding length tolerance too low) may cause the boolean object to be malformed.

3: Mode: The mode controls the tradeoff between set-up time and ray tracing speed. Use mode 1 for fast set up time and slower ray tracing, mode 2 for medium set up time and medium ray tracing, and mode 3 for slow set up time and fast ray tracing. Generally use mode 1 during set-up of the analysis, and mode 3 for analysis tracing a large number of rays.

4-6: # X Y, Z Voxels: Voxels is a name derived from "volume pixels". A voxel is a 3D rectangle that defines some portion of the total volume occupied by imported solids. Voxel technology allows for fast ray tracing by pre-computing which objects, or portions of objects; lie within a given voxel. A ray entering a voxellated space may only intersect some subset of the total number of voxels; and therefore only these voxels need to be checked for possible ray-object intersections. The greater the number of voxels, the longer the set-up time but the faster the ray tracing. It generally takes some experimenting to determine the optimum number of voxels. Use 5 for all three values if no other values seem obviously superior.

7: Chord Tolerance: This setting only affects the rendering of the solid after it is created. To render the solid, ZEMAX converts the solid to a list of triangles which approximate the shape. The tolerance is the maximum allowed distance in lens units between a single triangle and the actual surface of the solid. More triangles are added if the tolerance is set smaller which yields more accurate rendering, at the expense of speed and a larger memory requirement. The default value of zero will use a chord tolerance related to the size of the object sufficient to generate a coarse approximation of the object shape that will render quickly.

8-12: Unused, reserved for future expansion of this feature.

13-22: The object numbers that define which object is A, B, C, etc. See the discussion below.

Parameters 13-22 are used to define which "parent" objects will be combined. Parameter 13 is used to define which object is Object A, parameter 14 defines which object is Object B, etc. There is no need to define more objects than are required for the desired Boolean operation. If only two objects are to be combined, then only Object A and B need to be defined. Any unused objects may be left with an object number of zero. All defined objects must precede the Boolean object in the NSC Editor. Parent objects may not be sources or detectors. Once all the parent objects are defined, the Boolean operation is defined by a "Control String" placed in the "Comments" column of the form:

```
object oper object [oper object ]...
```

The "object" argument is a single letter, such as A, B, or C; which corresponds to the desired object. The "oper" is a single character operator. The following operators (between objects A and B in the logical example) are supported:

- + Combines the two objects together (logical A OR B).
- Subtracts the second object from the first (logical A AND NOT B).
- & Computes the intersection of the two objects (logical A AND B).
- ^ Yields the portion of the object that is part of one or the other, but not both objects (logical A XOR B).
- \$ Subtracts the first object from the second (logical NOT A AND B).

Here are some example Control Strings:

```
A+B  
E+A-B-C
```

All operations are carried out in left-to-right order. No order precedence, such as parentheses, is allowed. Each Boolean operation is performed on the result of previous operations to the left of the operator, with the object listed just to the right of the operator being the second argument. Note the \$ operator can be used to subtract the first object from the second. This is useful for combining objects and then subtracting the combination from a third object. For example, to create the object (A - B) and then subtract this combined object from C, use the string "A - B \$ C".

The local coordinate system of the resulting object will be the same as the local coordinate system of the first object listed in the Control String. The global object position and orientation is set by the object position and tilt parameters in the same manner as all NSC objects.

If the parent objects are not needed, use the "Rays Ignore This Object" and "Do Not Draw Object" settings on the parent object (See "Type tab" on page 426 and "Draw tab" on page 435). The Boolean object will still be created even if the parent objects are not drawn or ray traced.

The Face numbers and properties of the Boolean object are defined by the face numbers and properties of the parent objects. All face properties from the parent objects are copied to the corresponding faces of the Boolean object. To view each face on any object, use the NSC Object Viewer (See "NSC Object Viewer" on page 126). There is a limit to the total number of unique faces on any single object; for details see "Object faces" on page 437.

One significant advantage of the Boolean object is that the resulting volume is still parametrically defined by the parent objects. Changing any of the parent object parameters will change the Boolean object as well. If this capability is not needed, the Boolean object may be exported as a IGES or STEP file (see "Export IGES/SAT/STEP Solid" on page 257) and then imported using the Import object (see "Imported" on page 423), and then the parent objects may be deleted.

Compound Parabolic Concentrator (CPC)

A CPC is defined by the following parameters:


- 1: The radial aperture at $z = 0$, a .
- 2: The maximum acceptance angle in degrees, θ .
- 3: The length along the local Z axis, L .
- 4-5: Unused.
- 6: The "Is Volume?" flag.
- 7: The volume index of refraction.

A CPC is used to concentrate light entering one end of the CPC to the other end. Only rays that make an angle less than the acceptance angle with respect to the local Z axis will pass through the CPC; other rays will reflect back out. This type of CPC is the "Basic CPC" as described in detail in "High Collection Nonimaging Optics" by W. T. Welford and R Winston, Academic Press (1989).

If the "Is Volume?" flag is zero, then the object is a hollow shell. Otherwise, the object is a closed solid volume. If the CPC is a closed volume, it may be made of a refractive material. The reference index of the material can be entered for parameter 7. This index will be used to scale the acceptance angle (via Snell's law) and the refracted acceptance angle will be used to determine the CPC shape and maximum length.

The maximum value for the CPC length is given by

$$L = \frac{a(1 + \sin\theta)}{\tan\theta \sin\theta},$$



Longer lengths will be truncated to this value. The radial coordinate of points on the CPC as a function of the z coordinate along the axis is given by the positive real root of this quadratic equation:

$$C^2 r^2 + 2(CS z + aP^2)r + (z^2 S^2 - 2aCQz - a^2PT) = 0, \text{ where}$$

$$C = \cos\theta, S = \sin\theta, P = 1 + S, Q = 1 + P, \text{ and } T = 1 + Q.$$

The acceptance angle must be between 0.0 and 89.0 degrees, inclusive.

Face Numbers: Front face Face 1, rear face Face 2, outer faces Face 0.

CPC Rectangular


This object is similar to the CPC (see "Compound Parabolic Concentrator (CPC)" on page 411), but with rectangular rather than radial symmetry. See that object description for detailed equations defining the CPC.

The CPC Rectangular is defined by the following parameters:

- 1: The x direction semi-aperture at z = 0.
- 2: The maximum x-direction acceptance angle in degrees.
- 3: The y direction semi-aperture at z = 0.
- 4: The maximum y-direction acceptance angle in degrees.
- 5: The length along the local Z axis.
- 6: The "Is Volume?" flag.
- 7: The volume index of refraction.

The acceptance angles must be between 0.0 and 89.0 degrees, inclusive.

Face Numbers: Front face Face 1, rear face Face 2, outer faces Face 0.




Cone

A cone is defined by 4 parameters:

- 1: The z coordinate of the first point.
- 2: The radial coordinate of the first point.
- 3: The z coordinate of the second point.
- 4: The radial coordinate of the second point.

The line segment defined by the points is rotated about the z axis to form a section of a cone. This object can be used to make an annular or circular shape (if the two z coordinates are identical) or a cylindrical shape (if the two r coordinates are identical). In this sense the cone is redundant with the annulus and cylinder pipe objects. The cone is used as a primitive for creating Fresnel lenses.

The reference coordinate is locally (0, 0, 0), and the points that define the cone may be placed anywhere relative to the reference point. Face Numbers: All faces Face 0.




Cylinder Pipe

A cylinder pipe is a rotationally symmetric surface defined by 3 parameters:

- 1: The radius of the front circular aperture.
- 2: The length along the local Z axis of the cylinder.
- 3: The radius of the rear circular aperture.

This object is normally used to define a reflective light pipe.

The reference coordinate is the center of the front aperture.
Face Numbers: All faces Face 0.




Cylinder Volume

A cylinder volume is a rotationally symmetric volume defined by 3 parameters:

- 1: The radius of the front circular face.
- 2: The length along the local Z axis of the cylinder.
- 3: The radius of the rear circular face.

This object is very similar to the cylinder pipe, except the front and rear faces are included to make the shape a closed volume. Because the object is a closed volume, it may be reflective, refractive, or absorbing.

The reference coordinate is the center of the front aperture.
Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.




Cylinder 2 Pipe

A cylinder 2 pipe is a rotationally symmetric surface defined by 4 parameters:

- 1: The radius of the cylinder.
- 2: The length along the local Z axis of the cylinder.
- 3: The tilt of the front face along the y direction in degrees.
- 4: The tilt of the rear face along the y direction in degrees.
- 5: The tilt of the front face along the x direction in degrees.
- 6: The tilt of the rear face along the x direction in degrees.

The reference coordinate is the center of the front face. Face Numbers: All faces Face 0.



Cylinder 2 Volume


A cylinder volume is a rotationally symmetric volume defined by 4 parameters:

- 1: The radius of the cylinder.
- 2: The length along the local Z axis of the cylinder.
- 3: The tilt of the front face along the y direction in degrees.
- 4: The tilt of the rear face along the y direction in degrees.
- 5: The tilt of the front face along the x direction in degrees.
- 6: The tilt of the rear face along the x direction in degrees.

This object is very similar to the cylinder 2 pipe, except the front and rear faces are included to make the shape a closed volume.

Because the object is a closed volume, it may be reflective, refractive, or absorbing.

The reference coordinate is the center of the front aperture. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.



Diffraction Grating


The diffraction grating is a Standard Lens with a diffraction grating on one face. For a rectangular diffraction grating, see "Lenslet Array 1" on page 426. The Diffraction Grating is defined by the following parameters:

1-9: See "Standard Lens" on page 436 for information on these parameters.

10: The grating line frequency in lines/micrometer on the front face. See parameter 12 below.

11: The diffraction order for the front face.

12: The grating formula. If zero, the grating is a constant spacing linear grating with a grating frequency defined by parameter 10. If the formula is "1", the grating frequency in lines/micrometer is defined by



$$t = t_0 + t_1y + t_2y^2 + t_3y^3 + t_4y^4,$$

where t_0 is defined by parameter 10 and the other coefficients are defined by parameters 13-16.

The grating is assumed to consist of lines parallel to the local x axis. The grating frequency is the lines per micrometer along the y direction; projected down on to the surface.

This object does not diffract rays correctly if the front diffractive surface is made hyperhemispheric.

For important information on diffractive objects, see "Diffraction from NSC objects" on page 426.

Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Dual BEF Surface

The Dual Brightness Enhancement Film (BEF) Surface is an idealized surface that can reflect and transmit rays with variable amplitude depending upon the polarization state of the ray. The surface is always a rectangular plane, and may not have any material, coatings, or scattering function applied. The surface may be imbedded in any isotropic, homogenous medium. If either ray splitting or polarization is off, the surface has no optical effect. Otherwise, any incident ray is split into a transmitted ray and a reflected ray. The transmission and reflection intensity coefficients are defined as parameters 3-6, and are different for X and Y polarized light, where the X and

Y directions are defined in the surface coordinate system. The transmission and reflection for X polarized light may not exceed 1, and the same limit applies to Y polarized light.

Note that both reflected and transmitted electric fields must remain orthogonal to the respective ray vectors, and so for some combinations of incident angle, polarization state, and coefficient values, some loss of energy is possible when rays are incident upon this surface. The 6 parameters used to define this object are:

1-2: The X and Y Half Widths.

3-6: Reflection X, Reflection Y, Transmission X, Transmission Y

Face Numbers: All faces Face 0.

Ellipse


The ellipse is a plane elliptical surface shape defined by 2 parameters:

1: The X Maximum Half Width.

2: The Y Maximum Half Width.

The ellipse resides entirely within the local XY plane. This object is a special case of the more general annulus.

The reference coordinate is the center of the ellipse. Face Numbers: All faces Face 0.



Elliptical Volume

An elliptical volume is a tapered volume or surface with an elliptical cross section defined by these parameters:

1: The x half width of the front face.

2: The y half width of the front face.

3: The x half width of the back face.

4: The y half width of the back face.


5: The length along the local Z axis of the cylinder.

6-7: Unused.

8: The "Is Volume?" flag. Use 1 to make a closed volume, else use 0 to make a hollow shell.

This object is similar to the cylinder pipe and volume. Note this object may be either a hollow shell or a closed volume, depending upon the Is Volume? flag.

The reference coordinate is the center of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.




Even Asphere Lens

The Even Asphere surface shape is defined by:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \alpha_1r^2 + \alpha_2r^4 + \alpha_3r^6 + \alpha_4r^8 + \alpha_5r^{10} + \alpha_6r^{12} + \alpha_7r^{14} + \alpha_8r^{16},$$

which is exactly the same sag formula as for the Even Asphere surface. The Even Asphere Lens object consists of two of these faces, separated by a thickness. The total object shape is defined by 24 parameters:

- 1: The Maximum Radial aperture.
- 2: The Thickness of the lens at the center.
- 3-4: Unused.
- 5: The front face radius of curvature.
- 6: The front face conic constant k.
- 7-14: The front face coefficients $\alpha_1 - \alpha_8$.
- 15: The back face radius of curvature.
- 16: The back face conic constant k.
- 17-24: The back face coefficients $\alpha_1 - \alpha_8$.



The reference coordinate is the center of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Extended Polynomial Lens


The Extended Polynomial surface shape is defined by:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=1}^N A_i E_i(x, y).$$

where N is the number of polynomial coefficients in the series, and A_i is the coefficient on the i^{th} extended polynomial term. The polynomials are a power series in x and y . The first term is x , then y , then $x*x$, $x*y$, $y*y$, etc. There are 2 terms of order 1, 3 terms of order 2, 4 terms of order 3, etc. The maximum order is 20, which makes a maximum of 230 polynomial aspheric coefficients. The position values x and y are divided by a normalization radius so the polynomial coefficients are dimensionless.

The Extended Polynomial Lens object consists of two of these faces, separated by a thickness. The total object shape is defined by these parameters:

- 1: The radial height of the lens object in lens units. This value is used for the y direction half height if the lens is rectangular.
- 2: The x half width of the lens object in lens units. If zero, the lens is circular; otherwise, the lens is rectangular.
- 3: The center thickness of the lens in lens units.
- 4-5: Unused.
- 6: The front face radius of curvature. If this value is zero, then the curvature is assumed to be zero.
- 7: The front face conic constant k .
- 8: The front face normalization radius.
- 9: The front face number of extended polynomial terms.
- 10: The rear face radius of curvature. If this value is zero, then the curvature is assumed to be zero.
- 11: The rear face conic constant k .
- 12: The rear face normalization radius.
- 13: The rear face number of extended polynomial terms.



14 and on: The front and rear face extended polynomial coefficients. The rear face coefficients follow the front face coefficients.

The reference coordinate is the center of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Extended Polynomial Surface

This object is very similar to the Extended Polynomial Lens object. Rather than being a solid, the object is a shell surface with only 1 set of extended asphere coefficients. The object is defined by these parameters:

1: The radial height of the surface. This object supports user defined aperture shapes, see "User defined apertures" on page 437.

2: The radius of curvature.

3: The conic constant k .


4-5: Unused.

6: The normalization radius.

7: The number of extended polynomial terms.

8 and on: The extended polynomial coefficients.

The reference coordinate is the center of the surface. Face Numbers: All faces Face 0.



Extruded

An Extruded object is a User Defined Aperture extruded to form a solid volume. Any User Defined Aperture file may be used to create the Extruded object. For information on User Defined Apertures see "User defined apertures and obscurations" on page 80. The extrusion is along the Z axis, and the front and rear faces of the object may be independently and anamorphically scaled. The rear face may also be decentered relative to the front face.

The User Defined Aperture file name is placed in the comment column. The following additional parameters are used:

1: The length along the Z axis between the front and rear faces.

2-3: The front face X and Y direction scaling factors.

4-5: The rear face X and Y direction scaling factors.

6-7: The rear face X and Y direction decenters.

The reference coordinate is the origin of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.



Freeform Z

A Freeform Z object is formed by drawing a cubic spline curve through a series of points in the YZ plane, and then sweeping this curve around the Z axis to form either a solid volume or shell. The Freeform Z object is defined by two parameters and a user selectable number of defining points:

1: The number of points. This is the number of coordinate pairs used to define the curve in the YZ plane. The minimum is 5, the maximum is 124.

2: A flag to indicate if the resulting object is a solid volume or a hollow shell. If the "Is Volume?" parameter is zero, then ZEMAX assumes the shape is hollow shell. No "end caps" are added in this case. If the object is a solid volume, then flat "end caps" are added if the first or last points do not have a y coordinate of zero to close the solid.

3-250: The coordinate values of the YZ curve, in pairs.



The Z coordinate of the first point must be zero. No Z value may be less than or equal to the Z value of the preceding point. The absolute value of the Y values is always taken when creating the spline curve.

There is a special optimization operand designed to provide boundary constraints on the shape of Freeform Z objects, see "FREZ" on page 500 for more information.

A tool is available that greatly simplifies inserting and deleting points in the spline curve; see "Insert/Delete Z Point Tools" on page 470.

The reference point is the center of the front face of the lens. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Fresnel 1

This object is a general radially symmetric or cylindrical solid Fresnel lens constructed by modeling the detailed faces of the Fresnel surface. For an idealized Fresnel lens (which ignores the detailed structure of the Fresnel surface and thus ray traces much faster) see the Fresnel 2 object.

The substrate shape is a flat disk (if radial) or a rectangle (if cylindrical). One face of the substrate consists of radial or rectangular faces defining the profile of the Fresnel that yields optical power. The profile is constructed of radially flat faces (or a series of flat faces if subsegments are used) whose endpoints are defined by a sag expression identical to the Even Asphere surface:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \alpha_1r^2 + \alpha_2r^4 + \alpha_3r^6 + \alpha_4r^8 + \alpha_5r^{10} + \alpha_6r^{12} + \alpha_7r^{14} + \alpha_8r^{16}.$$

However, to make a Fresnel lens, each face is offset just enough along the z axis so all faces start at the same z coordinate as the center vertex point. This yields a lens that is "collapsed" into a relatively small volume. The Fresnel faces are automatically generated by ZEMAX using these 16 parameters:

1: Radial Height: This is the maximum radial aperture of the lens if radially symmetric, or the y half height if cylindrically symmetric.



2: X Half Width: This is the half width of the lens if cylindrically symmetric. If this parameter is zero, then a rotationally symmetric lens is generated.

3: +Depth/-Frequency: If this parameter is positive, then it corresponds to the depth of each groove in lens units. If negative, then it corresponds to the frequency of the grooves. For example, a value of -2.0 will yield 2 grooves per radial lens unit. If the groove depth is defined, the radial positions of the grooves will generally vary; if the groove frequency is defined; the groove depth will vary. For the case of a defined groove depth ZEMAX automatically computes the exact radial coordinate at which the sag has changed by the specified depth. This is done using an iterative search.

4: Pitch (degrees): The pitch is the angle the "inactive" faces (those faces nominal parallel to the local z axis) make with respect to the z axis. The pitch is always radially outward, whether the pitch angle is positive or negative. A pitch angle of a few degrees is typically added to Fresnel molds to make extraction of the molded part easier.

5: Thick: The thickness of the Fresnel in lens units. This value may be positive or negative; but should be chosen such that the absolute value of the thickness exceeds the deepest groove depth; or a non-physical Fresnel may be generated without warning or error message.

6: Radius: The base radius of curvature. This is one over the value "c" in the sag expression above.



7: Conic: The conic constant "k" in the sag expression above.

8-15: The coefficients on the even radial powers of r. Note these coefficients have units as r is not normalized.

16: The number of sub segments. The greater the number of sub segments, the smoother the approximation to a curved surface between the grooves. A value of 1 yields flat grooves, larger values yield progressively more smooth curved faces at the expense of ray tracing speed. The maximum number of sub segments is 20.

Because the object is a closed volume, it may be reflective, refractive, or absorbing.

The reference coordinate is the center vertex of the side of the lens with the Fresnel faces. If the radius or aspheric terms yield negative sag values, then it is important to offset the position of the Fresnel such that the entire solid resides inside the non-sequential group. If the entry port is placed inside the lens; incorrect ray tracing will result.

Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Fresnel 2

This object is an idealized Fresnel lens. Unlike the Fresnel 1, this Fresnel lens uses the approximation that the Fresnel faces are infinitesimally small, and may be ignored for purposes of computing the ray-object intersection point. Unlike the idealized Fresnel surface in sequential ray tracing (see "Fresnel" on page 303), this object does not add or subtract any phase to the ray.

The substrate shape is either a flat disk or a rectangle. The front face of the substrate consists of a radial or cylindrical Fresnel lens. The radial profile is defined by a sag expression identical to the Even Asphere surface:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \alpha_1r^2 + \alpha_2r^4 + \alpha_3r^6 + \alpha_4r^8 + \alpha_5r^{10} + \alpha_6r^{12} + \alpha_7r^{14} + \alpha_8r^{16} .$$

If the Fresnel surface is cylindrical, then the profile is described by an identical expression in the y coordinate alone:

$$z = \frac{cy^2}{1 + \sqrt{1 - (1+k)c^2y^2}} + \alpha_1y^2 + \alpha_2y^4 + \alpha_3y^6 + \alpha_4y^8 + \alpha_5y^{10} + \alpha_6y^{12} + \alpha_7y^{14} + \alpha_8y^{16} .$$

The Fresnel lens may be either radial or cylindrical independent of the substrate shape. It is possible to define a radial Fresnel on a rectangular substrate or a cylindrical Fresnel on a radial substrate if desired.


The Fresnel 2 object is defined using these 14 parameters:

1: Radial Height: This is the maximum radial aperture of the lens if radially symmetric, or the y half height if cylindrically symmetric.

2: X Half Width: This is the half width of the lens if cylindrically symmetric. If this parameter is zero, then a rotationally symmetric lens is generated.

3: Thick: The thickness of the Fresnel in lens units. This value must be positive.

4: Is Cylinder: If zero, the Fresnel surface will be radial, otherwise, the surface will be cylindrical.



5: Radius: The base radius of curvature. This is one over the value "c" in the sag expressions above.


6: Conic: The conic constant "k" in the sag expressions above.

7-14: The coefficients on the even radial powers of r if "Is Cylinder" is zero, otherwise these are the coefficients on the even radial powers of y. Note these coefficients have units as y and r are not normalized.

Because the object is a closed volume, it may be reflective, refractive, or absorbing.

The reference coordinate is the center vertex of the side of the lens with the Fresnel faces.

Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.



Grid Sag Lens

The Grid Sag Lens is a solid volume object with a Standard surface front face, and a Grid Sag surface rear face. The Standard surface shape is described in "Standard" on page 325. The Grid Sag surface shape is described in "Grid Sag" on page 313.

Grid Sag Lens objects use the following parameters:


- 1: The x half width in lens units. See the discussion below.
- 2: The y half width in lens units. See the discussion below.
- 3: The front surface radius of curvature in lens units.
- 4: The front surface conic in lens units.
- 5: The axial lens thickness in lens units.
- 6: Interpolate: Use 0 for bicubic, and 1 for linear. See "Bicubic spline vs. linear interpolation" on page 315.
- 7: The back surface radius of curvature in lens units.
- 8: The back surface conic constant.
- 9-16: The back surface α_j aspheric coefficients.

17-18: The X- and Y-decenter values in lens units. This is the decenter of the grid data relative to the aspheric substrate. The decenter values are added to any decenter values defined in the GRD file, described below.

If the x and y half widths are zero, the lens shape is rectangular and the size is defined by the data in the GRD file. If the x half width is zero, and the absolute value of the y-half width is not zero but is less than the size of the data in the GRD file, then the part is square (if the y half width is greater than zero) or circular (if the y half width is less than zero). If both the the x half width and the y half width are not zero, and the absolute values of both half widths are less than the size of the data in the GRD file, then the part is rectangular (if both half widths are positive) or elliptical (if either half width is negative).

The actual point by point sag data is defined by placing the data in a properly formatted file ending in the extension GRD. The file format is defined in "Importing grid data" on page 314. The grid data file must be placed in the folder <objects>\Grid Files (see "Folders" on page 68). The file name, without the path, is then placed in the comment column of the NSC Editor.

The reference coordinate is the center of the front face. Face Numbers: Side faces Face 0, front face Face 1, back face Face 2.




Grid Sag Surface

The Grid Sag Surface is an object whose shape is defined by a rectangular array of points listed in a file. This surface shape is similar to the sequential Grid Sag surface, although no Zernike terms are supported. For a detailed description of the sequential Grid Sag surface shape, see “Grid Sag” on page 313.

Grid Sag Surface objects use the following parameters:

- 1: Interpolate: Use 0 for bicubic, and 1 for linear. See “Bicubic spline vs. linear interpolation” on page 315.
- 2: The radius of curvature in lens units.
- 3: The conic constant.
- 4-11: The α_i aspheric coefficients.



12-13: The X- and Y-decenter values in lens units. This is the decenter of the grid data relative to the aspheric substrate. The decenter values are added to any decenter values defined in the GRD file, described below.


The actual point by point sag data is defined by placing the data in a properly formatted file ending in the extension GRD. The file format is defined in “Importing grid data” on page 314. The grid data file must be placed in the <objects>\Grid Files folder (see “Folders” on page 68). The file name, without the path, is then placed in the comment column of the NSC Editor.

Face Numbers: The entire face is Face 0.

Hexagonal Lenslet Array

The hexagonal lenslet array object models a rectangular volume with a front face consisting of a hexagonal array of even aspheric lenslets. To make the part rectangular, partial hexagonal lenses are added along the sides as required. The even asphere surface sag expression is identical to that of the Even Asphere Lens object (see “Even Asphere Lens” on page 415). The back face is plano. The object uses the following parameters:

- 1: The number of columns (must be an odd positive integer).
- 2: The number of rows (must be an odd positive integer).
- 3: Unused.
- 4: Width: This is the full width of a single hexagon lenslet in lens units.



5: Thickness: The thickness of the lenslet, measured along the +z axis from the center lenslet vertex to the back plano face.

6: Radius: The radius of curvature of each lenslet in lens units.

7: Conic: The conic constant of each lenslet.

8-15: The aspheric coefficients on the powers of r.

The reference coordinate is the center of the front face. Face Numbers: Side faces Face 0, front face Face 1, back face Face 2.

Hologram Lens

This object is an ideal optically fabricated hologram similar to the Hologram 1 and Hologram 2 sequential surface models (see “Hologram 1” on page 316 for a description of these surfaces). The hologram is a solid, and may be circular or rectangular in shape. The front and back faces may be plane, spheres, or conic aspheres. The hologram surface is on the front face. See also “Hologram Surface” on page 422.

The Hologram object is defined using these 18 parameters:

1: The radial height of the lens object in lens units. This value is used for the y direction half height if the lens is rectangular.

2: The x half width of the lens object in lens units. If zero, the lens is circular; otherwise, the lens is rectangular.

3: The center thickness of the lens in lens units.

4-5: Unused.

6: Radius 1: The radius of the front face.

7: Conic 1: The conic constant of the front face.

8: Radius 2: The radius of the back face.

9: Conic 2: The conic constant of the back face.

10: Holo Type: Use 1 for Hologram type 1 (both sources converging/diverging) or 2 for Hologram type 2 (one source converging, one diverging) See "Hologram 1" on page 316.

11: Order: The diffraction order to use. Multiple orders may be specified, see "Diffraction tab" on page 430.

12: Construction Wavelength: The wavelength in micrometers used to fabricate the hologram.

13-18: The X, Y, and Z coordinates in lens units of the construction points relative to the vertex of the front face of the Hologram.

Because the object is a closed volume, it may be air, reflective, refractive, or absorbing. If the order is zero or if the ray totally internally reflects at the hologram boundary, no hologram diffraction is computed.

For important information on diffractive objects, see "Diffraction from NSC objects" on page 426.

The reference coordinate is the center of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.


Hologram Surface

This object is an ideal optically fabricated hologram similar to the Hologram 1 and Hologram 2 sequential surface models (see "Hologram 1" on page 316 for a description of these surfaces). The hologram is a surface, and may be circular or user defined in shape. The surface shape may be plane, sphere, or a conic and/or polynomial asphere. See also "Hologram Lens" on page 421.

The surface shape is defined by the following sag equation:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1 + k)c^2r^2}} + \sum_{i=1}^N \alpha_i r^i,$$

where c is the curvature of the surface, k is the conic constant, r is the radial coordinate, and the α terms are aspheric coefficients. The surface supports specification of both a minimum and a maximum radial aperture; so annular surfaces may be defined. Note that both even and odd terms are defined; up to approximately 230 coefficients may be used. This is the same shape as the Aspheric Surface object, see "Aspheric Surface" on page 402.



The following parameters are used to define the hologram surface:

1: The radius of curvature. If this value is zero, then the curvature is assumed to be zero.

2: The conic constant k .

3: The maximum radial aperture in lens units.

4: The minimum radial aperture in lens units. This value may be zero.

5-6: Unused.

7: Holo Type: Use 1 for Hologram type 1 (both sources converging/diverging) or 2 for Hologram type 2 (one source converging, one diverging) See "Hologram 1" on page 316.

8: Order: The diffraction order to use. Multiple orders may be specified, see "Diffraction tab" on page 357.

9: Construction Wavelength: The wavelength in micrometers used to fabricate the hologram.

10-15: The X, Y, and Z coordinates in lens units of the construction points relative to the vertex of the front face of the Hologram.

16: The number of terms to use in the aspheric expansion. Ray tracing will be faster if this term is no larger than the highest order non-zero coefficient.

17-250: The α coefficients on the polynomial expansion.

Face Numbers: All faces Face 0. This object supports user defined apertures, see "User defined apertures" on page 364.

Imported

Objects may be imported in either IGES, SAT, or STEP format. These file formats are commonly supported by CAD programs. Imported objects may also be in ZOF (ZEMAX's native format, see "Imported objects and ZOF files" on page 424). Once imported, an object may have optical properties such as coatings, glass, and scattering functions applied, and then be ray traced like any other optical component.

To import an object, set the object type to "Imported" and select the file name from the drop-down box, or place the file name in the comment column. The file must reside in the <objects>\CAD Files folder (see "Folders" on page 68) and end in the extension IGS, IGES, SAT, STP, STEP, or ZOF. The extension must be included with the file name on the comment column.


The following parameters are used by the Import object type.

1: Scale: This dimensionless parameter scales the entire imported solid. Upon import, ZEMAX will automatically attempt to scale the dimensions of the imported solid to match the current dimensions in ZEMAX; this scale factor is applied after that conversion.

2: Mode: The mode controls the tradeoff between set-up time and ray tracing speed. Use mode 1 for fast set up time and slower ray tracing, mode 2 for medium set up time and medium ray tracing, and mode 3 for slow set up time and fast ray tracing. Generally use mode 1 during set-up of the analysis, and mode 3 for analysis tracing a large number of rays.

3-5: # X Y, Z Voxels: Voxels is a name derived from "volume pixels". A voxel is a 3D rectangle that defines some portion of the total volume occupied by imported solids. Voxel technology allows for fast ray tracing by pre-computing which objects, or portions of objects; lie within a given voxel. A ray entering a voxellated space may only intersect some subset of the total number of voxels; and therefore only these voxels need to be checked for possible ray-object intersections. The greater the number of voxels, the longer the set-up time but the faster the ray tracing. It generally takes some experimenting to determine the optimum number of voxels. Use 5 for all three values if no other values seem obviously superior.

6: Chord Tolerance: This setting only affects the rendering of the object. To render the object, ZEMAX generates a list of triangles which approximate the object shape. The tolerance is the maximum allowed distance in lens units between a single triangle and the actual surface of the object. More triangles are added if the



tolerance is set smaller which yields more accurate rendering, at the expense of speed and a larger memory requirement. The default value of zero will use a chord tolerance related to the size of the object sufficient to generate a coarse approximation of the object shape that will render quickly.

7: Face Mode: The Face Mode setting determines how ZEMAX assigns face numbers to the various surfaces on the imported object. Some CAD programs create data files that have many more small surfaces than is useful for optical analysis. For example, a simple cylinder may be described in the CAD file by hundreds of small surfaces, while for optical analysis only two or three different optical properties are applied to the entire object. Rather than assign optical properties to each of the many surfaces, it is usually more convenient to group surfaces by assigning a single face number to all surfaces that form a continuous smooth portion of the object. The following Face Modes are supported to automate this numbering:

Face Mode = 0: All surfaces are assigned face number 0. The entire object will have just one face.

Face Mode = 1: All surfaces whose edges meet along a non-zero length curve, and whose normal vectors along the curve of contact are parallel within a user defined angle tolerance are assigned a common face number. The angle tolerance is defined by the Face Angle (parameter 8). This mode allows control over how finely the faces are numbered. If the Face Angle is set to a large value (such as 180) then all faces that touch will share a common number. Larger Face Angles yield fewer unique faces.

Face Mode = 2: All surfaces are uniquely numbered. This mode yields the largest number of unique faces.

Face Mode = 3: Retains the face numbers defined in the imported file. Some CAD files, such as those created by ZEMAX, have face numbers already defined. If ZEMAX recognizes the face numbers, they will be used. If ZEMAX does not detect the face numbers, the surfaces will be numbered as for Face Mode = 2.

Face Mode = 4: All surfaces on each separate object defined in the CAD file are assigned a common face number. This option is useful for applying one property to all surfaces on each object when more than one object is defined with a single CAD file.

In all cases, the number of unique face numbers cannot exceed the maximum number of faces ZEMAX supports. Any faces which exceed the maximum limit will be renumbered to be face 0. ZEMAX also supports assignment of face numbers to imported object surfaces using the "Face" tab of the NSC Object Viewer, see "Face tab" on page 355. For visual confirmation of the face numbering, see "NSC Object Viewer" on page 126.

8: Face Angle: The angle in degrees used by Face Mode = 1 described above.

Imported objects and ZOF files

When importing CAD objects from IGES, STEP, or other format files, the data must be translated into the ZEMAX internal object representation. This conversion need only be done once, and thereafter the converted file can be used. This greatly speeds up the loading of imported objects. The converted file name is the original file name with the extension ZOF (ZEMAX Object Format) appended. For example, if the original imported file was MyObject.IGS, the ZOF file created will be MyObject.IGS.ZOF. The name of the file used by the imported object will be changed to the converted file name. ZEMAX will automatically convert imported files to ZOF format files if the original file requires more than 2.0 seconds to read and convert, *and if the scale factor is 1.0*. If the scale factor needs to be something other than 1.0, first set the scale factor to 1.0, then after the file has been converted to ZOF format, the scale factor may be adjusted to the desired value. This restriction prevents unexpected scaling of the ZOF object when more than one ZMX file is using the same ZOF object.

It is recommended that the original CAD file be retained, even though the ZOF file will be used by ZEMAX. This allows regeneration of the ZOF file should that be required. This feature may be disabled, see "Convert Imported Files To ZOF" under "Type tab" on page 353.

Comments about imported objects

The advantage to using imported objects is that solids of any shape may be ray traced within ZEMAX. There are no limits on the shape, complexity, or number of objects that may be imported and ray traced. Multiple objects may be imported in a single file. The disadvantage to using imported objects is ray tracing speed, and in some cases, ray tracing accuracy. See the discussion below for important considerations regarding imported objects. An alternative to imported objects exists, see "User Defined Object" on page 444.

Ray tracing accuracy for imported objects

Not all types of surface shapes may be ray traced with adequate accuracy using representations supported by CAD file formats, such as IGES, SAT, and STEP. For planes, spheres, and cylinders, the CAD representation, if done correctly, is of very high precision suitable for optical accuracy ray tracing. However, higher order shapes do not usually have a native representation in CAD formats. For example, an aspheric surface with a polynomial term of the form r^{16} may have no equivalent representation in the chosen CAD format. A CAD program will generally approximate this shape using a segmented spline, which is in general a piece-wise fit of the surface using multiple lower order polynomials. Typically, multiple third or fourth order polynomials are used to approximate the surface. This is probably adequate for mechanical design, but not for optical precision ray tracing, where surfaces must be known to tiny fractions of the wavelength of light.

This problem often arises when a high optical precision surface is modeled in ZEMAX, then exported as a CAD file, then imported as an CAD file for subsequent ray tracing. The optical precision of the part is lost upon exporting the native ZEMAX asphere as a CAD spline. See “Limitations of exported data” on page 259 for more information.

For non-imaging optics, the precision of the CAD representation is usually adequate, but for imaging systems, great care must be taken to verify that the imported CAD part is a suitably accurate description of the desired shape. Note ZEMAX uses a relative internal optical precision of about $1E-12$ for ray tracing. Most CAD representations of objects are many orders of magnitude more coarse.

Ray tracing speed for imported objects

Simple objects such as spherical lenses typically ray trace slower when imported in CAD format than the native ZEMAX object of identical shape. Ray tracing speed for imported objects is critically dependent upon the efficient representation of the solid shape within the imported file. The identical object may be represented in a nearly infinite number of ways using the various solid and surface entity types supported by the CAD formats ZEMAX can import. For example, an efficient representation of an object may use only a few spline surfaces; while an inefficient representation of the object may use hundreds of smaller spline surfaces. Although from a mechanical modeling perspective the two representations may both be valid and the resulting solids identical, the representation with the larger number of spline surfaces will ray trace dramatically slower. The only remedy is to return to the source of the CAD file and see if a more efficient representation may be generated.

Limitations on imported objects

Not all possible valid CAD format files may be imported, as ZEMAX may only import solids. No lines or surfaces are allowed. Shells must be converted to thin solids before being imported into ZEMAX. Solids must be simply closed with a continuous exterior surface without holes or gaps. No internal surfaces or faces are allowed. Solids composed of multiple solid volumes either in contact or overlapping cannot be ray traced. Multiple volumes not in contact are allowed. Files containing valid solids which do not import may be sent to technical support for review; however no guarantee is offered that ZEMAX can be made to accommodate all possible CAD format entity types and files.

Jones Matrix


The Jones Matrix is a plane elliptical surface whose shape is defined by 2 parameters, and polarization transmission/reflection properties by 8 "ABCD" parameters:

- 1: The X Maximum Half Width.
- 2: The Y Maximum Half Width.
- 3-10: Ar, Ai, Br, Bi, Cr, Ci, Dr, Di

The ellipse resides entirely within the local XY plane.

The Jones Matrix may be used to model neutral density filters, polarizers, and rotators. For a description of the Jones Matrix parameters, see “Defining polarizing components” on page 623.

The reference coordinate is the center of the ellipse. Face Numbers: All faces Face 0. This object supports user defined apertures, see “User defined apertures” on page 364.



Lenslet Array 1



The array aspects of this object are provided for back compatibility with older versions of ZEMAX. A more efficient method that ray-traces more quickly is to make a single lenslet and then create an array using the Array Object - see "Array" on page 399 for details.

A Lenslet Array 1 object consists of an array of rectangular volumes, each with a flat front face and a curved back face. The back face may be plane, sphere, conic, or polynomial asphere; or a spherical, conic, or polynomial aspheric toroid. The shape may be decentered with respect to the center of the lenslet. A toroidal surface is defined by a curve in the YZ plane which is then rotated about an axis parallel to the Y axis but displaced by a distance R; called the radius of rotation.

If the back face is rotationally symmetric, the radial profile is defined by a sag expression identical to the Even Sphere surface:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \alpha_1r^2 + \alpha_2r^4 + \alpha_3r^6 + \alpha_4r^8 + \alpha_5r^{10} + \alpha_6r^{12} + \alpha_7r^{14} + \alpha_8r^{16}.$$

If the back face is a toric, then the profile is described by an identical expression in the y coordinate alone:

$$z = \frac{cy^2}{1 + \sqrt{1 - (1+k)c^2y^2}} + \alpha_1y^2 + \alpha_2y^4 + \alpha_3y^6 + \alpha_4y^8 + \alpha_5y^{10} + \alpha_6y^{12} + \alpha_7y^{14} + \alpha_8y^{16},$$

where c is the reciprocal of the radius of curvature in YZ plane.

Either the front or the back surface may be a diffraction grating. The grating is assumed to consist of equally spaced lines parallel to the local x axis. The grating frequency is the lines per micrometer along the y direction; projected down on to the surface.



The Lenslet Array 1 is defined by these parameters:

- 1: The X Half-Width in lens units of each lenslet.
- 2: The Y Half-Width in lens units of each lenslet.
- 3: The thickness along the local Z axis of each lenslet.
- 4-5: Unused.
- 6: The radius of curvature of each lenslet, use zero for a plane.
- 7: The conic constant of each lenslet.
- 8: The "Is Toric" flag; if zero the lenslet surface is rotationally symmetric, otherwise it is toroidal.
- 9: The radius of rotation of each lenslet if the lenslet surface is toric.
- 10: The grating line frequency in lines/micrometer on either the front or back face, see parameter 24.
- 11: The diffraction order. Multiple orders may be specified, see "Diffraction tab" on page 357.
- 12-19: The aspheric coefficients 1-8.
- 20-21: The decenter x and decenter y of the curved back face.
- 22-23: The number of lenslet elements in x and y. If the total number of lenslets is large, it is generally much faster to use a single lenslet (set parameters 22 and 23 both equal to 1) and then use the Array object to create the array of lenslets. The Array object method is much more powerful, uses less memory, and will generally load and trace much faster than using a large number of lenslet elements. See "Array" on page 399.
- 24: The face to make diffractive. This value must be 1 (front face) or 2 (back face).

Note a cylinder lens results if the radius of rotation is set to zero. The reference coordinate is the center of the front face. Face Numbers: Outside edge side faces are Face 0, flat front face Face 1, curved back face Face 2, inside side faces Face 3 (these are the side faces formed in the interior of an array of lenslets when decentered lens faces meet).

Lenslet Array 2




The array aspects of this object are provided for back compatibility with older versions of ZEMAX. A more efficient method that ray-traces more quickly is to make a single lenslet and then create an array using the Arry Object - see "Array" on page 399 for details.

A Lenslet Array 2 object consists of an array of rectangular volumes, each with curved front and back faces. The faces may be plane, sphere, or conic asphere.

The Lenslet Array 2 is defined by these parameters:

- 1: The X Half-Width in lens units of each lenslet.
- 2: The Y Half-Width in lens units of each lenslet.
- 3: The thickness along the local Z axis of each lenslet.
- 4-5: Unused.
- 6: The radius of curvature of the front face of each lenslet, use zero for a plane.
- 7: The conic constant of the front face of each lenslet.
- 8: The radius of curvature of the back face of each lenslet, use zero for a plane.
- 9: The conic constant of the back face of each lenslet.




10-11: The number of lenslet elements in x and y. If the total number of lenslets is large, it is generally much faster to use a single lenslet (set parameters 10 and 11 both equal to 1) and then use the Array object to create the array of lenslets. The Array object method is much more powerful, uses less memory, and will generally load and trace much faster than using a large number of lenslet elements. See "Array" on page 399.

The reference coordinate is the center of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Micro Electro Mechanical System (MEMS)

This object simulates a Micro Electro Mechanical System (MEMS). The MEMS consists of an array of small rectangular mirrors (commonly called pixels). The mirrors may tip at any of three angles, each rotated about an axis to point the mirror in any direction. The mirrors may be turned on and off by rows, by columns, or by individual mirrors if desired to model any state the MEMS can be in. This type of device is sometimes called a Digital Mirror Device (DMD). The following parameters are used to define the MEMS:

- 1: The number of X pixels across each row.
- 2: The number of Y pixels down each column.
- 3: The full X width of the array in lens units. The X pixel width is this number divided by the number of X pixels.
- 4: The full Y width of the array in lens units. The Y pixel width is this number divided by the number of Y pixels.
- 5, 6, 7: The tip angle of the pixel when in state 0, 1, or 2.
- 8: The rotation angle around the Z axis about which the pixel will tip. This angle is measured clockwise from the +Y axis.



9: P-Flag. If this value is 0, then the pixels are addressed by rows. If 1, then the pixels are addressed by columns, and if 2, the pixels are addressed individually.

10 and above: Integer values which define the state of the rows, columns, or pixels, as described below.

The pixels are set to either state 0, 1, or 2 using a series of base-3 integer values. To determine the values for any logic state of the MEMS, construct a table similar to the one below, which to save space just shows the values for 3 rows or columns or pixels (r/c/p):

```
r/c/p: 3 2 1 X
        0 0 0 0
        0 0 1 1
        0 0 2 2
        0 1 0 3
        0 1 1 4
        0 1 2 5
        0 2 0 6
```

etc..

where the value X is entered into parameter 10 to control the first 15 r/c/p; parameter 11 to control r/c/p 16-30, and so on for as many r/c/p exist. Because of the way ZEMAX addresses the individual r/c/p; the total number of addressable pixels cannot exceed approximately 3750. This limits the number of x pixels, y pixels, or the product of x and y pixels, depending upon the address mode defined by the P-flag (parameter 9).

Note that X for any given state of the MEMS is given by:

$$X = M1*(3^0) + M2*(3^1) + M3*(3^2) + \dots$$

where M1 is the logic state (0, 1, or 2) of the first row/column/pixel and M2 is the logic state of the second row/column/pixel, etc. Up to 15 r/c/p values are defined by each parameter value in the editor. The pixels are numbered starting in the bottom left corner with pixel #1, and incrementing across the columns and then up along the rows.

Face Numbers: All faces Face 0.

Null Object

This is a non-existent object. It can be used as a place holder or a reference point for other objects.

The reference coordinate is locally (0, 0, 0).


Odd Asphere Lens

The Odd Asphere surface shape is defined by:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=1}^{12} \alpha_i r^i,$$

which is very similar to the Odd Asphere surface (there are 4 additional terms). The Odd Asphere Lens object consists of two of these faces, separated by a thickness. The total object shape is defined by 32 parameters:

- 1: The Maximum Radial aperture.
- 2: The Thickness of the lens at the center.
- 3-4: Unused.
- 5: The front face radius of curvature.
- 6: The front face conic constant.
- 7-18: The front face coefficients $\alpha_1 - \alpha_{12}$.
- 19: The back face radius of curvature.
- 20: The back face conic constant.
- 21-32: The back face coefficients $\alpha_1 - \alpha_{12}$.



The reference coordinate is the center of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Paraxial Lens

The Paraxial Lens is an ideal thin lens surface with separate X and Y focal lengths. The defining parameters are:

1/2: The X/Y Half Widths. Make positive for a rectangular shape, negative for an elliptical shape.


3/4: The X/Y direction focal length, measured in air, in lens units.

This object may be placed in air, or be fully or partially immersed within a homogenous or gradient index medium. No adjustment of the optical path length or phase of the ray is made during refraction, so no attempt at estimating the correct phase of the ray is made. No attempt is made to compute the polarization effects of refraction through this idealized part. If polarization effects are important, this object should be used with caution, especially if the ray angles on either side of the lens are large with respect to the normal.

Polygon Object

The polygon object is a very general user-defined object. It can be used to define an open polygon surface or a closed polygon volume with some portions reflective and others refractive or absorptive. The Polygon Object is based on a collection of 3D triangles whose vertices are placed in a file with the POB extension. See the "Defining Polygon Objects" on page 429 for more details. Any Polygon Object may be used as a detector as described in "Objects as detectors" on page 470.

There are no fixed limits to the total number of vertices or polygons. The POB file name is referenced in the "comments" column of the Polygon Object row. For example, if the POB file myobject.POB is placed in the <objects>\Polygon Objects folder (see "Folders" on page 68), then specify "myobject.POB" in the comment column of the Polygon Object type row in the NSC Editor. Polygon objects require two parameters:



1: A scale factor. All vertices in the POB file are multiplied by this parameter.

2: A flag to indicate if the POB file defines a volume or a surface. If the "Is Volume?" parameter is zero, then ZEMAX assumes the POB file defines an open surface. If the "Is Volume?" parameter is any non-zero value, then ZEMAX assumes the POB file defines a closed volume.

See "Special considerations for faceted objects" on page 391 for information on limitations of ray tracing through faceted objects. A tool for creating POB files for use as detectors is described in "Create Polygon Object" on page 96.

The reference coordinate is locally (0, 0, 0), and the polygons that compose the object may be placed anywhere relative to the reference point. Face Numbers: Each face has a face number as defined in the POB file, see below for details.

Defining Polygon Objects

Polygon objects are user-defined objects that are composed of a collection of triangles in 3D space. The triangles are defined in a text file ending in the extension POB (for polygon object). Each face, or a group of faces, on the polygon may be assigned a unique face number if desired. There is a limit to the total number of unique faces on any single object; for details see "Object faces" on page 364.

Polygon objects that form closed volumes must not contain any interior faces. Every face must lie on the exterior of the object, defining the boundary between the inside and the outside of the volume. To make an object that has a cemented boundary, such as a beam splitter, use two or more separate polygon objects and then place them in contact. For more information on placing objects in contact, see "Placing objects inside of, adjacent to, or overlapping other objects" on page 351.

The file format may be created and edited with any text editor. The file consists of a series of rows of data. Each row begins with a single letter or symbol, followed by data for that symbol. The supported symbols and their meanings are defined as follows.

The comment symbol: !

The symbol "!" is used to define a comment line.

Syntax:

```
! Any comment here
```

Example:

```
! A dove prism
```

The face name symbol: C

Names for the faces may be defined using the C command.

Syntax:

```
C facenumber "any name here"
```

Example:

```
C 0 "All faces"
```

The face names are used to display in the object dialog box the names for the various faces of an object. Note the name must be in quotes. This command is optional.

The invisible symbol: I

Lines which connect vertices that are not to be drawn (such as lines between two adjacent triangles that lie in the same plane) can be marked as invisible. To mark two vertices as invisible, use the symbol "I" followed by the two vertex numbers. Any line connecting the marked vertices will not be drawn.

This only affects the 3D Layout plot, and is strictly a cosmetic enhancement to the drawing of polygon objects.

Syntax:

```
I v1 v2
```

Example:

```
I 7 9
```

Note the vertex numbers must be integers. The numbers are separated by spaces. The I command must precede any triangles or rectangles that reference the invisible vertices.

The vertex symbol: V

Vertices are defined by the symbol "V" followed by the vertex number and the x, y, and z coordinates of the vertex.

Syntax:

```
V number x y z
```

Example:

```
V 1 -1.0 -1.0 0.0
```

Note the vertex number must be an integer, and the x, y, and z coordinates are floating point numbers. The numbers are separated by spaces.

The triangle symbol: T

Triangles are defined by the connection of 3 vertices.

Syntax:

```
T vertex1 vertex2 vertex3 isreflective face
```

The vertex numbers must be integer vertices previously defined in the file. If the triangle is part of an object that forms a closed volume, the vertices should be ordered so that the vector cross product of the vector between vertices 1 and 2, and the vector between vertices 1 and 3, points "out" of the volume.

The "isreflective" flag is -1 if the surface absorbs, 1 if the surface reflects, or 0 if the surface refracts. Using this flag allows some triangles to be reflective, and others to be refractive or absorptive, within the same Polygon

Object or face. However, the isreflective flag is not considered in the object produced by Boolean operations with the Polygon Object.

The value for face defines which face the triangle belongs to. If the face number is omitted; face 0 is assumed.

Example:

```
T 1 2 3 0 2
```

This would define a refractive triangle connecting vertices 1, 2, and 3, belonging to face 2.

The rectangle symbol: R

Rectangles are defined by the connection of 4 vertices, otherwise, they are very similar to triangles. Internally, ZEMAX converts rectangles into 2 triangles.

Syntax:

```
R vertex1 vertex2 vertex3 vertex4 isreflective face
```

The vertex numbers must be integer vertices previously defined in the file. The order of the vertices is not arbitrary, they must be listed in a continuously clockwise or counter clockwise order, not a "bow tie" order. Not all four vertex polygons form a closed rectangle; for more complex shapes use multiple triangle commands instead. If the rectangle is part of an object that forms a closed volume, the vertices should be ordered so that the vector cross product of the vector between vertices 1 and 2, and the vector between vertices 1 and 3, points "out" of the volume.

The "isreflective" flag is -1 if the surface absorbs, 1 if the surface reflects, or 0 if the surface refracts. Using this flag allows some rectangles to be reflective, and others to be refractive or absorptive, within the same Polygon Object or face. However, the isreflective flag is not considered in the object produced by Boolean operations with the Polygon Object.

The value for face defines which face the rectangle belongs to. If the face number is omitted; face 0 is assumed.

Example:

```
R 1 2 3 4 1 0
```

This would define a reflective rectangle connecting vertices 1, 2, 3, and 4, belonging to face 0.

Maximum triangles in the polygon object

There is no fixed upper limit on the number of triangles a polygon object may contain. The ultimate limit is determined by the amount of real or virtual RAM the computer has available. Each triangle takes up about 100 bytes of memory. However, ZEMAX often maintains multiple copies of the lens data simultaneously, and so a good rule of thumb is that ZEMAX will need 500 bytes of RAM for each triangle. A 2000 triangle object will require about 1 Mb of free RAM. A more practical limit is computer speed; ZEMAX will slow down noticeably if the number of triangles becomes very large.

Example POB file

Here is the complete text of a POB file defining a cube of refractive material. The file is included as the example CUBE.POB. All 8 sides of the cube belong to face 0.

```
! A cube
! front face vertices
V 1 -1 -1 0
V 2 1 -1 0
V 3 1 1 0
V 4 -1 1 0
! back face vertices
V 5 -1 -1 2
V 6 1 -1 2
V 7 1 1 2
V 8 -1 1 2
! Front
R 1 2 3 4 0 0
! Back
R 5 6 7 8 0 0
! Top
```

```

R 4 3 7 8 0 0
! Bottom
R 1 2 6 5 0 0
! Left side
R 1 4 8 5 0 0
! Right side
R 2 3 7 6 0 0

```

The 8 "V" commands define the vertex coordinates of the 8 corners of the cube. The 6 "R" commands define each of the 6 faces of the cube. Note the width of the faces is 2 units, and the z coordinates of the back face vertices is defined to be 2 units; so the shape is a perfect cube. All the coordinates are relative to the object reference point; which in this case will be the center of the front face. To place the reference point at the center of the cube, change the vertex definitions to:

```

V 1 -1 -1 -1
V 2 1 -1 -1
V 3 1 1 -1
V 4 -1 1 -1
V 5 -1 -1 1
V 6 1 -1 1
V 7 1 1 1
V 8 -1 1 1

```

For more examples see the POB files provided with ZEMAX in the <objects>\Polygon Objects folder (see "Folders" on page 68).

Ray Rotator

The Ray Rotator is an imaginary device that rotates all incident rays. The object is a plane surface with either an elliptical or a rectangular shape. At the point a ray intersects the rotator, a "ray" coordinate system is created. The X, Y, and Z directions of this ray coordinate system initially are parallel to the local X, Y, and Z directions of the rotator, but the ray coordinate system is centered on the ray intercept point rather than the surface vertex.

The incident ray direction cosines are then rotated around the Z axis of the ray coordinate system. The rotation direction uses the right hand rule. This Z rotation also rotates the ray coordinate Y axis by the same angle. The rays are then rotated around the *new* ray coordinate Y axis resulting from the first Z direction rotation. If the rays are polarized, the polarization vector is also rotated. Note the ray's position is never modified; only the ray direction cosines.

The rotator object will not refract, absorb, attenuate, alter optical path or phase, or otherwise modify any property of the ray other than its direction and the direction of any associated electric field data. If the material is set to MIRROR, this object will reflect rays, and the rotation is applied after the reflection. Because the phase and optical path length are not modified, wavefront phase is not preserved by this object and thus only incoherent computations may be meaningful on the beam subsequent to the ray rotation.

The defining parameters are:

1/2: The X/Y Half Widths. Make positive for a rectangular shape, negative for an elliptical shape.

3/4: The Z/Y direction rotation angles in degrees.


Rectangular Corner

Rectangular corners are defined by 1 parameter:

1: The full width of the squares.

Rectangular corners are composed of 3 squares which meet at 90 degrees. The squares are aligned in the positive XY, XZ, and YZ planes. Each square is of dimension X by X.

The reference point is the point where all 3 squares touch. Face Numbers: All faces Face 0.




Rectangle

A rectangle is a flat surface defined by 2 parameters:

- 1: The X half width.
- 2: The Y half width.

Rectangles are flat and reside entirely within the local XY plane and are placed at the local Z = 0 coordinate.

The reference point is the center of the rectangle. Face Numbers: All faces Face 0.




Rectangular Pipe

Rectangular pipes are defined by 9 parameters:

- 1: The X half width of the open front face.
- 2: The Y half width of the open front face.
- 3: The Z length of the pipe along the local Z axis.
- 4: The X half width of the open rear face.
- 5: The Y half width of the open rear face.
- 6-7: The X, Y angle tilt in degrees of the front face.
- 8-9: The X, Y angle tilt in degrees of the rear face.

Rectangular pipes are 4 sided boxes. The front and rear faces are open. This object is typically used as a rectangular light pipe. The reference point is the center of the front open face. Face Numbers: All faces Face 0.



See “Comment about tilted faces” on page 435, those comments apply to this object as well.

Rectangular Pipe Grating

This object is the same shape as the Rectangular Pipe and uses the same first 9 parameters. However, the grating version adds a linear diffraction grating on all four sides. There are two additional parameters:

- 10: The grating line frequency in lines/micrometer on the side faces.
- 11: The diffraction order. Multiple orders may be specified, see “Diffraction tab” on page 357.

The grating consists of equally spaced lines perpendicular to the local z axis, lying on each of the four faces. The grating frequency is the lines per micrometer along the z direction; projected down on to either the XZ or YZ plane. Note the grating exists on all four sides of the pipe. The reference point is the center of the front open face.

For important information on diffractive objects, see “Diffraction from NSC objects” on page 353.

Face Numbers: All faces Face 0.


Rectangular Roof

There are 3 parameters used to define a rectangular roof:

- 1: The X half width.
- 2: The Y half width.
- 3: The angle between the two rectangles.

A rectangular roof is composed of two rectangles which meet at an angle.

The reference point is the midpoint of the line joining the two rectangles. Face Numbers: All faces Face 0.



Rectangular Torus Surface

A rectangular torus is a surface formed by rotating a rectangle about a displaced axis. The rotation about the displaced axis may be over a full 360 degrees; or just some subset of that angular range. See also the discussion of the Rectangular Torus Volume for modeling refractive solid torus shapes.

The rectangular torus surface is defined by 6 parameters:

- 1: The outer radius of the torus, R_{out} .
- 2: The inner radius of the torus, R_{in} .
- 3: The start angle of the torus, θ_1 .
- 4: The stop angle of the torus, θ_2 .
- 5: The thickness of the torus, T_y .


The rectangle lies in the YZ plane with the center at $x = 0$, $y = 0$, $z = (R_{out} + R_{in})/2$. This position of the rectangle corresponds to the rotation angle $\theta = 0$. The angles of rotation are about the Y axis and must meet this condition:

$$0 \leq \theta_1 \leq \theta_2 \leq 360.$$

There is also the restriction that $R_{out} > R_{in} > 0$ and $T_y > 0$. The reference coordinate is the center of the axis of rotation. Face Numbers: All faces Face 0.

Rectangular Torus Volume


This object is essentially identical to the Rectangular Torus Surface, except the ends of the torus are closed to make a solid volume. This allows the object to be made of a refractive material. See the description of the Rectangular Torus Surface for details. Face Numbers: Start angle end cap Face 1, stop angle end cap Face 2, all other faces Face 0.



Rectangular Volume

There are 9 parameters used to define the rectangular volume:


- 1: The X half width of the front face.
- 2: The Y half width of the front face.
- 3: The Z length of the volume along the local Z axis.
- 4: The X half width of the rear face.
- 5: The Y half width of the rear face.
- 6: The front face tilt angle in degrees along X.
- 7: The front face tilt angle in degrees along Y.
- 8: The rear face tilt angle in degrees along X.
- 9: The rear face tilt angle in degrees along Y.



Rectangular volumes are usually 6 sided solids. However, shapes such as pyramids and wedges may be created by setting one or more of the parameters to zero as shown in the picture on the right.

When performing Boolean operations with rectangular volumes, or exporting rectangular volumes in CAD format, any face with a width of less than the export tolerance will be removed.

The reference point is the center of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.



Comment about tilted faces

Note that if the front and rear face X half widths are not the same, and there is a tilt along the Y direction on either the front or rear face, then the four corners on the sides of the "box" are no longer coplanar. This yields an odd shape that ZEMAX models as two triangles, with a crease along the diagonal. A similar condition occurs if the front and rear Y half widths are different and a tilt along X is defined on either face. Although no error message is issued in these cases, the resulting object model should be carefully studied to confirm the desired shape is generated.

Rectangular Volume Grating

This object is the same shape as the Rectangular Volume and uses the same first 9 parameters. However, the grating version adds a linear diffraction grating on four faces of the object: the top, bottom, left, and right faces. There is no grating on the front or back faces. There are two additional parameters:

- 10: The grating line frequency in lines/micrometer on the side faces.
- 11: The diffraction order. Multiple orders may be specified, see "Diffraction tab" on page 357.

The grating consists of equally spaced lines perpendicular to the local z axis, lying on each of the four faces. The grating frequency is the lines per micrometer along the z direction; projected down on to either the XZ or YZ plane. Note the grating exists on all four sides of the volume. The reference point is the center of the front face.

For important information on diffractive objects, see "Diffraction from NSC objects" on page 353.

Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Slide

Slides are color RGB transparencies defined by either a BMP or JPG format graphics file.

Slides use the following parameters:

- 1: The width in lens units of the image.
- 2: The pixel aspect ratio, defined as height/width. Most bitmaps use a value of 1.0 for the pixel aspect ratio. The height of the slide will be determined by number of pixels, the slide width, and the pixel aspect ratio.

This object can be used to model radiant scenes by placing a source behind the slide. For example, to make a Lambertian radiant scene, place a collimated rectangular source (see "Source Rectangle" on page 459) behind the slide, then set the slide to have a Lambertian scattering property (see "Scattering" on page 364).

Face Numbers: All faces Face 0.

Sphere


Spheres are defined by two parameters:

1: The radius of the sphere.

2: A flag to indicate if the sphere is a solid volume or a hollow shell. If the "Is Volume?" parameter is zero, then ZEMAX assumes the sphere defines a hollow shell. Rays must either reflect or absorb from such a surface; refraction is not allowed. If the "Is Volume?" parameter is any non-zero value, then ZEMAX assumes the sphere is a closed volume. The volume may be reflecting, refracting, or absorbing.

This object can be used to model bubbles by placing a sphere within a glass volume and setting the material type to blank (air) or the name of a defined glass which describes a gas.

The reference point is the center of the sphere. Face Numbers: All faces Face 0.



Standard Lens

The standard lens is a lens composed of standard ZEMAX surfaces. Standard surfaces may be planes, spheres, conic aspheres, or hyperhemispheres. The standard lens is composed of 5 separate sections:


1) A standard shape front face.

2) A standard shape rear face.

3) An annular ring between the clear aperture of the front face and the edge of the front face.

4) An annular ring between the clear aperture of the rear face and the edge of the rear face.

5) A possibly tapered cylindrical surface connecting the edges of the front and rear faces of the lens.



9 parameters are used to define a standard lens:

1: The radius of curvature of the front face. Use zero for infinity (flat).

2: The conic constant of the front face.

3: The semi-diameter to the clear aperture of the front face. Use a negative value to yield the hyperhemispheric sag point.

4: The radial aperture to the edge of the front side of the lens.

5: The center-to-center thickness of the lens.

6: The radius of curvature of the rear face. Use zero for infinity (flat).

7: The conic constant of the rear face.

8: The semi-diameter to the clear aperture of the rear face. Use a negative value to yield the hyperhemispheric sag point.

9: The radial aperture to the edge of the rear side of the lens.

All 5 surfaces may refract, reflect, or absorb light, depending upon the material properties.


The reference point is the center of the front face of the lens. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Standard Surface

The standard surface object is very much like a standard surface in sequential ZEMAX. Standard surfaces include planes, spheres, conic aspheres, and hyperhemispheric spheres and aspheres.

The standard surface requires 4 parameters:

- 1: The radius of curvature. Use zero for infinity (flat).
- 2: The conic constant.
- 3: The semi-diameter to the maximum clear aperture. Use a negative value to yield the hyperhemispheric sag point.
- 4: The semi-diameter to the minimum clear aperture. This value must be positive. If greater than zero, this value creates a "hole" in the surface.



The surface is rotationally symmetric about the local z axis.

If the radius is negative, then the surface is concave towards the negative z axis. If the radius is positive, the surface is concave towards the positive z axis.

If the semi-diameter is negative, then the surface will become hyper-hemispheric, with the radial aperture of the open end equal to the absolute value of the semi-diameter.

Surfaces may reflect or absorb rays.

The reference point is the center of the surface. Face Numbers: All faces Face 0. This object supports user defined apertures, see "User defined apertures" on page 364. For a similar surface shape with more complex aperture shapes supported, see "Aspheric Surface 2" on page 402.

STL Object

The STL object is a very general user-defined object. It can be used to define an open polygon surface or a closed polygon volume such as a prism or other solid. The STL Object format is based upon a collection of 3D triangles. This format is widely supported by mechanical CAD programs. Both the text and the binary variations of the STL file format are supported. See the section "Defining STL objects" which follows for details. See also "User Defined Object" on page 444, "Polygon Object" on page 429, and "Imported" on page 423.

There are no fixed limits to the total number of vertices or polygons.


The STL file name is referenced in the "comments" column of the STL Object row. For example, if the STL file myobject.STL is placed in the <objects>\CAD Files folder (see "Folders" on page 68), then specify "myobject.STL" in the comment column of the STL Object type row in the NSC Editor.

The reference coordinate is locally (0, 0, 0), and the polygons that compose the object may be placed anywhere relative to the reference point. Some STL export implementations only allow objects to be placed such that all vertex coordinates are positive. ZEMAX does not require this to be the case, and will import triangle vertices anywhere in 3D space.

Using an STL object as a detector is supported, see "Objects as detectors" on page 470. When the STL object is used as a detector, If "Fast Ray Trace" (see "Type tab" on page 353) is checked, ZEMAX will alter the order of the triangles defined in the STL file to speed up the ray trace. If "Fast Ray Trace" is unchecked, the order of the triangles is unmodified, which slows the ray tracing but preserves the original order of the triangles defined in the STL file.

STL objects require two parameters:

- 1: A scale factor. All vertices in the STL file are multiplied by this parameter.



2: A flag to indicate if the STL file defines a volume or a surface. If the "Is Volume?" parameter is zero, then ZEMAX assumes the STL file defines an open surface. Rays must either reflect or absorb from such a surface; refraction is not allowed. If the "Is Volume?" parameter is any non-zero value, then ZEMAX assumes the STL file defines a closed volume. The volume may be reflecting, refracting, or absorbing.

Face Numbers: All faces Face 0.

Defining STL Objects

The STL format is commonly supported by mechanical CAD programs to describe arbitrary objects. Objects are modeled as a collection of triangles, and the global vertices of the triangles are written out to a file. STL is a good format for faceted objects. For smoothly curved objects, such as lenses, STL is an approximation which may be of acceptable accuracy for some non-imaging systems.

There are binary and text formats for STL files; and ZEMAX supports both formats.

To use an STL object, choose the object type to be an "STL Object" and enter the name of the file (without the STL extension) in the comments column of the STL Object row. The STL file must be placed in the folder <objects>\CAD Files (see "Folders" on page 68).

Maximum triangles in STL objects

There is no fixed upper limit on the number of triangles an STL object may contain. The ultimate limit is determined by the amount of real or virtual RAM the computer has available. Each triangle takes up about 100 bytes of memory. However, ZEMAX often maintains multiple copies of the lens data simultaneously, and so a good rule of thumb is that ZEMAX will need 500 bytes of RAM for each triangle. A 2000 triangle object will require about 1 Mb of free RAM. A more practical limit is computer speed; ZEMAX will slow down noticeably if the number of triangles becomes very large.

Example STL files

Several example STL files may be found in the <objects>\CAD Files folder (see "Folders" on page 68).

Swept Object

A swept object is created by sweeping a cross section of a parent object around an arbitrary axis to form a new shape. The parent object may be any surface or solid volume. A cross section of the parent object is formed by the intersection of the parent and a "slicing" plane. The slicing plane is defined by the axis of rotation, and a vector that points away from this axis of rotation. The resulting cross section may be swept in angle from -360 to +360 degrees, with the restriction that the absolute value of the angle must be at least 0.1 degrees. Swept objects are defined by the following parameters:

1: The parent object number. This number must precede the swept object number in the editor.


2,3,4: The center x, y, and z coordinates in the unrotated, unshifted coordinate system of the parent object. These data define a point that lies on the axis of rotation.

5,6,7: The axis x, y, and z direction cosines in the coordinate system of the parent object. These cosines are normalized to determine the vector pointing along the axis of rotation. The axis of rotation must not intersect the object, or an invalid solid will be generated. Generally this condition cannot be automatically detected and no warning will be issued.

8,9,10: The plane x, y, and z direction cosines in the coordinate system of the parent object. These cosines are normalized to determine a vector pointing away from the axis of rotation. The slicing plane is defined as the plane that contains both the axis of rotation and the plane vector. Note the plane vector is not orthogonal to the slicing plane, but lies within the plane. The plane vector must not be parallel to the axis vector.

11: The sweep angle in degrees. The sweep begins at the plane.

12: The end caps option. Use 0 for no end caps, 1 for first end cap only, 2 for second end cap only, and 3 for both end caps. If the sweep angle is 360 degrees no end caps are formed, regardless of this setting. If the sweep



angle is less than 360 degrees, and the "Is Volume" flag is set to a non-zero value, then the end caps setting is ignored, and both end caps are automatically added.

13: Chord Tolerance: This setting only affects the rendering of the object. To render the object, ZEMAX generates a list of triangles which approximate the object shape. The tolerance is the maximum allowed distance in lens units between a single triangle and the actual surface of the object. More triangles are added if the tolerance is set smaller which yields more accurate rendering, at the expense of speed and a larger memory requirement. The default value of zero will use a chord tolerance related to the size of the object sufficient to generate a coarse approximation of the object shape that will render quickly.

14: Face Mode: The Face Mode setting determines how ZEMAX assigns face numbers to the various surfaces on the swept object. The following Face Modes are supported:

Face Mode = 0: All surfaces are assigned face number 0. The entire object will have just one face.

Face Mode = 1: All surfaces whose edges meet along a non-zero length curve, and whose normal vectors along the curve of contact are parallel within a user defined angle tolerance are assigned a common face number. The angle tolerance is defined by the Face Angle (parameter 15). This mode allows control over how finely the faces are numbered. If the Face Angle is set to a large value (such as 180) then all faces that touch will share a common number. Larger Face Angles yield fewer unique faces.

Face Mode = 2: All surfaces are uniquely numbered. This mode yields the largest number of unique faces.

In all cases, the number of unique face numbers cannot exceed the maximum number of faces ZEMAX supports. Any faces which exceed the maximum limit will be renumbered to be face 0.

15: Face Angle: The angle in degrees used by Face Mode = 1 described above.

16: Is Volume?: If the resulting shape is an enclosed volume, set this parameter to 1, otherwise, the object is a shell and this parameter should be set to zero.

Tabulated Faceted Radial

A tabulated object is based upon coordinates defined in a file. The file must reside in the <objects>\Tabulated Objects folder (see "Folders" on page 68) and end in the extension TOB. The coordinates represent the starting and ending points of facets. A figure of revolution is generated by replicating a specified number of facets over some angular range. The axis of revolution is the local Z axis.



The TOB file format is two columns of data separated by either one or more spaces or one or more tab characters. A sample TOB file might look like this:

```
1.5 3.5
2.2 4.5
3.0 5.5
3.0 6.0
```

The first number of each pair is the local Y coordinate. This value must be zero or positive. The second value is the local Z coordinate. Each pair of values after the first pair represent a "zone". If there are 6 pairs of numbers, then there are 5 zones to the object. There is a maximum of 246 zones per object. Multiple objects may be used if more zones are required.

ZEMAX will generate facets which approximate a smooth surface for each zone. The facets can cover any fraction of a full circle, defined by a start and stop angle. To make a full figure of revolution, the start angle should be set to 0.0 degrees, and the stop angle to 360.0 degrees. Both angles must be zero or positive and be less than or equal to 360.0 degrees.

The number of facets generated over this angular range can be specified independently for each zone; so there may be 40 facets in the first zone, 80 in the second, 50 in the



third, etc. If the number of facets changes from zone to zone, small gaps in the part may form and these gaps will generate geometry errors. If the gaps are small, the rays which trigger these errors may be ignored; see "The Ray Trace Control" on page 377 for more information.

The parameters used to define the object are:

1: A scale factor. All vertices in the TOB file are multiplied by this parameter.

2: A flag to indicate if the TOB file defines a volume or a surface. If the "Is Volume?" parameter is zero, then ZEMAX assumes the TOB file defines an open surface. Rays must either reflect or absorb from such a surface; refraction is not allowed. If the "Is Volume?" parameter is any non-zero value, then ZEMAX assumes the TOB file defines a closed volume. The volume may be reflecting, refracting, or absorbing.

3: Start angle. The angle in degrees to begin the revolution of the TOB defined coordinates.

4: Stop angle. The angle in degrees to end the revolution of the TOB defined coordinates.

5: Zone 1 facets. The number of facets between the start and stop angles for the first zone.

4+n: Zone n facets. The number of facets between the start and stop angles for the nth zone.

If the "Is Volume" flag is set, the TOB file must define an object that will be a closed volume upon rotation. This requires that the object be rotated fully around 360.0 degrees. Fully closed volumes defined by TOB files may be used to model faceted approximations to true Fresnel lenses.

Face Numbers: All faces Face 0.

Tabulated Faceted Toroid

A tabulated object is based upon coordinates defined in a file. The file must reside in the <objects>\Tabulated Objects folder (see "Folders" on page 68) and end in the extension TOB. The coordinates represent the starting and ending points of facets. A figure of rotation is generated by replicating a specified number of facets over some angular range. The axis of rotation is an axis parallel to the local Y axis offset by a specified radius. If the radius is set to zero, then a surface with cylindrical symmetry is generated instead of a torus.

The TOB file format is two columns of data separated by either one or more spaces or one or more tab characters. A sample TOB file might look like this:



```
1.5 3.5
2.2 4.5
3.0 5.5
3.0 6.0
```

The first number of each pair is the local Y coordinate. This value may be negative, zero, or positive. The second value is the local Z coordinate. Each pair of values after the first pair represent a "zone". If there are 6 pairs of numbers, then there are 5 zones to the object. There is a maximum of 246 zones per object. Multiple objects may be used if more zones are required.

ZEMAX will generate facets which approximate a smooth surface for each zone. The facets can cover any fraction of a full circle, defined by a start and stop angle. To make a full figure of revolution, the start angle should be set to -180.0 degrees, and the stop angle to 180.0 degrees. Both angles must have an absolute value less than or equal to 180.0 degrees.

The number of facets generated over this angular range can be specified independently for each zone; so there may be 40 facets in the first zone, 80 in the second, 50 in the third, etc.

The parameters used to define the object are:



1: A scale factor. All vertices in the TOB file are multiplied by this parameter.

2: Radius of rotation. If positive, then the axis of rotation is in the positive local Z direction, parallel to the local Y axis, in the YZ plane. If negative, then the axis of rotation is in the negative local Z direction, parallel to the local Y axis, in the YZ plane. If zero, then a cylinder results. In the special case of zero radius, the number of facets parameters are ignored (since a single facet perfectly models a flat plane) and the start and stop angle are interpreted as start and stop x coordinate in lens units.

3: Start angle. The angle in degrees to begin the rotation of the TOB defined coordinates, unless the radius of rotation is zero; in this case the start angle defines the starting x coordinate in lens units.

4: Stop angle. The angle in degrees to end the rotation of the TOB defined coordinates, unless the radius of rotation is zero; in this case the stop angle defines the stopping x coordinate in lens units.

5: Zone 1 facets. The number of facets between the start and stop angles for the first zone.

4+n: Zone n facets. The number of facets between the start and stop angles for the nth zone.

Face Numbers: All faces Face 0.

Tabulated Fresnel Radial

This object is nearly identical to the tabulated faceted radial object. The key difference is that the radially symmetric surfaces are smooth, rather than faceted. The maximum number of points is also much larger at 50,000. See “Tabulated Faceted Radial” on page 439 for a description of this type of object and the tabulated object file format. Face Numbers: All faces Face 0.

Toroidal Hologram

A Toroidal Hologram consists of a rectangular, circular, or elliptical lens with possibly aspheric toroidal surfaces on the front and back faces and an optically fabricated hologram on the front face. A toroidal surface is defined by a curve in the YZ plane which is then rotated about an axis parallel to the Y axis but displaced by a distance R; the radius of rotation. The curve in the YZ plane is defined by:

$$z = \frac{cy^2}{1 + \sqrt{1 - (1+k)c^2y^2}} + \alpha_1y^2 + \alpha_2y^4 + \alpha_3y^6 + \alpha_4y^8 + \alpha_5y^{10} + \alpha_6y^{12},$$

where c is the reciprocal of the radius of curvature in YZ plane. The hologram on the front face is defined in a manner similar to the Hologram 1 and Hologram 2 sequential surface models (see “Hologram 1” on page 316 for a description of these surfaces). The Toroidal Hologram is defined by 32 parameters:

1: The Radial Height of the lens in the y direction.

2: The X Half-Width. If this parameter is zero, then the outer boundary of the lens will be a circle with a radial size equal to the Radial Height. If this parameter is positive, the outer boundary of the lens will be rectangular. If this parameter is negative, the outer boundary of the lens will be elliptical.

3: The Thickness of the lens along the local Z axis.

4-5: Unused.

6, 7, 8: The radius of rotation, radius of curvature, and conic for the front face.


9-14: The coefficients on the powers of y for the front face.

15, 16, 17: The radius of rotation, radius of curvature, and conic for the back face.

18-23: The coefficients on the powers of y for the back face.

24: Holo Type: Use 1 for Hologram type 1 (both sources converging/diverging) or 2 for Hologram type 2 (one source converging, one diverging) See “Hologram 1” on page 316.

25: Order: The diffraction order to use. Multiple orders may be specified, see “Diffraction tab” on page 357.



26: Construction Wavelength: The wavelength in micrometers used to fabricate the hologram.

27-32: The X, Y, and Z coordinates in lens units of the construction points relative to the vertex of the front face of the Hologram.

To make any of the 4 radii flat; use a value of zero. Note a cylinder lens results if the radius of rotation is set to zero. If the order is zero or if the ray totally internally reflects at the hologram boundary, no hologram diffraction is computed.

For important information on diffractive objects, see “Diffraction from NSC objects” on page 353.

The reference coordinate is the center of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Toroidal Lens

A toroidal lens consists of a rectangular, circular, or elliptical lens with possibly aspheric toroidal surfaces on the front and back faces. A toroidal surface is defined by a curve in the YZ plane which is then rotated about an axis parallel to the Y axis but displaced by a distance R; the radius of rotation. The curve in the YZ plane is defined by:

$$z = \frac{cy^2}{1 + \sqrt{1 - (1+k)c^2y^2}} + \alpha_1y^2 + \alpha_2y^4 + \alpha_3y^6 + \alpha_4y^8 + \alpha_5y^{10} + \alpha_6y^{12},$$

where c is the reciprocal of the radius of curvature in YZ plane. The toroidal lens is defined by 23 parameters:

1: The Radial Height of the lens in the y direction.

2: The X Half-Width. If this parameter is zero, then the outer boundary of the lens will be a circle with a radial size equal to the Radial Height. If this parameter is positive, the outer boundary of the lens will be rectangular. If this parameter is negative, the outer boundary of the lens will be elliptical.

3: The Thickness of the lens along the local Z axis.

4-5: Unused.

6, 7, 8: The radius of rotation, radius of curvature, and conic for the front face.

9-14: The coefficients on the powers of y for the front face.

15, 16, 17: The radius of rotation, radius of curvature, and conic for the back face.


18-23: The coefficients on the powers of y for the back face.

To make any of the 4 radii flat; use a value of zero. Note a cylinder lens results if the radius of rotation is set to zero. The reference coordinate is the center of the front face. Face Numbers: Front face Face 1, back face Face 2, all other faces Face 0.

Toroidal Surface

A toroidal surface consists of a rectangular surface with a possibly aspheric toroidal shape. A toroidal surface is defined by a curve in the YZ plane which is then rotated about an axis parallel to the Y axis but displaced by a distance R; the radius of rotation. The curve in the YZ plane is defined by:


$$z = \frac{cy^2}{1 + \sqrt{1 - (1+k)c^2y^2}} + \alpha_1y^2 + \alpha_2y^4 + \alpha_3y^6 + \alpha_4y^8 + \alpha_5y^{10} + \alpha_6y^{12},$$



where c is the reciprocal of the radius of curvature in YZ plane. The toroidal lens is defined by 23 parameters:

- 1: The X Half-Width in lens units.
- 2: The Y Half-Width in lens units.
- 3-4: Unused.
- 5, 6, 7: The radius of rotation, radius of curvature, and conic for the surface.
- 8-13: The coefficients on the powers of y for the surface.

To make either the radius of rotation or radius of curvature flat; use a value of zero. Note a cylindrical surface results if the radius of rotation is set to zero. The reference coordinate is the center of the front face. Face Numbers: All faces Face 0.




Torus Surface

A torus is a circle rotated about a displaced axis. The rotation about the displaced axis may be over a full 360 degrees; or just some subset of that angular range. This object can be used to model optical fibers or curved light pipes. See also the discussion of the Torus Volume for modeling refractive solid torus shapes.

The torus surface is defined by 6 parameters:

- 1: The radius of rotation about the Y axis of the circle, R .
- 2: The radius of the circle, r .
- 3: The start angle of the torus, θ_1 .
- 4: The stop angle of the torus, θ_2 .



The circle lies in the YZ plane with the center of the circle at $x = 0$, $y = 0$, and $z = R$. This position of the circle corresponds to the rotation angle $\theta = 0$. The angles of rotation are about the Y axis and must meet this condition:

$$0 \leq \theta_1 \leq \theta_2 \leq 360.$$

There is also the restriction that $R > r$; otherwise, a closed volume or a smooth surface will not result.

The reference coordinate is the center of the axis of rotation. Face Numbers: All faces Face 0.

Torus Volume

This object is essentially identical to the Torus Surface, except the ends of the torus are closed to make a solid volume. This allows the object to be made of a refractive material. See the description of the Torus Surface for details. Face Numbers: Start angle end cap Face 1, stop angle end cap Face 2, all other faces Face 0.


Triangular Corner

Triangular corners are defined by 1 parameter:

1: The full width of the short side of the triangle.

Triangular corners are defined by 3 triangles which meet at 90 degrees. The triangles are aligned in the positive XY, XZ, and YZ planes. Each surface is a 45-45-90 triangle, where the two short sides are of a length specified by parameter 1.

The reference point is the point where all 3 triangles touch. Face Numbers: All faces Face 0.




Triangle

Triangles are defined by 3 points in the XY plane, which is a total of 6 parameters:

- 1: The X coordinate of vertex 1.
- 2: The Y coordinate of vertex 1.
- 3: The X coordinate of vertex 2.
- 4: The Y coordinate of vertex 2.
- 5: The X coordinate of vertex 3.
- 6: The Y coordinate of vertex 3.

The reference point is the (0, 0, 0) coordinate, which may not be part of the triangle at all; depending upon the values give for the vertex points. Face Numbers: All faces Face 0.



User Defined Object

There are several different ways to create user defined objects. For simple polygon based objects, such as prisms and other shapes with all flat faces, see “Polygon Object” on page 429. Objects may also be defined in an external CAD program and then imported into ZEMAX, see “Imported” on page 423 and “STL Object” on page 437. There are also methods for user definitions for gradient index media (“Defining GRIN media for non-sequential ray tracing” on page 373), surface scattering properties (“Scattering” on page 364), bulk scattering properties (“Bulk scattering” on page 370), and diffraction (“Defining DLLs for ray splitting at diffractive surfaces” on page 375).

The method of creating a user defined object described in this section involves the use of an external, user supplied program called a Dynamic Link Library or DLL. The advantages to defining an object using a DLL, rather than the other methods listed above, are:

- DLL defined objects generally ray trace much faster, and with much higher numerical precision, than objects imported from CAD programs.
- Any number of complex curved shapes may be combined in a single object, unlike the polygon object which only has flat faces.
- Objects may have a mixture of reflective and refractive curved faces, with user definable face names.
- The DLL description is inherently parametric, which means the object is dynamically regenerated when any defining property is modified. This allows interactive design, modification, and even optimization.
- User defined coating data, including detailed control over the complex amplitude reflection and transmission coefficients is supported. Coating data may be ray position, cosine, or object coordinate dependent.

The entire definition for the object is contained within the DLL. The DLL computes the following data:

- The computation of basic data about the object, such as whether it is a solid volume or a hollow shell, and

the names of all the used parameters.

-The determination of a list of triangles which approximate the shape of the object. These triangles are used to render the object as well as provide a "close first guess" for ray tracing purposes.

-The exact computation of the ray-surface intercept. Most of this code is actually inside ZEMAX; the DLL only needs to find the exact ray-surface intercept point given a very close guess. ZEMAX handles all the logic for non-sequential tracing, nesting, refraction, reflection, diffraction, optical path, scattering, etc.

-The user defined coating data used by the object, if any. The DLL or ZEMAX can compute coating data.

-The starting "safe" data to use for the various free parameters when the object is first created.

The DLL must include two functions:

UserObjectDefinition

UserParamNames

The UserObjectDefinition function is called to create a list of triangles which approximate the shape of the object. ZEMAX traces rays non-sequentially to the triangles to determine close intercepts, if any. Once a potential close intercept is found, the UserObjectDefinition function is called again to determine the exact intercept. This is implemented using an iterative routine provided in any of the sample object DLL source files. A single object may contain multiple complex curved faces, each with it's own defining function and iteration routines.

The best way to learn the use of Object DLLs is to study an existing DLL and modify it as needed. The sample DLLs provided with ZEMAX include extensive documentation and comments on the data formats and DLL functions; see any of the sample object code files for examples.



All Object DLLs must be placed in the <program>\DLL\Objects folder.

Face Numbers: All faces have user definable Face numbers.

If you require a User Defined Object DLL, but do not wish to write the DLL yourself, contact ZEMAX Technical Support for a quote on developing a custom DLL to meet your requirements. We have considerable experience in developing ray tracing algorithms, and can generally write user defined object DLL code for a small fee, often on very short notice.

Sample Object DLLs

The following sample object DLLs and source code are included with ZEMAX.

SAMPLE OBJECT DLLS

DLL name	Description
CoatingSample	A rectangular surface, with user definable dimensions in x, y. This sample shows how to implement user defined coating data. The transmission of this rectangle is 1.0 outside of a circle of radius "a", and a user defined value inside of the circle.
EllipticalVolume	An elliptical solid volume, with user definable dimensions in x, y, and z. This DLL demonstrates the creation of a list of triangles to approximate the object shape, naming object faces, and exact ray-surface iteration to arbitrary surface description functions.
HalfCylinder	A half cylinder, with user definable radius and length. This DLL demonstrates the creation of a list of triangles to approximate the object shape, naming object faces, and exact ray-surface iteration to arbitrary surface description functions.

Wolter Surface

The Wolter Surface is a radial symmetric shell object defined by the following parameters:

- 1: The length along the Z axis.
- 2: The radial aperture at $z = 0$, R_0 .
- 3: The angle θ in degrees.
- 4-7: The polynomial coefficients on Z^2 , Z^3 , Z^4 , and Z^5 .

The object starts at $Z = 0$ and ends at the specified length.

A Wolter Surface is commonly used in X-Ray telescope optics. The surface is a hollow shell and may be nested within other shells. The shell has a radius that is a function of the z coordinate given by the following expression:

$$R^2(z) = R_0^2[1 + 2z \tan \theta + a_2 z^2 + a_3 z^3 + a_4 z^4 + a_5 z^5].$$

Face Numbers: All faces Face 0.

Zernike Surface

A Zernike surface is defined by the following sag equation:

$$z = \frac{cr^2}{1 + \sqrt{1 - (1+k)c^2r^2}} + \sum_{i=1}^8 \alpha_i r^{2i} + \sum_{i=1}^N A_i Z_i(\rho, \varphi),$$

where c is the curvature of the surface, k is the conic constant, the α terms are aspheric coefficients, N is the number of Zernike coefficients in the series, A_i is the coefficient on the i^{th} Zernike Standard polynomial, r is the radial coordinate in lens units, ρ is the (optionally decentered) normalized radial coordinate, and φ is the (optionally decentered) angular coordinate. The Zernike Standard polynomials are defined in the table given in "Zernike Standard Coefficients" on page 199. The surface supports specification of both a minimum and a maximum radial aperture; so annular surfaces may be defined.

The following parameters are used to define the Zernike surface:



- 1: The radius of curvature. If this value is zero, then the curvature is assumed to be zero.
- 2: The conic constant k .
- 3: The maximum radial aperture in lens units. This aperture is also used to normalize the radial coordinates used in the Zernike polynomial expansion, unless the minimum radial aperture is negative, see below.

4: The minimum radial aperture in lens units. This value may be zero. If this value is negative, then the minimum radial aperture is set to zero, and the absolute value of the parameter is used as the normalization radius. This feature allows a normalization radius distinct from the maximum radial aperture.

5-6: The decenter of the Zernike terms with respect to the base asphere in lens units.

7-14: The even radial aspheric coefficients α .

15: The number of terms to use in the Zernike expansion.



16-247: The Zernike coefficients.

Face Numbers: All faces Face 0. This object supports user defined apertures, see “User defined apertures” on page 364.

Sources

Sources include points, ellipses, rectangles, volumes, data files, and user defined types. Any source may be placed inside of any object, or not in any object, but not both (a source may not straddle an object boundary).

Placing sources inside objects

By default, ZEMAX assumes a source is placed in the global medium of the non-sequential group. However, sources may be placed entirely inside any solid object, or inside any number of nested solids. The object the source is placed within must precede the source in the NSC Editor. Two steps are required to properly place a source inside of a solid object:

1) The source location and size must be correctly set such that all rays initiate from inside the object. This means the source cannot extend outside of the object in which it is placed.

2) The object in which the source is placed must be defined using the "Inside Of" data column of the source.

There are two ways to define the object in which the source is placed: either by absolute or relative reference. If a positive number greater than zero is specified in the Inside Of column, then the source is inside the specified object. This is an "absolute" reference. If the "Inside Of" object number is negative, then the source is inside of the object whose number is determined by adding the current object number to the negative "Inside Of" object number. This is a "relative" reference object. For example, if the "Inside Of" value is -3 on source object 8, the source must be placed inside of object 5 because $8 - 3 = 5$. Relative references are particularly useful when copying and pasting groups of sources and objects; this is easiest if all the sources and objects in the group use relative references to the first object in the group. If the source is placed inside of an object which is inside of a second object; then the first object must define the second object number in the "Inside Of" data column for the first object. If there are more levels of nested objects, each object must define the object it has been placed inside of on the "Inside Of" data column. All of these nested definitions may be absolute, relative, or a mixture of both.

The source must be listed after all objects that the source is placed inside of in the NSCE.

Adding new source types

If a source is required that is not listed, please contact technical support to suggest the new type be added to ZEMAX. Some simple sources, such as Lambertian planes, may be simulated by irradiating a Lambertian scattering surface with one of the other source types. Many objects may also act as detectors, for example any polygon object may be a detector. See the discussion for information on the Source DLL and Source File for user defined sources.

SUMMARY OF NSC SOURCES

Object Name	Description	Page
Source Diode	An array of diodes with separate X/Y distributions	450
Source DLL	A source defined by an external user supplied program.	452
Source Ellipse	An elliptical surface that emits light from a virtual source point.	453
Source EULUMDAT File	A source defined by lamp data in an EULUMDAT format file.	454
Source Filament	A source in the shape of a helical filament.	454
Source File	A user defined source whose rays are listed in a file.	454
Source Gaussian	A source with a Gaussian distribution.	457
Source IESNA File	A source defined by lamp data in an IESNA format file.	457
Source Object	A source defined by the shape of another object.	458

Object Name	Description	Page
Source Point	A point source that radiates into a cone. The cone may be of zero width or be extended up to a full sphere if desired.	459
Source Radial	A radial symmetric source based upon a spline fit of arbitrary intensity vs. angle data.	459
Source Ray	A point source aligned with direction cosines.	459
Source Rectangle	A rectangular surface that emits light from a virtual source point.	459
Source Tube	A source in the shape of a cylindrical tube.	459
Source Two Angle	A rectangular or elliptical surface that emits light into a cone with distinct angles in the X and Y directions.	459
Source Volume Cylinder	A volume source in the shape of a cylinder with an elliptical cross section.	460
Source Volume Ellipse	A source in the shape of an elliptical volume.	460
Source Volume Rectangle	A volume source in the shape of a rectangle.	460

The individual sources are described in the following sections.

Parameters common to all source objects

All source type objects have the same definition for parameters 1 through 5. These parameters are:

- 1: # Layout Rays: Defines how many random rays to launch from the source when creating layout plots.
- 2: # Analysis Rays: Defines how many random rays to launch from the source when performing analysis.
- 3: Power (units): Power is the total power over the defined range of the source. The power units are specified by the system source units. See "Source Units" on page 105 for details.
- 4: Wavenumber: The wavenumber to use when tracing random rays. Zero means polychromatic; which chooses ray wavelengths randomly with the weighting defined on the wavelength data editor.
- 5: Color #: The pen color to use when drawing rays from this source. If zero, the default color will be chosen.

The other parameters have source type specific meanings as described in the follow sections.

Source Diode

The source diode model can be used to define one diode, a 1D array of diodes, or a 2D array of diodes. Each diode has an intensity distribution given by:

$$I(\theta_x, \theta_y) = I_o e^{-2\left(\left(\frac{\theta_x}{\alpha_x}\right)^{2G_x} + \left(\frac{\theta_y}{\alpha_y}\right)^{2G_y}\right)},$$

where α_x is the XZ divergence angle in degrees, G_x is the "supergaussian" factor for the X direction, with similar definitions for the Y subscripted values. Note if G_x is 1.0, then a typical Gaussian distribution results. Both G_x and G_y must be greater than or equal to 0.01 and less than or equal to 50.0.

Most laser diode manufacturers specify the far field divergence angles as the full width of the distribution between the half power points, θ_{fwhm} . For a true Gaussian distribution ($G_x = 1$), setting the left hand side of the equation to $\frac{1}{2}I_o$, setting θ_y to zero, substituting for θ_x the value of $\frac{1}{2}\theta_{fwhm}$, then solving for α_x gives

$$\alpha_x = \frac{\theta_{fwhm}}{\sqrt{2 \ln(2)}}, \text{ or}$$

$$\alpha_x = (0.8493218)\theta_{fwhm}.$$

For example, a diode with a θ_{fwhm} in the x direction of 10 degrees, the value for α_x would be 8.493218 degrees. A similar conversion applies in the y direction.

An astigmatism term may also be defined. This value must be positive, and represents the distance along the local Z axis from which the XZ distribution is measured. At the local XY plane at Z = 0, the resulting ray pattern is a line oriented along the local X axis.

If more than 1 diode is desired, the number x, number y, delta x, and delta y spacings may be defined. The diodes will be symmetrically placed in the x and y directions about the local coordinate origin.

If the astigmatism term is zero, the rays may emanate from a point, or from a line, or from a rectangular region. The spatial distribution of the ray source points is given by:

$$I(x, y) = I_o e^{-2\left(\left(\frac{x}{s_x}\right)^{2H_x} + \left(\frac{y}{s_y}\right)^{2H_y}\right)}, \text{ and}$$

$$-W_x < x < W_x,$$

where s_x is measured in lens units, H_x is the "supergaussian" factor for the X direction, with similar definitions for the Y subscripted values. Note if H_x is 1.0, then a typical Gaussian distribution results. Both H_x and H_y must be greater than or equal to 0.01. The spatial distribution terms are all ignored if the astigmatism term is not zero.

The parameters used to define this source are:

1-5: See "Parameters common to all source objects" on page 450.

6: Astigmatism: The distance to offset the XZ distribution in lens units.

7: X divergence in degrees (α_x).

8: X supergaussian factor (G_x).

9: Y divergence in degrees (α_y).

10: Y supergaussian factor (G_y).

11-12: Number of diodes in X/Y.

13-14: Spacing of diodes in X/Y in lens units.

15: The X half width of the rectangular region from which the rays emanate in lens units (W_x).

16: The Gaussian width of the spatial distribution in the X direction (s_x).

17: The X direction spatial distribution super Gaussian factor (H_x).

18: The Y half width of the rectangular region from which the rays emanate in lens units (W_y).

19: The Gaussian width of the spatial distribution in the Y direction (s_y).

20: The Y direction spatial distribution super Gaussian factor (H_y).

Source DLL

Although many types of built-in sources are included with ZEMAX, there are times when the best way to model a source is to define an algorithm to generate rays with the desired properties.

ZEMAX also supports user-defined sources as tables of rays; see the Source File discussion on page 454.

To define a source using a program, the algorithm which generates random rays must be written and compiled into a Windows Dynamic Link Library, or DLL. Numerous DLLs are provided with ZEMAX with source code. New DLLs may be easily created with a suitable compiler. See also "Comments about DLLs" on page 471.

Source DLL parameters

Each DLL may use between zero and 30 user defined data values as parameters in the computation of the source properties. These values are defined by the DLL and are only used by the DLL.

Creating a new Source DLL

The DLL must include two functions:

UserSourceDefinition

UserParamNames

When launching a ray from a source modeled by a DLL, ZEMAX passes to the UserSourceDefinition function the source parameters, the wavelength, and other data. The UserSourceDefinition function then is required to compute the following values:

x, y, z: the starting coordinates for the ray

l, m, n: the starting direction cosines for the ray

i: the initial relative intensity of the ray (see "Comments on Source DLL relative intensity" below)

These values are returned to ZEMAX and are used to initiate the ray trace. The function UserParamNames is used to define the names of all used parameters. These names appear in the parameter columns of the source DLL object in the NSC Editor.

The best way to learn the use of Source DLLs is to study an existing DLL and modify it as needed. The sample DLLs provided with ZEMAX include extensive documentation and comments on the data format; see any of the sample source code files for examples.



All Source DLLs must be placed in the <program>\DLL\Sources folder.

Comments on Source DLL relative intensity

When ZEMAX traces a ray generated by a DLL, the initial power of the ray is given by the source total power, divided by the number of analysis rays, scaled by the relative weight. For this reason, the relative intensity of the ray as computed by the DLL must be set so that the average of all the relative intensities of all the analysis rays requested is 1.0. Note the number of analysis rays is passed as a parameter to the DLL. This condition is required because ZEMAX cannot normalize the entire collection of analysis rays to be traced without generating all the rays before tracing any.

The preferred way to normalize the rays is for the DLL to set the relative intensity of all rays to 1.0, and then choose rays according to the desired probability distribution of the source. This method implies that more rays will be selected with positions and direction cosines that correspond to more energy. The method for doing this is straightforward, and involves integrating the source over all spatial and direction space and using the normalized integral and a uniform random number generator to determine rays. The provided sample FIBER1 DLL illustrates the method.

An alternate way is to use non unity ray weights. This method requires that the DLL compute all the desired rays first, then internally normalize the rays to yield an average weight of 1.0. This method is conceptually simpler but is usually less efficient.

Sample Source DLLs

The following sample source DLLs and source code are included with ZEMAX.

SAMPLE SOURCE DLLS

DLL name	Description
FIBER1	<p>A planar fiber model. The rays emanate from a disk of radius R. The intensity profile is given by</p> $I(r) = A + Br^2 + Cr^4,$ <p>where</p> $0.0 \leq r \leq R.$ <p>Once the ray starting position is determined, the ray emanates from the surface into a cone whose numerical aperture depends upon the radial coordinate:</p> $NA(r) = D + Er^2 + Fr^4.$ <p>Within the angular cone defined by the numerical aperture, the ray distribution is uniform. The parameters R, A, B, C, D, E, and F are all defined on the NSC Editor.</p>

Source Ellipse

The Source Ellipse is a flat elliptical (and optionally annular) surface which emits rays. Although the origin of each ray launched lies on the surface of the ellipse with a uniform distribution, the distribution of the ray directions may be any one of these:

- a) All rays appear to emit from a point placed anywhere relative to the source. The location of this point is defined in the parameter list below. When used in this mode, the source acts like an imaging point source.
- b) Rays emit in a cosine distribution of the form

$$I(\theta) \approx I_o(\cos\theta)^{Cn},$$

where the exponent Cn may be any value between zero and 100, and need not be an integer. The larger Cn, the narrower the distribution becomes. Note this distribution is rotationally symmetric about the local Z axis.

- c) Rays emit in a Gaussian distribution of the form

$$I(l, m) \approx I_o e^{-(G_x l^2 + G_y m^2)},$$

where l and m are the direction cosines of the ray in the X and Y axis directions and Gx and Gy are constants. This form may be used to define a far field pattern that is different in the X or Y directions. The larger Gx and Gy are, the narrower the distribution becomes in the respective directions.

The nature of the ray distribution is defined by the parameter values. If Cn, Gx, and Gy, are all zero, then all rays will appear to emit from a virtual point source. If Cn is 1.0 or greater, then the cosine distribution will result (no matter how the source distance or Gx or Gy are set). If Cn is zero, but either Gx and Gy is non-zero, then a Gaussian distribution will result.

The parameters used to define this source are:

- 1-5: See "Parameters common to all source objects" on page 450.
- 6: X Half Width: The x half width in lens units.
- 7: Y Half Width: The y half width in lens units.

8: Source Distance: The distance along the local z axis from the apparent source point to the location of the source object. This value may be positive or negative. If zero, the rays are collimated. If positive, the apparent source point is behind the object. Considered only if Cn, Gx, and Gy are all zero.

9: Cosine Exponent: The power on the cosine term. This is Cn in the cosine distribution expression above. The ray cosines do not support Sobol sampling when this parameter is not zero.

10: Gaussian Gx: The X term in the Gaussian distribution. Ignored if Cn is not zero. The ray cosines do not support Sobol sampling when this parameter is not zero.

11: Gaussian Gy: The Y term in the Gaussian distribution. Ignored if Cn is not zero. The ray cosines do not support Sobol sampling when this parameter is not zero.

12: Source X: The X coordinate of the point that emits the rays. If Source Distance is zero, then the Source X parameter is the X direction cosine of the collimated ray bundle. Considered only if Cn, Gx, and Gy are all zero.

13: Source Y: The Y coordinate of the point that emits the rays. If Source Distance is zero, then the Source Y parameter is the Y direction cosine of the collimated ray bundle. Considered only if Cn, Gx, and Gy are all zero.

14: Min X Half Width: The minimum x half width in lens units. Use a value greater than zero and less than the X Half Width to define an annular region. The ray coordinates do not support Sobol sampling when this parameter is not zero.

15: Min Y Half Width: The minimum y half width in lens units. Use a value greater than zero and less than the Y Half Width to define an annular region. The ray coordinates do not support Sobol sampling when this parameter is not zero.

Source EULUMDAT File

This source is defined by measured photometry data for a real lamp, properly formatted in a file commonly referred to as EULUMDAT format. The file extension must be LDT and the file must be placed in the <data>\Sources\EULUMDAT folder (see "Folders" on page 68). Although LDT files usually contain a specification of the total flux of the lamp, to remain consistent with other ZEMAX source models the total lamp flux is a user definable parameter.

The parameters are:

1-5: See "Parameters common to all source objects" on page 450.

Source Filament

The Source Filament can be thought of as a thin wire coiled in a helix shape. The wire turns "N" times along the full Z coordinate length given by "L". The radius of the turns is defined by "R". Rays emanate from a randomly chosen point along the helix in a random direction.

The parameters used to define this source are:

1-5: See "Parameters common to all source objects" on page 450.

6: Length "L" in lens units.

7: Radius "R" in lens units.

8: Number of turns "N" (dimensionless). N may be fractional or even negative to reverse the rotation direction of the helix.

The axis of the helix, and the length L, are oriented along the Z axis.

Source File

The Source File is a source whose ray coordinates, cosines, and intensity are defined in a user supplied file. The name of the file containing the ray data must be placed in the comment column of the object. The file extension may be either DAT or SDF and the file must be placed in the <data>\Objects\Sources\Source Files folder (see "Folders" on page 68). The extension SDF is recommended and DAT is supported only for backward compatibility. The file format may be either text or binary, both formats are described below.

The parameters are:

1-5: See "Parameters common to all source objects" on page 450.

6: Randomize?: If set to zero, the rays will normally be traced in the order listed in the file. If non-zero, then the ray order is randomized once, when the file is read or any parameter in the NSC Editor changes for the source object and the source is updated. The randomize feature is only available if the total number of rays in the file is less than the maximum number of source rays to hold in memory. Files with more than the maximum number of source rays to hold in memory are too large to hold in memory and cannot be randomized. These large files are partially read as required as described below. The maximum number of source rays to hold in memory is a user definable parameter; see "Maximum Source File Rays In Memory" on page 114.

7: Power/Lumens In File: The value for this parameter is for reference, and is set by ZEMAX upon reading the file and should not be altered or set by the user. The behavior of this field is different for "flux only" and "spectral color" formats, as defined in "Binary Source File format" on page 456.

For flux only format files:

This reference value is the total power in current source units defined by the rays in the file. Only the number of rays used for analysis is considered for making this power computation. If the current source units are photometric (lumens) then the power in watts in the source file are converted to photometric units at the monochromatic wavelength defined for the source file in the parameter 4 value. Note there is no way to convert power in watts to lumens if the source is polychromatic, since the wavelengths are randomly assigned as described in "Parameters common to all source objects" on page 450. If the source is polychromatic, this parameter will be set to zero. The actual power of each ray is defined by parameter 3 above, the number of rays being traced, and the relative power of each ray as defined in the source file.

For spectral color format files:

The wavelength of each ray is defined in the file, and those wavelengths will be used for tracing no matter what the wavelength parameter in the NSC Editor is set to. The Power/Lumens value will be valid even if this source is polychromatic, and the actual flux and wavelengths defined in the file are used to determine the total power in current source units.

8/9: These values are for reference only, and are set by ZEMAX upon reading the file and should not be altered or set by the user. These references values are the minimum and maximum wavelength in micrometers for the source. For spectral color format files, these values are read from the file. For flux only format files these values are determined by the system wavelengths or color spectrum defined for the source, if any. See "Defining the color and spectral content of sources" on page 433.

Text vs. binary format files

The maximum number of rays currently allowed for text files is 1,000,000. If there are more than 1,000,000 rays defined in a text source file, ZEMAX will offer to convert the file to binary format automatically. The original file will be renamed with the ".OLD" extension, and the new binary SDF file will be created using the same name as the original file. After the conversion is complete, an option will be presented to save or delete the old file. This conversion is done to save disk space (binary files are about 30% as large as equivalent text files), load time (binary files read about 20x faster than text files), and system memory (large binary files can be left on disk and thus require very little RAM) even if billions of rays are being traced.

Memory requirements for source file objects

For both text and binary files, if the total number of rays to be traced is less than the maximum number of source rays in memory, ZEMAX stores the rays in system memory for maximum speed. This requires about 32 bytes of memory for each ray, or up to 32 Mb of memory for 1,000,000 rays. If the total number of rays is greater than the maximum number of source rays in memory, ZEMAX leaves the ray data on disk, and reads the file as required. This requires less memory, and allows ZEMAX to trace a huge number of rays, limited only by the systems file storage capacity. The maximum number of source rays to hold in memory is a user definable parameter; see "Maximum Source File Rays In Memory" on page 114.

Restrictions on the number of rays selected

There are restrictions on the values for the number of layout and analysis rays when using the Source File:

- The number of analysis rays must be equal to or larger than the number of layout rays. The number of analysis rays will be set to the larger of the two numbers entered in the NSC Editor.

- The number of analysis and layout rays may not exceed the number of rays defined in the file.

Binary Source File format

The Binary Source File consists of a header structure of the form:

```
typedef struct
{
  int Identifier; // Format version ID, current value is 1010
  unsigned int NbrRays; // The number of rays in the file
  char Description[100]; // A text description of the source
  float SourceFlux; // The total flux in watts of this source
  float RaySetFlux; // The flux in watts represented by this Ray Set
  float Wavelength; // The wavelength in micrometers, 0 if a composite
  float AzimuthBeg, AzimuthEnd; // Angular range for ray set (Degrees)
  float PolarBeg, PolarEnd; // Angular range for ray set (Degrees)
  long DimensionUnits; // METERS=0, IN=1, CM=2, FEET=3, MM=4
  float LocX, LocY, LocZ; // Coordinate Translation of the source
  float RotX, RotY, RotZ; // Source rotation (Radians)
  float ScaleX, ScaleY, ScaleZ; // Scale factor to expand/contract source
  float unused1, unused2, unused3, unused4;
  int ray_format_type, flux_type;
  int reserved1, reserved2;
} NSC_RAY_DATA_HEADER;
```

The data types float, int, and unsigned int are all 32 bit types. Identifier defines the file format, currently 1010. Older formats normally can be read by ZEMAX as well, although ZEMAX only writes files in the current format. ZEMAX only uses the NbrRays, DimensionUnits, ray_format_type, and flux_type parameters. The other data is ignored.

The maximum valid number of rays in the file as defined by NbrRays is 4 billion (4,000,000,000 rays). Exceeding this number may cause problems with the reading of the data file by ZEMAX. The ray_format_type must be either 0 for flux only format, or 2 for the spectral color format (both are defined below). Other format types are not supported. If and only if the ray_format_type is 0, then the flux_type is 0 for watts, and 1 for lumens. For the spectral color format, the flux must be in watts, and the wavelength in micrometers. The flux and wavelength values and units are used in the ray structures defined below.



Note that when the spectral color format is used, ZEMAX will always trace rays at the wavelengths defined in the file, no matter what system wavelengths are defined.

After the header follow NbrRays ray structures. The ray structure format depends upon the ray format type.

For the flux only format, the ray structure is:

```
typedef struct
{
  float x, y, z;
  float l, m, n;
  float flux;
} FLUX_ONLY;
```

For the color spectral color format, the ray structure is:

```
typedef struct
{
  float x, y, z;
  float l, m, n;
  float flux, wavelength;
} SPECTRAL_COLOR;
```

ZEMAX can read either the flux only or the spectral color format. The spectral color format only is used by ZEMAX when generating source files from ray trace data (see the discussion of "Save Rays on Object n" in the section "Ray Database Viewer" on page 455).

Text Source File format

The text Source File consists of a single line of header data with just two integer numbers of the form:

```
number_of_rays dimension_flag
```

The number_of_rays indicates the total number of rays in the file. The dimension_flag is 0 for meters, 1 for inches, 2 for centimeters, 3 for feet, and 4 for millimeters.

The remaining lines in the file are of the format:

```
x y z l m n intensity wavelength
```

Any line starting with the "!" symbol is assumed to be a comment line and is ignored. Any number of rays may be defined in a single file using this text format. However, if the number of rays exceeds 1,000,000, the file will automatically be converted to binary format by ZEMAX the first time the file is opened for a source file object. For a discussion of this conversion see "Text vs. binary format files" on page 455. The wavelength argument is optional, however, if present, the ray will always be traced at the defined wavelength in micrometers, and the source file will be treated as described by the spectral color format in the discussion on binary file formats.

Intensity normalization in source files

Each ray may have a different relative intensity. If the relative intensity value is not 1.0 for each ray, then normalization occurs as described below.

When the source file is first loaded into memory, the intensity of each ray is summed and then normalized to the average intensity. If the total flux of the source is then defined to be some number of watts, a subset of the rays can be traced and their intensities will yield approximately, but not exactly, the total flux. Normalization is required for an arbitrary subset of the rays to yield approximately the total desired power.

Source Gaussian

The Source Gaussian has a Gaussian distribution of rays which appear to emanate from a point source. The parameters are:

1-5: See "Parameters common to all source objects" on page 450.

6: Beam Size: The beam radius at the 1 over e² point in intensity in lens units.

7: Position: The distance from the apparent point of divergence of the rays to the source plane location. If zero, the rays are collimated.

Source IESNA File

This source is defined by measured photometry data for a real lamp, properly formatted in a file according to the Illumination Engineering Society of North America (IESNA) Standard File Format for Electronic Transfer of Photometric Data (IESNA91, IESNA LM-63-95, or IESNA LM-63-02). The file extension must be IES and the file must be placed in the <objects>\Sources\IESNA folder (see "Folders" on page 68). Although IES files usually contain a specification of the total flux of the lamp, to remain consistent with other ZEMAX source models the total lamp flux is a user definable parameter.

The parameters are:

1-5: See "Parameters common to all source objects" on page 450.

Source Imported

This source has a shape defined by an imported data file, in either IGES, STEP, or SAT format. The source volume has no properties other than to define the starting location and orientation of the rays. Rays are spatially distributed uniformly in area over the object. The angular distribution of the rays is given by

$$P(\theta) = (\cos\theta)^x,$$

where θ is the angle measured from the local surface normal pointing outside the volume, and x is a user defined parameter between 0.0 and 100.0. If $x = 0.0$, the rays emanate in a hemisphere with equal probability in all directions. If $x = 1.0$, the distribution is Lambertian. If $x \geq 100.0$, the rays are always normal to the surface.

The parameters are:

1-5: See “Parameters common to all source objects” on page 450.

6: Scale. This dimensionless parameter scales the entire imported solid. Upon import, ZEMAX will automatically attempt to scale the dimensions of the imported solid to match the current dimensions in ZEMAX; this scale factor is applied after that conversion.

7: Chord Tolerance: The chord tolerance determines the position accuracy of the starting point for the rays, and affects the rendering of the source. To render or trace rays from the source, ZEMAX converts the imported file to a list of triangles which approximate the shape. The tolerance is the maximum allowed distance in lens units between a single triangle and the actual surface of the source. More triangles are added if the tolerance is set smaller which yields more accurate rendering, at the expense of speed and a larger memory requirement. The default value of zero will use a chord tolerance related to the size of the object sufficient for most purposes.

8: Cosine Factor: The cosine power exponent "x" described above.

Source Object

The source object creates a source with a size and shape based upon any other "parent" object. Any prior object in the NSC Editor may be used to define the shape of the Source Object, including User Defined and Boolean objects. Any changes in the properties of the parent object will dynamically affect the distribution of the rays on the Source Object. The parent object shape is used to determine the starting location and orientation of the rays. Rays are spatially distributed uniformly in area over the object. The angular distribution of the rays is given by

$$P(\theta) = (\cos\theta)^x,$$

where θ is the angle measured from the local surface normal pointing outside the volume, and x is a user defined parameter between 0.0 and 100.0. If $x = 0.0$, the rays emanate in a hemisphere with equal probability in all directions. If $x = 1.0$, the distribution is Lambertian. If $x \geq 100.0$, the rays are always normal to the surface.

For best results, the parent object should be a solid volume. If the parent object is a surface, the rays may radiate from either one or both sides of the surface, depending upon exactly how the parent object is defined.

There are two ways to prevent rays leaving the source from immediately striking the parent object. The first way is to set the Source Object to ignore the parent object from which it is defined. This will allow the ray to ignore the parent object until after the ray leaves the source (See “Defining an Ignore Objects list” on page 437). The second way is to define a pre-propagation distance that is large compared to the chord tolerance but small compared to the distance the ray might legitimately travel before it interacts with the source itself (See “Sources tab” on page 431). This latter method is useful if the interaction of the rays with the source body is desired and the source is hollow or concave so that rays leaving one part of the object may strike another part before striking any other object.

The Source Object may be placed independently from the parent object. To superimpose the Source Object on the parent object, set the reference object of the Source Object to the parent object (see “Reference objects” on page 423) and leave the position and tilt values at zero.

The parameters are:

1-5: See “Parameters common to all source objects” on page 450.

6: Parent Object #. The integer object number of the object used to define the source shape. This object number must *precede* the object number of the Source Object in the NSC Editor.

7: Chord Tolerance: The chord tolerance determines the position accuracy of the starting point for the rays, and affects the rendering of the source. To render or trace rays from the source, ZEMAX converts the imported file to a list of triangles which approximate the shape. The tolerance is the maximum allowed distance in lens units between a single triangle and the actual surface of the source. More triangles are added if the tolerance is set smaller which yields more accurate rendering, at the expense of speed and a larger memory requirement. The default value of zero will use a chord tolerance related to the size of the object sufficient for most purposes.

8: Cosine Factor: The cosine power exponent "x" described above.

Source Point

The Source Point is a point which emits rays into a cone. The cone angle may be any value between 0 and 180 degrees (which would radiate into a full sphere). The parameters are:

- 1-5: See "Parameters common to all source objects" on page 450.
- 6: Cone angle: The semi-cone angle in degrees.

Source Radial

The Source Radial is a flat, rectangular or elliptical shape that emits rays into a hemisphere. The distribution of the rays in this hemisphere is symmetric about the local z axis and is defined by a cubic spline fit to arbitrary intensity vs. angle data. The number of points can be any integer between 5 and 181, inclusive. The angular range is evenly divided among the specified number of points. The data for each point is the relative intensity measured in the far field of the source at the angle corresponding to that point. The provided data is fit to a cubic spline. Rays generated from the spline fit will follow the correct distribution in the far field of the source.

The parameters are:

- 1-5: See "Parameters common to all source objects" on page 450.
- 6: X Half Width: The x half width in lens units. If less than zero, the emitting region will be elliptical.
- 7: Y Half Width: The y half width in lens units. If less than zero, the emitting region will be elliptical.
- 8: Minimum Angle in degrees. The value must be between 0.0 and 89.9.
- 9: Maximum Angle in degrees. The value must be between the minimum angle plus 0.1 and 90.0.
- 10: The number of points, must be between 5 and 181, inclusive.
- 11+: The relative intensity data at each angle from the normal.

Source Ray

The Source Ray is a point which emits rays along specified direction cosines. This is useful for debugging, for example, the direction cosines may be used to recreate a ray that causes geometry errors. The parameters are:

- 1-5: See "Parameters common to all source objects" on page 450.
- 6-8: Cosines: The local X, Y, and Z direction cosines of the ray. These are automatically normalized.
- 9: Random Seed: If not zero, the random number generator is reseeded with the specified value. This allows identical rays to be generated which subsequently have random but reproducible properties, such as when bulk or surface scattering is on.

Source Rectangle

The Source Rectangle is a flat rectangular surface which emits rays from a virtual source point. The parameters are identical to the Source Ellipse, but the shape of the source is a rectangle rather than an ellipse.

Source Tube

The Source Tube is similar to the Source Volume Cylinder, except the rays only emanate from the surface of the tube rather than the full volume.

The parameters used to define this source are:

- 1-5: See "Parameters common to all source objects" on page 450.
- 6: Length "L" in lens units.
- 7: Radius "R" in lens units.

The axis of the tube, and the length L, are oriented along the Z axis.

Source Two Angle

The Source Two Angle is a flat surface, with either an elliptical or a rectangular shape, which emits rays. The origin of each ray launched lies on the surface of the source with a uniform distribution. The angular distribution of the rays is either uniform in angle space or uniform as projected onto a distant plane. The angular distribution may also be either rectangular or elliptical, with distinct maximum half angles for each direction.

Using just two orthogonal angles to describe a source distribution underconstrains the possible energy distribution, since there is no clear definition of what happens at intermediate azimuthal angles. For this reason, this source model should only be used at angles less than roughly 60 degrees, and only when the aspect ratio between the X and Y angles is no more than about 4 to 1. It is a good idea to test the resulting energy distribution using a polar detector to confirm the approximations used in this model yield the desired distribution.

The parameters used to define this source are:

1-5: See “Parameters common to all source objects” on page 450.

6: X Half Width: The x half width in lens units.

7: Y Half Width: The y half width in lens units.

8: X Half Angle in degrees: The half angle of the cone of rays in the XZ plane. Maximum 89 degrees.

9: Y Half Angle in degrees: The half angle of the cone of rays in the YZ plane. Maximum 89 degrees.

10: Spatial Shape: Use 0 for rectangular, 1 for elliptical.

11: Angular Shape: Use 0 for rectangular, 1 for elliptical.

12: Uniform Angle: Use 0 for uniform irradiance on a distant projected plane, or 1 for uniform in angle space.

Source Volume Cylinder

The Source Volume Cylinder is a 3D volume formed by an elliptical shape in the XY plane that is symmetrically extruded along the Z axis. The center of the source volume is at the local origin of the object. Rays are emitted in random directions from anywhere inside the volume, with uniform probability in both position and ray direction. Points within the cylinder volume satisfy these relations:

$$\left(\frac{X}{W_x}\right)^2 + \left(\frac{Y}{W_y}\right)^2 \leq 1, \text{ and } \left(\frac{Z}{W_z}\right)^2 \leq 1,$$

where W refers to the half widths in X, Y, and Z. The parameters are:

1-5: See “Parameters common to all source objects” on page 450.

6: X Half Width: The x half width in lens units.

7: Y Half Width: The y half width in lens units.

8: Z Half Width: The z half width in lens units.

Source Volume Ellipse

The Source Volume Ellipse is a 3D volume formed by an elliptical shape in the XY, XZ, and YZ planes. The center of the source volume is at the local origin of the object. Rays are emitted in random directions from anywhere inside the volume, with uniform probability in both position and ray direction. Points within the ellipse satisfy this relation:

$$\left(\frac{X}{W_x}\right)^2 + \left(\frac{Y}{W_y}\right)^2 + \left(\frac{Z}{W_z}\right)^2 \leq 1,$$

where W refers to the half widths in X, Y, and Z. The parameters are otherwise identical to that for the Source Volume Cylinder.

Source Volume Rectangle

The Source Volume Rectangle is a 3D volume formed by a rectangular shape in the XY, XZ, and YZ planes. The center of the source volume is at the local origin of the object. Rays are emitted in random directions from anywhere inside the volume, with uniform probability in both position and ray direction. Points within the rectangle satisfy these relations:

$$|X| \leq W_x, |Y| \leq W_y, \text{ and } |Z| \leq W_z,$$

where W refers to the half widths in X , Y , and Z . The parameters are otherwise identical to that for the Source Volume Cylinder.

Detectors

The available types of detectors in ZEMAX are summarized in the table below.

SUMMARY OF NSC DETECTORS

Object Name	Description	Page
Detector Color	A flat rectangular detector with an arbitrary number of pixels. This detector can record and display incoherent illumination data defined by tristimulus response. This detector can accurately record and display the color of illumination.	463
Detector Polar	A section of a sphere, or a full sphere, that collects angular (far field) intensity data. Data collected by this detector may be exported into source data files such as IESNA and EULUMDAT.	464
Detector Rectangle	A flat rectangular detector with an arbitrary number of pixels. This detector can record and display incoherent, coherent, point spread function, polarization, and other data. This is the most powerful detector in terms of the analysis it provides, but it is limited to a flat rectangular shape.	466
Detector Surf	A circular or annular detector with an arbitrary number of pixels in the radial and angular directions. The surface may follow a plane, sphere, conic asphere, or aspherical shape. The surface detector can only record incoherent irradiance data.	468
Detector Volume	A rectangular volume with an arbitrary number of voxels in the local x, y, and z directions. The detector volume may be nested within or straddle any other object. Multiple detector volumes may also be superimposed and all will be illuminated by rays passing through the individual voxels.	470
Objects as detectors	Most objects of arbitrary shape may be used as a detector that records incoherent irradiance data.	470

Detailed descriptions of these detector types follow.

Detector Color object

The Detector Color object stores power and tristimulus (visual light color) data from NSC source rays that strike it. The resulting data distributions may be viewed for incoherent light in spatial or angular domains. Detector Color objects may be reflective, transparent, or absorbing if the material is set to "MIRROR", blank, or "ABSORB", respectively.

The defining parameters are:

- 1: X Half Width. The x width in lens units.
- 2: Y Half Width. The y width in lens units.
- 3: # X Pixels. The number of pixels along the x direction. The maximum number of pixels is 5000.
- 4: # Y Pixels. The number of pixels along the y direction. The maximum number of pixels is 5000.

Parameters 5-9 are only used to define how the detector is displayed on shaded model plots:

- 5: Data type: Select 0 for irradiance, 1 for intensity. If photometric units are selected, the options are the equivalent photometric quantities. For more information on selecting units, see "Analysis Units" on page 105.
- 6: Color: Select for 0 grey scale, 1 for inverse grey scale, 2 for false color, 3 for inverse false color, and 4 for true color.

7: Smoothing: The amount of smoothing. The algorithm is described in "The Detector Viewer" on page 384.

8: Scale: Choose 0 for linear, 1 for log -5, 2 for log -10, and 3 for log -15. This option is ignored if Color is set to 4 for true color.

9: Plot Scale: choose a maximum value to normalize the color display to; this is useful for setting a common scale across multiple detectors. This value is ignored if it does not exceed the maximum displayed value.

10: Front Only. If 0, then rays may strike the detector on the front or the back side. If this flag is 1, then rays coming from the back will be ignored, and will miss the detector. The "back" side is the side facing towards the positive local z axis.

11: Unused.

12-15: The minimum and maximum x and y direction angles in degrees. These settings are used to control the sensitivity of the detector to rays incident upon the detector at an angle when displaying intensity. The default setting is -90.0 degrees to +90.0 degrees in both x and y directions, which allows display of angle data for all rays striking the detector. If the angular range is set to cover a subset of the possible incident angles, rays that strike the detector at angles outside the defined range are ignored ONLY for intensity displays.

16: Polarization flag. If 0, the detector ignores polarization of the incident rays and considers all incident rays. If the flag is 1, 2, or 3, the detector will only consider the component of the rays polarized along the local X, Y, and Z directions, respectively. If the flag is 4, 5, or 6, the detector will only consider the component of the rays polarized along X and Y, X and Z, and Y and Z, respectively.

17: Mirroring flag. Mirroring allows the detector to exploit symmetry in the incident rays. Any ray that lands on the detector will be divided into parts, and each part of the ray will be copied to the symmetric mirror images of the detector. If mirroring is set to 0, the detector does not use the mirroring feature. If mirroring is set to 1, then the detector will assume left-right symmetry exists in the incident rays. If mirroring is set to 2, then the detector will assume top-bottom symmetry. If mirroring is set to 3, then both left-right and top-bottom mirroring will be done. The mirroring may also be set to 4, 5, or 6; these values are similar to 1, 2, and 3, respectively, except the intensity of the rays is not scaled down to conserve energy. For example, using a value of 1 would divide a single 1-watt ray into two 0.5 watt rays, while a value of 4 would create two 1-watt rays. Which value is correct depends upon whether or not the sources generating the illumination are also scaled down to reflect symmetry or not, and only the user can make this determination. Note that if mirroring is used, and the illumination and optics are not actually symmetric in the expected way, erroneous data may result.

The information stored on any detector may be viewed by selecting Detectors, Detector Viewer on the NSC Editor, or when the program mode is Non-sequential, on the Analysis menu. The Detector Viewer also has an option for displaying data from any previously saved ray data base. If "none" is selected for the ray database, then the data currently stored in the detector will be displayed, otherwise, the data stored in the selected ray database will be used to regenerate the displayed image. For more information see "The Detector Viewer" on page 384.

For more information on the definitions and units for the terms irradiance, intensity, and flux, see "Analysis Units" on page 105.

Comments on Detector Color pixel numbering

For the Detector Color, the pixel numbers start at 1 in the lower left (-x, -y) corner of the rectangle. The pixel numbers increase along the +x direction first, and then move up one row in the +y direction, and start over at the -x edge. For a detector with n_x by n_y pixels, the pixel numbers are 1 through n_x on the bottom row, and then n_x+1 through $2n_x$ on the next row above the bottom, until the last pixel (number n_x*n_y) is reached in the upper right (+x, +y) corner. Another way of stating the pixel numbering is that the x index changes the fastest, and then the y index.

Comments on Pixel Interpolation

See "Pixel Interpolation" under "Type tab" on page 353.

Detector Polar object:

The Detector Polar object stores power and tristimulus (visual light color) data from NSC source rays that strike it. The resulting data distributions may be viewed for incoherent light in the angular domains. Detector Polar objects may be reflective, transparent, or absorbing if the material is set to "MIRROR", blank, or "ABSORB", respectively.

The object shape is a faceted approximation to a portion of a sphere centered on the object coordinates. Although the detector only records angular data, the detector must be large enough to encompass all the rays of interest, and the size of the detector is a user-definable parameter. The polar object spans an angular region, with the zero polar angle centered around the object local Z axis. The polar angle may span any angle between 1 and 180 degrees (which would form a complete sphere). The detector only considers the angle the ray striking the detector makes, and not the position the ray strikes the detector. The azimuthal angle is defined by

$$\phi = \text{atan}\left(\frac{y}{x}\right),$$

where the zero angle is along the +x axis and the angle increases in a counter-clockwise direction all the way to 360 degrees. The azimuthal angle always spans 360 degrees. The defining parameters are:

- 1: Maximum Angle. The maximum polar angle in degrees. Must be between 1 and 180 degrees.
- 2: Radial Size. The radial size of the detector in lens units.
- 3: # P Pixels. The number of pixels along the polar angle direction. Must be between 10 and 721.
- 4: # A Pixels. The number of pixels along the azimuth angle direction. Must be between 12 and 720, and must be a multiple of 4.
- 5: Mirroring flag. Mirroring allows the detector to exploit symmetry in the incident rays. Any ray that lands on the detector will be divided into parts, and each part of the ray will be copied to the symmetric mirror images of the detector. If mirroring is set to 0, the detector does not use the mirroring feature. If mirroring is set to 1, then the detector will assume left-right symmetry exists in the incident rays. If mirroring is set to 2, then the detector will assume top-bottom symmetry. If mirroring is set to 3, then both left-right and top-bottom mirroring will be done. The mirroring may also be set to 4, 5, or 6; these values are similar to 1, 2, and 3, respectively, except the intensity of the rays is not scaled down to conserve energy. For example, using a value of 1 would divide a single 1-watt ray into two 0.5 watt rays, while a value of 4 would create two 1-watt rays. Which value is correct depends upon whether or not the sources generating the illumination are also scaled down to reflect symmetry or not, and only the user can make this determination. Note that if mirroring is used, and the illumination and optics are not actually symmetric in the expected way, erroneous data may result.

The information stored on any detector may be viewed by selecting Detectors, Detector Viewer on the NSC Editor, or when the program mode is Non-sequential, on the Analysis menu. The Detector Viewer also has an option for displaying data from any previously saved ray data base. If "none" is selected for the ray database, then the data currently stored in the detector will be displayed, otherwise, the data stored in the selected ray database will be used to regenerate the displayed image. For more information see "The Detector Viewer" on page 384.

For more information on the definitions and units for the terms irradiance, intensity, and flux, see "Analysis Units" on page 105.

Exporting Detector Polar Data

Intensity data stored on a Detector Polar maybe exported to a source file. See "Export Source Data" on page 259.

Comments on Detector Polar Pixel Numbering

For the Detector Polar, the pixel numbers start at 1 in the polar center at an azimuthal angle of zero. The pixel numbers increase along the polar direction until pixel np, where np is the number of polar pixels. The pixel np+1 is at zero polar angle at the first azimuthal increment counter-clockwise from 0 degrees, and the pixel number increase along this azimuthal arm until 2*np. The next pixel number is 2*np+1, which starts again at a polar angle of zero and proceeds along the next azimuthal arm. Another way of stating the pixel numbering is that the polar index changes the fastest, and then the azimuthal index.

Comments on Pixel Interpolation

See "Pixel Interpolation" under "Type tab" on page 353.

Detector Rectangle object

The detector rectangle object stores energy data from NSC source rays that strike it. The resulting data distributions may be viewed for incoherent light in spatial or angular domains, or spatially coherent irradiance or phase, or coherent irradiance as a point spread function. Detector rectangles may be reflective, transparent, or absorbing if the material is set to "MIRROR", blank, or "ABSORB", respectively.

The defining parameters are:

1: X Half Width. The x width in lens units.

2: Y Half Width. The y width in lens units.

3: # X Pixels. The number of pixels along the x direction. The maximum number of pixels is 5000.

4: # Y Pixels. The number of pixels along the y direction. The maximum number of pixels is 5000.

Parameters 5-9 are only used to define how the detector is displayed on shaded model plots:

5: Data type: Select 0 for incoherent irradiance, 1 for coherent irradiance, 2 for coherent phase, 3 for radiant intensity, 4 for radiance in position space, and 5 for radiance in angle space. If photometric units are selected, the options are the equivalent photometric quantities. For more information on selecting units, see "Analysis Units" on page 105.

6: Color: Select for 0 grey scale, 1 for inverse grey scale, 2 for false color, and 3 for inverse false color.

7: Smoothing: The amount of smoothing. The algorithm is described in "The Detector Viewer" on page 384.

8: Scale: Choose 0 for linear, 1 for log -5, 2 for log -10, and 3 for log -15.

9: Plot Scale: choose a maximum value to normalize the color display to; this is useful for setting a common scale across multiple detectors.

10: Front Only. If 0, then rays may strike the detector on the front or the back side. If this flag is 1, then rays coming from the back will be ignored, and will miss the detector. The "back" side is the side facing towards the positive local z axis.

11: PSF Mode Wavenumber. If 0, then the coherent irradiance display is the coherent data detected for each pixel individually. This is the default and most appropriate choice if the detector is large and the intent is to look at the beam far from focus, such as fringes created by two overlapping beams from an interferometer. If the wavenumber is equal to a wavelength number, then the coherent data for each pixel is the Huygens integral of all rays incident upon the detector, resulting in a diffraction Huygens Point Spread Function. When using PSF Mode, the width of the detector should be a small number of wavelengths, to allow adequate pixel resolution to see the diffraction structure, and to avoid aliasing artifacts inherent to the Huygens model when applied to large areas. The irradiance is normalized so that the total power is equal to the total incoherent power on the detector. This gives accurate results as long as the rays incident upon the detector cover a broad enough range of angles so that both constructive and destructive interference occur over the size of the detector. This normalizing will not give accurate results if the incident rays are all nearly parallel and PSF Mode should not be used for this purpose. If PSF Mode is used and the detector is placed inside of another object, such as a block of glass, then the "Inside Of" value for the detector should be used to indicate which object the detector is immersed within.

12-15: The minimum and maximum x and y direction angles in degrees. These settings are used to control the sensitivity of the detector to rays incident upon the detector at an angle when displaying radiant intensity or radiance in angle space. The default setting is -90.0 degrees to +90.0 degrees in both x and y directions, which allows display of angle data for all rays striking the detector. If the angular range is set to cover a subset of the possible incident angles, rays that strike the detector at angles outside the defined range are ignored ONLY for radiant intensity and radiance in angle space displays.

16: Polarization flag. If 0, the detector ignores polarization of the incident rays and considers all incident rays. If the flag is 1, 2, or 3, the detector will only consider the component of the rays polarized along the local X, Y, and Z directions, respectively. If the flag is 4, 5, or 6, the detector will only consider the component of the rays polarized along X and Y, X and Z, and Y and Z, respectively. Setting the polarization flag to 1, 2, or 3 will also consider the phase of the electric field as well as the phase of the ray (due to optical path, index, and phase objects) when computing the coherent amplitude and phase displayed by the detector.

17: Mirroring flag. Mirroring allows the detector to exploit symmetry in the incident rays. Any ray that lands on the detector will be divided into parts, and each part of the ray will be copied to the symmetric mirror images of the detector. If mirroring is set to 0, the detector does not use the mirroring feature. If mirroring is set to 1,

then the detector will assume left-right symmetry exists in the incident rays. If mirroring is set to 2, then the detector will assume top-bottom symmetry. If mirroring is set to 3, then both left-right and top-bottom mirroring will be done. The mirroring may also be set to 4, 5, or 6; these values are similar to 1, 2, and 3, respectively, except the intensity of the rays is not scaled down to conserve energy. For example, using a value of 1 would divide a single 1-watt ray into two 0.5 watt rays, while a value of 4 would create two 1-watt rays. Which value is correct depends upon whether or not the sources generating the illumination are also scaled down to reflect symmetry or not, and only the user can make this determination. Note that if mirroring is used, and the illumination and optics are not actually symmetric in the expected way, erroneous data may result.

The information stored on any detector may be viewed by selecting Detectors, Detector Viewer on the NSC Editor, or when the program mode is Non-sequential, on the Analysis menu. The Detector Viewer also has an option for displaying data from any previously saved ray data base. If "none" is selected for the ray database, then the data currently stored in the detector will be displayed, otherwise, the data stored in the selected ray database will be used to regenerate the displayed image. For more information see "The Detector Viewer" on page 384.

Four data items are stored for every ray that strikes the detector:

The incoherent intensity of the ray. This energy is stored by incrementing the counter corresponding to the pixel the ray struck. The sum of the ray intensity divided by the pixel area yields the irradiance in flux per area.

The incoherent intensity of the ray in angular space. This energy is stored by incrementing the counter corresponding to the pixel the ray struck in angle space. The sum of the ray intensity divided by the solid angle the pixel represents yields the intensity in flux per solid angle.

The coherent real and imaginary amplitudes of the ray. The amplitude is the square root of the ray intensity, while the phase is determined by the total optical path length from the source referenced to the center of the spatially distributed pixel. By separating the real and imaginary parts of the ray, interference between many rays may be simulated. Phase due to polarization effects is also considered if the polarization flag is 1, 2, or 3; see the discussion for parameter 16 above.

For more information on the definitions and units for the terms irradiance, intensity, and flux, see "Analysis Units" on page 105.

Comments on coherent data computations

The propagation and interference of light generally has properties of both particles and waves. Rays can be thought of as the particle representation, and diffraction interference (such as for a diffraction PSF) can be thought of as a wave representation.

For NSC analysis, ZEMAX uses ray tracing to determine optical paths and energy distributions. ZEMAX accounts for the phase along the ray, and this allows for computation of some interference and diffraction effects. However, it is important that the user understand what assumptions the model makes and how these assumptions affect the accuracy of the results.

When a ray strikes a detector, ZEMAX computes the real and imaginary parts of the electric field by using the intensity and phase of the ray referenced to the center of the pixel struck. The real and imaginary parts may then be summed for many rays that strike the same pixel. ZEMAX also sums the intensity (amplitude squared) for each pixel.

Because the phase is accounted for, some rays will constructively interfere with other rays while other rays will destructively interfere. This allows ZEMAX to simulate effects such as fringes in interferometers (shearing or otherwise) or interference from various orders of a diffraction grating. However, computing the coherent irradiance involves some assumptions. Physically, destructive interference means the energy would have propagated somewhere other than where the ray went. In a similar way, when constructive interference occurs, the squaring of the amplitude sum of many rays artificially and unphysically increases the energy in the beam. ZEMAX cannot determine where the energy went (or came from), and therefore cannot account for conservation of energy in coherent irradiance calculations without making assumptions. ZEMAX can make one of two different assumptions: either to normalize the coherent power on the detector to match the incoherent power, or to determine the coherent power using the incoherent power on each pixel and scaling by the degree of coherence of rays landing on each pixel. Which method is used depends upon the number of pixels on the detector, and on the setting of the Normalize Coherent Power checkbox (see "The object properties dialog box" on page 353).

If there is only one pixel defined for the detector, or if the Normalize Coherent Power checkbox is off, then the coherent irradiance for each pixel is computed by summing the real and imaginary parts of every ray incident upon that pixel, computing the magnitude of this sum, dividing by the square of the sum of the coherent amplitude of all incident rays upon that pixel, then finally multiplying this ratio by the incoherent irradiance on the pixel. This method allows the computed coherent irradiance to vary between a value of zero and the incoherent irradiance. However, there is no way to accurately determine the true coherent irradiance in this case because of the limitation of the ray model in the presence of constructive and destructive interference as described above. Specifically, it is unknowable where energy lost in this computation would have propagated to.

If there is more than one pixel in the detector, and the Normalize Coherent Power checkbox is on, then the coherent irradiance is computed from the ray data by squaring and summing the real and imaginary parts of the amplitude pixel by pixel, and then re-normalizing the total coherent irradiance for the entire detector to be equal to the incoherent irradiance incident upon the detector.

With two methods available, the question naturally arises - which method is correct? If one or the other method was always correct, there would be no need for the Normalize Coherent Power switch. The general rule is simple. If all the incoherent power from one or more sources falls on the same detector, then the Normalize Coherent Power switch should be checked on. This is the default setting. This would be case, for example, of two plane waves interfering to form a sinusoidal fringe pattern with many fringes across a detector. All of the energy is present on the detector, but the wavefront interference causes ray energy that would have landed in a dark fringe to be redirected to a neighboring bright fringe on the same detector. Conversely, when two parallel, perfectly collimated beams are combined, yielding a single bright or dark (or something in between) fringe, then the Normalize Coherent Power switch should be checked off. This would be the case where an interferometer was set up to yield a single dark fringe on one detector and a single bright fringe on another detector - physically the energy never would propagate along the dark path, but the rays cannot know this in advance. Note that in this case (even when using the recommended setting), the bright fringe will only show half the actual power - the rays cannot create power, and the coherent power computation is still capped at the incoherent energy. The energy lost in the dark fringe physically would have been directed to the bright fringe, but again, the rays are incapable of determining that.

There may be cases where neither of these coherent power computation methods yields the correct results. These cases are simply beyond the scope of the ray model, and ZEMAX should not be used for critical analysis of the physical optics aspect of these systems.

It is important to understand that ZEMAX considers ALL sources to be coherent with respect to one another, and for the phase of the ray to be zero at the starting coordinate of the ray, wherever it may be. This generally limits the usefulness of the interference analysis to monochromatic sources. Polychromatic sources cannot be modeled correctly for coherent analysis because all rays, regardless of their wavelength, are coherently summed. The initial phase of the ray and the coherence length of the source may be defined, see "Coherence length modeling" on page 373 and "Sources tab" on page 358.

Comments on Detector Rectangle pixel numbering

For the Detector Rectangle, the pixel numbers start at 1 in the lower left (-x, -y) corner of the rectangle. The pixel numbers increase along the +x direction first, and then move up one row in the +y direction, and start over at the -x edge. For a detector with n_x by n_y pixels, the pixel numbers are 1 through n_x on the bottom row, and then n_x+1 through $2n_x$ on the next row above the bottom, until the last pixel (number n_x*n_y) is reached in the upper right (+x, +y) corner. Another way of stating the pixel numbering is that the x index changes the fastest, and then the y index.

Comments on Pixel Interpolation

See "Pixel Interpolation" under "Type tab" on page 353.

Detector Surface object

The detector surface object stores energy data from NSC source rays that strike it. The resulting data distributions may be viewed as irradiance for incoherent light. Detector surfaces may be reflective, transparent, or absorbing if the material is set to "MIRROR", blank, or "ABSORB", respectively.

The surface shape of a detector surface object is identical to the Aspheric Surface object. See "Aspheric Surface" on page 402 for information on the defining formulas for the surface shape.

The defining parameters are:

1: Radius: The radius of curvature in lens units. If zero, a plane base surface will result.

2: Conic: The conic constant. See "Aspheric Surface" on page 402 for the surface shape formula.

3: Max Aperture: The maximum radial size in lens units.

4: Min Aperture: The minimum radial size in lens units. If this value is greater than zero, an annular detector will be created.

5: # A Zones: The number of angular zones, Na. See "Comments about detector triangles" below.

6: # R Zones: The number of radial zones, Nr. See "Comments about detector triangles" below.

Parameters 7-11 are only used to define the appearance of the Detector Surface on Shaded Model Plots:

7: Data type: Select 0 for incoherent irradiance, 1 for incoherent flux.

8: Color: Select for 0 grey scale, 1 for inverse grey scale, 2 for false color, and 3 for inverse false color.

9: Unused.

10: Scale: Choose 0 for linear, 1 for log -5, 2 for log -10, and 3 for log -15.

11: Max Scale: choose a maximum value to normalize the color display to; this is useful for setting a common scale across multiple detectors.

12: Front Only. If 0, then rays may strike the detector on the front or the back side. If this flag is 1, then rays coming from the back will be ignored, and will miss the detector. The "back" side is the side facing towards the positive local z axis.

13: The number of terms to use in the aspheric expansion. Ray tracing will be faster if this term is no larger than the highest order non-zero coefficient.

14-250: The α coefficients on the polynomial expansion. See "Aspheric Surface" on page 402 for the surface shape formula.

The information stored on any detector may be viewed by selecting Detectors, Detector Viewer on the NSC Editor, or when the program mode is Non-sequential, on the Analysis menu. The Detector Viewer also has an option for displaying data from any previously saved ray data base. If "none" is selected for the ray database, then the data currently stored in the detector will be displayed, otherwise, the data stored in the selected ray database will be used to regenerate the displayed image. For more information see "The Detector Viewer" on page 384.

One data item is stored for every ray that strikes the detector: The incoherent intensity of the ray. This energy is stored by incrementing the counter corresponding to the pixel the ray struck. The sum of the ray intensity divided by the pixel area yields the irradiance in flux per area.

For more information on the definitions and units for the terms irradiance, intensity, and flux, see "Analysis Units" on page 105.

Comments about detector triangles

ZEMAX uses a collection of triangles to conform to the shape of the curved surface. There are Nr radial zones, and Na angular zones. Together these two parameters define $Nr * Na$ polygonal regions. Each region is generally divided into two triangles, so the total number of triangular pixels is $2 * Nr * Na$. However, if the minimum aperture (parameter 4 above) is zero, then the innermost radial zone has only Na triangles rather than $2 * Na$. In this case the total number of triangles is $2 * Nr * Na - Na$. The algorithm that selects the triangle coordinates attempts to make the area of each triangle approximately the same, however the area from triangle to triangle can not be held exactly constant. The corners of each triangle will lie on the aspheric surface defined by the radius, conic, and aspheric coefficients. The triangles are flat and do not lie exactly on the surface if the surface is not flat.

Comments on Detector Surface pixel numbering

The triangles are the individual pixels used to record irradiance and flux data. The pixels are numbered in consecutive radial zones. Pixel 1 starts at angle zero in the innermost radial zone. The first $2 * Na$ pixels (or just Na pixels if the minimum aperture is zero) are numbered going counter clockwise around z to complete the first

annular zone. The pixel numbering continues at angle zero in the second radial zone for the next $2 * N_a$ pixels, until all radial zones are complete.

If "Fast Ray Trace" (see "Type tab" on page 353) is checked, ZEMAX will alter the order of the triangles to speed up the ray trace. If "Fast Ray Trace" is unchecked, the order of the triangles is unmodified, which slows the ray tracing but preserves the ordering described above.

Detector Volume object

The detector volume object stores incident and absorbed energy data from NSC source rays that pass through the individual voxels. The resulting data distributions may be viewed as incident flux, absorbed flux, or absorbed flux per unit volume. Detector volumes may be transparent, or be made of any valid material. Detector volumes may also be nested within or straddle other objects. Objects within detector volumes may be refractive, reflective, absorbing, gradient index, or have surface and or bulk scatter properties defined.

The shape of a detector volume object is a rectangular solid.

The defining parameters are:

1-3: The X, Y, and Z half widths of the volume. The center of the volume is at the local coordinate origin.

4-6: The number of voxels in the X, Y, and Z directions. The total number of voxels is $n_x * n_y * n_z$.

The information stored on any detector may be viewed by selecting Detectors, Detector Viewer on the NSC Editor, or when the program mode is Non-sequential, on the Analysis menu. Data for detector volume objects is not available when using previously saved data from a ray data base. Pressing the left and right arrow keys while viewing detector volume data on the NSC detector viewer will cycle through all the Z-planes. For more information see "The Detector Viewer" on page 384.

Two data items are stored for every ray that passes through each voxel: The incident flux of the ray in source units, and the absorbed flux in the voxel. The incident flux is the flux of the ray as it enters the voxel. A ray segment that starts in a particular voxel will not increase the incident flux in that voxel, because the incident flux is only incremented for rays entering the voxel. The absorbed flux is a function of the path length of the ray through the voxel, as well as the coefficient of absorption of the material the ray is passing through.

Detector volumes may be nested within other NSC object types. For example, to measure the absorbed flux in a cylindrical glass rod, place the rod within the detector volume, and set the glass type of the rod on the rod object. If the glass type has an absorption coefficient associated with the material, the portion of the flux of rays passing through the glass rod will be recorded by the detector volume.

For more information on the definitions and units for the terms irradiance, intensity, and flux, see "Analysis Units" on page 105.

Comments on Detector Volume pixel numbering

For the Detector Volume, the pixel numbers start at 1 in the lower left corner of the rectangle formed by the first z-plane (on the -z side of the volume). The pixel numbers in this z-plane increase along the +x direction first, and then move up one row in the +y direction, and start over at the -x edge. For a detector with n_x by n_y by n_z pixels, the pixel numbers in the first z-plane are 1 through n_x on the bottom row, and then n_x+1 through $2n_x$ on the next row above the bottom, until the pixel n_x*n_y is reached in the upper right (+x, +y) corner. The numbering then proceeds to the next z-plane, starting again at the bottom left corner. Another way of stating the pixel numbering is that the x index changes the fastest, then the y index, and then finally the z index.

Objects as detectors

Most objects may act like a detector. This includes objects with flat faces, such as the polygon, STL, and rectangular volume objects. Most objects that approximate the shape of a curved surface or solid for rendering purposes by using the drawing resolution setting (see "Draw tab" on page 362), such as the aspheric surface, toroidal lens, biconic lens, and other complex shapes, may also be detectors.

To make any of these objects a detector, choose "Object Is A Detector" from the Type tab of the Object Properties dialog box. Once the object is set to be a detector, the intensity detected will be displayed on the Shaded Model display, and on the text listing of the Detector Viewer window. The color scheme the detector will display on the Shaded Model may also be selected on the Object Properties dialog box. The colors chosen may represent total flux or may be normalized to be flux per unit area.

Objects acting as detectors can only display incoherent intensity. Each individual triangle used to draw the object becomes a single pixel. All rays which strike the pixel will be summed to yield the total intensity. The object detectors are cleared whenever all detectors are cleared from the detector control dialog box.

A tool for creating POB files for use as detectors is described in “Create Polygon Object” on page 96.

Comments on general objects as detectors pixel numbering

The pixels are numbered in the order the polygons are defined, and there is no general way to describe the pixel ordering. For these objects, some experimentation may be required to determine the exact pixel ordering method.

Introduction

Solves are functions which actively adjust specific values. Solves can be specified on curvatures, thicknesses, glasses, semi-diameters, conics, parameters, TCE values, and Extra Data Editor values. Solves are set by double clicking on the cell you wish to place the solve on. The solve type control window appears. Select the solve type from the drop-down list. Some solves require additional parameters. The available solves and required parameters are summarized in the following table. Parameters may also have "fixed" or "variable" status; these are not solves, however, they are listed for completeness. To turn solves off, select the "fixed" option from the drop-down list. Note that this turns off the solve, it does not change the last value. Pressing Ctrl-Z twice with the cursor on the solved parameter will also remove any solve.

SUMMARY OF SOLVES

Solve type	Param 1	Param 2	Param 3	Param 4	Code	Integer
Curvature: Fixed						0
Curvature: Variable					V	1
Curvature: Marginal ray angle	Angle				M	2
Curvature: Chief ray angle	Angle				C	3
Curvature: Pickup	Surface	Scale factor		Column	P	4
Curvature: Marginal ray normal					N	5
Curvature: Chief ray normal					N	6
Curvature: Aplanatic					A	7
Curvature: Element power	Power				X	8
Curvature: Concentric with surface	Surface to be concentric to				S	9
Curvature: Concentric with radius	Surface to be concentric with				R	10
Curvature: F/#	Paraxial F/#				F	11
Curvature: ZPL Macro	Macro name				Z	12
Thickness: Fixed						0
Thickness: Variable					V	1
Thickness: Marginal ray height	Height	Pupil zone			M	2
Thickness: Chief ray height	Height				C	3
Thickness: Edge thickness	Thickness	Radial height (use zero for semi-diameter)			E	4
Thickness: Pickup	Surface	Scale factor	Offset	Column	P	5
Thickness: Optical path difference	OPD	Pupil zone			O	6
Thickness: Position	Surface	Length from surface			T	7

Solve type	Param 1	Param 2	Param 3	Param 4	Code	Integer
Thickness: Compensator	Surface	Sum of surface thicknesses			S	8
Thickness: Center of curvature	Surface to be at the COC of				X	9
Thickness: Pupil Position	-	-	-	-	U	10
Thickness: ZPL Macro	Macro name				Z	11
Glass: Fixed						0
Glass: Model	Index Nd	Abbe Vd	Dpgf			1
Glass: Pickup	Surface				P	2
Glass: Substitute	Catalog name				S	3
Glass: Offset	Index Nd offset	Abbe Vd offset			O	4
Semi-Diameter: Automatic						0
Semi-Diameter: Fixed					U	1
Semi-Diameter: Pickup	Surface	Scale factor		Column	P	2
Semi-Diameter: Maximum					M	3
Semi-Diameter: ZPL Macro	Macro name				Z	4
Conic: Fixed						0
Conic: Variable					V	1
Conic: Pickup	Surface	Scale factor		Column	P	2
Conic: ZPL Macro	Macro name				Z	3
Parameter: Fixed						0
Parameter: Variable					V	1
Parameter: Pickup	Surface	Scale factor	Offset	Column	P	2
Parameter: Chief ray	Field	Wavelength			C	3
Parameter: ZPL Macro	Macro name				Z	4
TCE: Fixed						0
TCE: Variable					V	1
TCE: Pickup	Surface				P	2
Extra Data Value: Fixed						0
Extra Data Value: Variable					V	1
Extra Data Value: Pickup	Surface	Scale factor	Offset	Column	P	2
Extra Data Value: ZPL Macro	Macro name				Z	3

Curvature: Marginal ray angle

The effective focal length of the lens can be controlled by placing a paraxial marginal ray angle solve on the curvature of the last surface before the image. For example, suppose a lens has a 20 mm entrance pupil diameter, and it is desired to constrain the effective focal length to 100 mm. This requires a marginal ray exit angle from the

last surface of -1 (this number is 20 divided by 2 divided by 100, the minus sign is because the ray is converging, or going down toward the image surface). See the F number solve.

Curvature: Chief ray angle

The chief ray angle solves work the same way as the marginal ray angle, except the paraxial chief ray is used for the calculations instead. The chief ray solve is useful for maintaining a particular magnification or for maintaining collimation. The letter "C" will appear in the curvature cell to indicate that a chief ray angle solve is active on that surface.

Curvature: Pickup

The curvature pickup solve uses a scaled value from another surface and column as the curvature on the target surface. Note the solve computes the correct curvature, although ZEMAX displays the radius (one over the curvature) in the Lens Data Editor. The pickup allows variables to be coupled together, and they will change under the influence of other solves, editing, or optimization. See also "Solve restrictions" on page 479. To set a pickup solve from a ZPL macro, see "Integer codes for column numbers" on page 480.

Curvature: Marginal ray normal

This solve will force the surface to be normal to the paraxial marginal ray. This is also called an image-centered surface. These special surfaces introduce no spherical aberration or coma. The letter "N" will appear in the curvature cell to indicate that a normal ray angle solve is active on that surface.

Curvature: Chief ray normal

This solve will force the surface to be normal to the paraxial chief ray. This is also called a pupil-centered surface. These special surfaces introduce no coma, astigmatism, or distortion. The letter "N" will appear in the curvature cell to indicate that a normal ray angle solve is active on that surface.

Curvature: Aplanatic

This solve will force the surface to be aplanatic with respect to the paraxial marginal ray. These special surfaces introduce no spherical aberration, coma, or astigmatism. The letter "A" will appear in the curvature cell to indicate that an aplanatic solve is active on that surface.

Curvature: Element power

The power of an element is given by:

$$\phi = c_1(n_2 - n_1) + c_2(n_3 - n_2) - c_1(n_2 - n_1) \cdot c_2(n_3 - n_2) \cdot \frac{t_2}{n_2}.$$

This solve adjusts the value of c_2 , to maintain the specified element power. The solve is assumed to be on the second of two adjacent surfaces. This solve is ignored if the surface number is less than 2, or if n_3 and n_2 are identical at the primary wavelength.

Curvature: Concentric with surface

This solve will force the curvature of the surface such that the surface is concentric about the specified surface. The specified surface must precede the surface on which the solve is placed. The letter "S" will appear in the cell to indicate that the surface concentric solve is active on that surface.

Curvature: Concentric with radius

This solve will force the curvature of the surface such that the surface is concentric about the same point that the specified surface is concentric to. The specified surface must precede the surface on which the solve is placed. The letter "R" will appear in the cell to indicate that the radius concentric solve is active on that surface.

Curvature: F number

This solve will force the curvature of the surface such that the marginal ray angle exiting the surface is $-1/2F$ where F is the paraxial $F/\#$. The letter "F" will appear in the cell to indicate that the $F/\#$ solve is active on that surface.

Curvature: ZPL Macro

For a complete description of ZPL Macro solves, see “Using ZPL Macro solves” on page 479. Note curvature solves must return data in units of curvature, which is the reciprocal of the radius of curvature.

Thickness: Marginal ray height

The most common thickness solve is the marginal ray height solve, which can be used to constrain the image surface to the paraxial focus. To move the image surface to the paraxial focus, double click on the thickness of the last surface before the image surface, and select "Marginal Ray Height" from the drop-down list which appears. Since the marginal ray has a height of zero (it crosses the axis) at the paraxial image surface (assuming a rotationally symmetric system), we would like to set this last thickness to bring the image surface to the point where the marginal ray height is zero. The default height value is zero, which is desired, and so click on OK to exit the window, and the solve will adjust the thickness appropriately.

In principle, the solve could have been set for some other marginal ray height by entering a value other than zero for height, the first optional parameter. Note that any surface can have a marginal ray height solve, not just the thickness before the image surface. The "Height" is the height of the marginal ray on the tangent plane of the next surface (again, not necessarily the image surface). Note the height is measured at the point the ray intercepts the tangent plane, not the actual curved surface of the next surface.

The third value, "Pupil Zone", allows the ray pupil coordinate to be defined. The default is zero, which indicates that a paraxial ray should be used. Any non-zero value indicates that a real marginal ray is to be used. The zone value must be between -1 and 1. This is the Py coordinate, or the normalized entrance pupil coordinate in the y direction. This solve can be used to constrain that a particular ray, such as the .7 zonal ray, to have zero transverse aberration on axis. The letter "M" will appear in the thickness cell to indicate that a marginal ray height solve is active on that surface.

Thickness: Chief ray height

This is similar to the marginal ray height solve, except the paraxial chief ray is used. This solve is useful for locating a surface at a pupil plane. The letter "C" will appear in the thickness cell to indicate that a chief ray height solve is active on that surface.

Thickness: Edge thickness

This solve dynamically adjusts the spacing between two surfaces to maintain a specified distance between the surfaces at a specified radial aperture. This is useful for preventing negative, or overly sharp edges on elements. The letter "E" will appear in the thickness cell to indicate that an edge thickness solve is active on that surface. If the radial aperture is set to zero, then the current semi-diameter will be used. See “Edge thickness” on page 52 for an important discussion on the definition of edge thickness.

Thickness: Pickup

The thickness pickup solve uses a scaled and offset value from another surface and column as the thickness on the target surface. The pickup thickness T is given by:

$$T = O + S \cdot V$$

where V is the value of the source data, S is the scale factor, and O is the offset. See also “Solve restrictions” on page 479. To set a pickup solve from a ZPL macro, see “Integer codes for column numbers” on page 480.

Thickness: Optical path difference

This solve will actively adjust the thickness to maintain a specific optical path difference at a specific pupil coordinate. The OPD is measured at the exit pupil, not at the surface the solve is placed on. The two parameters to be set are the OPD in primary waves, and the pupil zone at which to evaluate the OPD. For example, to maintain the focal position so that the real marginal ray has the same optical path length as the real chief ray, define an OPD solve on the last thickness before the image surface. Set the OPD parameter to zero, and the pupil zone to 1.0. The letter "O" will appear to indicate the OPD solve is active. Perform an OPD plot and verify that the OPD is in fact zero at the pupil edge. Only the primary wavelength is used, and only the on-axis field is considered.

More complicated OPD constraints are available by using optimization, described in a separate chapter. The letter "O" will appear in the thickness cell to indicate that an OPD solve is active on that surface.

Thickness: Position

The position solve maintains the "z" distance from a specified reference surface. If the reference surface precedes the surface on which the solve is placed, then the sum of the thicknesses from the reference surface to the surface following the solve surface will be maintained at the specified value. If the reference surface follows the surface on which the solve is placed, then the sum of the thicknesses from the solve surface to the reference surface will be maintained at the specified value. If the reference surface is the same as the solve surface, then the thickness of the solve surface will be set at the solve length value.

The position length solve is particularly useful for maintaining the length of a portion of a zoom lens to a fixed value. The solve can also be used to meet a total lens length constraint. In either case, the solve may eliminate optimization variables and operands, enhancing optimization convergence and speed. The letter "T" will appear in the thickness cell to indicate that a total length solve is active on that surface.

This solve assumes that all the surfaces in the affected range are in the same coordinate system.

Thickness: Compensator

The thickness compensator solve is very similar to the position solve. This solve maintains the thickness on a surface so that the sum of the surface thickness and the thickness on another "reference" surface equals a constant. In equation form, this solve maintains the condition $T = S - R$, where S is the sum of the thicknesses of the two surfaces, and R is the thickness of the reference surface. The reference surface must precede the surface with the solve.

Thickness: Center of curvature

This solve will adjust the thickness of the surface to place the following surface at the center of curvature of any prior surface. The letter "X" will appear in the cell to indicate that the center of curvature solve is active on that surface.

Thickness: Pupil Position

This solve will place the next surface at the pupil position for the optical space following refraction from the current surface. The pupil position is determined by tracing real, differential rays about the central field chief ray.

Thickness: ZPL Macro

For a complete description of ZPL Macro solves, see "Using ZPL Macro solves" on page 479.

Glass: Model

The model glass state is not a solve. The model glass is used to idealize a glass by describing the dispersion in the visible wavelength region using three parameters: the index of refraction at d-light, the Abbe number at d-light, and a term describing the partial dispersion. For more information on model glasses, see "Using model glasses" on page 601.

Glass: Pickup

The glass pickup solve uses a value from another surface as the glass on the target surface. See also "Solve restrictions" on page 479.

Glass: Substitute

The substitute state is not a solve. If the glass solve type is set to "Substitute", then the global optimization algorithms are permitted to change the glass type during optimization. This is similar to a variable status in that the glass type can change.

If no catalog name is given (i.e. the catalog field is blank) then glasses may be selected from all catalogs selected for use on the General System Data dialog box. If a catalog name is given (i.e. Hoya) then only glasses from that one catalog will be selected.

To prevent only certain glasses from being chosen during optimization, choose "Exclude Substitution" for the glass that should be avoided on the glass catalog dialog box, or use a glass substitution template; see "Using glass substitution" on page 542 for more information.

Glass: Offset

The offset solve allows a small change in index and/or Abbe number to be added to the index as computed by the dispersion formulas and the glass catalog dispersion data. The primary use of this solve is for tolerancing.

There are several conditions for using the index and Abbe numbers in computing the change in index as a function of wavelength:

-The minimum wavelength in use is greater than 0.3 micrometers and less than 2.5 micrometers

-The dispersion data in the glass catalog must span the wavelength range of 0.4861327 to 0.6562725 micrometers.

If these conditions are met, then the change in index is given by the difference in the model glass index computed at the base N_d and V_d as compared to the offset N_d and V_d values. See the discussion of the "Model" glass in this chapter for a discussion of the model glass properties. In equation form the index is:

$$n = n_{base} + n(N_d + \Delta N_d, V_d + \Delta V_d, P_d) - n(N_d, V_d, P_d)$$

where the function $n()$ is the model glass function. Note this model will add a different index offset to the index at each wavelength. If the wavelength band criteria listed above is not met, then the index offset is added to all the base index values:

$$n = n_{base} + \Delta N_d$$

The Abbe offset is ignored in this case.

Semi-Diameter: Pickup

The semi-diameter pickup solve uses a scaled value from another surface and column as the semi-diameter on the target surface. See also "Solve restrictions" on page 479. To set a pickup solve from a ZPL macro, see "Integer codes for column numbers" on page 480.

Semi-Diameter: Maximum

The maximum solve is used for setting the semi-diameter value to the maximum required across multiple configurations. For example, a zoom lens which has three configurations will generally have three different semi-diameter values for every surface that is using "automatic" semi-diameters. The maximum solve will compute the semi-diameter for each configuration, and then use the largest of the values.

Semi-Diameter: ZPL Macro

For a complete description of ZPL Macro solves, see "Using ZPL Macro solves" on page 479.

Conic: Pickup

The conic pickup solve uses a scaled value from another surface and column as the conic on the target surface. See also "Solve restrictions" on page 479. To set a pickup solve from a ZPL macro, see "Integer codes for column numbers" on page 480.

Conic: ZPL Macro

For a complete description of ZPL Macro solves, see "Using ZPL Macro solves" on page 479.

Parameter: Pickup

The parameter pickup solve uses a scaled and offset value from another surface and column as the parameter on the target surface. The pickup parameter is given by:

$$P' = O + S \cdot V$$

where V is the value of the source data, S is the scale factor, and O is the offset. See also "Solve restrictions" on page 479. To set a pickup solve from a ZPL macro, see "Integer codes for column numbers" on page 480.

Parameter: Chief ray

The chief ray solve works only on coordinate break surfaces, and only on the first four parameters. When placed on a decenter x or y parameter, the solve will set the decenter to place the selected wavelength (use zero for primary wavelength) real chief ray from the selected field position at an x or y coordinate of zero on the coordinate break surface, respectively. If set on the tilt x and tilt y parameters, the tilt angles are adjusted to make the exiting chief ray angle zero in the y and x directions, respectively. The order flag on the coordinate break surface will determine the order in which the parameters are fit, and the solution may not be unique. The output chief ray coordinates can be checked on the Analysis, Calculations, Ray Trace feature; and it is a good idea to verify the solve works correctly. This solve is not allowed when the stop surface follows the surface the solve is placed on and ray aiming (see "Ray Aiming" on page 106) is on.

There are times when the resulting values will yield an output chief ray angle or coordinate that is small, but not zero. This is caused by the non-orthogonal rotations not always being able to completely "undo" an arbitrary rotation in a single set of matrix rotations. There are two solutions to this problem:

- 1) Try changing the "order" flag from 1 to 0 or 0 to 1, then update the system, to see which set of coordinate transforms works better.
- 2) Use two adjacent coordinate break surfaces with identical chief ray solves. This is easy to do by copying and pasting the entire coordinate break surface. The second coordinate break will usually bring the chief ray many orders of magnitude closer to the desired zero coordinates or angles.

Parameter: ZPL Macro

For a complete description of ZPL Macro solves, see "Using ZPL Macro solves" on page 479.

TCE: Pickup

The TCE pickup solve uses a value from another surface as the TCE on the target surface. See also "Solve restrictions" on page 479.

Extra Data Value: Pickup

The extra data value pickup solve uses a scaled and offset value from another surface and column as the extra data value on the target surface. The extra data value E is given by:

$$E = O + S \cdot V$$

where V is the value of the source data, S is the scale factor, and O is the offset. See also "Solve restrictions" on page 479. To set a pickup solve from a ZPL macro, see "Integer codes for column numbers" on page 480.

Extra Data Value: ZPL Macro

For a complete description of ZPL Macro solves, see "Using ZPL Macro solves" on page 479.

Solve restrictions

Solves are evaluated sequentially from the first surface to the image surface, in this order: curvature, thickness, glass, semi-diameter, conic, parameters (in order left to right), TCE, and extra data values (in order left to right). Because of this order, pickup solves may only pickup values from prior surfaces, or from the same surface if the source column precedes the target column.

Because solves can affect the entrance pupil position, ZEMAX does not allow certain types of solves, such as marginal ray height solves or macro solves, to be placed before the stop surface. This will avoid ambiguous or incorrect settings for the solved parameters. There is no way to predict the results of solves placed before the stop. A similar problem arises when a marginal ray solve is used on the last surface to control the focal length - the aperture definition must be entrance pupil diameter (rather than F/#) to keep the solution unique.

Using ZPL Macro solves

ZPL Macro solves call a user defined ZPL Macro to determine the solve value. For information about the ZPL Macro language, see "ZEMAX PROGRAMMING LANGUAGE" on page 651 and specifically "SOLVERRETURN" on page 734.

Macro solves work by invoking a user defined ZPL Macro to compute the solve. Any numerical value that can be computed in a macro can be returned to the editor calling the solve macro. The macro may make use of the data that exists elsewhere in the editors, such as previous surfaces. Once the data is computed, the macro uses the keyword SOLVERETURN to pass the data back to the editor.

For a simple example, here is a macro that computes the first -order optical power of the interface between surfaces 1 and 2, and returns this power as the solve value:

```
n1 = INDX(1)
n2 = INDX(2)
c2 = CURV(2)
SOLVERETURN (n2-n1)*c2
```

There is a disadvantage to referring to specific surfaces directly by surface number as this macro does. If new surfaces are inserted or deleted, the macro will need to be edited to reflect the new surface numbers. Also, certain features, such as mirror substrate drawing, will not work correctly if macro solves refer to fixed surface numbers. The solution is to use the ZPL function SURC to find surfaces with specific text in the comment column of the Lens Data Editor. Suppose the comment "My Surface" was placed on the desired first surface in the Lens Data Editor. The macro could support the power computation above with this revised macro:

```
A$ = "My Surface"
SURF = SURC(A$)
n1 = INDX(SURF)
n2 = INDX(SURF+1)
c2 = CURV(SURF+1)
SOLVERETURN (n2-n1)*c2
```

The surface or object number from which the macro solve is being called can also be extracted using the SOS function described on page 666.

Use surface comments and the SURC function in solve macros whenever possible.

Important considerations for ZPL Macro solves

Macro solves are very general, allowing virtually any computation to be used to determine a solve value. All functions and keywords supported by ZPL are available. There are no differences between ZPL macros executed from the macro menu or executed as a solve. However, many ZPL keywords and some functions should not be used in a macro solve. For example, calling "UPDATE" from the solve will update all solves, which will invoke the macro again, resulting in an infinite loop. Other keywords, such as INPUT, will require the user to enter data every time the solve is called, which may be an enormous number of times. Other keywords that set values in the editors or create merit functions should never be used. In general, macro solves should be kept short and simple, avoid lengthy computations, and should not make modifications of any lens data. Macro solves should also never depend upon data in the editors that is subsequent to the solve; as this may also create erroneous data if the solve is called first and then the source data is modified by a separate, subsequent solve. ZEMAX makes *no attempt to qualify or validate* the solve macro. This feature offers great power and flexibility, but must be used carefully.

If an error or invalid condition occurs in the macro, and the solve value cannot be computed, then the macro should return without calling SOLVERETURN. The absence of the SOLVERETURN call will indicate to ZEMAX that the solve cannot be computed and the optical system is in an error condition. This is particularly important during optimization.

Integer codes for column numbers

To set a Pickup solve from the ZPL macro language, the column number is required. The column number is defined as follows:

- 0: The current column, which is the column the solve is placed in. This is the default for all pickup solves.
- 1-4: Radius, Thickness, Conic, and Semi-Diameter, respectively.
- 5-17: Parameter 0 through 12, respectively.
- 18-259: Extra Data value 1 through 242, respectively.

Suggestions for use

Solves are highly efficient, and should be used when possible instead of variables during optimization. For example, it is usually better to place a curvature solve on the last surface to control effective focal length than to optimize for focal length explicitly.

Introduction

The optimization feature provided by ZEMAX is quite powerful, and is capable of improving lens designs given a reasonable starting point and a set of variable parameters. Variables can be curvatures, thicknesses, glasses, conics, parameter data, extra data, and any of the numeric multi-configuration data. ZEMAX uses either an actively damped least squares or an orthogonal descent algorithm. The algorithms are capable of optimizing a merit function composed of weighted target values; these target values are called "operands". ZEMAX has several different default merit functions, described in a subsequent section. These merit functions can be changed easily using the Merit Function Editor. For details on this procedure see "Modifying the merit function" on page 488.

Optimization requires three steps: 1) a reasonable system which can be traced, 2) specification of the variables, and 3) a merit function. A reasonable system is a rather loose concept which means that poorly conceived designs are not likely to be transformed into exceptional designs by the optimization algorithm (although there are exceptions). The variables, and there must be at least one for the optimization algorithm to be able to make any progress, are specified on the various editors, as described in the next section. To reach the optimization window, select Tools, Optimization. You must specify all variables before using optimization.

The algorithm used by the optimization feature described in this chapter is designed to find the "local" minimum of the specified merit function. However, ZEMAX also has the capability to search for a "global" minimum of the merit function. The global minimum is the lowest possible value for the merit function, and if the merit function is selected appropriately, this implies the best possible solution to the problem. The global optimization feature is not for novice users, and is not appropriate for interactive designing. For details see "GLOBAL OPTIMIZATION" on page 539.

Selecting variables

Variables for optimization are specified by pressing Ctrl-Z when the highlighted bar is on the parameter to be varied in the Lens Data Editor. Note that Ctrl-Z is a toggle. The Multi-Configuration and Extra Data Editors also contain numeric data that may be made variable by using Ctrl-Z. Glasses cannot be made variable directly because they are discrete. To optimize glasses, see "Optimizing glass selection" on page 528.

Defining the default merit function

The merit function is a numerical representation of how closely an optical system meets a specified set of goals. ZEMAX uses a list of operands which individually represent different constraints or goals for the system. Operands represent goals such as image quality, focal length, magnification, and many others.

The merit function is proportional to the square root of the weighted sum of the squares of the difference between the actual and target value of each operand in the list. The merit function is defined this way so a value of zero is ideal. The optimization algorithm will attempt to make the value of this function as small as possible, and so the merit function should be a representation of what you want the system to achieve. You do not have to use the default merit function, you may construct your own as described in a later section.

The easiest way to define a merit function is to select the Tools, Default Merit Function option on the Merit Function Editor menu bar. A dialog box will appear which allows selection of options for the default merit function. Each option is explained in the following paragraphs.

Selecting the type of optimization

Several different types of merit functions are available. The default merit function is constructed using four key choices: The optimization type, data type, reference point, and integration method. The choices are described in the following tables.

OPTIMIZATION TYPES

Name	Description
RMS	RMS is an abbreviation for Root-Mean-Square. This type is by far the most commonly used. The RMS is the square root of the average value of the squares of all the individual errors.
PTV	PTV is an abbreviation for Peak-To-Valley. There are rare cases where the RMS is not as important as the maximum extent of the aberrations. For example, if all the rays need to land within a circular region on a detector or fiber. In these cases, the PTV may be a better indicator of performance. This merit function type attempts to minimize the PTV extent of the errors. However, the numerical value of the merit function is not the actual peak-to-valley error.

OPTIMIZATION DATA

Name	Description
Wavefront	Wavefront is the aberration measured in waves. See “Comments about RMS wavefront computations” on page 155.
Spot Radius	The radial extent of the transverse ray aberrations in image space.
Spot X	The x extent of the transverse ray aberrations in image space.
Spot Y	The y extent of the transverse ray aberrations in image space.
Spot X and Y	Both the x and y extent of the transverse ray aberrations in image space. The x and y components are considered separately, and both are optimized together. This is similar to Spot Radius, except the signs of the aberrations are retained, which yields better derivatives. Note that in computing the radius of an aberration, the sign information is lost.
Angular Radius	The radial extent of the angular aberrations in image space. Intended for afocal systems. See “Afocal Image Space” on page 103.
Angular X	The x extent of the angular ray aberrations in image space. Intended for afocal systems. See “Afocal Image Space” on page 103.
Angular Y	The y extent of the angular ray aberrations in image space. Intended for afocal systems. See “Afocal Image Space” on page 103.

OPTIMIZATION REFERENCE POINTS

Name	Description
Centroid	The RMS or PTV computation of the data is referenced to the centroid of all the data coming from that field point. Centroid reference is generally preferred, especially for wavefront optimization. For wavefront optimization, reference to the centroid subtracts out the mean wavefront over the entire pupil, x-tilt, and y-tilt, none of which degrade image quality. Centroid reference also yields more meaningful results when coma is present, since coma shifts the image centroid away from the chief ray location.

Name	Description
Chief	The RMS or PTV computation of the data is referenced to the chief ray at the primary wavelength. For wavefront optimization, reference to the chief ray subtracts out the mean wavefront over the entire pupil, but not x-tilt or y-tilt. Note the exact point at which the OPD is defined to be zero is arbitrary; this is the reason the chief ray reference subtracts out the mean wavefront.
Unreferenced	This option is only available if the data is wavefront. If the wavefront is unreferenced, then the OPD data with respect to the chief ray is used without subtracting the mean wavefront or tilt.

Physically significant RMS merit functions

Note that the numerical value of the merit function is physically significant when using RMS as the optimization type. If the merit function is RMS-Wavefront-Centroid, then the numerical value of the merit function is the RMS wavefront error in waves. If the merit function is RMS-Spot Radius-Chief, then a value of 0.145 means the RMS spot radius is 0.145 lens units. If the lens units were millimeters, this would correspond to 145 micrometers RMS. If more than one field or wavelength is defined, then the merit function numerical value is the weighted average of the RMS values for the various fields and wavelengths.

Note that optimization using the RMS spot radius merit function will in general yield a different optimum design than the RMS wavefront merit function. The basic reason they are different is that ray aberrations are proportional to the derivative of the wave aberrations. Therefore, it is unreasonable to expect that the minimum of one corresponds to the minimum of the other. A general rule of thumb to use is that if the system is close to diffraction limited (say a PTV wavefront error of less than two waves) then use wavefront error. Otherwise, use spot radius.

Generally speaking, the merit functions which use the centroid as a reference are superior to those that use the chief ray. Most diffraction based performance measures, such as MTF or encircled energy, improve when the RMS wavefront error referenced to the centroid decreases. However, it is always best to reoptimize a final design with the various merit functions to verify which one provides the best performance for the system being designed. For example, the RMS wavefront centroid reference often yields better low frequency MTF response, but worse medial frequency response, than the RMS chief ray reference optimization does.

For PTV merit functions, the merit function numerical value is proportional but not equal to the actual PTV of the aberration. The merit function will however minimize the PTV when the merit function is minimized.

Selecting the pupil integration method

There are two different pupil integration methods used to construct the merit function: Gaussian quadrature (GQ) or rectangular array (RA). The GQ algorithm is VASTLY superior for almost all cases of practical interest. The GQ algorithm uses a carefully selected and weighted ray set to accurately compute the RMS or PTV error over the entrance pupil (strictly speaking, the PTV algorithm is not a GQ algorithm, but it is very similar). The weighting for all rays is applied according to the weights set on the wavelength and field data dialog boxes, any pupil apodization function, and by the GQ merit function algorithm. For RMS merit functions, the weighting and ray set selection used is based on a method described in G. W. Forbes, "Optical system assessment for design: numerical ray tracing in the Gaussian pupil", J. Opt. Soc. Am. A, Vol. 5, No. 11, p1943 (1988). For the PTV merit functions, the ray set is based on solutions to the Chebyshev polynomials, described in Numerical Recipes, Cambridge University Press (1989). If you are interested in detailed information on the basis and accuracy of these methods, see these references. GQ is much, much more accurate than any other known method, and requires fewer rays. Therefore, you get the best of both worlds: greater speed and greater accuracy. The GQ algorithm requires specification of the number of "Rings" and the number of "Arms", and these terms are defined in subsequent paragraphs. The only drawback to GQ is that the algorithm assumes the pupil is a circle, or more generally, an ellipse. For non-elliptical pupils, GQ does not work accurately. For example, if there are surface apertures in the optical system that vignette enough rays to alter the effective shape of the pupil significantly, GQ should not be used. One notable exception is when using circular pupils with modest central obscurations, such as a Newtonian telescope. Modest central obscurations do not usually affect the RMS significantly because the aberrations tend to be smaller in the central zone of the pupil. Note also that GQ works fine when used with vignetting factors, since the ray pattern is redistributed from a circle to an ellipse.

The RA algorithm traces a grid of rays through the pupil. The "Grid" size determines the number of rays traced, and is described in a subsequent paragraph. The "Delete Vignetted" option (also described later) allows the vignetted rays to be deleted from the ray set. Vignetted rays in this context are those rays clipped by surface apertures, not rays which have been altered by the use of vignetting factors (see "CONVENTIONS AND DEFINITIONS" on page 51). The advantage to the RA algorithm is the ability to accurately account for the effects of vignetting in the merit function. This is useful in systems such as obscured telescopes and camera lenses which intentionally clip troublesome rays. The disadvantage to the RA algorithm is speed and accuracy. Usually, more rays are required to achieve a given degree of accuracy than the GQ algorithm. The bottom line: don't use RA unless you are using surface apertures.

Rings

The "Rings" setting is only used in the GQ algorithm. It determines how many rays are traced at each field and at each wavelength. For on-axis fields (zero degrees field angle in a rotationally symmetric system), the number of rays is equal to the number of rings. For all other fields in symmetric systems, the number of rays traced per ring is equal to half the number of "arms" (defined in the next paragraph). Only half the rays are traced because the left-right symmetry of the system is exploited. Each set of rays is traced for each defined wavelength. For example, if you have one on-axis field, two off-axis fields, three wavelengths, four rings, and six arms selected, the number of rays traced is $3 * (4 + 4*3 + 4*3) = 84$. For systems without rotational symmetry, the number of rays per ring is the number of arms independent of field. In the prior example, this means $3 * 3 * 4 * 6 = 216$ rays.

Arms

The "Arms" setting is also only used in the GQ algorithm. It determines how many radial arms of rays in the pupil are traced. By default six equally spaced (in angle) arms are traced (or three if the system is rotationally symmetric). This number may be changed to eight, ten, or twelve. For most common optical systems, six is sufficient.

Sampling considerations for Rings and Arms

You should select the number of rings and the number of arms according to the order of aberrations present in your system. The highest order of aberration the GQ algorithm can accurately integrate is $(2*n - 1)$ where n is the number of rings. If an optical design will be limited by 5th order ray aberrations n should be at least 3. A simple way of determining the correct number of rings is to select the minimum number, one. Then go to the optimization dialog box and note the merit function. Now go back to the default merit function tool, and select two rings. If the merit function changes by more than a few percent, go back and select three, and so on until the merit function does not change significantly (perhaps 1%). Repeat the procedure for the number of arms (six arms is almost always plenty). Selecting more rings or arms than required will not improve the optimization performance, it will only slow the algorithm down needlessly. Tracing more rays than required will not help you find better solutions!



Selecting more rings or arms than required will not improve the optimization performance, it will only slow the algorithm down needlessly.

Grid

The "Grid" is only used by the RA algorithm, and the value determines the number of rays to be used. The grid size can be 4x4 (16 rays per field per wavelength), 6x6 (36 rays per field per wavelength) etc. Rays on the grid are automatically omitted if they fall outside the entrance pupil, so the actual number of rays used will be lower than the grid size squared. Selecting a larger grid size generally yields more accurate results at the expense of slower execution. However, there may be an advantage in choosing a large grid density, and then selecting the "Delete Vignetted" checkbox (described in the next paragraph). The reason is that a large grid density will fill the pupil with rays, and then the operands which are vignetted will be deleted. The result is a reasonable number of rays which accurately reflect the aperture of the system.

Delete Vignetted

The "Delete Vignetted" checkbox option is only used by the RA algorithm. If selected, then each ray in the merit function will be traced through the system, and if it is vignetted by a surface aperture, if it misses any surface, or if it is total internal reflected at any surface, the ray is deleted from the merit function. This keeps the total

number of rays in the merit function to a minimum. The disadvantage is that if the vignetting changes as the design is optimized, then the merit function may have to be regenerated. It is always a better choice to use the vignetting factors and then use the GQ algorithm than to delete vignetted rays if possible. Vignetting factors can be adjusted, if required, during optimization using SVIG in the merit function.

Note that ZEMAX will attempt to trace any ray defined in the merit function regardless if that ray is vignetted or not. For example, if the chief ray height is targeted using REAY, and there is a central obscuration that vignettes the chief ray, ZEMAX will still trace the ray and use the operand results as long as the ray can be traced. ZEMAX does not check to see if defined rays are vignetted, because this introduces substantial overhead during optimization.

In general, avoid vignetting of rays by surface apertures, and use vignetting factors to shape the beam size when possible. To optimize on the fraction of unvignetted rays, a macro must be defined to perform the required computations. However, this method is very prone to stagnation during optimization because small changes in lens parameters lead to discrete changes in the merit function as rays abruptly jump from being vignetted to not being vignetted.

Setting thickness boundary values

Boundary constraints may be automatically generated and included in the default merit function by checking the air and/or glass boundary values on. If selected, then MNCG, MXCG, and MNEG operands will be added to the merit function to constrain the minimum center thickness, maximum center thickness, and minimum edge thickness for glass surfaces, respectively. MNCA, MXCA, and MNEA operands will be added to the merit function to constrain the minimum center thickness, maximum center thickness, and minimum edge thickness for air spaces, respectively.

The automatic boundary constraint feature is meant to save some manual entry of routine boundary constraints on optical systems with or without mirrors. More complex lenses, such as those with complex coordinate breaks, or multi-configurations usually require additional boundary constraints to be added to the merit function manually.

Start At

The "Start At" option is used to add the default merit function at a specific position within the Merit Function Editor operand list. To control what the default start at value will be, see the DFMS operand definition. Use a value of 0 to start at the end of the current merit function.

Assume Axial Symmetry

If selected, then the default merit function will exploit the left-right and rotational symmetry of the lens when constructing and evaluating the merit function. Fewer rays will be traced, accelerating the optimization with no loss of accuracy. In systems with coordinate breaks or non-rotationally symmetric systems, the default is unselected, which means symmetry will not be exploited. Overriding the default system symmetry is useful, however, if you are designing a lens that ZEMAX thinks is non-symmetric, but the lack of symmetry does not affect the aberrations. For example, if tilted but flat fold mirrors are present, these mirrors do not eliminate the left-right symmetry of the system, but ZEMAX will by default assume symmetry does not exist. Some gradient index surfaces also use non-symmetric index variation terms which are often zero (they are used only for tolerancing). Select the checkbox to accelerate the optimization in these cases. See also the "USYM" operand description.

If the merit function will be used for tolerancing, check this option off, because even symmetric lenses generally become non-symmetric during tolerancing.

Ignore Lateral Color

By default, ZEMAX references all RMS or PTV computations to a common reference point for each field. All of the rays are traced for all wavelengths for each field point, and the primary wavelength chief ray or the centroid of all the rays is used as the reference point. If "Ignore Lateral color" is selected, then ZEMAX computes an independent reference point for each wavelength as well. This is useful for designing systems that intentionally divide the beam by wavelength, such as a prism or spectrometer system. This option will cause the merit function to optimize each color spot independently.

Relative X Weight

The relative X weight is the additional weighting to be placed on the X component of the transverse aberrations when computing the PTV or RMS SPOT X + Y merit function. This setting has no effect on the other merit functions. If the relative X weight is less than unity, then the Y components are weighted more heavily; if the relative weight is greater than unity, then the X components are weighted more heavily. If left at the default value of unity; then the components are equally weighted. This control is useful for systems which intentionally form slit images, such as spectrometers.

Overall Weight

When switching between different default merit function types, such as RMS spot radius and RMS wavefront error, the numerical magnitude of the default merit function can change dramatically. This may make it tedious to manually adjust the weights of operands that are not part of the default merit function. The overall weight is a factor that scales all the weights in the default merit function. Under most circumstances, this weight may be left at one.

Pitfalls with the default merit function

The default merit function is easy to set up, numerically efficient, and suitable for a large number of optimization problems. However, most optical designs require extensions or modifications to the default as the design progresses. ZEMAX offers significant flexibility in the definition of the merit function, as described in the following sections.

Note that if the field or wavelength values or weights are changed, you must reconstruct the default merit function. If you are using the RA algorithm, reconstruct the default merit function if the vignetting influence changes appreciably during optimization.



If field or wavelength values or weights are changed, reconstruct the merit function.

Optimization with apodized beams

If no pupil apodization has been specified (see “Apodization Type” on page 102) then ZEMAX assumes uniform illumination when constructing the default merit function. If the illumination is not uniform, then the rays in the default merit function are weighted according to the apodization factor. Since the rays selected may be insufficient to adequately represent an apodized beam, use a larger number of rays (described previously) when using apodization factors.

Modifying the merit function

The merit function can be modified by the user. To change the merit function, select Editors, Merit Function from the main menu bar. New operands can be added to the list, or others deleted, using the insert and delete keys. The current merit function value and the value of each operand can be updated by selecting Tools, Update.

Operands are set by typing the name in the first column and then filling in the remaining data fields. There are multiple fields that may be required to define an operand. The fields are called Int1 and Int2 for the two integer values, and Data1 through Data6 for up to 6 double precision values. Not all of the operands use all of the fields provided.

Hx, Hy, Px, and Py

Many of the operands use Data1 through Data4 for the values of Hx, Hy, Px, and Py; these are the normalized field and pupil coordinates (see “Normalized field coordinates” on page 55). Note that ZEMAX does not check to see if the specified Hx, Hy, Px, and Py coordinates are within the unit circle. For example a pupil coordinate of (1, 1) is actually outside the entrance pupil, but you will not get an error message when tracing these rays unless the rays cannot physically be traced.

Target, Value, and Weight

The target is the desired value of the specified parameter. The difference between the target and the value of the operand is squared, and summed over all operands to yield the merit function. The value of the target and

the operand itself is unimportant in optimization, only the difference between the two. The larger the difference, the greater the contribution to the merit function.

The weight is the relative importance of that parameter. The weight can be any number, positive or negative. However, the optimizer will act somewhat differently if the weight is negative, zero, or positive.

Percent Contribution

This column lists the relative contribution of each operand to the overall merit function as a percentage. This column is not user editable.

Operand weights less than zero

When the weight is negative, the operand will be treated as a Lagrangian multiplier. The Lagrangian multipliers force the optimization algorithm to find a solution which exactly meets the specified constraint, regardless of the effect on the other operands. This is sometimes useful to exactly meet an optimization target, such as focal length or magnification. In some respects, this is similar to a weight of "infinity", however it is implemented in a way that is numerically more stable.

Because there is generally a non-linear relationship between the variables and the operand targets, ZEMAX may not converge to the exact target value in a single optimization cycle; however, multiple cycles will usually converge to the Lagrangian targets with extremely high precision in a few cycles if a solution exists. It is possible to define Lagrangian targets that cannot be met with the variables provided, especially if there is more than one Lagrangian target defined. For purposes of computing the overall merit function value, ZEMAX will use the absolute value of the weight.

For best results, use Lagrangian multipliers sparingly, if at all. Better optimization and adequate accuracy is usually just as easily achieved using heavier weights on those operands which require exact (or nearly so) values. If a Lagrangian multiplier operand does not exactly converge with a modest weight, such as -1, try a larger magnitude negative weight, such as -1000. This will allow the merit function to decrease while meeting the target exactly. Lagrangian multipliers should not be used with boundary operands, such as CTGT, because appropriate derivatives cannot always be computed for these types of operands.

Operand weights equal to zero

When the weight is zero, the optimization algorithm calculates but ignores the operand. This is very useful for computing a result that does not have a specific target, but might be used elsewhere in the merit function; or if the value is used as a check or monitored parameter.

Operand weights greater than zero

If the weight is greater than zero, then the operand will be treated as an "aberration" to be minimized along with the merit function. The vast majority of operands should have positive weights.

Merit function definition

The merit function is defined as:

$$MF^2 = \frac{\sum W_i (V_i - T_i)^2}{\sum W_i},$$

where W is the absolute value of the weight of the operand, V is the current value, T is the target value, and the subscript i indicates the operand number (row number in the spreadsheet). The sum index "i" is normally over all operands in the merit function, however the merit function listing feature (see "Merit Function Listing" on page 240) sums the user defined and default operands separately.

Optimization operands

The following tables describe the available operands. The first table is a "quick reference" guide which categorizes the operands by general subject. The second table provides a detailed description of each operand

(listed alphabetically) and states which operands use which data fields. Names shown in **bold** define data values used to define the operand. These same names will appear in the editor column headings.

OPTIMIZATION OPERANDS BY CATEGORY

Category	Related Operands
First-order optical properties	AMAG, ENPP, EFFL, EFLX, EFLY, EPDI, EXPD, EXPP, ISFN, LINV, OBSN, PIMH, PMAG, POWF, POWP, POWR, SFNO, TFNO, WFNO
Aberrations	ABCD, ANAC, ANAR, ANAX, ANAY, ANCX, ANCY, ASTI, AXCL, BIOG, BIOD, BSER, COMA, DIMX, DISA, DISC, DISG, DIST, FCGS, FCGT, FCUR, LACL, LONA, OPDC, OPDM, OPDX, OSCD, PETC, PETZ, RSCE, RSCH, RSRE, RSRH, RWCE, RWCH, RWRE, RWRH, SMIA, SPCH, SPHA, TRAC, TRAD, TRAE, TRAI, TRAR, TRAX, TRAY, TRCX, TRCY, ZERN
MTF data	GMTA, GMTS, GMTT, MSWA, MSWS, MSWT, MTFA, MTFS, MTFT, MTHA, MTHS, MTHT
PSF/Strehl Ratio data	STRH
Encircled energy	DENC, DENF, ERF, GENC, GENF, XENC, XENF
Constraints on lens data	COGT, COLT, COVA, CTGT, CTLT, CTVA, CVGT, CVLT, CVVA, DMGT, DMLT, DMVA, ETGT, ETLT, ETVA, FTGT, FTLT, MNCA, MNCG, MNCT, MNCV, MNEA, MNEG, MNET, MNP, MXCA, MXCG, MXCT, MXCV, MXEA, MXEG, MXET, MNSD, MXSD, TGTH, TTGT, TTHI, TTLT, TTVA, XNEA, XNET, XNEG, XNEA, XNEG, XXET, ZTHI
Constraints on lens properties	CVOL, MNDT, MXDT, SAGX, SAGY, SSAG, STHI, TMAS, TOTR, VOLU, NORX, NORZ, NORD, SCUR
Constraints on parameter data	PMGT, PMLT, PMVA.
Constraints on extra data	XDGT, XDLT, XDVA
Constraints on glass data	GCOS, GTCE, INDX, MNAB, MNIN, MNP, MXAB, MXIN, MXP, RGLA
Constraints on paraxial ray data	PANA, PANB, PANC, PARA, PARB, PARC, PARR, PARX, PARY, PARZ, PATX, PATY, YNIP
Constraints on real ray data	CENX, CENY, CNAX, CNAY, CNPX, CNPY, DXDX, DXDY, DYDX, DYDY, HHCN, IMAE, OPTH, PLEN, RAED, RAEN, RAGA, RAGB, RAGC, RAGX, RAGY, RAGZ, RAID, RAIN, RANG, REAA, REAB, REAC, REAR, REAX, REAY, REAZ, RENA, RENB, RENC, RETX, RETY
Constraints on element positions	GLCA, GLCB, GLCC, GLCR, GLCX, GLCY, GLCZ
Changing system data	CONF, IMSF, PRIM, SVIG, WLEN, CVIG, FDMO, FDRE
General math operands	ABSO, ACOS, ASIN, ATAN, CONS, COSI, DIFF, DIVB, DIVI, EQUA, LOGE, LOGT, MAXX, MINN, OPGT, OPLT, OPVA, OSUM, PROB, PROD, QSUM, RECI, SQRT, SUMM, SINE, TANG, ABGT, ABLT
Multi-configuration (zoom) data	CONF, MCOL, MCOG, MCOV, ZTHI
Gaussian beam data	GBPD, GBPP, GBPR, GBPS, GBPW GBSD, GBSP, GBSR, GBSS, GBSW
Gradient index control operands	DLTN, GRMN, GRMX, InGT, InLT, InVA, LPTD
Foucault analysis	FOUC
Ghost focus control	GPIM, GPRT, GPRX, GPRY, GPSX, GPSY

Category	Related Operands
Fiber coupling operands	FICL, FICP, POPD
Relative illumination operand	RELI, EFNO
Optimization with ZPL macros	ZPLM
User defined operands	UDOP
Merit function control operands	BLNK, DMFS, ENDX, GOTO, OOFF, SKIN, SKIS, USYM
Constraints on non-sequential object data.	FREZ, NPGT, NPLT, NPVA, NPXG, NPXL, NPXV, NPYG, NPYL, NPYV, NPZG, NPZL, NPZV, NSRM, NTXG, NTXL, NTXV, NTYG, NTYL, NTYV, NTZG, NTZL, NTZV
Non-sequential ray tracing and detector operands.	NSDC, NSDD, NSDE, NSDP, NSRA, NSST, NSTR
Constraints on construction optics for optically fabricated holograms	CMFV
Constraints on optical coatings, polarization ray trace data	CMGT, CMLT, CMVA, CODA, CEGT, CELT, CEVA, CIGT, CILT, CIVA
Physical Optics Propagation (POP) results	POPD, POPI
Best Fit Sphere data	BFSD
Tolerance sensitivity data	TOLR
Thermal Coefficient of Expansion data	TCGT, TCLT, TCVA

OPTIMIZATION OPERANDS AND DATA FIELD USAGE

NAME	Description
ABCD	The ABCD values used by the grid distortion feature to compute generalized distortion. See “Grid Distortion” on page 184. The reference field number is defined by Ref Fid . The wavelength number is defined by Wave . Data is 0 for A, 1 for B, 2 for C, and 3 for D. See also “DISA” on page 496.
ABGT	Absolute value of operand greater than. This is used to make the absolute value of the operand defined by Op# greater than the target value.
ABLT	Absolute value of operand less than. This is used to make the absolute value of the operand defined by Op# less than the target value.
ABSO	Absolute value of the operand defined by Op# .
ACOS	Arccosine of the value of the operand defined by Op# . If Flag is 0, then the units are radians, otherwise, degrees.
AMAG	Angular magnification. This is the ratio of the image to object space paraxial chief ray angles at the wavelength defined by Wave . Not valid for non-paraxial systems.
ANAC	Angular aberration radial direction measured in image space with respect to the centroid at the wavelength defined by Wave . This quantity is defined as: $\varepsilon = \sqrt{(l - l_c)^2 + (m - m_c)^2}$ where <i>l</i> and <i>m</i> are the x and y direction cosines of the ray and the <i>c</i> subscript indicates the centroid. See “Hx, Hy, Px, and Py” on page 488.

NAME	Description
ANAR	Angular aberration radius measured in image space at the wavelength defined by Wave with respect to the primary wavelength chief ray. This quantity is defined as: $\varepsilon = \sqrt{(l - l_c)^2 + (m - m_c)^2}$ where <i>l</i> and <i>m</i> are the x and y direction cosines of the ray and the <i>c</i> subscript indicates the chief ray. See “Hx, Hy, Px, and Py” on page 488.
ANAX	Angular aberration x direction measured in image space at the wavelength defined by Wave with respect to the primary wavelength chief ray. This quantity is defined as: $\varepsilon = l - l_c$ where <i>l</i> is the x direction cosine of the ray and the <i>c</i> subscript indicates the chief ray. See “Hx, Hy, Px, and Py” on page 488.
ANAY	Angular aberration y direction measured in image space at the wavelength defined by Wave with respect to the primary wavelength chief ray. This quantity is defined as: $\varepsilon = m - m_c$ where <i>m</i> is the y direction cosine of the ray and the <i>c</i> subscript indicates the chief ray. See “Hx, Hy, Px, and Py” on page 488.
ANCX	Angular aberration x direction measured in image space at the wavelength defined by Wave with respect to the centroid. This quantity is defined as: $\varepsilon = l - l_c$ where <i>l</i> is the x direction cosine of the ray and the <i>c</i> subscript indicates the centroid. ANCX has the same restrictions that TRAC does; see TRAC for a detailed discussion. See “Hx, Hy, Px, and Py” on page 488.
ANCY	Angular aberration y direction measured in image space at the wavelength defined by Wave with respect to the centroid. This quantity is defined as: $\varepsilon = m - m_c$ where <i>m</i> is the y direction cosine of the ray and the <i>c</i> subscript indicates the centroid. ANCY has the same restrictions that TRAC does; see TRAC for a detailed discussion. See “Hx, Hy, Px, and Py” on page 488.
ASIN	Arcsine of the value of the operand defined by Op# . If Flag is 0, then the units are radians, otherwise, degrees.
ASTI	Astigmatism in waves contributed by the surface defined by Surf at the wavelength defined by Wave . If Surf is zero, the sum for the entire system is used. This is the third order astigmatism calculated from the Seidel coefficients, and is not valid for non-paraxial systems.
ATAN	Arctangent of the value of the operand defined by Op# . If Flag is 0, then the units are radians, otherwise, degrees.
AXCL	Axial color, measured in lens units for focal systems and diopters for afocal systems. This is the image separation between the two wavelengths defined by Wave1 and Wave2 . If Zone is zero, paraxial rays are used to determine the paraxial image locations. If Zone is greater than 0.0 and less than or equal to 1.0, real marginal rays are used to determine the image locations. In this case, Zone corresponds to the Py coordinate of the real marginal ray. See “Hx, Hy, Px, and Py” on page 488.

NAME	Description
BFSD	<p>Best Fit Sphere (BFS) data. For a complete description, see “Sag Table” on page 208. The data is computed for the surface defined by Surf. The value returned depends upon the value of Data as follows:</p> <ul style="list-style-type: none"> 0 - BFS curvature in inverse lens units 1 - BFS radius of curvature in lens units 2 - Vertex offset from BFS in lens units 3 - Maximum depth of material to remove in lens units 4 - Total volume of material to use in cubic lens units 5 - The maximum slope (dz/dr) difference between the surface and BFS (dimensionless) <p>The MinR and MaxR are the minimum and maximum radial coordinate over which the BFS data is computed. If both values are zero, the maximum radial coordinate used is the semi-diameter of the surface.</p>
BIOC	<p>Biocular Convergence. Returns the convergence between two eye configurations in milliradians. The left and right eye configurations are specified using the Left and Right values. The other parameters are:</p> <p>Wave: The wavelength number to use.</p> <p>UseCos: If 0 field units are degrees, otherwise field is in direction cosine units.</p> <p>Xang/Yang: The X and Y angle or cosines at which to compute the convergence.</p> <p>If the chief rays from both configurations at the specified angles do not pass through to the image without vignetting, an error is reported. See “Dipvergence/Convergence” on page 181 for more information and important assumptions.</p>
BIOD	<p>Biocular Dipvergence. Returns the dipvergence between two eye configurations in milliradians. See BIOC above for details.</p>
BIPF	<p>Unused.</p>
BLNK	<p>Does nothing. Used for separating portions of the operand list. A comment line may optionally be typed in the space to the right of the operand name; this comment will be displayed in the editor as well as in the merit function listing.</p>
BSER	<p>Boresight error. Boresight error is defined as the radial chief ray coordinate traced for the on axis field and wavelength defined by Wave divided by the effective focal length. This definition yields a measure of the angular deviation of the image.</p>
CEGT	<p>Boundary operand that constrains the coating extinction offset of the coating layer defined by Layr for the surface defined by Surf to be greater than the target value. Use Surf = 0 for all surfaces, Layr = 0 for all layers.</p>
CELT	<p>Boundary operand that constrains the coating extinction offset of the coating layer defined by Layr for the surface defined by Surf to be less than the target value. Use Surf = 0 for all surfaces, Layr = 0 for all layers.</p>
CENX	<p>Centroid X position. This operand uses a grid of rays to determine the x coordinate of the centroid of all rays from a single field point. The centroid accounts for apodization and apertures, and optionally polarization. The parameters for this operand are:</p> <p>Surf: The surface number to use (use 0 for the image surface).</p> <p>Wave: The wavelength number to use (0 for polychromatic, otherwise use the monochromatic wavelength number).</p> <p>Field: The integer field number to use.</p> <p>Pol?: Set to 0 to ignore polarization and 1 to consider it.</p> <p>Samp: The grid size. A value of 10 would yield a 10 x 10 grid of rays.</p> <p>When CENX is followed by a CENY with identical settings, both are computed with the same ray set to save time.</p> <p>See also CNPX, CNPY, CNAX, and CNAY.</p>
CENY	<p>Centroid Y position. See CENX.</p>

NAME	Description
CEVA	Boundary operand that constrains the coating extinction offset of the coating layer defined by Layr for the surface defined by Surf to be equal to the target value.
CIGT	Boundary operand that constrains the coating index offset of the coating layer defined by Layr for the surface defined by Surf to be greater than the target value. Use Surf = 0 for all surfaces, Layr = 0 for all layers.
CILT	Boundary operand that constrains the coating index offset of the coating layer defined by Layr for the surface defined by Surf to be less than the target value. Use Surf = 0 for all surfaces, Layr = 0 for all layers.
CIVA	Boundary operand that constrains the coating index offset of the coating layer defined by Layr for the surface defined by Surf to be equal to the target value.
CMFV	Construction merit function value. This operand calls the merit functions defined in either of the two construction systems used to define an optically fabricated hologram. The Cons# is either 1 or 2, for the first or second construction system, respectively. The Op# is either 0, which will return the entire merit function value from the construction system; or it is an integer which defines the operand row # from which to extract the value from. For example, if Cons# is 2 and Op# is 7, CMFV will return the value of merit function operand 7 in construction file 2. If there are more than one optically fabricated hologram surfaces in the playback system being optimized, the Cons# may be incremented by 2 to specify the second surface parameters be used, or by 4 to indicate the third hologram surface construction optics be used, and so on. For example, a Cons# of 7 would indicate construction system 1 on the fourth optically fabricated hologram surface present.
CMGT	Boundary operand that constrains the coating multiplier of the coating layer defined by Layr for the surface defined by Surf to be greater than the target value. Use Surf = 0 for all surfaces, Layr = 0 for all layers.
CMLT	Boundary operand that constrains the coating multiplier of the coating layer defined by Layr for the surface defined by Surf to be less than the target value. Use Surf = 0 for all surfaces, Layr = 0 for all layers.
CMVA	Boundary operand that constrains the coating multiplier of the coating layer defined by Layr for the surface defined by Surf to be equal to the target value.
CNAX	Centroid angular x direction. This operand computes the X angle in radians relative to the local Z axis of the centroid of rays from any field point. The centroid accounts for apodization and apertures, and optionally polarization. Surf : The surface number to use (use 0 for the image surface). Wave : The wavelength number to use (0 for polychromatic, otherwise use the monochromatic wavelength number). Hx/Hy : The normalized field coordinates to use. See “Hx, Hy, Px, and Py” on page 488. Pol? : Set to 0 to ignore polarization and 1 to consider it. Samp : The grid size. A value of 10 would yield a 10 x 10 grid of rays. When CNAX is followed by a CNAY with identical settings, both are computed with the same ray set to save time. See also CNAY, CNPX, CNPY, CENX, CENY.
CNAY	Centroid angular y direction. See CNAX.
CNPX	Similar to CNAX, but computes the centroid position rather than angle.
CNPY	Similar to CNAY, but computes the centroid position rather than angle.

NAME	Description
CODA	<p>Coating Data. This feature traces a polarized ray using the system global polarization state (see "Polarization (ZEMAX-EE only)" on page 108). Even if the global polarization state is set to "unpolarized", the defined polarization state is used because the data computed by CODA is polarization state specific.</p> <p>The polarized ray may be traced from any field point in object space defined by Field to any point in the pupil, to the surface defined by Surf. If Surf is zero, the ray is traced to the image surface. The wavelength is defined by Wave while Px and Py define the normalized pupil coordinate. The absolute value of Data determines what data is returned as follows:</p> <p>0: The relative transmitted polarized intensity (see 8) 1, 2, 3: The intensity reflectance R, transmittance T, absorption A 4,5: The field amplitude transmittance real, imaginary 6,7: The field amplitude reflectance real, imaginary 8: The relative transmitted unpolarized intensity (see 0) 101, 102: E field out X real, imaginary 103, 104: E field out Y real, imaginary 105, 106: E field out Z real, imaginary 110: The phase difference between Ex and Ey; Pxy 111, 112, 113: The E field phase Px, Py, Pz</p> <p>If the data is related to the coating (1-7), and Data is negative, the data is for the "S" polarization, otherwise the data is for the "P" polarization.</p> <p>Other coating and polarization data may be added to this list upon request. See "Hx, Hy, Px, and Py" on page 488.</p>
COGT	Boundary operand that constrains the conic of the surface defined by Surf to be greater than the specified target value.
COLT	Boundary operand that constrains the conic of the surface defined by Surf to be less than the specified target value.
COMA	Coma in waves contributed by the surface defined by Surf at the wavelength defined by Wave . If Surf is zero, the sum for the entire system is used. This is the third order coma calculated from the Seidel coefficients, and is not valid for non-paraxial systems.
CONF	Configuration. This operand is used to change the configuration number to the configuration defined by Cfg# during merit function evaluation. This permits optimization across multiple configurations with a single merit function. This operand does not use the target or weight columns.
CONS	Constant value. This is used to enter in constant values for use in other operand computations. The value will be identical to the target value.
COSA	Unused.
COSI	Cosine of the value of the operand defined by Op# . If Flag is 0, then the units are radians, otherwise, degrees.
COVA	Conic value. Returns the conic constant of the surface defined by Surf .
CTGT	Center thickness greater than. This boundary operand constrains the center thickness of the surface defined by Surf to be greater than the specified target value. See also MNCT.
CTLT	Center thickness less than. This boundary operand constrains the center thickness of the surface defined by Surf to be less than the specified target value. See also MXCT.
CTVA	Center thickness value. Constrains the center thickness of the surface defined by Surf to be equal to the specified target value.
CVGT	Curvature greater than. This boundary operand constrains the curvature of the surface defined by Surf to be greater than the target value.
CVIG	Clears the vignetting factors. This operand clears the current vignetting factors for the remainder of the operands in the current configuration. See also "SVIG" on page 520.

NAME	Description
CVLT	Curvature less than. This boundary operand constrains the curvature of the surface defined by Surf to be less than the target value.
CVOL	Cylinder volume. This operand computes the volume in cubic lens units of the smallest cylinder that will contain the range of surfaces defined by Surf1 and Surf2 . Only the vertex positions and semi-diameters are used in the calculation, not the sag. The range of surfaces should not include any coordinate breaks.
CVVA	Curvature value. This operand constrains the curvature of the surface defined by Surf to be equal to the specified target value.
DENC	<p>Diffraction Encircled Energy (distance). This operand computes the distance to the fraction of diffraction encircled, ensquared, x only, or y only (enslitted) energy defined by Frac. For focal mode, the units are micrometers. For afocal mode, the units are afocal mode units. The other parameters are:</p> <p>Samp: The pupil sampling, where 1 yields 32 x 32, 2 yields 64 x 64 etc.</p> <p>Wave: The wavelength number to use (use 0 for polychromatic).</p> <p>Field: The field number.</p> <p>Type: 1 for encircled, 2 for x only, 3 for y only, and 4 for ensquared.</p> <p>Refp: The reference point/algorithm to use. For FFT encircled energy, use Refp = 0 for chief ray, 1 for centroid, and 2 for vertex. For Huygens encircled energy, use Refp = 3 for chief ray, 4 for centroid, and 5 for vertex. When using the Huygens method, I Samp refers to the Huygens image sampling (1 for 32 x 32, 2 for 64 x 64, etc.) and I Delta refers to the Huygens image delta. For a detailed description of the Huygens image sampling and image delta, see “Encircled Energy” on page 158.</p> <p>If the sampling is too low, the radius returned is 1e+10. See also DENC, GENC and XENC.</p>
DENF	<p>Diffraction Encircled Energy (fraction). This operand computes the fraction of diffraction encircled, ensquared, x only, or y only (enslitted) energy at a given distance from the reference point defined by Dist. For focal mode, the distance units are micrometers. For afocal mode, the distance units are afocal mode units.</p> <p>The options and settings are identical to DENC, except Dist, which here is used as the distance at which the fraction of energy is desired. See also DENC, GENC, GENF, and XENC.</p>
DIFF	Difference of two operands (Op#1 - Op#2).
DIMX	<p>Distortion maximum. This is similar to DIST, except it specifies only an upper bound for the absolute value of the distortion. Field can be zero, which specifies that the maximum field coordinate be used, or any valid field number. Note the maximum distortion does not always occur at the maximum field coordinate. The distortion is calculated at the wavelength defined by Wave.</p> <p>If Absolute is 0, the value returned is in units of percentage. If Absolute is 1, the distortion is given as an absolute length rather than a percentage.</p> <p>This operand may not be valid for non-rotationally symmetric systems.</p>
DISA	<p>Distortion, ABCD. This operand computes the Radial, X, or Y direction distortion relative to the reference field (Ref Fld), for the chief ray at the wavelength defined by Wave. Data is 0 for radial distortion, 1 for X direction distortion, and 2 for Y direction distortion. The distortion is computed for the chief ray at the field point defined by Field. The A, B, C, D values are user defined. The distortion is computed in the same manner as the grid distortion feature (see “Grid Distortion” on page 184). The key difference between this operand and DISG is that the ABCD values are user defined. See also “ABCD” on page 491 and “DISG” on page 497.</p>
DISC	<p>Distortion, calibrated. This operand computes the calibrated f-theta distortion across the y-field of view at the wavelength defined by Wave, and returns the absolute value of the maximum deviation from linearity of the f-theta condition.</p> <p>If Absolute is 0, the value returned is in units of percentage. If Absolute is 1, the distortion is given as an absolute length (for focal systems) or an absolute deviation in cosine space (for afocal systems) rather than a percentage. This operand is useful for designing f-theta lenses.</p>

NAME	Description
DISG	<p>Generalized distortion, either in percent or as an absolute distance. This operand computes the distortion for any ray in the pupil, from anywhere in the field, at the wavelength defined by Wave, using the field point defined by Field as a reference. The method used and assumptions made are the same as for the grid distortion plot, see “Grid Distortion” on page 184. DISG cannot be calculated if the field units are angles and the maximum angle equals or exceeds 90 degrees. DISG assumes the predicted magnification is not symmetric.</p> <p>If the field is defined in terms of angles, the normalized field coordinates Hx and Hy are defined differently than for other ZEMAX features. In this case only, Hx and Hy correspond to divisions on a uniform projected grid in object space, centered on the reference field coordinate. These normalized field coordinates are defined as</p> $H = \frac{\tan(\theta)}{\tan(\theta_M)}$ <p>where θ is the angle relative to the reference field chief ray and θ_M is the maximum field angle (see “Maximum field” on page 54). These definitions are needed to support the reference field feature.</p> <p>If Wave is a positive number, DISG returns the distortion as a percentage. If Wave is a negative number, the absolute value of Wave is used to define the wavelength and the returned value is in units of absolute length rather than percentage.</p> <p>As with all distortion concepts, the best way to avoid confusion and misleading results is to use finite object distances and object heights to define fields rather than field angles.</p> <p>See “Hx, Hy, Px, and Py” on page 488.</p>
DIST	<p>Distortion in waves contributed by the surface defined by Surf at the wavelength defined by Wave. This is the third order distortion calculated from the Seidel coefficients (see “Seidel Coefficients” on page 194), and is not valid for non-paraxial systems.</p> <p>If Surf is zero, the distortion is given in percent instead (see “Field Curvature/Distortion” on page 182 for a detailed definition).</p> <p>If Absolute is set to 1, and the surface number is zero, the distortion is given as an absolute length rather than a percentage.</p> <p>See also DISG.</p>
DIVB	<p>Divides the value of any prior operand defined by Op# by any factor defined by Factor.</p>
DIVI	<p>Division of first by second operand (Op#1 / Op#2). See also “RECI” on page 518.</p>
DLTN	<p>Delta N. Computes the difference between the maximum and minimum index of refraction on axis for a gradient index surface defined by Surf. The wavelength used is defined by Wave. The minimum and maximum z coordinates used are computed by accounting for the sag of both ends of the surface. See the section “Using gradient index operands”.</p>
DMFS	<p>Default merit function start. This operand is a flag to indicate where the default merit function should be appended, if one is subsequently created. The row number after this operand will appear as the default “Start At” row on the default merit function dialog box.</p>
DMGT	<p>Diameter greater than. This boundary operand constrains the diameter of the surface defined by Surf to be greater than the specified target value. The diameter is twice the semi-diameter value displayed on the main spreadsheet.</p>
DMLT	<p>Diameter less than. This boundary operand constrains the diameter of the surface defined by Surf to be less than the specified target value. The diameter is twice the semi-diameter value displayed on the main spreadsheet.</p>
DMVA	<p>Diameter value. This operand constrains the diameter of the surface defined by Surf to be equal to the specified target value. The diameter is twice the semi-diameter value displayed on the main spreadsheet.</p>
DXDX	<p>Derivative of transverse x-aberration with respect to x-pupil coordinate. This is the slope of the ray fan plot at the specified pupil coordinate at the wavelength defined by Wave.</p> <p>See “Hx, Hy, Px, and Py” on page 488.</p>

NAME	Description
DXDY	Derivative of transverse x-aberration with respect to y-pupil coordinate. This is the slope of the ray fan plot at the specified pupil coordinate at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.
DYDX	Derivative of transverse y-aberration with respect to x-pupil coordinate. This is the slope of the ray fan plot at the specified pupil coordinate at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.
DYDY	Derivative of transverse y-aberration with respect to y-pupil coordinate. This is the slope of the ray fan plot at the specified pupil coordinate at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.
EFFL	Effective focal length in lens units. The wavelength used is defined by Wave . This is paraxial, and may not be accurate for non-paraxial systems.
EFLX	Effective focal length in the local x plane of the range of surfaces defined by Surf1 and Surf2 at the primary wavelength.
EFLY	Effective focal length in the local y plane of the range of surfaces defined by Surf1 and Surf2 at the primary wavelength.
EFNO	Effective F/#. This operand computes the Effective F/# of the field point defined by Field . For a discussion of Effective F/# see "Effective F/#" on page 188. The other parameters are: Samp : The grid size. A value of 10 would yield a 10 x 10 grid of rays. Wave : The wavelength number to use. Pol? : Set to 0 to ignore polarization and 1 to consider it. See also RELI.
ENDX	End execution. Terminates the computation of the merit function. All remaining operands are ignored.
ENPP	Entrance pupil position in lens units, with respect to the first surface. This is the paraxial pupil position, valid only for centered systems.
EPDI	Entrance pupil diameter in lens units.
ERFP	Edge Response Function Position. This operand computes the x or y position of the point at which the edge response function reaches a certain relative value. For details on the edge response function calculation, see "Geometric Line/Edge Spread" on page 160. The Sampling value is 1 for 32x32, 2 for 64x64, etc. Wave is the wavelength number or 0 for polychromatic. Field is the field number. Type determines that data to be returned. If Type is 0 or 1, the x position (for edges parallel to the local y axis) or y position (for edges parallel to the local x axis) relative to the chief ray in lens units is returned. If Wave is zero, the primary wavelength chief ray is the reference point. If Type is 2 or 3, the x or y position in lens units relative to the surface vertex is returned. Fraction is the relative value of the edge response curve, and must be between 0.01 and 0.99. Max Radius is the maximum radial size of the integration window in micrometers. If Max Radius is zero a default value is used, and this is the recommended setting in most cases. Note that the edge response function is normally defined with the "bright" side of the edge being on the + side of the integration coordinate, which means the edge response goes to 1 as the coordinate becomes more positive. To compute results for the reversed edge orientation, with the bright side on the negative side of the coordinate, use the value (1-fraction) instead of (fraction). For example, to get the 80% response coordinate for a reversed edge, use Fraction = 0.20. If afocal mode is used, all returned values are in afocal analysis units.
EQUA	Equal operand. This operand constrains all operands within the range of operands defined by Op#1 and Op#2 to have the same value within the tolerance specified by the target. The value of this operand is computed by finding the average of the range of values, and then summing the absolute value of the errors between each operand and the average if the error exceeds the target value. See SUMM and OSUM.

NAME	Description
ETGT	Edge thickness greater than. This boundary operand constrains the edge thickness of the surface defined by Surf to be greater than the specified target value. The edge thickness is calculated along the +y axis if Code is zero, the +x axis if Code is 1, the -y axis if Code is 2, and the -x axis if Code is 3. See also MNET.
ETLT	Edge thickness less than. This boundary operand constrains the edge thickness of the surface defined by Surf to be less than the specified target value. The edge thickness is calculated along the +y axis if Code is zero, the +x axis if Code is 1, the -y axis if Code is 2, and the -x axis if Code is 3. See also MXET.
ETVA	Edge thickness value. Constrains the edge thickness of the surface defined by Surf to be equal to the specified target value. The edge thickness is calculated along the +y axis if Code is zero, the +x axis if Code is 1, the -y axis if Code is 2, and the -x axis if Code is 3. See also MNET.
EXPD	Exit pupil diameter in lens units. This is the paraxial pupil diameter, valid only for centered systems.
EXPP	Exit pupil position in lens units, with respect to the image surface. This is the paraxial pupil position, valid only for centered systems.
FCGS	Generalized field curvature, sagittal. The field curvature value for any field point, at the wavelength defined by Wave . The value is generalized to return reasonable results even for non-rotationally symmetric systems; see "Field Curvature/Distortion" on page 182. See "Hx, Hy, Px, and Py" on page 488.
FCGT	Generalized field curvature, tangential; see FCGS.
FCUR	Field curvature in waves contributed by the surface defined by Surf at the wavelength defined by Wave . If Surf is zero, the sum for the entire system is used. This is the third order field curvature calculated from the Seidel coefficients, and is not valid for non-paraxial systems.
FDMO	Field Data Modify. This operand allows temporary modification of field position data. The operand allows new field coordinates and vignetting factors (except for vignetting angle, which cannot be modified by this operand) for any field coordinate. All subsequent operands will use the modified field data. The original field data is restored when a FDRE operand is reached with the same field number argument, or a CONF operand is reached (regardless of the configuration number referenced by the CONF operand), or after the last operand in the merit function is reached. Field : The field number to modify. Hx, Hy : The normalized field coordinates for the new field position. VDX, VDY, VCX, VCY : The vignetting factor data for the new field position. See "Vignetting factors" on page 115.
FDRE	Field Data Restore. See FDMO. Field : The field number to restore.
FICL	Fiber coupling efficiency for single mode fibers. The calculated value is the total coupled energy efficiency, relative to unity. The parameters of this operand are: Samp : The pupil sampling, where 1 yields 32 x 32, 2 yields 64 x 64 etc. Wave : The wavelength number to use. Field : The integer field position number. IgSrc : If this parameter is zero, then the object source fiber is considered; otherwise, the object source fiber is ignored. Sna : The source fiber NA. Rna : The receiver fiber NA. Data : This controls what data is returned as well as the algorithm to use. If Data is 0, the fast algorithm is used and the power coupling is returned. If Data is 1, the Huygens algorithm is used and the power coupling is returned. If Data is 2 or 4, the fast algorithm is used and the real or imaginary part of the amplitude coupling is returned, respectively. If Data is 3 or 5, the Huygens algorithm is used and the real or imaginary part of the amplitude coupling is returned, respectively. Pol? : Set to 0 to ignore polarization and 1 to consider it. See "Fiber Coupling Efficiency" on page 206 for details. See also FICP.

NAME	Description
FICP	<p>Fiber coupling as computed using the Physical Optics Propagation (POP) algorithm, using whatever the current default settings are for the POP feature.</p> <p>Data: This controls what data is returned. If Data is 0 the total power coupling is returned. If Data is 1, the amplitude coupling for the Ex field is returned. If Data is 2 or 3, the real or imaginary part of the amplitude coupling for the Ex field is returned, respectively. If Data is 4, the amplitude coupling for the Ey field is returned. If Data is 5 or 6, the real or imaginary part of the amplitude coupling for the Ey field is returned, respectively. If the beam is unpolarized the Ey values will be zero.</p> <p>To use this operand, first define the settings on the POP analysis feature as desired, then press Save on the settings box. The operand FICP will return the same efficiency as computed by the POP feature.</p> <p>This operand is redundant with the more general POPD. See "Computing Fiber Coupling" on page 645. See also FICL.</p>
FOUC	<p>Foucault analysis. This operand returns the RMS difference between the computed and reference shadowgram as computed by the Foucault analysis feature, using whatever the current default settings are. To use this operand, first define the settings on the Foucault analysis feature as desired, then press Save on the settings box. The operand FOUC will return the RMS difference between the computed and reference shadowgrams, regardless of how the "data" option is set in the saved configuration file. Using this operand, the optical system wavefront aberrations may be optimized to produce the reference shadowgram.</p>
FREZ	<p>Freeform Z object constraints. This operand computes various values from the shape of a Freeform Z object described in "Freeform Z" on page 506. Surf is the surface number of the NSC group, use 1 for nonseq program mode. Object is the object number, which must be a Freeform Z object for this operand to compute any data. The Data value determines what value is computed by the operand as follows:</p> <ol style="list-style-type: none"> 1: Maximum z value. Note minimum z value is always zero. 2: Maximum z increment. This is the biggest difference between adjacent z control points. 3: Minimum z increment. This is the smallest difference between adjacent z control points. 4: Minimum y value. This value is for the entire length of the solid, not just at control points. 5: Maximum y value. This value is for the entire length of the solid, not just at control points. 6: Maximum y increment. This is the biggest difference between adjacent y control points. 7: Minimum y increment. This is the smallest difference between adjacent y control points. 8: Volume of the object in lens units cubed. 9: Monotonic deviation. This is the largest change in adjacent y values whose sign deviates from the sign of the first y change going from 0 to positive z coordinates. 10: Minimum slope. 11: Maximum slope. <p>Once the Data has been computed, the Mode value determines how the operand behaves. If Mode is 1, the FREZ operand returns the data. If Mode is 2, the operand returns the data value only if the data value is less than the target value, otherwise the target value is returned. This mode is used to enforce a "greater than" boundary on the data. If Mode is 3, the operand returns the data value only if the data value is greater than the target value, otherwise the target value is returned. This mode is used to enforce a "less than" boundary on the data.</p>
FTGT	<p>Full thickness greater than. This boundary operand constrains the full thickness of surface Surf to be greater than the specified target value. The full thickness is computed at 200 points between the vertex and edge along the +y radial direction, including the sag of the surface and the sag of the next surface. The operand is useful for constraining surfaces which do not have their minimum or maximum thickness at the center or edge, but at some intermediate zone. See FTLT.</p>
FTLT	<p>Full thickness less than. See FTGT.</p>

NAME	Description
GBPD	Gaussian beam (paraxial) divergence in the optical space following the surface defined by Surf at the wavelength defined by Wave . The other parameters are: UseX : If this parameter is non-zero, then the computation is for the x-direction beam, otherwise, it is for the y-direction. W0 : The input beam waist size in lens units. S1toW : The distance from Surface 1 to the waist location in lens units. See the Gaussian beam feature for details.
GBPP	Gaussian beam (paraxial) position, which is the distance from the waist to the surface, in the optical space following the specified surface. See GBPD.
GBPR	Gaussian beam (paraxial) radius of curvature in the optical space following the specified surface. See GBPD.
GBPS	Gaussian beam (paraxial) size in the optical space following the specified surface. See GBPD.
GBPW	Gaussian beam (paraxial) waist in the optical space following the specified surface. See GBPD.
GBSD	Gaussian beam (skew) divergence in the optical space following the specified surface. The input beam is aligned along the chief ray of the field defined by Field . The other parameters are: In# : The surface number to start the propagation. Out# : The surface at which to compute the Skew Gaussian Beam data. Wave : The wavelength number to use. StoW : The starting surface to waist distance in lens units. W0 : The input beam waist size in lens units. If W0 is positive, then the computation is for the y-direction beam, otherwise, it is for the x-direction. For details on the Skew Gaussian Beam feature, see “Skew Gaussian Beam” on page 225.
GBSP	Gaussian beam (skew) position, which is the distance from the waist to the surface, in the optical space following the specified surface. See GBSD.
GBSR	Gaussian beam (skew) radius in the optical space following the specified surface. See GBSD.
GBSS	Gaussian beam (skew) size in the optical space following the specified surface. See GBSD.
GBSW	Gaussian beam (skew) waist in the optical space following the specified surface. See GBSD.
GCOS	Glass cost. This operand returns the relative cost factor as listed in the glass catalog for the glass on the surface defined by Surf .
GENC	Geometric Encircled Energy (distance). This operand computes the distance to the fraction of geometric encircled, ensquared, x only, or y only (enslitted) energy defined by Frac . For focal mode, the units are micrometers. For afocal mode, the units are afocal mode units. The other parameters are: Samp : The pupil sampling, where 1 yields 32 x 32, 2 yields 64 x 64 etc. Wave : The wavelength number to use (use 0 for polychromatic). Field : The field number. Type : 1 for encircled, 2 for x only, 3 for y only, and 4 for ensquared. Refp : The reference point to use. Use 0 for chief ray, 1 for centroid, 2 for vertex, and 3 for middle of the spot. GENC always scales by the diffraction limited data. See also GENF, DENC, DENF, and XENC.
GENF	Geometric Encircled Energy (fraction). This operand computes the fraction of geometric encircled, ensquared, x only, or y only (enslitted) energy at a given distance from the reference point defined by Dist . The options and settings are identical to GENC, except Dist , which here is used as the distance at which the fraction of energy is desired. See also GENC, DENC, DENF, and XENC.
GLCA	Global x-direction orientation vector component of the surface defined by Surf .

NAME	Description
GLCB	Global y-direction orientation vector component of the surface defined by Surf .
GLCC	Global z-direction orientation vector component of the surface defined by Surf .
GLCR	Global Coordinate Rotation Matrix component at the surface defined by Surf . The 3 x 3 R matrix has 9 components. If Data is 1, GLCR returns R[1][1], if Data is 2, GLCR returns R[1][2], etc... through Data = 9 returning R[3][3].
GLCX	Global vertex x-coordinate of the surface defined by Surf .
GLCY	Global vertex y-coordinate of the surface defined by Surf .
GLCZ	Global vertex z-coordinate of the surface defined by Surf .
GMTA	<p>Geometric MTF average of sagittal and tangential response. The parameters are:</p> <p>Samp: Higher sampling yields higher accuracy at the expense of computation time. To confirm the computation has acceptable accuracy, start at 1, and increment the sampling until the results change by less than the desired accuracy. Note that extreme precision is not required for good optimization results; three significant figures is usually adequate.</p> <p>Wave: The wavelength number to use (use 0 for polychromatic).</p> <p>Field: The field number.</p> <p>Freq: The spatial frequency in MTF units (see "MTF Units" on page 106).</p> <p>!Sci: If zero, then the diffraction limit will be used to scale the results (recommended) otherwise, no scaling is done.</p> <p>If Grid is zero, a fast, sparse sampling integration method is used to compute the MTF. The fast Geometric MTF algorithm is only accurate for systems with circular or elliptical pupils with modest or no apodization. For systems that violate this assumption, set Grid equal to 1. The fast sampling method used by GMTA, GMTS, and GMTT is not directly related to the Geometric MTF Analysis feature. Because only a single spatial frequency is required, the method of computation used by the MTF operands is different, and generally much faster, than the algorithm used by the analysis feature. To select the alternate grid-based algorithm used by the Geometric MTF analysis feature, set Grid equal to 1. The grid-based algorithm is usually slower than the default algorithm if the MTF is reasonably good (greater than 5%), but the grid algorithm converges faster if the aberrations are very large and the resulting MTF is very low.</p> <p>If both the tangential and sagittal MTF are needed; place the GMTT and GMTS operands on adjacent lines and they will be computed simultaneously. The Geometric MTF, though approximate, can usually be computed more quickly than the Diffraction MTF, and is therefore useful for optimization. See "Using MTF operands" on page 525.</p>
GMTS	Geometric MTF sagittal response. See GMTA.
GMTT	Geometric MTF tangential response. See GMTA.
GOTO	Skips all operands between the GOTO operand line and the operand number defined by Op# . Execution of the merit function will start again at the Op# line.

NAME	Description
GPIM	<p>Ghost pupil image. GPIM controls the location of ghost pupils (and optionally ghost images) relative to the image surface. Double-bounce ghosts form images of the pupil, and if these images are formed near the focal plane they will contaminate the image with unwanted light. This is the cause of the familiar “sun flare” images of the pupil seen through camera lenses pointed near the sun.</p> <p>The operand computes any one specific or all possible ghost pupil image locations and returns one over the absolute value of the distance from the image surface to the closest pupil ghost. The operand is defined in this manner so it can be targeted to zero and weighted and optimized to reduce ghost pupil affects. If the Surf1 and Surf2 parameters are set to specific surface numbers, that specific ghost path is computed, if either or both of the Surf1 and Surf2 values are -1, then all possible surface combinations are considered. For example, if Surf1 is 12 and Surf2 is -1, then all double bounces that first bounce off surface 12 and then 11, 10, 9, etc. are considered. If both numbers are negative, all possible ghosts are considered.</p> <p>This same operand also can be used for detecting and controlling image ghosts (which are distinct from pupil ghosts) by changing Mode from 0 to 1, or to control ghost pupil magnification (the ratio of the ghost exit pupil diameter to entrance pupil diameter), by setting Mode to 2.</p> <p>The WFB and WSB columns will list the worst combination found for reference and possible further analysis. Only surfaces with index changes are considered as possible ghost generators. First bounces off mirrors are ignored.</p> <p>See also GPRT, GPRX, GPRY, GPSX, and GPSY.</p>
GPRT	<p>Ghost ray transmission. This operand computes the unpolarized transmission of any ray from object to image surface for a double-bounce ghost path. The Surf1 and Surf2 values define the first and second surface numbers to define the ghost path. Note the first bounce surface number must be larger than the second bounce surface number. The primary wavelength is always used.</p> <p>GPRT makes no changes to the field or aperture specifications, and therefore may return meaningless results if the field or aperture data are defined in a way that changes the ghost path. No warning is issued in these cases. For example, using GPRT while using image height for field points or using image space F/# for the system aperture type will likely return useless results. The recommended approach is to use field angles or object heights, use entrance pupil diameter, and place the stop prior to the first bounce surface. To study the double bounce system in detail see “Ghost Focus Generator” on page 263.</p> <p>See “Hx, Hy, Px, and Py” on page 488.</p>
GPRX	<p>Ghost ray real x coordinate. This operand is similar to GPRT, see that discussion for assumptions and limitations. The returned value is the x coordinate of the real ray on the image surface.</p>
GPRY	<p>Ghost ray real y coordinate. This operand is similar to GPRT, see that discussion for assumptions and limitations. The returned value is the y coordinate of the real ray on the image surface.</p>
GPSX	<p>Ghost ray paraxial x coordinate. This operand is similar to GPRT, see that discussion for assumptions and limitations. The returned value is the x coordinate of the paraxial ray on the image surface.</p>
GPSY	<p>Ghost ray paraxial y coordinate. This operand is similar to GPRT, see that discussion for assumptions and limitations. The returned value is the y coordinate of the paraxial ray on the image surface.</p>
GRMN	<p>Gradient index minimum index. This boundary operand sets the minimum allowable index of refraction value for the gradient index surface defined by Surf at the wavelength defined by Wave. The index is checked at six places: the front vertex, front +y top, front +x side, rear vertex, rear +y top, and the rear +x side. See also InGT, InLT and GRMX.</p>
GRMX	<p>Gradient index maximum index. This boundary operand sets the maximum allowable index of refraction value for the gradient index surface defined by Surf at the wavelength defined by Wave. The index is checked at six places: the front vertex, front +y top, front +x side, rear vertex, rear +y top, and the rear +x side. See also InGT, InLT and GRMN.</p>
GTCE	<p>Glass TCE. This operand returns the Thermal Coefficient of Expansion Alpha1 as listed in the glass catalog for the glass on the surface defined by Surf. For non-glass surfaces, see “TCVA” on page 520.</p>
HACG	<p>Unused.</p>

NAME	Description
HHCN	<p>Test for the hyperhemisphere condition. ZEMAX traces the specified ray at the wavelength defined by Wave to the surface defined by Surf, and computes the x, y, and z intercept coordinates. Then, the x and y coordinates only are used in the sag expression for that surface to see what z coordinate results. If the z coordinates are not the same, then HHCN returns 1, otherwise it returns zero. This operand can be used to prevent optimizations from reaching solutions that require hyperhemispheric surface shapes.</p> <p>See "Hx, Hy, Px, and Py" on page 488.</p>
IMAE	<p>Image analysis data. This operand returns fractional efficiency, centroid, or RMS data as computed by the geometric image analysis feature, using whatever the current default settings are, except for "Show" which is always set to spot diagram for this computation, "Surface", which may be selected using Surf, "Field", which may be selected using Field, and "Field Size", which may be selected using Field Size. To use this operand, first define the settings on the geometric image analysis feature as desired, then press Save on the settings box. The operand IMAE will return the desired data as computed by the image analysis feature.</p> <p>If Surf is 0, the surface specified by the saved settings will be used. If Surf is greater than zero, then the efficiency will be computed at the specified surface.</p> <p>If Field is 0, the field number specified by the saved settings will be used. If Field is greater than zero, then the efficiency will be computed at the specified field.</p> <p>If Field Size is 0, the Field Size specified by the saved settings will be used. If Field Size is greater than zero, then the efficiency will be computed using the specified Field Size.</p> <p>Data is 0 for efficiency, 1 for X-centroid, 2 for Y-centroid, or 3-5 for X-, Y- or Radial-direction RMS, respectively. All values other than efficiency are in lens units.</p> <p>See the discussion "Optimizing with the IMAE operand" on page 532.</p>
IMSF	<p>Image Surface. This operand dynamically changes which surface is interpreted as the "image surface" for all subsequent operands. The new image surface is defined by Surface. The primary usage for this operand is to optimize image quality at intermediate surfaces. Set Surface to 0 to restore the image surface back to the original surface.</p> <p>If Refocus? is 0, the specified surface becomes the new image surface and is not modified in any way. If Refocus? is 1, then a dummy plane surface will follow the specified surface, and a paraxial focus solve will be used to place this dummy surface at the focus of the specified surface. This option is most useful for optimizing aberrations at a virtual focus for intermediate surfaces that are not already at focus.</p> <p>Note this operand only temporarily changes a copy of the lens data use for evaluating the merit function and has no impact on the original lens data. Care should be taken when the selected surface precedes the original stop surface. For information on how this operand is implemented see "Evaluating results at intermediate surfaces" on page 128.</p>
INDX	<p>Index of refraction. Returns the current index at the surface defined by Surf at the wavelength defined by Wave.</p>
InGT	<p>Index "n" greater than. This boundary operand constrains the index of refraction at the wavelength defined by Wave of the gradient index surface defined by Surf at one of six points inside the gradient index lens. For n=1, the point is the front vertex, n=2 is the front +y top, n=3 is the front +x side, n=4 is the rear vertex, n=5 is the rear +y top, and n=6 is the rear +x side. In all cases the operand bounds the index at the specified point to be greater than the specified target value. For example, "I4GT" constrains the minimum index at the rear vertex of the surface (i.e. the vertex of the next surface) of the gradient index lens. In all cases the +y top and +x side distance is defined by the larger of the front and rear semi-diameters set on the main spreadsheet. See also GRMN and GRMX, which are similar operands that are easier to use.</p>
InLT	<p>Index "n" less than. This operand is similar to InGT except it constrains the maximum value of the index of refraction rather than the minimum. See InGT for a complete description of the parameter "n".</p>
InVA	<p>This operand is similar to InGT except it constrains the current value of the index of refraction. See InGT for a complete description of the parameter "n".</p>
ISFN	<p>Image space F/#. This operand is the paraxial infinite conjugate F/#. See WFNO.</p>

NAME	Description
ISNA	Image space numerical aperture. This operand is the paraxial image space na at the defined conjugates. See ISFN.
LACL	Lateral color. For focal systems, this is the y-distance between the paraxial chief ray intercepts of the two extreme wavelengths defined by Minw and Maxw , measured in lens units. For afocal systems, this is the angle in afocal mode units between the paraxial chief rays of the two extreme wavelengths defined by Minw and Maxw .
LINV	Lagrange (or optical) invariant of system in lens units at the wavelength defined by Wave . The paraxial marginal and chief ray data are used to compute this value.
LOGE	Log base e of an operand. Op# is the row number of the operand value to take the log of. If the value is less than or equal to zero, the value zero is returned.
LOGT	Log base 10 of an operand. Op# is the row number of the operand value to take the log of. If the value is less than or equal to zero, the value zero is returned.
LONA	Longitudinal aberration, measured in lens units for focal systems and diopters for afocal systems. This is the defocus from the current image to the image at the wavelength defined by Wave and pupil zone defined by Zone . If Zone is zero, paraxial rays are used to determine the paraxial image locations. If Zone is greater than 0.0 and less than or equal to 1.0, real marginal rays are used to determine the image locations. In this case, Zone corresponds to the Py coordinate of the real marginal ray. See AXCL.
LPTD	This boundary operand constrains the slope of the axial gradient index profile from changing signs within a gradient index component for the surface defined by Surf . See the section "Using gradient index operands".
MAXX	Returns the largest value within the indicated range of operands defined by Op#1 and Op#2 . See MINN.
MCOG	Multi-configuration operand greater than. This is used to constrain values in the multi-configuration editor. Op# defines which multi-configuration operand is used. Cfg# defines which configuration is used.
MCOL	Multi-configuration operand less than. This is used to constrain values in the multi-configuration editor. See MCOG.
MCOV	Multi-configuration operand value. This is used to directly target or compute values in the multi-configuration editor. See MCOG.
MINN	Returns the smallest value within the indicated range of operands. See MAXX.
MNAB	Minimum Abbe number. This boundary operand constrains the Abbe number of surfaces from Surf1 to Surf2 to be greater than the specified target value. See also MXAB. This operand controls multiple surfaces simultaneously.
MNCA	Minimum center thickness air. This boundary operand constrains each of the center thicknesses of surfaces from Surf1 to Surf2 which have air (i.e. no glass) as the glass type to be greater than the specified target value. See also MNCT and MNCG. This operand controls multiple surfaces simultaneously.
MNCG	Minimum center thickness glass. This boundary operand constrains each of the thicknesses of surfaces from Surf1 to Surf2 which have a non-air glass type to be greater than the target value. See also MNCT and MNCA. This operand controls multiple surfaces simultaneously.
MNCT	Minimum center thickness. This boundary operand constrains each of the center thicknesses of surfaces from Surf1 to Surf2 to be greater than the specified target value. See also MNCG and MNCA. This operand controls multiple surfaces simultaneously.

NAME	Description
MNCV	Minimum curvature. This boundary operand constrains each of the curvatures of surfaces from Surf1 to Surf2 to be greater than the specified target value. See also MXCV. This operand controls multiple surfaces simultaneously.
MNDT	Minimum diameter to thickness ratio. Controls the minimum allowable value on the ratio of surface diameter to center thickness of surfaces from Surf1 to Surf2 . Only surfaces with non-unity index of refraction are considered. See also MXDT. This operand controls multiple surfaces simultaneously.
MNEA	Minimum edge thickness air. This boundary operand constrains each of the edge thicknesses of surfaces from Surf1 to Surf2 which have air (i.e. no glass) as the glass type to be greater than the specified target value. See also MNET, MNEG, ETGT, and XNEA. This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XNEA for constraining non-rotationally symmetric surfaces. Zone , if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.
MNEG	Minimum edge thickness glass. This boundary operand constrains each of the edge thicknesses of surfaces from Surf1 to Surf2 which have a non-air glass type to be greater than the specified target value. See also MNET, MNEA, ETGT, and XNEG. This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XNEG for constraining non-rotationally symmetric surfaces. Zone , if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.
MNET	Minimum edge thickness. This boundary operand constrains each of the edge thicknesses of surfaces from Surf1 to Surf2 to be greater than the specified target value. See also MNEG, MNEA, ETGT, and XNET. This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XNET for constraining non-rotationally symmetric surfaces. Zone , if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.
MNIN	Minimum index at d-light. This boundary operand constrains the Nd value of surfaces from Surf1 and Surf2 to be greater than the specified target value. See also MXIN. This operands controls multiple surfaces simultaneously.
MNPD	Minimum $\Delta P_{g,F}$. This boundary operand constrains the deviation of the partial dispersion of surfaces between Surf1 and Surf2 to be greater than the specified target value. See also MXPDP. This operand controls multiple surfaces simultaneously.
MNSD	Minimum semi-diameter. Constrains the semi-diameter to be larger than the specified target over the surface range between Surf1 and Surf2 . This operand controls multiple surfaces simultaneously.
MSWA	Modulation square-wave transfer function, average of sagittal and tangential. See MTF for details.
MSWS	Modulation square-wave transfer function, sagittal. See MTF for details.
MSWT	Modulation square-wave transfer function, tangential. See MTF for details.

NAME	Description
MTFA	<p>Diffraction modulation transfer function, average of sagittal and tangential. The parameters are:</p> <p>Samp: Higher sampling yields higher accuracy at the expense of computation time. To confirm the computation has acceptable accuracy, start at 1, and increment the sampling until the results change by less than the desired accuracy. Note that extreme precision is not required for good optimization results; three significant figures is usually adequate. There are two algorithms available for computing the MTF. If Grid is zero (recommended), a fast, sparse sampling integration method is used to compute the MTF. The fast sampling method used by MTFA, MTFS, and MTFT is not directly related to the MTF Analysis feature. Because only a single spatial frequency is required, the method of computation used by the MTF operands is different, and generally much faster, than the algorithm used by the analysis feature. To select the alternate grid-based algorithm used by the MTF analysis feature, set Grid equal to 1. The grid-based algorithm is usually slower than the default algorithm if the MTF is reasonably good (greater than 5%), but the grid algorithm converges faster if the aberrations are very large and the resulting MTF is very low.</p> <p>Wave: The wavelength number to use (use 0 for polychromatic).</p> <p>Field: The field number.</p> <p>Freq: The spatial frequency in MTF units (see “MTF Units” on page 106). If the sampling is set too low for accurate computation of the MTF, then the MTF operands all return zero.</p> <p>If both the tangential and sagittal MTF are needed; place the MTFT and MTFS operands on adjacent lines and they will be computed simultaneously. See “Using MTF operands” on page 525.</p>
MTFS	Modulation transfer function, sagittal. See MTFA for details.
MTFT	Modulation transfer function, tangential. See MTFA for details.
MTHA	<p>Huygens Modulation transfer function, average of sagittal and tangential. This operand computes the diffraction MTF using the Huygens method (see “Huygens MTF” on page 137). The parameters are:</p> <p>Samp: The pupil sampling, where 1 yields 32 x 32, 2 yields 64 x 64 etc. The sampling is assumed to be the same for both pupil and image.</p> <p>Wave: The wavelength number to use (use 0 for polychromatic).</p> <p>Field: The field number.</p> <p>Freq: The spatial frequency in MTF units (see “MTF Units” on page 106). If the sampling is set too low for accurate computation of the MTF, then the MTF operands all return zero.</p> <p>Pol?: Set to 0 to ignore polarization and 1 to consider it.</p> <p>All Conf?: Set to 0 to use the current configuration (defined by the last CONF operand preceding this operand), and 1 to sum over all configurations. See “Huygens MTF” on page 137 for a full discussion of this option.</p> <p>Ima Delta: The image delta in micrometers used for the computation. If zero, the default image delta is used.</p> <p>If both the tangential and sagittal MTF are needed; place the MTHT and MTHS operands on adjacent lines and they will be computed simultaneously. See “Using MTF operands” on page 525.</p>
MTHS	Huygens Modulation transfer function, sagittal. See MTHA for details.
MTHT	Huygens Modulation transfer function, tangential. See MTHA for details.
MXAB	Maximum Abbe number. This boundary operand constrains the Abbe number of surfaces from Surf1 to Surf2 to be less than the specified target value. See also MNAB. This operand controls multiple surfaces simultaneously.
MXCA	Maximum center thickness air. This boundary operand constrains each of the thicknesses of surfaces from Surf1 to Surf2 which have air (i.e. no glass) as the glass type to be less than the target value. See also MXCT and MXCG. This operand controls multiple surfaces simultaneously.
MXCG	Maximum center thickness glass. This boundary operand constrains each of the center thicknesses of surfaces from Surf1 to Surf2 which have a non-air glass type to be less than the target value. See also MXCT and MXCA. This operand controls multiple surfaces simultaneously.
MXCT	Maximum center thickness. This boundary operand constrains each of the center thicknesses of surfaces from Surf1 to Surf2 to be less than the specified target value. See also MXCG and MXCA. This operand controls multiple surfaces simultaneously.

NAME	Description
MXCV	Maximum curvature. This boundary operand constrains each of the curvatures of surfaces from Surf1 to Surf2 to be less than the specified target value. See also MNCV. This operand controls multiple surfaces simultaneously.
MXDT	Maximum diameter to thickness ratio. Controls the maximum allowable value on the ratio of surface diameter to center thickness from Surf1 to Surf2 . Only surfaces with non-unity index of refraction are considered. See also MNDT. This operand controls multiple surfaces simultaneously.
MXEA	Maximum edge thickness air. This boundary operand constrains each of the edge thicknesses of surfaces from Surf1 to Surf2 which have air (i.e. no glass) as the glass type to be less than the specified target value. See also MXET, MXEG, ETLT, and XAEA. This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XAEA for constraining non-rotationally symmetric surfaces. Zone , if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.
MXEG	Maximum edge thickness glass. This boundary operand constrains each of the edge thicknesses of surfaces from Surf1 to Surf2 which have a non-air glass type to be less than the target value. See also MXET, MXEA, ETLT, and XSEG. This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XSEG for constraining non-rotationally symmetric surfaces. Zone , if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.
MXET	Maximum edge thickness. This boundary operand constrains each of the edge thicknesses of surfaces from Surf1 to Surf2 to be less than the specified target value. See also "MXEG", "MXEA", "ETLT", and "XSET". This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XSET for constraining non-rotationally symmetric surfaces. Zone , if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.
MXIN	Maximum index at d-light. This boundary operand constrains the Nd value of surfaces between Surf1 and Surf2 to be less than the specified target value. See also MNIN. This operand controls multiple surfaces simultaneously.
MXPD	Maximum $\Delta P_{g,F}$. This boundary operand constrains the deviation of the partial dispersion of surfaces between Surf1 and Surf2 to be less than the specified target value. See also MNPD. This operand controls multiple surfaces simultaneously.
MXSD	Maximum semi-diameter. Constrains the semi-diameter to be less than the specified target over the surface range from Surf1 to Surf2 .
NORD	Normal distance to the next surface. This operand computes the surface normal vector at the coordinate defined by X and Y on the surface defined by Surf , then returns the distance to the next surface measured along the normal vector.
NORX	Normal vector x component. This operand returns the x component of the surface normal vector at the coordinate defined by X and Y on the surface defined by Surf . If Global is zero, the vector is in surface local coordinates, if Global is 1, the vector is in global coordinates.
NORY	Normal vector y component. This operand returns the y component of the surface normal vector at the coordinate defined by X and Y on the surface defined by Surf . If Global is zero, the vector is in surface local coordinates, if Global is 1, the vector is in global coordinates.
NORZ	Normal vector z component. This operand returns the z component of the surface normal vector at any coordinate defined by X and Y on the surface defined by Surf . If Global is zero, the vector is in surface local coordinates, if Global is 1, the vector is in global coordinates.
NPGT	Non-sequential parameter greater than. Surf defines the surface number of the NSC group (always 1 in pure NSC systems). Object defines the object number in the NSC group. The parameter number is defined by Param .

NAME	Description
NPLT	Non-sequential parameter less than. See NPGT.
NPVA	Non-sequential parameter value. See NPGT.
NPXG	Non-sequential object position x greater than. Surf defines the surface number of the NSC group (always 1 in pure NSC systems). Object defines the object number in the NSC group. If Ref? is 0, the coordinates are relative to the reference object. If Ref? is 1, the coordinates are relative to the origin of the NSC coordinate system or entry port. If Ref? is 2, then the coordinates are relative to the global coordinate reference surface.
NPXL	Non-sequential object position x less than. See NPXG.
NPXV	Non-sequential object position x value. See NPXG.
NPYG	Non-sequential object position y greater than. See NPXG.
NPYL	Non-sequential object position y less than. See NPXG.
NPYV	Non-sequential object position y value. See NPXG.
NPZG	Non-sequential object position z greater than. See NPXG.
NPZL	Non-sequential object position z less than. See NPXG.
NPZV	Non-sequential object position z value. See NPXG.
NSDC	Non-sequential coherent data. Surf defines the surface number of the NSC group (always 1 in pure NSC systems). Det# refers to the object number of the desired detector. If Pix# is a positive integer, then the data from the specified pixel is returned. If Pix# is zero, then the sum of the data for all pixels for that detector object is returned. Data is 0 for real, 1 for imaginary, 2 for amplitude, and 3 for power. See “Optimizing with sources and detectors in non-sequential mode” on page 531 for complete details.
NSDD	Non-sequential incoherent intensity data. Surf defines the surface number of the NSC group (always 1 in pure NSC systems). Det# refers to the object number of the desired detector. If Det# is zero, then all detectors are cleared. If Det# is less than zero, then the detector defined by the absolute value of Det# only is cleared. <u>For Detector Rectangles, Detector Surfaces, and all faceted detectors:</u> If Pix# is a positive integer greater than zero, then the data from the specified pixel is returned. Otherwise, the following data is returned depending upon the value of Pix# : 0: Sum of the total flux in position space, average flux/area in position space, or total flux in angle space for all pixels for that detector object, for Data = 0, 1, or 2, respectively. -1: Maximum flux or flux/area. -2: Minimum flux or flux/area. -3: Number of rays striking the detector. -4: Standard deviation (RMS from the mean) of all the <i>non-zero</i> pixel data. -5: The mean value of all the <i>non-zero</i> pixel data. -6, -7, -8: The x, y, or z coordinate of the position or angle Irradiance or Intensity centroid, respectively. -9,-10, -11, -12, -13: The RMS radius, x, y, z, or xy cross term distance or angle of all the pixel data with respect to the centroid. These are the second moments r^2 , x^2 , y^2 , z^2 , and xy , respectively. Data is 0 for flux, 1 for flux/area, and 2 for flux/solid angle pixel. Only values 0 and 1 (for flux and flux/area) are supported for faceted detectors. See “Optimizing with sources and detectors in non-sequential mode” on page 531 for complete details. <u>For Detector Volumes:</u> Pix# is interpreted as the voxel number. For Data values of 0, 1, or 2, the returned value is incident flux, absorbed flux, or absorbed flux per unit volume, respectively. If Pix# is zero, the value returned is the sum for all pixels. For Detector Color objects, use NSDE instead. For Detector Polar objects, use NSDP.

NAME	Description
NSDE	<p>Non-sequential Detector Color object data. Surf defines the surface number of the NSC group (always 1 in pure NSC systems). Det# refers to the object number of the desired detector. If Det# is zero, then all detectors are cleared. If Det# is less than zero, then the detector defined by the absolute value of Det# only is cleared.</p> <p>If Pix# is a positive integer greater than zero, then the data from the specified pixel is returned. Otherwise, the following data is returned depending upon the value of Pix#:</p> <p>0: Sum of all data for all pixels on the detector. For chromaticity coordinates this is the average over all pixels in the detector that have non-zero lumen values. If Pix# is zero the <area> values described below are the area of the entire detector, measured in analysis units for position space data or solid angle in steradians for angle space data.</p> <p>-1: Maximum data. -2: Minimum data. -3: Number of rays striking the detector.</p> <p>Angle? is 0 for position space data, 1 for angle space data.</p> <p>Data is an integer code that defines the type of data to return, defined as follows:</p> <p>1: power in <sprefix> watts where <sprefix> is the source units prefix. 2: power/area in <aprefix> watts per <area> where <aprefix> is the analysis units prefix. If Angle? is 0 <area> is the area in analysis units. If Angle? is 1 then <area> is solid angle in steradians. 3: power in <sprefix> lumens where <sprefix> is the source units prefix. 4: power/area in <aprefix> lumens per <area> where <aprefix> is the analysis units prefix. If Angle? is 0 <area> is the area in analysis units. If Angle? is 1 then <area> is solid angle in steradians. 5/6: CIE 1931 chromaticity coordinate x/y. 7/8: CIE 1976 chromaticity coordinate u'/v'. 9/10/11: CIE 1931 tristimulus X/Y/Z value in lumens per <area>. If Angle? is 0 <area> is the area in analysis units. If Angle? is 1 then <area> is solid angle in steradians.</p> <p>If the source units are joules, then the data 1 and 2 are in joules and data 3 and 4 are in talbots. See "Units" on page 104. See also NSDD and NSDP.</p>
NSDP	<p>Non-sequential Detector Polar object data. Surf defines the surface number of the NSC group (always 1 in pure NSC systems). Det# refers to the object number of the desired detector. If Det# is zero, then all detectors are cleared. If Det# is less than zero, then the detector defined by the absolute value of Det# only is cleared.</p> <p>If Pix# is a positive integer greater than zero, then the data from the specified pixel is returned. Otherwise, the following data is returned depending upon the value of Pix#:</p> <p>0: Sum of all data for all pixels on the detector. For chromaticity coordinates this is the average over all pixels in the detector that have non-zero lumen values.</p> <p>-1: Maximum data. -2: Minimum data. -3: Number of rays striking the detector. -4: The radial RMS in degrees of the distribution with respect to 0 degrees.</p> <p>Data is an integer code that defines the type of data to return, defined as follows:</p> <p>1: power in <sprefix> watts where <sprefix> is the source units prefix. 2: power/solid angle in <sprefix> watts per steradian where <sprefix> is the source units prefix. 3: power in <sprefix> lumen where <sprefix> is the source units prefix. 4: power/solid angle in <sprefix> lumens per steradian where <sprefix> is the source units prefix. 5/6: CIE 1931 chromaticity coordinate x/y. 7/8: CIE 1976 chromaticity coordinate u'/v'. 9/10/11: CIE 1931 tristimulus X/Y/Z value in lumens/steradian.</p> <p>If the source units are joules, then the data 1 and 2 are in joules and data 3 and 4 are in talbots. See "Units" on page 104. See also NSDD and NSDE.</p>

NAME	Description
NSRA	<p>Non-sequential single ray trace. Src# refers to the object number of the desired source. This source must be defined to trace just a single analysis ray. If Split? is non-zero, then splitting is on. If Pol? is non-zero then polarization will be used. If splitting is on polarization is automatically selected. Scattering is always turned off for this feature, since scattering introduces random paths for the child rays, which is not suitable for optimization. Errors are always considered.</p> <p>If multiple NSRA operands trace the same ray are adjacent in the merit function editor, the ray will only be traced once for efficiency.</p> <p>Seg# refers to the segment number that contains the data to be returned. Use -1 for the last segment.</p> <p>Data refers to the data type for the specified segment.</p> <p>Use Data values 1-9 for x-coordinate, y-coordinate, z-coordinate, x-cosine, y-cosine, z-cosine, x-normal, y-normal, and z-normal, respectively. These values are relative to the entry port, see 31-39 below.</p> <p>Use Data values 10-15 for path-to, intensity, phase of, phase at, index, and starting phase, respectively.</p> <p>Use Data values 16-17 for the sum of the path, or optical path, respectively, in lens units from the source to the end of the specified segment. These values do not include the phase of diffractive surfaces.</p> <p>Use Data values 21-26 for Ex real, Ex imaginary, Ey real, Ey imaginary, Ez real, and Ez imaginary, respectively. Polarization must be on for real and imaginary data to be returned.</p> <p>Use Data value 27-29 for the phase of Ex, Ey, and the phase difference between Ex and Ey, respectively. All values are in radians. Polarization must be on for valid data to be returned.</p> <p>Use Data values 31-39 for the coordinate data defined by Data values 1-9 converted to coordinates relative to the global coordinate reference surface.</p> <p>Source # refers to the element number to use if the source is an array; see “Modeling arrays of sources” on page 563 for the numbering scheme.</p> <p>For more information on these data items see “Ray database (ZRD) files” on page 583.</p>
NSRM	<p>Non-sequential Rotation Matrix component. Surf defines the surface number of the NSC group (always 1 in pure NSC systems). Object defines the object number in the NSC group.</p> <p>If Ref? is 0, the coordinates are relative to the reference object. The rotation matrix will always be the identity matrix when using this reference.</p> <p>If Ref? is 1, the coordinates are relative to the origin of the NSC coordinate system or entry port.</p> <p>If Ref? is 2, then the coordinates are relative to the global coordinate reference surface.</p> <p>The 3 x 3 R matrix has 9 components. If Data is 1, NSRM returns R[1][1], if Data is 2, NSRM returns R[1][2], etc... through Data = 9 returning R[3][3].</p>
NSST	<p>Non-sequential single ray trace. This operand traces a single sequential ray through the system to any specified Non-Sequential surface, and returns various data about the ray within the Non-Sequential surface. Surf is the surface number of the Non-Sequential surface. Wave is the wavelength number. For a definition of Hx, Hy, Px, and Py, see “Hx, Hy, Px, and Py” on page 488.</p> <p>Data determines what data NSST will compute and return as follows:</p> <ul style="list-style-type: none"> 0, 1, 2: The x, y, or z coordinate at the ray-object intercept point. 3, 4, 5: The x, y, or z direction cosines of the ray after reflection/refraction from the surface. 6, 7, 8: The x, y, or z direction cosines of the ray up to the ray-object intercept point. 9, 10, 11: The surface normal at the ray-object intercept point. 12: The face number of the object the ray struck. <p>Object specifies which object the desired data is for.</p> <p>In the general case, a ray may strike the same object multiple times. By default, NSST will return the data for the <i>last</i> intercept on the specified object. To select a specific intercept, add 1000 to the Data value for each time the ray strikes the object. For example, to compute the y coordinate value of the ray on the third time the ray strikes an object, use a Data value of 3001. If the ray does not strike the object, or does not strike the object the specified number of times, the operand returns a value of zero, but no warning or error is issued.</p> <p>All coordinates and cosines are in the coordinate system of the Non-Sequential surface. See also NSTR.</p>

NAME	Description
NSTR	Non-sequential trace. Src# refers to the object number of the desired source. If Src# is zero, all sources will be traced. If Split? is non-zero, then splitting is on. If Scat? is non-zero, then scattering is on. If Pol? is non-zero then polarization will be used. If splitting is on polarization is automatically selected. If IgEr? is non-zero, then errors will be ignored. See “Optimizing with sources and detectors in non-sequential mode” on page 531 for complete details. See also NSST.
NTXG	Non-sequential object tilt about x greater than. See NPXG.
NTXL	Non-sequential object tilt about x less than. See NPXG.
NTXV	Non-sequential object tilt about x value. See NPXG.
NTYG	Non-sequential object tilt about y greater than. See NPXG.
NTYL	Non-sequential object tilt about y less than. See NPXG.
NTYV	Non-sequential object tilt about y value. See NPXG.
NTZG	Non-sequential object tilt about z greater than. See NPXG.
NTZL	Non-sequential object tilt about z less than. See NPXG.
NTZV	Non-sequential object tilt about z value. See NPXG.
OBSN	Object space numerical aperture. This is only useful for finite conjugate systems, and is calculated on axis at the primary wavelength.
OOFF	This operand indicates an unused entry in the operand list. OOFF operands are automatically converted to BLNK operands upon evaluation of the merit function. OOFF is only used to indicate that the merit function operand type was not recognized.
OPDC	Optical path difference with respect to chief ray in waves at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
OPDM	Optical path difference with respect to the mean OPD over the pupil at the wavelength defined by Wave . OPDM has the same restrictions that TRAC does; see TRAC for a detailed discussion. See “Hx, Hy, Px, and Py” on page 488.
OPDX	Optical path difference with respect to the mean OPD over the pupil with tilt removed at the wavelength defined by Wave . OPDX has the same restrictions that TRAC does; see TRAC for a detailed discussion. See “Hx, Hy, Px, and Py” on page 488
OPGT	Operand greater than. This is used to make the value of the operand defined by Op# greater than the target value.
OPLT	Operand less than. This is used to make the value of the operand defined by Op# less than the target value.
OPTH	Optical path length. This is the distance, in lens units, the specified ray travels to the surface defined by Surf at the wavelength defined by Wave . The distance is measured from the object for finite conjugates; otherwise the distance is referenced to the first surface. The optical path accounts for the index of refraction of the media, and for phase adding surfaces such as gratings and binary optics. See PLEN. See “Hx, Hy, Px, and Py” on page 488.
OPVA	Operand value. This operand constrains the value of the operand defined by Op# to be equal to the target value.

NAME	Description
OSCD	Offense against the sine condition (OSC) at the wavelength defined by Wave . There are two definitions for OSC supported. The first definition is as described in Welford, Aberrations of Optical Systems (see "REFERENCES ON LENS DESIGN" on page 37). This definition is used if Zone is zero. An alternate definition due to Prof. Roland Shack which supports computation of OSC as a function of pupil zone and uses only real rays is available. This definition is used if Zone is not zero. In this case, Zone corresponds to the Py coordinate of the real marginal ray. The two methods will give very similar results for systems with modest F/#'s and aberrations when Zone is 1.0 for the alternate definition. This operand has no meaning if the system is not axially symmetric.
OSUM	Sums the values of all operands between the two operands defined by Op#1 and Op#2 . See SUMM.
PANA	Paraxial ray x-direction surface normal at the ray-surface intercept at the wavelength defined by Wave . This is the x component of the surface normal vector at the intersection point of the specified paraxial ray and the surface defined by Surf , in the local coordinate system. See "Hx, Hy, Px, and Py" on page 488.
PANB	Paraxial ray y-direction surface normal at the ray-surface intercept at the wavelength defined by Wave . This is the y component of the surface normal vector at the intersection point of the specified paraxial ray and the surface defined by Surf , in the local coordinate system. See "Hx, Hy, Px, and Py" on page 488.
PANC	Paraxial ray z-direction surface normal at the ray-surface intercept at the wavelength defined by Wave . This is the z component of the surface normal vector at the intersection point of the specified paraxial ray and the surface defined by Surf , in the local coordinate system. See "Hx, Hy, Px, and Py" on page 488.
PARA	Paraxial ray x-direction cosine of the ray after refraction from the surface defined by Surf at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.
PARB	Paraxial ray y-direction cosine of the ray after refraction from the surface defined by Surf at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.
PARC	Paraxial ray z-direction cosine of the ray after refraction from the surface defined by Surf at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.
PARR	Paraxial ray radial coordinate in lens units at the surface defined by Surf at the wavelength defined by Wave . This is the radial distance from the local axis to the intersection of the surface defined by Surf and the specified paraxial ray, in the local coordinate system. See "Hx, Hy, Px, and Py" on page 488.
PARX	Paraxial ray x-coordinate in lens units at the surface defined by Surf at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.
PARY	Paraxial ray y-coordinate in lens units at the surface defined by Surf at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.
PARZ	Paraxial ray z-coordinate in lens units at the surface defined by Surf at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.
PATX	Paraxial ray x-direction ray tangent. This is the tangent of the angle the paraxial ray makes in the X-Z plane after refraction from surface defined by Surf at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.

NAME	Description
PATY	Paraxial ray y-direction ray tangent. This is the tangent of the angle the paraxial ray makes in the Y-Z plane after refraction from surface defined by Surf at the wavelength defined by Wave . See "Hx, Hy, Px, and Py" on page 488.
PETC	Petzval curvature in inverse lens units at the wavelength defined by Wave . Not valid for non-paraxial systems.
PETZ	Petzval radius of curvature in lens units at the wavelength defined by Wave . Not valid for non-paraxial systems.
PIMH	Paraxial image height at the paraxial image surface at the wavelength defined by Wave . Not valid for non-paraxial systems.
PLEN	Path length. This operand computes the total optical path length (including index of refraction and phase surfaces) between surfaces Surf1 and Surf2 for the specified ray, which is always traced at the primary wavelength. PLEN is essentially the difference between two OPTH operands. See OPTH. See "Hx, Hy, Px, and Py" on page 488.
PMAG	Paraxial magnification. This is the ratio of the paraxial chief ray height on the paraxial image surface to the object height at the wavelength defined by Wave . Only useful for finite conjugate systems. Note the paraxial image surface is used even if the system is not at paraxial focus.
PMGT	Parameter greater than. This boundary operand constrains the value of the parameter defined by Param for the surface defined by Surf to be greater than the target value. The parameter values have different meanings depending upon the surface type. See the Chapter "Surface Types" for a description of the parameter values.
PMLT	Parameter less than. This boundary operand constrains the value of the parameter defined by Param for the surface defined by Surf to be less than the target value. The parameter values have different meanings depending upon the surface type. See the Chapter "Surface Types" for a description of the parameter values.
PMVA	Parameter value. This boundary operand constrains the value of the parameter defined by Param for the surface defined by Surf to be equal to the target value. The parameter values have different meanings depending upon the surface type. See the Chapter "Surface Types" for a description of the parameter values.
PnGT	This operand is obsolete, use PMGT instead.
PnLT	This operand is obsolete, use PMLT instead.
PnVA	This operand is obsolete, use PMVA instead.

NAME	Description
POPD	<p>Physical Optics Propagation Data. For important details see “PHYSICAL OPTICS PROPAGATION” on page 627.</p> <p>To use this operand, first define the settings on the POP analysis feature as desired, then press Save on the settings box. The operand will return data based upon the selected settings.</p> <p>If Surf is zero, then the saved ending surface number will be used; otherwise, the specified surface will be used as the ending surface. If Wave is zero, then the saved wavelength number will be used; otherwise, the specified wavelength number will be used. If Field is zero, then the saved field number will be used; otherwise, the specified field number will be used.</p> <p>Data determines what data the POP feature will compute and return as follows:</p> <p>0: The total fiber coupling. This is the product of the system efficiency and the receiver efficiency.</p> <p>1: The system efficiency for fiber coupling.</p> <p>2: The receiver efficiency for fiber coupling.</p> <p>3: The total power.</p> <p>4: The peak irradiance.</p> <p>5, 6, 7: The pilot beam position, Rayleigh range, beam waist (x).</p> <p>8, 9, 10: The pilot beam position, Rayleigh range, beam waist (y).</p> <p>11, 12, 13: The local X, Y, Z coordinates of the center of the beam array on the end surface (this is a reference point and is not related to the amplitude of the beam).</p> <p>21, 22: The X, Y coordinates of the centroid of the intensity distribution in local coordinates relative to the center of the beam.</p> <p>23, 24, 25, 26: The X, Y beam width and X, Y M-squared values, respectively. See “Beam width and M-squared” on page 646.</p> <p>27, 28, 29: The second moment of the beam for x^2, y^2, and xy, respectively.</p> <p>30, 31, 32: The mean, RMS and PTV irradiance variation, respectively, of the non-zero amplitude portion of the beam. These operands should only be used when the beam has nearly uniform irradiance and has just been clipped by a surface aperture.</p> <p>33, 34, 35: The mean, RMS and PTV phase variation in radians, respectively, of the non-zero amplitude portion of the beam. These operands should only be used when the beam has nearly uniform irradiance and has just been clipped by a surface aperture.</p> <p>40, 41, 42: The fraction of the total power enclosed within a circle of radius specified by the Xtr1 value in lens units, referenced to the beam centroid (40), chief ray (41), or surface vertex (42).</p> <p>50, 51, 52: The radius in lens units of the circle at which the fraction of the total power enclosed is equal to the value specified by the Xtr1 value. The circle is centered on the beam centroid (50), chief ray (51), or surface vertex (52).</p> <p>60, 61, 62, 63: The fiber coupling receiver efficiency amplitude and phase in radians for the Ex field (60 and 61) and the Ey field (62 and 63). These values do not consider system efficiency (see Data type 1 above).</p> <p>Undefined Data values will return 0.</p> <p>The Xtr1 and Xtr2 values are only used by selected data numbers which are reserved for future expansion of this feature.</p> <p>If adjacent POPD operands all have the same Surf, Wave, Field, Xtr1, and Xtr2 values, then the POP analysis is done just once and all data returned at one time. Note the POPD operands must be on adjacent rows in the MFE for this efficiency to be implemented.</p>

NAME	Description
POPI	<p>Physical Optics Propagation Data. For important details see “PHYSICAL OPTICS PROPAGATION” on page 627.</p> <p>To use this operand, first define the settings on the POP analysis feature as desired, then press Save on the settings box. The operand will return data based upon the selected settings.</p> <p>If Surf is zero, then the saved ending surface number will be used; otherwise, the specified surface will be used as the ending surface. If Wave is zero, then the saved wavelength number will be used; otherwise, the specified wavelength number will be used. If Field is zero, then the saved field number will be used; otherwise, the specified field number will be used.</p> <p>Data determines what data the POP feature will compute and return as follows: 0, 1, 2: The total (Ex + Ey), Ex only, or Ey only Irradiance. 3, 4, 5, 6: Ex-real, Ex-Imaginary, Ey-real, Ey-imaginary. 7, 8: Ex, Ey phase in radians. Undefined Data values will return 0.</p> <p>The Pix# refers to the desired pixel of the beam. The pixel number is greater than or equal to zero, and less than $n_x \times n_y$, where n_x and n_y are the number of columns and rows, respectively. The pixel number is generally defined as $p = x + y \times n_x$ where x is the integer row number and y is the integer column number, and $0 \leq x < n_x$ and $0 \leq y < n_y$.</p> <p>If adjacent POPI operands all have the same Surf, Wave, and Field values, then the POP analysis is done just once and all data returned at one time. Note the POPI operands must be on adjacent rows in the MFE for this efficiency to be implemented.</p>
POWF	<p>Power at a field point. Computes the power or effective focal length (EFL) after refraction from the surface defined by Surf at the wavelength defined by Wave for any point in the field. For the Data parameter, use 0 for spherical, 1 for cylinder, 2 for max, 3 for min, 4 for tangential, 5 for sagittal, 6 for y, 7 for x, and 8 for astigmatic power in diopters. To get the EFL values, add 9 to these codes. For example, for tangential EFL use Data = 13. For cylinder, the EFL value is the difference between the maximum and minimum focal lengths. For astigmatic, the EFL is the difference between the x and y direction focal lengths (x-y).</p> <p>For a full description of this type of analysis, see “Power Field Map” on page 191. See “Hx, Hy, Px, and Py” on page 488.</p>
POWP	<p>Power at a point in the pupil. Computes the power or effective focal length (EFL) after refraction from the surface defined by Surf at the wavelength defined by Wave for any point in the field, at any point in the pupil. For the Data parameter, use 0 for spherical, 1 for cylinder, 2 for max, 3 for min, 4 for tangential, 5 for sagittal, 6 for y, 7 for x, and 8 for astigmatic power in diopters. To get the EFL values, add 9 to these codes. For example, for tangential EFL use Data = 13. For cylinder, the EFL value is the difference between the maximum and minimum focal lengths. For astigmatic, the EFL is the difference between the x and y direction focal lengths (x-y).</p> <p>For a full description of this type of analysis, see “Power Pupil Map” on page 193. See “Hx, Hy, Px, and Py” on page 488.</p>
POWR	<p>The surface power (in inverse lens units) of the surface defined by Surf at the wavelength defined by Wave. This operand only works for standard surfaces.</p>
PRIM	<p>Primary wavelength. This is used to change the primary wavelength number to the wavelength defined by Wave during merit function evaluation. This operand does not use the target or weight columns.</p>
PROB	<p>Multiplies the value of the operand defined by Op# by the factor defined by Factor.</p>
PROD	<p>Product of two operands (Op#1 X Op#2). See PROB.</p>
QOAC	<p>Unused.</p>
QSUM	<p>Quadratic sum. This operand squares and then adds all operands between the two operands defined by Op#1 and Op#2, then takes the square root of the sum. See also SUMM, OSUM, EQUA.</p>

NAME	Description
RAED	Real ray angle of exitance. This is the angle in degrees between the surface normal and the ray after refraction or reflection for the surface defined by Surf at the wavelength defined by Wave . See also RAID. See “Hx, Hy, Px, and Py” on page 488.
RAEN	Real ray angle of exitance. This is the cosine of the angle between the surface normal and the ray after refraction or reflection at the surface defined by Surf at the wavelength defined by Wave . See also RAIN. See “Hx, Hy, Px, and Py” on page 488.
RAGA	Global ray x-direction cosine. This is the direction cosine of the ray in the global coordinate system at the surface defined by Surf at the wavelength defined by Wave . The origin of the global coordinate system is at the global reference surface. See “Hx, Hy, Px, and Py” on page 488.
RAGB	Global ray y-direction cosine. See RAGA.
RAGC	Global ray z-direction cosine. See RAGA.
RAGX	Global ray x-coordinate. This is the coordinate in lens units in the global coordinate system at the surface defined by Surf at the wavelength defined by Wave . The origin of the global coordinate system is at the global reference surface. See “Hx, Hy, Px, and Py” on page 488.
RAGY	Global ray y-coordinate. See RAGX.
RAGZ	Global ray z-coordinate. See RAGX.
RAID	Real ray angle of incidence. This is the angle in degrees between the surface normal and the incident ray at the surface defined by Surf at the wavelength defined by Wave . Note the angle of incidence is always positive. See also RAED. See “Hx, Hy, Px, and Py” on page 488.
RAIN	Real ray angle of incidence. This is the cosine of the angle between the surface normal and the ray before refraction at the surface defined by Surf at the wavelength defined by Wave . See also RAEN. See “Hx, Hy, Px, and Py” on page 488.
RANG	Ray angle in radians with respect to z axis. The angle is measured with respect to the local Z axis of the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
REAA	Real ray x-direction cosine of the ray after refraction from the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
REAB	Real ray y-direction cosine of the ray after refraction from the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
REAC	Real ray z-direction cosine of the ray after refraction from the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
REAR	Real ray radial coordinate in lens units at the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
REAX	Real ray x-coordinate in lens units at the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.

NAME	Description
REAY	Real ray y-coordinate in lens units at the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
REAZ	Real ray z-coordinate in lens units at the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
RECI	Returns the reciprocal of the value of operand Op#1 . See also “DIVI” on page 497.
RELI	Relative illumination. This operand computes the relative illumination of the field point defined by Field relative to the (0, 0) field point. Note that in some systems, the illumination increases off axis, and for these systems the RI may be greater than one since the RELI operand uses (0, 0) as a reference field point. The other parameters are: Samp : The grid size. A value of 10 would yield a 10 x 10 grid of rays. Wave : The wavelength number to use. Pol? : Set to 0 to ignore polarization and 1 to consider it. See also EFNO.
RENA	Real ray x-direction surface normal at the ray-surface intercept at the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
RENB	Real ray y-direction surface normal at the ray-surface intercept at the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
RENC	Real ray z-direction surface normal at the ray-surface intercept at the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
RETX	Real ray x-direction ray tangent (slope) at the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
RETY	Real ray y-direction ray tangent (slope) at the surface defined by Surf at the wavelength defined by Wave . See “Hx, Hy, Px, and Py” on page 488.
RGLA	Reasonable glass. This operand restricts the deviation the index, Abbe, and deviation of the partial dispersion values may take from actual glasses in the currently loaded glass catalogs. See “Optimizing glass selection” on page 528 for a complete discussion including a description of Wn , Wa and Wp . The constraint is active over the surface range specified by Surf1 and Surf2 .
RSCE	RMS spot radius with respect to the centroid in lens units. This operand uses a Gaussian quadrature method that is accurate for systems with unvignetted circular pupils. Ring is used to specify the number of rings of rays traced. If Wave is zero; then a wavelength weighted polychromatic calculation is performed; otherwise, the specified wavelength number will be used. See “Hx, Hy, Px, and Py” on page 488.
RSCH	RMS spot radius with respect to the chief ray in lens units. This operand uses a Gaussian quadrature method that is accurate for systems with unvignetted circular pupils. Ring is used to specify the number of rings of rays traced. If Wave is zero; then a wavelength weighted polychromatic calculation is performed; otherwise, the specified wavelength number will be used. See “Hx, Hy, Px, and Py” on page 488.
RSRE	RMS spot radius with respect to the centroid in lens units. This operand uses a rectangular grid of rays to estimate the RMS. This operand considers vignetting. A Samp value of n will trace an n x n grid per pupil quadrant. If Wave is zero; then a wavelength weighted polychromatic calculation is performed; otherwise, the specified wavelength number will be used. See “Hx, Hy, Px, and Py” on page 488.

NAME	Description
RSRH	RMS spot radius with respect to the chief ray in lens units. This operand uses a rectangular grid of rays to estimate the RMS. This operand considers vignetting. A Samp value of n will trace an n x n grid per pupil quadrant. If Wave is zero; then a wavelength weighted polychromatic calculation is performed; otherwise, the specified wavelength number will be used. See "Hx, Hy, Px, and Py" on page 488.
RWCE	RMS wavefront error with respect to the centroid in waves. This operand uses a Gaussian quadrature method that is accurate for systems with unvignetted circular pupils. Ring is used to specify the number of rings of rays traced. If Wave is zero; then a wavelength weighted polychromatic calculation is performed; otherwise, the specified wavelength number will be used. See "Hx, Hy, Px, and Py" on page 488, and "OPTIMIZATION REFERENCE POINTS" on page 484.
RWCH	RMS wavefront error with respect to the chief ray in waves. This operand uses a Gaussian quadrature method that is accurate for systems with unvignetted circular pupils. Ring is used to specify the number of rings of rays traced. If Wave is zero; then a wavelength weighted polychromatic calculation is performed; otherwise, the specified wavelength number will be used. See "Hx, Hy, Px, and Py" on page 488, and "OPTIMIZATION REFERENCE POINTS" on page 484.
RWRE	RMS wavefront error with respect to the centroid in waves. This operand uses a rectangular grid of rays to estimate the RMS. This operand considers vignetting. A Samp value of n will trace an n x n grid per pupil quadrant. If Wave is zero; then a wavelength weighted polychromatic calculation is performed; otherwise, the specified wavelength number will be used. See "Hx, Hy, Px, and Py" on page 488, and "OPTIMIZATION REFERENCE POINTS" on page 484.
RWRH	RMS wavefront error with respect to the chief ray in waves. This operand uses a rectangular grid of rays to estimate the RMS. This operand considers vignetting. A Samp value of n will trace an n x n grid per pupil quadrant. If Wave is zero; then a wavelength weighted polychromatic calculation is performed; otherwise, the specified wavelength number will be used. See "Hx, Hy, Px, and Py" on page 488, and "OPTIMIZATION REFERENCE POINTS" on page 484.
SAGX	The sag in lens units of the surface defined by Surf at X = the semi-diameter, and Y = 0. See also SSAG.
SAGY	The sag in lens units of the surface defined by Surf at Y = the semi-diameter, and X = 0. See also SSAG.
SCUR	Surface curvature. Computes curvature and related data of the surface defined by Surf at the coordinate defined by X and Y . In all cases below, the "maximum" data is determined by calculation of the curvature at 50 points equally spaced from the vertex of the surface to the specified X and Y coordinate, and the largest absolute value found is returned. If data is 0-3, the returned value is the tangential, sagittal, tangential-sagittal, or maximum tangential-sagittal curvature, respectively. If data is 4-7, the returned value is the x, y, x-y, or maximum x-y curvature, respectively. If data is 8-9, the returned value is the absolute value of (R*Sc), where R is the radial coordinate and Sc is the sagittal curvature, or the maximum of this value, respectively.
SFNO	Sagittal working F/#, computed at the field point defined by Field and the wavelength defined by Wave . See TFNO.
SINE	Sine of the value of the operand defined by Op# . If Flag is 0, then the units are radians, otherwise, degrees.
SKIN	Skip if not symmetric. See SKIS.
SKIS	Skip if symmetric. If the lens is rotationally symmetric, then computation of the merit function continues at the operand defined by Op# .
SMIA	SMIA-TV Distortion. Field is the zero distortion reference field position, or use zero to indicate the field position (0, 0). Wave is the wavelength number, or use zero for the primary wavelength. X-Width and Y-Width are the full field width in field units. For more information see "SMIA-TV Distortion" on page 186.

NAME	Description
SPCH	Spherochromatism in lens units. This is the difference between the real marginal axial color and the paraxial axial color of the two extreme wavelengths defined by Minw and Maxw . The distance is measured along the Z axis. Zone defines the zone for which the real marginal axial color is computed. Zone corresponds to the Py coordinate of the real marginal ray. Not valid for non-paraxial systems.
SPHA	Spherical aberration in waves contributed by the surface defined by Surf at the wavelength defined by Wave . If Surf is zero, the sum for the entire system is used.
SQRT	Square root of the operand defined by Op# .
SSAG	The sag in lens units of the surface defined by Surf at the coordinate defined by X and Y . See also SAGX, SAGY.
STHI	Surface Thickness. This operand computes the thickness of the surface defined by Surf to the next surface at the coordinate defined by X and Y on the surface. The calculation accounts for the sag and center thickness of the surface and the sag of the next surface, but not any tilts and decenters between the surfaces.
STRH	Strehl Ratio. This operand computes the Strehl Ratio using the Huygens PSF computation (see "Huygens PSF" on page 146). The parameters are: Samp : The pupil sampling, where 1 yields 32 x 32, 2 yields 64 x 64 etc. The sampling is assumed to be the same for both pupil and image. Wave : The wavelength number to use (use 0 for polychromatic). Field : The field number. Pol? : Set to 0 to ignore polarization and 1 to consider it. All Conf? : Set to 0 to use the current configuration (defined by the last CONF operand preceding this operand), and 1 to sum over all configurations. See "Huygens PSF" on page 146 for a full discussion of this option. The default image delta is used for all Huygens PSF operands.
SUMM	Sum of two operands (Op#1 + Op#2). See OSUM.
SVIG	Sets the vignetting factors for the current configuration. The Precision is either 0, 1, or 2; for high, medium, or low precision, respectively. The SVIG operand can take a significant time to execute at high precision, although the resulting optimization is generally superior. Note that the operands SVIG and CVIG only modify the vignetting factors for evaluating subsequent optimization operands. When the end of the merit function is reached, the vignetting factors will be restored to their original values. See also "CVIG" on page 495.
TANG	Tangent of the value of the operand defined by Op# . If Flag is 0, then the units are radians, otherwise, degrees.
TCGT	Thermal Coefficient of expansion greater than. This boundary operand constrains the TCE of the surface defined by Surf to be greater than the specified target value.
TCLT	Thermal Coefficient of expansion less than. This boundary operand constrains the TCE of the surface defined by Surf to be less than the specified target value.
TCVA	Thermal Coefficient of expansion value. This boundary operand constrains the TCE of the surface defined by Surf to be equal to the specified target value. For glass surfaces, see "GTCE" on page 503.
TFNO	Tangential working F/#, computed at the field point defined by Field and the wavelength defined by Wave . See SFNO.
TGTH	Sum of glass thicknesses from Surf1 to Surf2 . Note that the sum is inclusive, it is not the thickness between the two surfaces. See TTHI.

NAME	Description
TMAS	Total mass. Computes the mass of the glass lenses within the range of surfaces from Surf1 to Surf2 . The mass of a surface considers the volume enclosed up to the following surface; therefore to compute the mass of a single element the first and last surface numbers should be the same. See “Comments on computing element volumes” on page 272 for a discussion of how element masses and volumes are computed.
TOLR	Tolerance data. Data is 0 for RSS estimated change in performance, 1 for nominal performance, and 2 for estimated performance (nominal plus estimated change). File is the integer number corresponding to the tolerance settings file to use. Config# is -2 for the configuration defined by the last CONF operand, -1 for all configurations, 0 to use the configuration set in the tolerance settings file and 1 or greater for a specific configuration number. For details, see “Optimizing tolerance sensitivity” on page 530.
TOTR	Total track (length) of lens in lens units. See “Total track” on page 60.
TRAC	Transverse aberration radial direction measured in image space with respect to the centroid for the wavelength defined by Wave . Unlike most other operands, TRAC critically depends upon the placement of other TRAC operands within the Merit Function Editor to work correctly. TRAC operands must be grouped together by field position and wavelength. ZEMAX traces all TRAC rays with a common field point together, and then uses the collective data to compute the centroid of all the rays. Each ray individually is then referenced to the computed centroid. This operand should only be entered into the Merit Function Editor by the Default Merit Function tool, and is not recommended for use directly by the user. See “Hx, Hy, Px, and Py” on page 488.
TRAD	The x component of the TRAR only. TRAD has the same restrictions that TRAC does; see TRAC for a detailed discussion.
TRAE	The y component of the TRAR only. TRAE has the same restrictions that TRAC does; see TRAC for a detailed discussion.
TRAI	Transverse aberration radius measured at the surface defined by Surf at the wavelength defined by Wave with respect to the chief ray. Similar to TRAR, except a surface other than the image surface may be specified. See “Hx, Hy, Px, and Py” on page 488.
TRAN	Unused.
TRAR	Transverse aberration radial direction measured in image space at the wavelength defined by Wave with respect to the chief ray. See ANAR. See “Hx, Hy, Px, and Py” on page 488.
TRAX	Transverse aberration x direction measured in image space at the wavelength defined by Wave with respect to the chief ray. See “Hx, Hy, Px, and Py” on page 488.
TRAY	Transverse aberration y direction measured in image space at the wavelength defined by Wave with respect to the chief ray. See “Hx, Hy, Px, and Py” on page 488.
TRCX	Transverse aberration x direction measured in image space with respect to the centroid. TRCX has the same restrictions that TRAC does; see TRAC for a detailed discussion.
TRCY	Transverse aberration y direction measured in image space with respect to the centroid. TRCY has the same restrictions that TRAC does; see TRAC for a detailed discussion.
TTGT	Total thickness greater than. This boundary operand constrains the total thickness, including surface sags of the surface defined by Surf and the immediately following surface at their respective semi-diameter values, to be greater than the specified target value. The thickness is calculated along the +y axis if Code is zero, the +x axis if Code is 1, the -y axis if Code is 2, and the -x axis if Code is 3. This operand automatically changes the sign on thicknesses in mirror spaces to always yield a positive value for physically possible lenses. See TTLT and TTVA.

NAME	Description
TTHI	Sum of thicknesses of surfaces from Surf1 to Surf2 . Note that the sum is inclusive, it is not the thickness between the two surfaces. See TGTH.
TTLT	Total thickness less than. See TTGT.
TTVA	Total thickness value. See TTGT.
UDOP	User defined operand. Used for optimizing numerical results computed in externally compiled programs. See "User defined operands" on page 533. See also ZPLM.
USYM	If present in the merit function, this operand instructs ZEMAX to assume radial symmetry exists in the lens even if ZEMAX detects symmetry does not exist. This speeds execution of the merit function in some special cases. See "Assume Axial Symmetry" on page 487.
VOLU	Volume of element(s) in cubic cm. Computes the volume of the lenses and air spaces for the range of surfaces defined by Surf1 and Surf2 . The volume of a surface includes the volume enclosed up to the following surface; therefore to compute the volume of a single element the first and last surface numbers should be the same. See "Comments on computing element volumes" on page 272 for a discussion of how element masses and volumes are computed.
WFNO	Working F/#. See "Working F/#" on page 62, and ISFN, SFNO, and TFNO.
WLEN	Wavelength. This operand returns the wavelength defined by Wave in micrometers.
XDGT	Extra data value greater than. The Edv# must be between 1 and 200 to indicate which of the extra data values is selected for the surface defined by Surf .
XDLT	Extra data value less than. The Edv# must be between 1 and 200 to indicate which of the extra data values is selected for the surface defined by Surf .
XDVA	Extra data value. The Edv# must be between 1 and 200 to indicate which of the extra data values is selected for the surface defined by Surf .
XENC	Extended source encircled energy (distance). This operand computes the distance in micrometers to the specified fraction of extended source geometric encircled energy, using whatever the current default settings are. To use this operand, first define the settings on the extended source encircled energy feature as desired, then press Save on the settings box. The only setting that is overwritten is set by Type , which is 1 for encircled, 2 for x only, 3 for y only, 4 for ensquared, 5 for x distribution, and 6 for y distribution. Frac is the fraction of energy desired, and must be between zero and 1, exclusive. Frac is ignored for type 5 or 6; for these types the returned valued is the full width half max independent of Frac . See also XENF, DENC, DENF, GENC, and GENF.
XENF	Extended source encircled energy (fraction). This operand computes the fraction of extended source geometric encircled, ensquared, x only, or y only (enslitted) energy at a given distance from the reference point. The Type is 1 for encircled, 2 for x only, 3 for y only, and 4 for ensquared. The options and settings are identical to XENC, except Dist , which here is used as the distance at which the fraction of energy is desired. See also XENC, GENC, GENF, DENC, and DENF.
XNEA	Minimum edge thickness for the range of air surfaces defined by Surf1 and Surf2 . This operand checks the edge thickness at numerous points around the perimeter of the surface, and tests if all points are at least the minimum specified thickness. This operand controls multiple surfaces simultaneously. See MNEA. Zone , if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.
XNEG	Minimum edge thickness for the range of glass surfaces defined by Surf1 and Surf2 . This operand checks the edge thickness at numerous points around the perimeter of the surface, and tests if all points are at least the minimum specified thickness. This operand controls multiple surfaces simultaneously. See MNEG. Zone , if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.

NAME	Description
XNET	<p>Minimum edge thickness for the range of surfaces defined by Surf1 and Surf2. This operand checks the edge thickness at numerous points around the perimeter of the surface, and tests if all points are at least the minimum specified thickness. This operand controls multiple surfaces simultaneously. See MNET.</p> <p>Zone, if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.</p>
XXEA	<p>Maximum edge thickness for the range of air surfaces defined by Surf1 and Surf2. This operand checks the edge thickness at numerous points around the perimeter of the surface, and tests if all points are no more than the maximum specified thickness. This operand controls multiple surfaces simultaneously. See MXEA.</p> <p>Zone, if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.</p>
XXEG	<p>Maximum edge thickness for the range of glass surfaces defined by Surf1 and Surf2. This operand checks the edge thickness at numerous points around the perimeter of the surface, and tests if all points are no more than the maximum specified thickness. This operand controls multiple surfaces simultaneously. See MXEG.</p> <p>Zone, if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.</p>
XXET	<p>Maximum edge thickness for the range of surfaces defined by Surf1 and Surf2. This operand checks the edge thickness at numerous points around the perimeter of the surface, and tests if all points are no more than the maximum specified thickness. This operand controls multiple surfaces simultaneously. See MXET.</p> <p>Zone, if non zero, scales the radial aperture at which the thickness is computed. A Zone value of 0.5 would compute the thickness at 0.5 times the semi-diameter.</p>
YNIP	<p>YNI-paraxial. This number is the product of the paraxial marginal ray height times the index times the angle of incidence at the surface defined by Surf at the wavelength defined by Wave. This quantity is related to the narcissus contribution of the specified surface. See Applied Optics, Vol. 21, 18, p3393.</p>
ZERN	<p>Zernike Fringe coefficient. The parameters are:</p> <p>Term: The Zernike term number (1 - 37 for fringe, 1 - 231 for standard or annular).</p> <p>Wave: The wavelength number.</p> <p>Samp: The pupil sampling, where 1 yields 32 x 32, 2 yields 64 x 64 etc.</p> <p>Field: The field number.</p> <p>Type: The Zernike type (0 for fringe, 1 for standard, 2 for annular).</p> <p>Epsilon: The obscuration ratio (for annular coefficients only).</p> <p>Vertex?: If 1, the OPD is referenced to the surface vertex.</p> <p>The Term value, if negative or zero, may also be used to return other data from the Zernike fitting as follows:</p> <ul style="list-style-type: none"> -8: Peak to Valley OPD (to centroid) -7: Peak to Valley OPD (to chief) -6: RMS to zero reference (unused by ZEMAX) -5: RMS to chief ray -4: RMS to centroid -3: Variance -2: Strehl Ratio -1: RMS fit error 0: Maximum single point fit error <p>Note that if you use multiple ZERN operands which only differ in the Term value, they should be placed on adjacent lines in the editor so ZEMAX only does the fitting once; otherwise, the computation is slower.</p>

NAME	Description
ZPLM	Used for optimizing numerical results computed in ZPL macros. See “User defined operands” on page 533. See also UDOP.
ZTHI	This operand controls the variation in the total thickness of the range surfaces defined by Surf1 and Surf2 over multiple configurations. It is similar to the TTHI operand, except it is an inequality operator. The target value specified is the maximum allowed difference between the TTHI at each defined configuration. For example, if there are 3 configurations where TTHI 3 8 would evaluate to 17, 19, and 18.5, respectively, ZTHI will return 2 (i.e. 19-17) if the target is less than 2. Otherwise, ZTHI returns the target value. To keep all zoom configurations the same length, use a target of 0.

The operational operands (SUMM, OSUM, DIFF, PROD, DIVI, SQRT) along with the parametric operands (CVGT, CVLT, CTGT, CTLT, etc..) can be used to define very general and complex optimization operands, as discussed in the section “Defining complex operands” on page 527.

Because of the dimensional differences between parameters such as effective focal length (tens of millimeters or more) and RMS spot radius (micrometers), usually a weighting of one is sufficient for quantities measured in lens units. However, the residual value of the effective focal length with this weighting is not likely to be zero. Increasing the weighting will bring the resulting system closer to the desired effective focal length. This effect is often noticeable when defining ETGT (edge thickness greater than) operands. Usually ETGT with a target of zero will often yield a value just slightly less than zero. Rather than increase the weight, it is much simpler to provide a target value of 0.1, or some such number.

After making changes to the operand list, the current values of each operand can be updated by selecting Tools, Update. This is also useful for checking to see what the current values of each operand are, and which has the greatest contribution to the merit function. The percent contribution is defined as

$$\% \text{ contrib}_i = 100 \times \frac{W_i(V_i - T_i)^2}{\sum_j W_j(V_j - T_j)^2}$$

where the index j indicates the sum over all operands.

The merit function is automatically saved with the lens file.

Understanding boundary operands

The boundary operands such as MNCT, CTGT, DIMX, and others behave somewhat differently than specific target operands such as TRAR and REAY. When you specify a boundary on a parameter, you specify the target value as the definition of the boundary. For example, to maintain a minimum center thickness on surface 5 of 10 mm, you might use a command such as CTGT 5 10 (where the 5 is in the Int1 column and 10 is in the target column). If you update the merit function and then observe the "value" column of that operand, there are two possibilities for the value: 1) if the boundary is violated, that is, the center thickness is less than 10, then the actual value of the thickness will be displayed, or 2) if the boundary is not violated, that is, the center thickness is greater than 10, then the value 10 will be displayed.

The rule is if the boundary is violated, the actual value is shown; if it is not violated, the value is set to the target and is therefore ignored by the optimization algorithm. If during optimization the boundary becomes violated, then the value will automatically be updated and the optimization algorithm will attempt to correct the offending parameter.

The boundary operands which constrain a range of surfaces are slightly more complicated. These multiple surface operands return values which represent the total effect of all violated boundaries within the specified surface range. For example, the operand MNCT 1 10 will constrain the minimum center thickness of surfaces 1 through 10. If the target is 3.0, which defines the boundary, then the difference between the value of the operand and the target is the sum of the difference between 3.0 and the thicknesses of all surfaces between 1 and 10 whose center thickness is less than 3.0. If only one surface in the range has a center thickness less than 3.0, say 2.5, then the operand has a value of 2.5. If a second surface is added that has a thickness less than 3.0, say 2.2,

then the operand will have a value of 1.7 (2.5 minus .8; the .8 is 3.0 - 2.2). The total difference between the target of the operand and the value is 3.0 - 1.7 or 1.3. This difference of 1.3 is due to the violation of 0.5 by the first surface and another 0.8 by the second surface.

If calculating the value of these boundary operands seems confusing, don't worry; ZEMAX does all the calculations for you. All you need to do is to specify the boundary type (such as MNCT or MNET) the boundary range (surface 1 through 10, or whatever) and the desired value (3 mm or whatever). If all the boundary constraints are met, then the operand value is equal to the target, otherwise, the value will be different and the merit function will increase. The increased merit function will cause the optimization algorithm to seek a reduction of the operand contribution.

If a boundary operand does not seem to work, there are several things to check:

1) Make sure the variables you have defined can have some effect on the boundary operands. A common mistake is to specify MNCT and have some "frozen" thickness within the surface range. If the thickness violates the boundary and it is not variable, ZEMAX can't fix it. The operands DO NOT ignore violated but frozen boundaries.

2) If there is a small residual error, try increasing the boundary value. For example, if MNCT is used with a target of 0.0, and the value is a small number (like -.001) the problem is not that the operand doesn't work, it is that the residual error is too small to increase the merit function significantly. It is usually better to increase the target to 0.1, or some other number, rather than to increase the weight. Increasing the weight will only lead to a smaller violation (like -.0000001) rather than meeting the boundary.

3) Check to see if there is a reasonable contribution to the merit function. You can easily check this with the percent contribution column. By looking at the percent contribution column, verify that the operand in question has enough influence on the total merit function. If it does not, increase the weight, or see the preceding paragraph for advice on changing the target.

Understanding the boundary operands is a crucial part of mastering ZEMAX optimization, and with a little practice you will find them to offer excellent control and flexibility.

Using MTF operands

The diffraction based MTF operands (MSWT, MSWS, MSWA, MTFT, MTFS, MTFA, MTHT, MTHS, and MTHA) and the geometric MTF operands GMTA, GMTS, and GMTT provide a capability to directly optimize the MTF. This is a powerful capability, however, using MTF operands requires some care.

The MTF cannot be computed accurately if the aberrations are too large. Any optical system for which the MTF plot produces invalid data will also produce meaningless data during optimization. Large aberrations also tend to produce MTF curves that drop to zero at some spatial frequencies, and then the MTF increases again at higher spatial frequencies. Local optimization of the MTF at frequencies beyond the first zero in the MTF curve is usually ineffective because the MTF must actually decrease before increasing as the aberrations improve, and the local optimizer is always looking for incremental improvement. The MTF operands should not be used if the starting MTF is beyond the first zero in the MTF curve at the desired spatial frequency. If the aberrations are large, MTF optimization is likely not required nor useful. An alternate approach is to first optimize using RMS wavefront error. Systems with low RMS wavefront error will have reasonable MTF performance. After the design is very nearly in the final form, then MTF optimization may be used for a final "touch-up".

MTF optimization is slower than RMS spot radius or RMS wavefront error optimization. Diffraction MTF is slower than Geometric MTF; however, the Geometric MTF is an approximation that is only reasonably valid for systems with circular or elliptical pupils and little or no apodization. Huygens based MTF can be many orders of magnitude slower than RMS optimization, and is thus not recommended for optimization. The slow execution speed may be noticeable when updating the Merit Function Editor display, and entering and exiting the optimization dialog box.

Note that if you use both MTFT and MTFS (or GMTT and GMTS) for the same field and wavelength data, they should be placed on adjacent lines in the editor; otherwise, the MTF is computed twice. If the sampling is too low, or the aberrations too large for accurate computation of the MTF, then the MTF operands return zero rather than a meaningless number.

Performing an optimization

To begin optimization, choose Tools, Optimization from the main menu bar. The optimization control dialog box will appear with the following options.

OPTIMIZATION OPTIONS

Item	Description
Automatic	Executes until ZEMAX determines the system is no longer improving significantly.
1 Cycle	Executes a single optimization cycle.
5 Cycles	Executes 5 optimization cycles.
10 Cycles	Executes 10 optimization cycles.
50 Cycles	Executes 50 optimization cycles.
Inf. Cycles	Executes optimization cycles in an infinite, continuous loop until "Terminate" is pressed.
Terminate	Terminates a running optimization, and returns control back to the dialog box.
Exit	Closes the optimization dialog box.
Algorithm	<p>There are two distinct algorithms supported for optimization.</p> <p>Damped Least Squares (DLS) uses numerically computed derivatives to determine a direction in solution space which produces a design with a lower merit function. This gradient method has been developed specifically for optical system design and is recommended for all optimization problems.</p> <p>Orthogonal Descent (OD) uses an orthonormalization of the variables and discrete sampling of solution space to reduce the merit function. The OD algorithm <i>does not</i> compute numerical derivatives of the merit function. For systems with inherently noisy merit functions, such as non-sequential systems, OD will usually outperform DLS. For a discussion of a similar algorithm, see S. Kudaev, P. Schreiber, "Automated optimization of non-imaging optics for luminaries", in <i>Optical Design and Engineering II</i>; Laurent Mazuray, Rolf Wartmann, Eds., Proc. SPIE 5962, p. 87-95 (2005).</p> <p>Note that both DLS and OD are local optimization algorithms whose final results depend upon the starting point. For global optimization, see "GLOBAL OPTIMIZATION" on page 539.</p>
Auto Update	If checked, ZEMAX will automatically update and redraw all open windows at the end of each optimization cycle. This allows monitoring of the optimization progress using any analysis feature.
# CPU's	Selects the number of CPU's over which to spread the optimization task. More than 1 may be selected, even on a single CPU computer, in which case the single CPU will time share the multiple simultaneous tasks. The default is the number of processors detected by the operating system.

Selecting automatic will cause the optimizer to run until no progress is being made. The other options will run the specified number of cycles. Automatic mode is highly recommended. The time required to run a given optimization cycle varies enormously with the number of variables, the complexity of the system, the number of solves, the number of operands, and of course the computer speed. If the cycle is taking too long, or if it appears to be hung up, or if you feel the design is not progressing adequately, click on Terminate to end the optimization run.

When the optimization begins, ZEMAX first updates the system merit function. If any of the operands cannot be computed, the optimization cannot begin, and an error message will be displayed. Operands cannot be computed if they require the tracing of rays which miss surfaces or which undergo total internal reflection (TIR) at an index boundary. If such an error message appears, usually the starting lens prescription is in error, or the ray targets are incorrectly defined (this will not happen with the default merit functions, but might happen with

user defined rays). ZEMAX can automatically recover if the merit function cannot be evaluated during the course of optimization; only the starting system need be adequate to compute all operands in the merit function.

Defining complex operands

Although the default merit function, coupled with a few predefined operands, is perfectly suitable for the majority of optical designs, there are times when an unusual constraint needs to be added to the merit function. Rather than define a very large number of very specific operands, ZEMAX allows you to build your own operands out of simple building blocks.

ZEMAX allows very general operand definitions. There are two tricks to creating these operands. First, use certain operands with zero-weighting to define the parameters you need, and second, use the operational operands to define relationships between them. For example, suppose you require that the thickness of surface 3 and the thickness of surface 4 sum to 10. There is an operand that does this, TTHI. The command structure would look like this:

Number	Type	Int1	Int2	Target	Weight
1	TTHI	3	4	10	1

However, for illustration only, note that there is an alternate way of calculating the same thing:

Number	Type	Int1	Int2	Target	Weight
1	CTVA	3		0	0
2	CTVA	4		0	0
3	SUMM	1	2	10	1

Operand 1 uses the Center Thickness VALUE (CTVA) command to extract the value of the thickness of surface 3. Similarly, operand 2 is used to extract the thickness of surface 4. The zero weighting on both operands means the optimization algorithm ignores the constraint; it is only used as an intermediate step. Operand 3 now sums two operands: number 1 and number 2. The result is the sum of the thicknesses of surfaces 3 and 4 is the value of operand 3, and this has a non-zero weight. The optimization algorithm will attempt to drive the sum to 10.

Why go to all the trouble of this three-step process if a single TTHI command would do the same thing? The reason is that this approach can be extended to develop very general operands. For example, suppose you wanted the radius of curvature of surface 5 to be centered on the vertex of surface 8. Study the following commands to see if you understand how this is done:

Number	Type	Int1	Int2	Target	Weight
1	CVVA	5		0	0
2	TTHI	5	7	0	0
3	PROD	1	2	1	1

The CVVA command extracts the curvature of surface 5, the curvature we want to control. TTHI 5 7 calculates the distance from surface 5 to surface 8 (note we only sum to surface 7 to get to surface 8, since the thickness of surface 8 gives the distance to surface 9). Since the curvature of the surface is the reciprocal of the radius, the product of the curvature and the distance must be one; hence the target is 1 for operand 3. Operand 3 is also the only weighted operand in the sequence.

Now consider the requirement that the thickness of surface 5 must be greater than twice the radius of curvature of surface 4 plus the conic constant of surface 2 (this is nonsensical, but illustrative of the flexibility in the approach):

Number	Type	Int1	Int2	Target	Weight
1	CTVA	5		0	0
2	CVVA	4		0	0
3	CONS			2	0
4	DIVI	3	2	0	0
5	COVA	2		0	0
6	SUMM	4	5	0	0
7	DIFF	1	6	0	0
8	OPGT	7		0	1

Operand 1 extracts the (center) thickness of surface 5. Operand 2 extracts the value of the curvature of surface 4. Operand 3 sets a constant of two, and operand 4 divides the value 2 by the curvature (yielding twice the radius of curvature). COVA extracts the conic, and SUMM adds operands 2 and 4. Operand 7 takes the difference of the thickness and twice the radius plus the conic. Since we want the former to exceed the latter, we set an operand greater than constraint; the only one to have a non-zero weighting.

Optimizing glass selection

Optimization of glasses is handled somewhat differently than other data. Optimizing the glass choice directly is a difficult and unpredictable process because there does not exist a continuum of glasses on the glass map. There are two methods for dealing with this problem: by using model glasses or by using glass substitution. Glass substitution is usually the superior method.

Using model glasses

The model glass method is to idealize the glass dispersion using a few numerical parameters, and then optimize these parameters while constraining either the parameter values or the computed index values to be similar to available glasses. This is the "model" glass method. Model glasses are described in detail in the chapter "Using Glass Catalogs". One disadvantage of the model glass method is that the optimized parameters and resulting index values may not correspond to any physically existing glass. Another disadvantage is that model glasses are only sufficiently accurate in the visible spectrum. This method is used by the conventional optimizer described in this chapter.

To optimize glasses then requires several steps. First, change the glass of the appropriate surface to a "Model" glass using the glass solve dialog box in the Lens Data Editor. For information on model glasses, see "Using model glasses" on page 601. When you change the glass from "Fixed" to "Model", ZEMAX will make a suitable guess for the index, Abbe number, and partial dispersion; you only need to change the values if you wish. The three values can be made variable by clicking on the "Vary" box next to each entry.

The Ctrl-Z shortcut on the glass column will also work; it makes the index, Abbe, and partial dispersion variable automatically. The model glass data values can now be optimized using the optimization feature in the usual way.

Unconstrained glass optimization usually will lead to very high index materials being selected. This is because surfaces with high refractivity (a large difference in index across the boundary) need less curvature than low refractivity surfaces to have the same optical power. Lower curvature surfaces introduce less aberration.

Unfortunately, high index materials are expensive, heavy, harder to fabricate, and may be brittle, delicate, or susceptible to stains and scratches. Also, very high index materials do not always exist; there are few glasses (for the visible spectrum) available with an N_d higher than about 1.9. The V_d value also is limited to the range of roughly 20 to 80. Therefore, it is essential to limit the N_d and V_d values to reasonable ranges during optimization. The partial dispersion deviation also must be limited in range.

There are two ways to limit the Nd, Vd, and $\Delta P_{g,F}$ values. The simplest way is to add the RGLA operand somewhere in the operand list. The RGLA operand measures the "distance" on the glass map between the index, Abbe number, and partial dispersion of the model glass to the closest glass in the currently loaded catalogs. For example, if you are optimizing the index and Abbe and you have specified that the Schott and Hoya catalogs are used (these are specified on the general data window), the RGLA operand computes the "distance" to each glass in these catalogs. If the smallest "distance" is less than the target value specified for the RGLA operand, then the boundary condition is met, and the value of the operand is equal to the target. If the closest glass is farther than the target value, then the RGLA value is the actual "distance". The "distance" is defined by the square root of the weighted sum of the squares of the difference between the index, Abbe, and partial dispersion terms for two glasses. The "distance" between any two glasses is given by

$$d = [W_n(Nd_1 - Nd_2)^2 + W_a(Vd_1 - Vd_2)^2 + W_p(\Delta P_{g,F1} - \Delta P_{g,F2})^2]^{\frac{1}{2}},$$

where the factors W_n , W_a , and W_p weight the various terms. The weighting factors may be user defined on the RGLA operand parameter list, or if left at zero, will default to 1.0, 1E-04, and 1E+02, respectively.

The best way to use RGLA is to specify the surface range that covers all of the surfaces you are optimizing. For a target value, start off with 0.05. This will allow the glasses to easily move all over the glass map, since the spacing between various glasses is usually less than 0.05. After optimization, decrease the target to roughly 0.02 and reoptimize. This will encourage the optimized system to choose index and Abbe numbers reasonably near actual glasses.

The other method for constraining index and Abbe values is to use the MNIN, MXIN, MNAB, and MXAB controls. These operands are mnemonics for Minimum and Maximum Index and Abbe values, and they are documented in the preceding tables. These operands can be used to restrict the optimization to specific rectangles on the glass map. It may be useful to use RGLA along with MXIN, for example, to restrict the glass selection to existing glasses with an index lower than some value.

At some point you will want to convert your variable index data back to a real glass. To convert from a model glass to the closest real glass, remove the model glass status using Ctrl-Z, or choose the glass solve type as "Fixed". There will usually not be a perfect match between the optimized Nd and Abbe values and those of an actual glass in the current catalog. However, ZEMAX will search through the catalog and find the "best fit" glass using a least-squares criterion similar to the RGLA definition above (the partial dispersion term is omitted). If a glass substitution template is being used, then the template will also be considered, and only glasses meeting the template specifications will be considered. For information on templates, see "Glass Substitution Template" on page 240. The glass in the catalog which differs the least from the variable index parameters is the glass selected. This glass is also reported on the "Surface Data Summary" feature (select Reports, Surface Data). The index of refraction data shown is that calculated from the Nd and Abbe values, not the best fit glass. After converting from a model glass to a real glass, another optimization run is generally required. For systems with delicate chromatic aberration balancing, the best glass choice might never be found using variable glasses, because the model glass dispersion is never identical to the dispersion of a real glass.

Using glass substitution

The glass substitution method is to directly alter the glass types, and then reoptimize to see if the new glasses yield a better solution. This method can be used manually, by changing the glass type and then reoptimizing, or the process can be automated using the global optimization techniques described in the next chapter, "Global Optimization". The global optimization method uses only actual glass catalog materials, and in this sense is the superior method. See "Using glass substitution" on page 542 for a discussion on this method for optimizing glass selection.

Optimizing zoom and multi-configuration lenses

Optimizing zoom lenses is virtually identical to optimizing conventional single-configuration lenses. See the Chapter "Multi-Configurations" for details.

Optimizing extra data



This discussion is only relevant for users of ZEMAX-EE.

Certain surface types supported by ZEMAX-EE, such as the Zernike, Zernike phase, extended polynomial, and binary optic surface use the extra data values. These extra data values can be edited, loaded from text files, and may be used as variables for optimization. For details on the editor, see “Extra Data” on page 93.

To make an extra data value a variable, open up the Extra Data Editor. When the extra data editor appears, move the cursor to the row and column with the value you want to optimize, and press Ctrl-Z (the same command used to set variables on the Lens Data Editor). The variable will now be optimized when the optimization is run.

There are also several boundary constraints for use with the extra data values. XDVA, XDGT, and XDLT are extra data value, greater than, or less than, respectively. The **Surf** value indicates the surface number to which the operand applies, and **Edv#** is used to specify which of the extra data values is to be used.

Optimizing tolerance sensitivity

The idea of optimizing a lens for reduced tolerance sensitivity is that a very low merit function is not useful if the as-built performance is significantly degraded by tolerance sensitivity. Therefore, the optimal design has a reasonable nominal performance but is relatively insensitive to manufacturing defects. The theory is there may be a potential trade-off between performance and tolerance sensitivity.

Tolerance sensitivity in an optical design comes from many sources, including angles of incidence of rays on surfaces, aberration balancing, and the nature of the potential fabrication defects. The interaction of multiple defects makes accurate tolerance prediction a difficult statistical problem. For a complete discussion of tolerance analysis, see the chapter “TOLERANCING” on page 545.

The optimization operand TOLR can in principle be used to optimize for reduced sensitivity to tolerances. To use TOLR, first optimize a design for reasonable starting performance. Then, define the relevant tolerance operands, limits, compensators, and criterion as described in the chapter on Tolerancing. Save the options of the tolerance dialog box; ZEMAX uses these saved options to compute the data returned by TOLR. TOP files are created by saving the tolerance options from the tolerance dialog, see “Other buttons” on page 564. To use the most recently saved tolerance settings, set the file number on the TOLR operand to zero. To use a specific saved settings file, use an integer value between 1 and 999. The saved settings file name must be of the format TOLRnnn.TOP, where nnn is the 3 digit integer specified on the TOLR operand file argument. For example, if the desired tolerance settings are saved in a file named TOLR005.TOP, the integer file value should be 5. Note only files with the correctly formatted name may be referenced and used by TOLR. The TOP files must be stored in the <data>\Configs folder (see “Folders” on page 68). The only criterion that is not allowed is “Merit Function” since that would create an infinite loop. If a script tolerance criterion is being used, do not load any merit functions within the script that contain TOLR operands, as this may also create an infinite loop that ZEMAX cannot detect.

As part of the tolerance sensitivity analysis, ZEMAX computes a nominal performance estimate, and predicts an RSS estimated change. The predicted total performance is the sum of nominal and estimated change. These values are computed and returned by the TOLR operand for optimization. TOLR values may be targeted and weighted as any other optimization operand.

TOLR takes a single “data” integer. If data is 0, TOLR returns the estimated change in performance. If data is 1, TOLR returns the nominal performance. If data is 2, TOLR returns the predicted performance, which is the sum of the nominal and estimated performance change. The first instance of TOLR (the first row in which TOLR appears) is the only instance for which the tolerance analysis is actually run, so that TOLR only runs once for each merit function evaluation. The other values computed by TOLR are saved and can be retrieved by using data values other than 0 on subsequent calls to TOLR. Other values of data are currently unused and may be used for future expansion of this feature.

The practical difficulty of optimization of tolerance sensitivity lies in the computation time. Complicated lenses may have hundreds of tolerance operands, and complex criterion may take considerable time to compute. When using TOLR, the sensitivity analysis may be computed thousands or even millions of times, which can make the use of TOLR impractical. The following tips should be used to improve the efficiency of the tolerance analysis and make TOLR efficient:

- Test and verify the tolerance sensitivity analysis outside of optimization before using TOLR.
- Use "Sensitivity" and not "Inverse Sensitivity". Inverse sensitivity is slower and meaningless for optimization.
- Use "Paraxial Focus" for the "Comp" control if possible... optimization of the compensator values requires an "optimization within an optimization" that is slow and memory intensive; although it is supported if required.
- ZEMAX will automatically ignore the Monte Carlo portion of the tolerance analysis.
- Do not use "Merit Function" as the tolerance criterion, as this will create an infinite loop.

The faster the tolerance sensitivity runs, the faster TOLR and thus the optimization will execute.

Optimizing objects in a non-sequential group with sequential rays

Optimizing variables within a non-sequential group is fundamentally no different from optimizing other numerical parameters. Variables are set in the same way as for parameters in the Lens Data Editor. However, optimizing non-sequential object properties is difficult because of the unpredictable way in which rays may (or may not) propagate through a non-sequential group. For non-sequential objects such as prisms, usually a small change in position or size of the prism does not dramatically affect the ray path. However, for objects such as light pipes, a small change in the object definition can dramatically affect the ray path. Rays that once propagated through an object may miss the object completely if the object position or angle changes slightly. This usually causes severe errors in the computation of derivatives, and the optimization performs either poorly or not at all.

Another problem with some non-sequential systems is the exit pupil may not be a reasonable image of the entrance pupil. For this reason, rectangular array rather than Gaussian Quadrature should be used if the system is a non-imaging system that does not form an image of the entrance pupil at the exit pupil.

For these systems, optimization may proceed more effectively using the global optimization algorithms, which do not rely exclusively on derivative computation.

Optimizing with sources and detectors in non-sequential mode

Optimizing illumination systems or other optical systems that use non-sequential sources and detectors is supported using the NSDC, NSDD, NSDE, NSDP and NSTR operands.

A typical merit function would consist of three groups of operands:

First, NSDD operands would be used to clear the data in the current detectors. Use NSDD with the detector number set to zero to clear all energy in all detectors. Usually a single NSDD at the top of the merit function is all that is needed. NSDD returns a value of zero and has no effect on the merit function value when used to clear detectors.

Second, NSTR operands are used to trace rays from NSC sources. NSTR *i* traces analysis rays from source *i*; NSTR 0 traces all analysis rays from all sources. Note the number of analysis rays on the NSC editor determines how many rays are traced and how long the evaluation of the NSTR operand will take. NSTR always returns a value of zero and has no effect on the merit function value. The NSTR operand supports options to split, scatter, and use polarization.

Third, a new group of NSDC, NSDD, NSDE, or NSDP operands are used to read out the detector data. NSDD has four arguments: surface, detector, pixel, and data. Surface is the surface number of the NSC group (use 1 if the program mode is non-sequential). Detector is the object number of the detector. Both detector objects and faceted detectors may be used as detectors. If pixel is an integer greater than zero, the flux, flux/area, or flux/solid angle is returned for that pixel. Which of the three is determined by the data argument, which should be 0, 1, or 2 for flux, irradiance, or intensity, respectively. If pixel is 0, the sum of all the flux or flux/area in the detector is returned. The units of the returned data is determined by the system units, see "Analysis Units" on page 105. NSDD can compute other data. For a full description of the pixel argument, see "NSDD" on page 509. If the object is a faceted detector, or if the pixel number is -1 or 0, only data options 0 and 1 are supported. Similar capability to that described for NSDD exists for coherent data using NSDC, for detector color data using NSDE, and for detector polar data using NSDP.

The practical difficulty in optimizing these systems is the difficulty of computing derivatives of detected energy with respect to variable parameters because of the relatively large uncertainty in computing detected energy. Many rays must be traced to determine illumination patterns approximately.

Comments about random numbers and NSTR

Note that when launching rays with the NSTR operand, ZEMAX seeds the random number generator with the identical value every time the merit function is evaluated. This means the identical set of random numbers are used on every evaluation of the merit function. If no changes are made to the optical system, the merit function will always evaluate to a consistent value; a desirable property for the optimization algorithm to function. However, if a change is made to the system, rays may take different paths and the identical set of random numbers may be used in different ways; for example, in the computation of scattering paths.

Pixel numbering for detectors

For a detailed discussion of the pixel numbering used by the various types of detectors, see “Detectors” on page 532.

Optimizing with the IMAE operand

The IMAE operand estimates the efficiency of an optical system by launching many rays into the entrance pupil; computing the fraction of rays that pass through all surface apertures to any surface. Optimization with this operand may not proceed smoothly if only hard-edged surface apertures, such as the circular aperture, are used. This is because ZEMAX estimates the derivative of operand values by making very small differential changes in the value of each variable, then computes a finite difference of the operand value. For the IMAE operand, a small change in the value of a variable may not change the efficiency estimate; if no rays are close enough to an aperture to change from being vignetted to unvignetted or vice-a-versa.

The solution is to replace hard edged apertures with soft-edged apertures placed on a user-defined surface. A soft-edged aperture has a transmission that is unity over most of the clear aperture, but near the edge the transmission drops to zero gradually over a small region, rather than abruptly.

Filter functions for doing this are included with ZEMAX as sample DLL files; see the Chapter "Surface Types" under "User Defined Surface" for details. See in particular the discussion of the US_FILT4.DLL sample.

IMAE uses the current saved settings for the image analysis feature; except for "show" which is always set to spot diagram for this computation. See “Geometric Image Analysis” on page 165.

Using gradient index operands

There are several optimization operands which are used to control the properties of gradient index materials during optimization. Some of them are described below.

DLTN

DLTN is used to control the maximum total change in index within a gradient index lens. **Surf** defines the surface number, and **Wave** defines the wavelength number. DLTN is defined as:

$$DLTN = n_{max} - n_{min}.$$

The min and max index values are computed at the extreme z coordinates, Z min and Z max. Z min and Z max are the Z coordinates of the minimum and maximum axial positions of the blank used to make the lens, before the shaping begins. For a convex surface, they correspond to the vertex. For a concave surface, they correspond to the maximum sag at that surface.

LPTD

LPTD is used to control the profile of the gradient within the material. **Surf** is used to define the surface number of the gradient index surface. LPTD is an acronym for LightPath Technology Delta, and the constraint is used to keep a nonlinear profile monotonically increasing or decreasing. It only needs to be used when the quadratic or cubic term of the axial gradient is variable. This operand only affects GRIN 5 surface types.

The LPTD operand should be used with a target of 0. The boundary constraint enforces the following conditions:

$$\frac{\partial n}{\partial z_{min}} > 0 \text{ and } \frac{\partial n}{\partial z_{max}} > 0, \text{ or } \frac{\partial n}{\partial z_{min}} < 0 \text{ and } \frac{\partial n}{\partial z_{max}} < 0.$$

Z min and Z max are the Z coordinates of the minimum and maximum axial positions of the blank used to make the lens, before the shaping begins. For a convex surface, they correspond to the vertex. For a concave surface, they correspond to the maximum sag at that surface. If the residual value of the operand is less than zero, then the target may be decreased slightly (try 0.1). Changing the target is usually more effective than increasing the weight. The value of the LPTD operand must be zero for the blank to be fabricated. Always check the gradient profile to make sure the slope does not change sign.

User defined operands

There are times when very complex calculations need to be performed, and the results of the computation need to be optimized. ZEMAX supports some of these calculations already, such as the MTF A operand which traces many rays, computes the MTF, and then returns a single resulting number to the Merit Function Editor "value" column. Some limited calculations can be performed within the merit function itself; see for example the discussion in the section "Defining complex operands" on page 527.

However, there are problems for which only the flexibility of a user defined program is sufficient for defining the data computed by an operand. There are two ways of achieving this:

- 1) Through the use of a ZPL macro
- 2) Through the use of an externally defined and compiled program

The use of ZPL macros is simpler, well integrated with ZEMAX, and requires very little programming experience. ZPL macro optimization is generally a better choice for simpler macros that execute fairly quickly.

Externally defined programs are more complex to program, require an external C or other language compiler, and at least some programming experience. However, externally defined programs can be vastly more complex than what is supported by the ZPL macro language, and usually run faster. Complex calculations will benefit more from being externally compiled than simple calculations.

Both the ZPL and the externally compiled methods of implementing UDO's are described in detail below.

Optimizing with ZPL macros

If the ZPL macro language (see "ZEMAX PROGRAMMING LANGUAGE" on page 651) is sufficient to perform the required computations, then the operand ZPLM may be used to call a ZPL macro from within the merit function. The macro performs the required computations, then returns the result using the ZPL OPTRETURN keyword.

ZPLM is simple to use. The **Mac#** and **Data** values are used to specify the macro number and data field number, respectively. The macro number is used to indicate which ZPL macro should be executed, while the data field number indicates which value computed by the macro should be optimized.

The macro number must be an integer between 0 and 99. If the **Mac#** value is 17, for example, then the macro to be executed must be named ZPL17.ZPL. The macro name must always use a two digit representation of the macro number. If the macro number was 6, then the macro to be executed would be ZPL06.ZPL. The ZPL macro file must reside in the folder for ZPL macros; see "Folders" on page 68.

The data field number may be any number between 0 and 50, inclusive. This number refers to a position in a global array associated with the lens in memory. During execution of the macro, the macro keyword OPTRETURN specifies which data field number stores the results of the macro calculation. There are 51 different data fields, so that a single macro call can be used to optimize up to 51 different values simultaneously. For example, suppose you needed a macro which computed the total length of the lens from surface 1 to the image surface (this is in effect a user-defined version of the TOTR operand). The macro might look like this:

```
n = NSUR()
x = 0
FOR i = 1, n, 1
    x = x + THIC(i)
NEXT
OPTRETURN 0, x
```

Note the use of the OPTRETURN keyword. This keyword stores the resulting value for "x" in the global array position 0. Suppose this macro was named ZPL15.ZPL. To optimize the resulting value for x, the ZPLM merit

function operand would be added to the Merit Function Editor, with **Mac#** = 15 and **Data** = 0. After updating the merit function, the "value" would be the same as that returned by TOTR, and it can be optimized in the same way.

ZPLM also permits the use of the data in the Merit Function Editor columns. These data fields can be read by the ZPL macro using the PVHX, PVHY, PVPX, and PVPY ZPL functions, respectively. "PV" is a mnemonic for "Pass Value".

There is one very important thing to know about the data field number. If it is zero, then the macro is executed and the value from OPTRETURN 0 is returned. However, if the data field number is not zero, then the macro is not executed, but any previous value stored from an earlier call to the macro is used instead. The advantage to this convention is substantial. If the macro computes many values, all of which need to be optimized, the macro only needs to be called once, yet multiple ZPLM operands can access the data. This is much more efficient than calling the macro multiple times.

For example, suppose a macro named ZPL11.ZPL computes three values, all of which require optimization. In the macro, the values are stored using OPTRETURN:

```
OPTRETURN 0, x
OPTRETURN 1, y
OPTRETURN 2, z
```

Then three ZPLM operands in the merit function can extract the data and perform the optimization with a single call to the macro:

```
ZPLM 11 0
ZPLM 11 1
ZPLM 11 2
```

The macro ZPL11.ZPL is only called during the evaluation of the ZPLM 11 0 operand. Note the data columns can only be used if the Int2 value is zero, since only in this case is the macro evaluated.

Changes made to the lens from within the ZPLM macro

The merit function is always evaluated using a temporary copy of the lens. After evaluation of the merit function, the copy of the lens, and any changes made to the lens, are discarded. For this reason, no changes should be made to the lens data from within the macro called by the ZPLM operand. These changes are not retained and may interfere with the computation of operands following the ZPLM operand in the same merit function evaluation. ZEMAX does not restore the lens being evaluated to the state it was in prior to the evaluation of the ZPLM specified macro. If however the macro is intentionally used to alter the lens data prior to evaluation of subsequent operands, two macros should be executed. The first should modify the data as required, and the second should restore the data to the original condition. Both macros can be listed in the merit function editor, with the intervening operands executing on the altered lens data.

ZPLM and the default merit function

ZPLM should not be used in the middle of a default merit function, but should instead be placed either prior to or after the portion of the merit function that ZEMAX defined by default.

Optimizing with externally compiled programs

The second method of creating a user defined operand (UDO) is to write an external Windows program which computes the data, then use Dynamic Data Exchange (DDE) to pass the data to and from ZEMAX. The DDE interface in ZEMAX is documented and described in "ZEMAX EXTENSIONS" on page 749. The material presented there is not duplicated here; this discussion assumes the material in that chapter is understood.

The operand UDOP is used to call an external client program from within the merit function. The client program performs the required computations, possibly by making multiple DDE calls back to the ZEMAX server, then returns the result to ZEMAX using the DDE interface. The computed data is then placed in the "value" column of the Merit Function Editor and thus may be optimized in the usual way.

UDOP is simple to use. The **Prog#** and **Data** values are used to specify the client program number and data field number, respectively. The client program number is used to indicate which client program should be executed, while the data field number indicates which value computed by the client program should be optimized.

The client program number must be an integer between 0 and 99. If **Prog#** is 17, for example, then the client program number to be executed must be named UDO17.EXE. The client program name must always use a two digit representation of the client program number. If the client program number was 6, then the client program to

be executed would be UDO06.EXE. The client program file must reside in the <data>\UDO folder (see "Folders" on page 68).

When reaching a UDOP operand with a data field number of zero (more on the data field number shortly), ZEMAX will call the client program. The client program is called with the following syntax, assuming the client program number was 17:

```
UDO17.EXE buffercode Data1 Data2 Data3 Data4
```

where Data1 through Data4 are the first four data columns from the editor. The buffercode is an integer value provided by ZEMAX to the client that uniquely identifies the correct lens. Because ZEMAX is capable of evaluating multiple lenses simultaneously, the buffercode is used as an identifier so that when the client requests or returns data, it is associated with the correct lens. Note that while optimizing ZEMAX may be evaluating dozens of lenses simultaneously, each slightly different as derivatives are evaluated and the optimization progresses. The client must compute data for the indicated lens.

Once execution in the client has begun, the client program must do the following critical steps:

- 1) Establish a DDE link with the ZEMAX server.
- 2) Load the correct lens in the ZEMAX server's memory.
- 3) Compute the required data.
- 4) Pass the data back to ZEMAX.
- 5) Clear the ZEMAX server's memory.
- 6) Terminate the DDE link and exit.

The DDE link is typically maintained by the ZCLIENT code, which is described in the DDE chapter (the user is of course free to write their own, if preferred). ZCLIENT calls the UserFunction defined by the user to compute the operand data.

To load the correct lens into the ZEMAX server's memory, a single item must be sent to the ZEMAX server, GetUDOSystem. The syntax is "GetUDOSystem, buffercode". This will cause ZEMAX to retrieve the correct lens from system memory, and all subsequent DDE calls will be for actions (such as ray tracing) on this lens.

The data is then computed, optionally using any of the DDE item calls defined in the DDE chapter. Once the data is computed, up to 1001 values may be sent back to the server, and ultimately to the optimizer within ZEMAX, with the SetUDOIItem command. The syntax is "SetUDOIItem, buffercode, data_field, data". All data items must be free formatted integer, exponential, or floating point numbers.

It is critical that the client program pass back a string using SetUDOIItem. ZEMAX will wait for the client program to pass this string back. ZEMAX has no way of knowing how long the computation may take; and so ZEMAX will "hang" until it receives the data. If the client program crashes, or never returns the data, ZEMAX will never complete execution of the operand, and will hang forever. Pressing the escape key from within ZEMAX will "break" the evaluation and cause ZEMAX to skip the operand evaluation.

Terminating the link is achieved by the client program UserFunction returning control to ZCLIENT.

A sample code that computes 3 data items, called a, b, and c, might look like this:

```
void UserFunction(char *szCommandLine)
{
double a, b, c;
char szBuffer[5000], szSub[256];
int buffer_code;

/* get the buffer code that identifies the lens */
buffer_code = atoi(GetString(szCommandLine, 0, szSub));

/* set the correct lens in the server's memory */
sprintf(szBuffer, "GetUDOSystem,%i", buffer_code);
PostRequestMessage(szBuffer, szBuffer);

/* Here is where we compute the data... these lines are omitted */
```

```

/* now return the data a, b, and c in data_field 0, 1, and 2 */
sprintf(szBuffer,"SetUDOIItem,%i, 0, %.9E", buffer_code, a);
PostRequestMessage(szBuffer, szBuffer);
sprintf(szBuffer,"SetUDOIItem,%i, 1, %.9E", buffer_code, b);
PostRequestMessage(szBuffer, szBuffer);
sprintf(szBuffer,"SetUDOIItem,%i, 2, %.9E", buffer_code, c);
PostRequestMessage(szBuffer, szBuffer);

/* close the communication, this tells ZEMAX we are done sending data. */
sprintf(szBuffer,"CloseUDOData,%i", buffer_code);
PostRequestMessage(szBuffer, szBuffer);

}

```

Note that data values may be returned to the client program using SetUDOIItem. This is where the data_field number comes into play. The data field number may be any number between 0 and 1000, inclusive. This number refers to a position in a global array associated with the lens in memory. During execution of the client program, the client DDE item "SetUDOIItem" is used to return numerical data values. These values are stored in the lens buffer for later use by UDOP.

There are 1001 different data fields, so that a single client program call can be used to optimize up to 1001 different values simultaneously. The data field number indicates which of the returned values should be placed in the "value" column for that UDOP operand.

UDOP also permits the use of the Data1 through Data4 fields. These data fields can be read by the client program, because they are passed on the command line after the buffercode.

There is one very important thing to know about the data_field number when used on the Merit Function Editor UDOP operand. If it is zero, then the client program is executed and the value from data position 0 is placed in the value column. However, if the data field number is not zero, then the client program is not executed, but any previous value stored from an earlier call to the client program is used instead. The advantage to this convention is substantial. If the client program computes many values, all of which need to be optimized, the client program only needs to be called once, yet multiple UDOP operands can access the data. This is much more efficient than calling the client program multiple times.

For example, suppose a client program named UDO25.EXE computes three values, all of which require optimization. In the client program, the values are passed back using SetUDOIItem. Then three UDOP operands in the merit function can extract the data and perform the optimization with a single call to the client program:

```

UDOP 25 0
UDOP 25 1
UDOP 25 2

```

The client program UDO25.EXE is only called during the evaluation of the UDOP 25 0 operand. Note the Data1 through Data4 values can only be used if the data field value is zero, since only in this case is the client program evaluated.

Unlike using ZPL macros, UDO's may freely change the lens data during evaluation since all DDE commands are executed on a copy of the lens, not the actual lens being optimized.

There is a sample UDO source code file called UDO_DEMO.C; it may be compiled and linked in with ZCLIENT. The executable needs to be renamed UDOPxx.EXE where xx is a two digit integer. The sample UDO returns 6 values: Data1, Data2, Data3, Data4, and two dummy constants, in data positions 0 through 5, respectively.

Suggestions for use

In preliminary design stages, it is rarely required to trace all of the rays for all of the wavelengths at each field position during optimization. For this reason, execution times may be substantially decreased by limiting the number of fields and wavelengths used during optimization. If the weight of selected fields and wavelengths is set to zero, then the default merit function algorithm will skip the zero weighted fields or wavelengths when constructing the merit function. This results in fewer rays being traced, speeding execution.

For example, if the lens is being evaluated at five field points, it is possible that only the first, third, and fifth field need be included in the merit function. Of course, later in the design process all fields may need to be included and the default merit function reconstructed.

There are a few other tricks to improve performance. Avoid setting boundary operands on variables unless the optimized solution persists on implausible designs. Boundary operands add computational overhead. Use solves instead of explicit operands whenever possible. For example, use a curvature solve to control the focal length rather than an operand if possible.

Optimization is inseparable from the art of modern lens design, and only practice will make a designer a proficient user of optimization algorithms. Users who are expert at other software optimization algorithms will probably find ZEMAX easier to use, and with a little practice, the mechanics of using the interface will slip into the subconscious, and the designer can concentrate upon the design itself. If you are new to computerized optimization of lenses, there is no better way to learn than to practice.

The global optimum

The design that yields the lowest possible value of the merit function is called the "global" optimum and is by definition the best possible design. However, there is no known optimization algorithm that can universally find the global optimum for an arbitrary design problem, unless you consider "direct search" an optimization algorithm (in other words, try all of the infinite number of possible solutions to see which is best). The art of optical design with computer assistance has two basic components. First, the designer must be able to determine a suitable starting design, and second, he or she must play the role of supervisor during the optimization process. A good supervisor knows when and how to back up and coax the program into a more fruitful direction.

Unfortunately, this often requires considerable experience, and even more often, excessive tedium. An experienced designer uses a combination of intuition, analysis, and luck in searching for new, better, design forms. ZEMAX provides an automated capability for performing this global optimum search; the feature is described in the next chapter.

Introduction



The material presented in this chapter critically depends upon the user having read and understood the material presented in the previous chapter on "Optimization".

The conventional means of optimizing lenses has for decades been the use of the damped least squares algorithm (DLS). DLS has many attractive features; it is efficient, and it is very good at finding the "local" minimum of the merit function. In this context, the word local means the lowest value of the merit function that can be reached from the current position in solution space without ever increasing the merit function (this is something of an idealization, in reality DLS can hop over small regions of increased merit function).

To visualize this problem, imagine you are on a hike, and are trying to find the bottom of a valley from a starting point on the side of a hill. You would like to find the lowest point in the valley. Suppose you cannot see the valley; it is very foggy and all you can see is the terrain very close to where you now stand. You can determine which way is downhill, and proceed in that direction until the slope becomes uphill again; at which point you find a new downhill direction. You might repeat this procedure over and over until a point is reached where all directions are uphill. This lowest point is then a local minimum, and the bottom of at least this valley.

The problem with this method is that once you have arrived at the local minimum, there is no known way (for the general optimization problem) to determine if there is not a better, lower minimum somewhere else. For example, if you were to walk uphill in any direction from this point until you reached a local peak, and then were to proceed onward downhill into the next valley, you will eventually reach a new local minimum. Is this new minimum lower or higher than the previous valley? The only way to find out is to take the hike!

You might ask, since computers are so fast, why not just try out every possible configuration to see which is best? To get a feel for the scope of the problem, consider a cemented doublet lens with six degrees of freedom (degrees of freedom manifest themselves as variables for optimization). If you assume that each variable can take on 100 possible values (a coarse sampling), then there are $1E+12$ different possible systems. If each system evaluation requires the tracing of 20 rays (a low estimate) and you can trace 1,000,000 rays per surface per second, then the time required is about $8E+07$ seconds, or about 2.5 years. For a four-element lens (16 variables) evaluated at three fields and three wavelengths, using 100 rays for evaluation would require $1E+32$ system evaluations or many billions of times the age of the universe.

There are approaches to the global optimization problem that (thankfully) do not require unreasonable computational effort. These algorithms include simulated annealing, multistart, expert systems, neural networks, and others. All of these algorithms have strengths and weaknesses which are beyond the scope of this chapter.

Capabilities of ZEMAX

There are two separate global optimization algorithms in ZEMAX, each with a different purpose. The first algorithm you are likely to use is called "Global Search", and it is used to find new design forms given only the merit function and a starting design. Global search uses a combination of genetic algorithms, multistart, conventional damped least squares, and some expert system heuristics to search for new design forms. The Global Search algorithm is very good at finding promising design forms, however, it does not usually produce "finished" designs. The second algorithm is used for this purpose.

The second algorithm is called "Hammer" optimization (lens designers often talk about hammering on a design to squeeze out the last bit of performance). The Hammer algorithm is used for exhaustively searching for the optimum solution once a reasonably good starting point is found, presumably by prior experience or the Global Search algorithm. The Hammer algorithm only requires a partially optimized lens and merit function in the form of a ZMX file.

Although the global optimization algorithms are extremely useful, it is important to realize that there is no guarantee that the true global optimum will always, or even occasionally be found. Of course, there is no way to even determine if any solution is the global optimum, even if it is the best you have ever found (remember the universe-lifetimes scale of the problem).



It is important to realize that there is no guarantee that the true global optimum will always, or even occasionally be found.

Both the Global Search and Hammer algorithms require extensive computational effort to be effective. These algorithms are not intended to be used interactively! (that is what DLS optimization is for). If you set up global optimization, and watch while the computer works, you are bound to be disappointed. Global optimization is highly effective when you set the problem up and let the computer run for many hours, or even several days, but not for ten minutes. The ideal situation is to set up the problem before stopping work for the night, and let Global Search (or Hammer, depending upon your requirements) work overnight. In the morning, you should have useful results to look at.

The Global Search algorithm

Before you begin the Global Search, you must come up with a very rough starting point. "Very rough" means the design has the correct number of surfaces, a defined stop surface, and initial glasses selected. The fields and wavelengths must be defined. You also need to define a merit function; see "OPTIMIZATION" on page 483 for details on this procedure. The very rough design can be parallel plates of glass with a curvature solve on the last surface to control the focal length. If a solve is not used to control the focal length, then the system should have at least the approximate focal length desired. Also, the variable parameters must all be defined. The lens must be saved before initiating Global Search. ZEMAX uses the starting focal length as a scaling parameter, so the initial design should have at least the approximately correct focal length!

From the main screen, select Tools, Global Search. There are four buttons on the dialog box labeled: Start, Stop, Resume, and Exit. There is also an option to select the number of "best" files to save. Select Start. ZEMAX copies the initial file into new ZEMAX format files called GLOPT_001.ZMX through GLOPT_nnn.ZMX where nnn is the maximum number of files to save. ZEMAX will then begin looking at various combinations of lens parameters extracted from the ranges you have defined. The optimization will proceed on a newly generated lens until ZEMAX concludes the new lens has been sufficiently optimized.

The algorithm control allows selection of either Damped Least Squares (DLS) or Orthogonal Descent (OD). The DLS algorithm is preferred for most imaging systems, while OD is preferred for systems with noisy, low-precision merit functions, such as illumination systems. For more information, see "Performing an optimization" on page 526.

As each new lens is generated, ZEMAX will compare the merit function of the new lens to the best lenses found so far, and will place it in the correct location in the best lenses list, renaming other lens files as required. If the lens has a higher merit function than all of the lenses on the best list, then it is discarded. The cycle repeats indefinitely. Each time a new lens is found which is better than the worst in the list of best lenses, it is placed in the correct place in the list. After hundreds of lenses have been replaced (which may require many tens of thousands of lenses to be evaluated) the resulting set will hopefully contain some very good designs, or at least some promising forms. The global search dialog box will also display the merit functions of the ten best lenses found so far. If the number of lenses to save is greater than ten, these files are stored on disk, but the merit functions are not displayed.

The algorithm also periodically returns to lenses in the best list to see if they can be improved upon. Occasionally some lenses will be improved and placed back in the list. If this happens, the older design being replaced is rejected if it has the same basic form as the new lens. This is done to keep some diversity in the best list, otherwise all the lenses will be of nearly identical form.

To terminate the search, select Stop. Depending upon what the algorithm is doing, it may exit immediately, or it may require several seconds. Once the algorithm has terminated, you can click on Exit. You can now open any of the GLOPT_xxx.ZMX files for further analysis.

The Resume button is very similar to the Start button, however, Resume will first load the existing GLOPT_xxx files and place their current merit functions in the best list. Thus, Resume begins the search from where a previous run ended. Resume does not erase the existing best files, whereas Start erases the files and begins the search anew entirely based upon the lens currently in the Lens Data Editor. If Resume is selected when a completely unrelated file is in the Lens Data Editor, ZEMAX will attempt to optimize the lens using the current data in the Lens Data and Merit Function Editors, but will use the old GLOPT files respective merit functions for comparison purposes.

Global search will rarely find the global optimum by itself. The reason is that the Global Search concentrates the effort on finding new, promising design forms rather than converging exactly on the best possible solution for each form. This latter job is left to a separate algorithm, called "Hammer Optimization", described in the following section.

The Hammer algorithm

After reviewing the design forms generated by Global Search, you will probably want to investigate one or two of them. The one with the lowest merit function is not always the best (although it should be if you designed your merit function well). For example, the second-best solution may be easier to fabricate. Whatever criterion you use to determine the most promising solution, you now want to find the best design possible using this selected lens as a starting point.

The following controls are on the Hammer dialog box:

Hammer: Begins the Hammer optimization loop. Hammer optimization will take the lens and exhaustively attempt to refine it by making adjustments and optimizations. Each time the lens is improved, it will be saved to disk in a temporary file.

Auto: Calls the conventional (local) damped least squares (DLS) or orthogonal descent (OD) optimizer and runs it in Automatic mode. Sometimes it is useful to optimize a lens before calling Hammer if the lens has not already been optimized.

Stop: Terminates the Hammer search.

Exit: Closes the dialog box.

Algorithm: Choose either Damped Least Squares (DLS) or Orthogonal Descent (OD). The DLS algorithm is preferred for most imaging systems, while OD is preferred for systems with noisy, low-precision merit functions, such as illumination systems. For more information, see "Performing an optimization" on page 526.

#CPU: The number of CPU's to use in implementing the Hammer search.

Auto Update: If checked, all open windows will be updated every time a lower merit function is found.

Hammer optimization only requires a ZMX file with variables and a merit function to get started. The Hammer optimization screen shows the starting merit function, and the best merit function found so far. Although good results can occasionally be had in several minutes, the algorithm should be allowed to run for several hours, and preferably overnight. To terminate the search, select Stop, then Exit.

If ZEMAX abnormally terminates, the last saved Hammer file can be found in the temporary file. The temporary file name is constructed from the starting lens file name. If the lens being optimized is stored in the file

```
C:\ZEMAX\SAMPLES\MYFILE.ZMX
```

then the temporary file will be called

```
C:\ZEMAX\SAMPLES\MYFILE_HAMMER.ZMX.
```

The Hammer algorithm can also be used effectively on partially optimized designs that were not generated with Global Search. Feel free to use the Hammer on any design!

Optimizing glass selection

If a glass is made into a "model" glass described by the index, Abbe, and partial dispersion deviation, then the parameters to the model may be made variable and optimized just like any other numerical parameters. However, the model glass method has one serious drawback. After a good solution is found using model glasses, a conversion must be made back from the model glass to a real glass. The design must then be reoptimized using the new glass selection.

Unfortunately, for many systems the newly optimized design will perform worse than the model glass design. Even more frustrating is that the optimal design using real glasses may have a different form than that found using the model glass.

Traditionally, to find a better combination of glass types, the designer would locate alternate glasses on the glass map, substitute in the new selections, and reoptimize. If the new solution was better, the glass choices would be retained; otherwise, a new set of glasses would be evaluated. The procedure would continue as long as the designer could bear to continue the search.

Using glass substitution

ZEMAX automates this procedure by allowing glasses to have a "substitute" status associated with them. If a glass is marked (using the glass solve dialog box) as a substitute glass, then the global optimization algorithms (both Hammer and Global Search) automatically perform iterative substitution of similar glass types during optimization. This allows ZEMAX to optimize not only the numerical prescription values such as radius and thickness, but also allows direct optimization of real glass selection, without resorting to idealizing the glass dispersion.

To use the substitution feature, set the status of each glass free to change to "substitute" on the glass solve dialog box (double clicking on any glass name is the quickest way to reach this dialog box).

Once the glass substitution status is defined, then invoke either the Hammer or Global Search optimizers. ZEMAX will automatically alter the glass types during the search for better designs. By default, ZEMAX will choose any glass from any of the currently active catalogs, which are specified on the System, General dialog box. Glass substitution does not use table glasses or MIL number glasses.

Restricting selected glasses

In practice, it is usually necessary to restrict the available glasses for substitution for these reasons:

Not all "glasses" in the catalogs are actually glasses; some are liquids, gases, crystals, or plastics that may not be acceptable for use.

Many glasses are very expensive, heavy, brittle, or have other undesirable mechanical properties.

Perhaps only certain catalogs are desired for the newly chosen glasses; while other catalogs must still be used for other surfaces in the optical system.

For these reasons, ZEMAX provides several optional methods to restrict the pool of available glasses used for substitution:

By defining a specific catalog only be used for any given surface.

By defining a "template" which places limits on the cost, acid resistance, stain resistance, and other properties of substituted glasses.

By "excluding" glasses on a case-by-case basis in the glass catalog.

By defining "penalties" in the merit function which discourage selection of glasses with undesirable properties.

The glass solve dialog box permits specification of a catalog name to use for candidate substituted glasses. If no catalog name is given (i.e. the catalog field is blank) then glasses may be selected from all catalogs selected for use on the System, General dialog box. If a catalog name is given (i.e. Hoya) then only glasses from that one catalog will be selected. This allows different surfaces to select from different catalogs, if required. It is also useful for restricting the number of glasses considered for substitution in a very general way. Note that user defined catalogs may be specified. No extension should be specified, ZEMAX automatically appends the .AGF extension.

A Glass Substitution Template may be defined to restrict which glasses from all catalogs are chosen for substitution based upon their cost, AR, SR, FR, CR, and PR code values. See "Glass Substitution Template" on page 240 for information on defining a template.

To prevent only certain glasses from being chosen during optimization, choose "Exclude Substitution" for the glass that should be avoided on the glass catalog dialog box; see "Description of catalog data" on page 588 for details. The advantage in excluding glasses is that other glasses in the catalog may still be used, without having to define a separate catalog.

Penalties are operands in the merit function such as GCOS, GTCE, and INDX which are used to increase the merit function if glasses have unacceptable properties. This is the least efficient method, because glasses may still be selected, and the subsequent lens optimized, even if the resulting merit function ends up being too high to be considered a good solution. However, it is useful for defining relationships between glasses, such as minimizing the difference of GTCE between two glasses which form a cemented surface (to prevent the lens from breaking under thermal stress).

Suggestions for use

There are several techniques for maximizing the performance of the global optimization algorithms (both Global Search and Hammer):

- 1) If possible, place the stop on the first surface. If your entrance pupil is imbedded in the system, you can model this by using a dummy first surface for the stop and then use a negative thickness to get to surface 2. This enhances performance because ZEMAX does not need to calculate where the entrance pupil is. You can make this thickness variable if appropriate. This technique does not work well for systems which inherently have entrance pupil distortion, such as wide angle lenses.
- 2) Use a marginal ray angle solve on the last glass surface curvature to control the effective focal length instead of the EFFL operand. Dropping one variable decreases the dimensionality of the problem significantly, and the EFFL operator does require one additional ray trace, slowing down the merit function evaluation.
- 3) Use a marginal ray height solve (for zero ray height) on the last thickness before the image surface. Most lenses are well corrected for the 0.7 pupil zone on axis. You may want to use another pupil zone if your intuition guides you to. This solve will ensure that every design generated by Global Search is in focus, just as the curvature solve ensures the correct focal length. These two solves together can improve the performance of the Global Search algorithm by several orders of magnitude. For Hammer optimization, replace the solve with a variable to allow for optimal defocus.
- 4) Use the MNCT and MNET operands. These are essential for avoiding negative center and edge thicknesses. ZEMAX uses these boundary constraints to determine the appropriate range of each variable in the search. The solutions will wander into these unacceptable regions of solution space unless you specifically forbid it.
- 5) Keep your merit function as simple as possible. Using 2 or 3 rings in the default merit function construction (see the chapter "Optimization") is often a good idea. You can always go to more rings for the Hammer optimization.
- 6) Be willing to run several long runs; perhaps one before starting a design problem to get ideas for design forms, a few more to explore various promising designs, and a last run to ensure the design you have settled on cannot be improved (Hammer is particularly good at this last task).
- 7) Use substitute glasses rather than model glasses, especially when using Hammer. There is not much point in using Hammer on a design that uses model glasses, because ultimately the glasses will need to be converted back to real glasses and then optimized again. Substitute glasses are highly useful, especially late in the design process when the design form has been determined.
- 8) The Hammer algorithm may be terminated and then restarted without significant loss of information. As the algorithm proceeds, it automatically saves any improved designs.
- 9) If a Global Search is terminated, the search may be resumed at a later time by choosing resume instead of start.
- 10) Some computes may be programmed to go into "sleep" mode if there is no user activity for a set amount of time. This feature may need to be disabled to allow ZEMAX to continue running, even if there is no input from the keyboard or mouse for long periods.

Summary

Finally, it is again stressed that the global optimum search feature may not ever find the global optimum to any particular problem. It is very good at finding alternative design forms that would have been tedious to discover by hand. It is a very powerful tool to add to your toolbox; it is not a substitute for optical design skills.

The Global Search algorithm does have a strong random component, and therefore no two runs will yield the exact same solutions every time. Sometimes the solutions will be worse, sometimes better, but usually they are just different for runs of similar duration.

Introduction

ZEMAX provides a flexible and powerful tolerance development and sensitivity analysis capability. The tolerances available for analysis include variations in construction parameters such as curvature, thickness, position, index of refraction, Abbe number, aspheric constants, and much more. ZEMAX also supports analysis of decentration of surfaces and lens groups, tilts of surfaces or lens groups about any arbitrary point, irregularity of surface shape, and variations in the values of any of the parameter or extra data. Since the parameter and extra data terms may describe aspheric coefficients, gradient index coefficients, and more, any of these values may also be made part of the tolerance analysis. The various tolerances may be used in any combination to estimate alignment and fabrication error effects on system performance.



ZEMAX always uses exact ray tracing for tolerance analysis; there are no approximations or extrapolations of first order results in the ZEMAX tolerance algorithms.

Tolerances are defined using operands, such as TRAD, which defines a tolerance on a radius. The tolerance operands are automatically saved with the lens file. Tolerance operands are edited on the Tolerance Data Editor available from the Editors group on the main menu bar.

Tolerances may be evaluated by several different criterion, including RMS spot radius, RMS wavefront error, MTF response, boresight error, user defined merit function, or a script which defines a complex alignment and evaluation procedure. Additionally, compensators may be defined to model allowable adjustments made to the lens after fabrication. ZEMAX also allows limits to be placed on the change of a compensator.

Tolerances may be computed and analyzed three ways:

Sensitivity Analysis: For a given set of tolerances, the change in the criterion is determined for each tolerance individually. Optionally, the criterion for each field and configuration individually may be computed.

Inverse Sensitivity: For a given permissible change in criterion, the limit for each tolerance is individually computed. Inverse sensitivity may be computed by placing a limit on the change in the criterion from nominal, or by a limit on the criterion directly. The criterion may be computed as an average over all fields and configurations, or on each field in each configuration individually.

Monte Carlo Analysis: The sensitivity and inverse sensitivity analysis considers the effects on system performance for each tolerance individually. The aggregate performance is estimated by a root-sum-square calculation. As an alternative way of estimating aggregate effects of all tolerances, a Monte Carlo simulation is provided. This simulation generates a series of random lenses which meets the specified tolerances, then evaluates the criterion. No approximations are made other than the range and magnitude of defects considered. By considering all applicable tolerances simultaneously and exactly, highly accurate simulation of expected performance is possible. The Monte Carlo simulation can generate any number of designs, using normal, uniform, parabolic, or user defined statistics.

The basic procedure

Tolerancing a lens consists of these steps:

- 1) Define an appropriate set of tolerances for the lens. Usually, the default tolerance generation feature described in this chapter is a good place to start. Tolerances are defined and modified on the Tolerance Data Editor, available on the Editors menu of the main menu bar.
- 2) Modify the default tolerances or add new ones to suit the system requirements.
- 3) Add compensators and set allowable ranges for the compensators. The default compensator is the back focal distance, which controls the position of the image surface. Other compensators, such as image surface tilt, may be defined. There is no limit to the number of compensators that may be defined.
- 4) Select an appropriate criterion, such as RMS spot radius, wavefront error, MTF, or boresight error. More complex criterion may be defined using a user defined merit function, or for comprehensive flexibility, a tolerance script, which is described later.

5) Select the desired mode, either sensitivity or inverse sensitivity. For inverse sensitivity, choose criterion limits or increments, and whether to use averages or compute each field individually.

6) Perform an analysis of the tolerances.

7) Review the data generated by the tolerance analysis, and consider the budgeting of tolerances. If required, modify the tolerances and repeat the analysis.

Details on this basic procedure will be provided in subsequent sections.

Tolerance operands

A tolerance operand has a four letter mnemonic, such as TRAD for Tolerance Radius. Three integer values, abbreviated Int1, Int2, and Int3, are associated with the mnemonic to identify the surface or surfaces of the lens to which the tolerance applies. Some tolerance operands use the Int numbers for purposes other than defining surface numbers as indicated in the following table.

Each tolerance operand also has a minimum and maximum value. These values refer to the maximum acceptable change from the nominal value. Each operand also has space for an optional comment to make the tolerance set easier to read. The available tolerance operands are listed in the following table, and are described in detail below.

TOLERANCE OPERANDS

Name	Int1	Int2	Int3	Description	Page
Surface Tolerances					
TRAD	Surf #	-	-	Tolerance on radius in lens units.	548
TCUR	Surf #	-	-	Tolerance on curvature in inverse lens units.	548
TFRN	Surf #	-	-	Tolerance on radius in fringes.	548
TTHI	Surf #	Adjust #	-	Tolerance on thickness or position in lens units, with optional adjuster.	548
TCON	Surf #	-	-	Tolerance on conic constant (dimensionless).	549
TSDI	Surf #	-	-	Tolerance on semi-diameter in lens units.	549
TSDX	Surf #	-	-	Tolerance on Standard surface x-decenter in lens units.	549
TSDY	Surf #	-	-	Tolerance on Standard surface y-decenter in lens units.	549
TSTX	Surf #	-	-	Tolerance on Standard surface tilt (TIR) in x in degrees.	549
TSTY	Surf #	-	-	Tolerance on Standard surface tilt (TIR) in y in degrees.	549
TIRX	Surf #	-	-	Tolerance on Standard surface tilt (TIR) in x in lens units.	549
TIRY	Surf #	-	-	Tolerance on Standard surface tilt (TIR) in y in lens units.	549
TIRR	Surf #	-	-	Tolerance on Standard surface irregularity.	550
TEXI	Surf #	Max Term#	Min Term#	Tolerance on surface irregularity using Fringe Zernikes.	550

Name	Int1	Int2	Int3	Description	Page
TEZI	Surf #	Max Term#	Min Term#	Tolerance on surface irregularity using Standard Zernikes.	550
TPAI	Surf #	Parameter #	-	Tolerance on the inverse of parameter value of surface.	552
TPAR	Surf #	Parameter #	-	Tolerance on parameter value of surface.	552
TEDV	Surf #	Extra Data #	-	Tolerance on extra data value of surface.	552
TIND	Surf #	-	-	Tolerance on index of refraction at d light (see notes).	552
TABB	Surf #	-	-	Tolerance on Abbe Vd number (see notes).	553
TCMU	Surf #	Layer #	-	Tolerance on coating multiplier.	553
TCIO	Surf #	Layer #	-	Tolerance on coating index offset.	553
TCEO	Surf #	Layer #	-	Tolerance on coating extinction offset.	553
Element Tolerances					
TEDX	First Surf	Last Surf	-	Tolerance on element x- decenter in lens units.	553
TEDY	First Surf	Last Surf	-	Tolerance on element y- decenter in lens units.	553
TETX	First Surf	Last Surf	-	Tolerance on element tilt about x in degrees.	554
TETY	First Surf	Last Surf	-	Tolerance on element tilt about y in degrees.	554
TETZ	First Surf	Last Surf	-	Tolerance on element tilt about z in degrees.	554
User Defined Tolerances					
TUDX	Surf #	-	-	Tolerance on user defined decenter x.	554
TUDY	Surf #	-	-	Tolerance on user defined decenter y.	554
TUTX	Surf #	-	-	Tolerance on user defined tilt x.	554
TUTY	Surf #	-	-	Tolerance on user defined tilt y.	554
TUTZ	Surf #	-	-	Tolerance on user defined tilt z.	554
Non-Sequential Component Tolerances					
TNMA	Surf #	Object #	Data	Tolerance on NSC material Nd and Vd. The data code is 1 for Nd, and 2 for Vd.	554
TNPS	Surf #	Object #	Data	Tolerance on NSC object position. The data code is 1-6 for x, y, z, tilt x, tilt y, tilt z, respectively.	554
TNPA	Surf #	Object #	Parameter	Tolerance on NSC object parameter.	554

Name	Int1	Int2	Int3	Description	Page
Multi-Configuration Value Tolerances					
TMCO	Row #	Config#	-	Tolerance on multi-configuration editor value.	555

For each tolerance, a minimum and a maximum value are specified on the Tolerance Data Editor. The tolerances are described in detail below.

TRAD: Tolerance on radius

Used to tolerance directly on the radius of curvature. Min and max are the extreme errors in lens units.

For example, if the nominal radius of a surface is 100 mm, and the min and max TRAD values for that surface are -.50 mm and +.50 mm, then the tolerance analysis will be performed with the radius of the surface set to 99.50 mm and 100.50 mm.

TCUR: Tolerance on curvature

Used to tolerance in units of curvature, which is directly related to power. Min and max are the extreme errors in inverse lens units.

For example, if the nominal radius of a surface is 100 mm, then the nominal curvature is 0.01 inverse mm. If the min and max TCUR values for that surface are -.001 and +.001 inverse mm, then the tolerance analysis will be performed with the radius of the surface set to 111.11 mm and 90.909 mm.

TFRN: Tolerance on fringes

Fringe tolerances are very useful when tolerancing flat or large radii surfaces. Min and max are the extreme errors in (dimensionless) fringes. The TWAV operand is used to define the test wavelength. The change in the sag of a surface for small changes in curvature is given approximately by

$$\Delta Z = \frac{r^2}{2} \Delta C,$$

where r is the semi-diameter of the surface. The change in the sag is related to the error in fringes by

$$\Delta Z = \frac{\lambda}{2} N,$$

where N is the number of fringes. The factor of one half assumes a double pass Newton's rings type test. For more information see Malacara, Optical Shop Testing.

TTHI: Tolerance on thickness

TTHI is used to tolerance both absolute positions of elements as well as thicknesses of lenses within element groups.

When ZEMAX creates the default tolerances, it is assumed that all variations in thickness affect only that surface and any surfaces in contact with that element. For example, if the first lens in a contact doublet has a +1.0 mm change in thickness, the front and rear vertex of the second lens both shift by +1.0 mm.

However, since ZEMAX defines the position of all surfaces by using an offset from the previous surface, adding 1.0 mm to the surface will shift all subsequent lenses in the system by +1.0 mm. What is more likely to occur in fabrication is that the +1.0 mm offset would be absorbed by the first air space after the lens group. TTHI can handle this case by allowing a "adjustment" surface to be specified. When ZEMAX creates the default tolerances, the adjustment surface is specified as the first air space which follows the surface being tolerated.

To illustrate, imagine a lens where surface 3 was made of BK7, and surface 4 was made of F2, and surface 5 was air. The nominal thicknesses are 3, 4, and 6 mm, respectively. If a TTHI operand was defined by the default tolerance algorithm for surface 3, an adjustment would be defined for surface 5. If the tolerance value was +.1

mm, then during analysis the thicknesses would be changed to 3.1, 4.0, and 5.9, respectively. Thus, the absolute positions of surfaces 6 through the image surface are unaffected by the change in thickness on surface 3.

The adjustment is optional; to disable it, set the adjustment to the same surface number as the tolerance, such as TTHI 3 3. For some lens systems, such as those that are assembled by stacking spacers in a tube, the adjustment may not be desired.

Int1 is used to define the surface number, Int2 in the adjustment surface number, unless Int2 is equal to Int1. Min and max are the extreme errors in lens units.

TCON: Tolerance on conic

TCON is used to define a tolerance on a conic constant. Min and max are the extreme errors, dimensionless.

TSDI: Tolerance on semi-diameter

TSDI is used to define a tolerance on the semi-diameter of a surface. Min and max are the extreme errors in lens units.

TSDX, TSDY: Tolerance on surface decenters

TSDX and TSDY are used to analyze decentration in X and Y respectively of a Standard surface type. The Int1 value indicates the number of the surface, and this surface must be a Standard surface type. Surfaces other than Standard surface types may be toleranced using the TEDX and TEDI operands described later. The min and max values are the decentering in lens units.

The analysis of TSDX and TSDY uses the Irregular surface type. See the discussion "Tolerancing with the Irregular surface type" on page 555.

TSTX, TSTY: Tolerance on surface tilts

TSTX and TSTY are used to analyze tilts of a Standard surface about the X and Y axes, respectively. The Int1 value indicates the number of the surface, and this surface must be a Standard surface type. Surfaces other than Standard surface types may be toleranced using the TETX and TETI operands described later.

The min and max values are the tilt in degrees about the local X and Y axis of the lens.

See the description of the related TIR operands TIRX and TIRY.

The analysis of TSTX and TSTY uses the Irregular surface type. See the discussion "Tolerancing with the Irregular surface type" on page 555.

TIRX, TIRY: Tolerance on surface TIR

TIRX and TIRY are used to analyze tilts of a Standard surface along the X and Y axes, respectively. The Int1 value indicates the number of the surface, and this surface must be a Standard surface type. Surfaces other than Standard surface types may be toleranced using the TETX and TETI operands described later.

TIRX and TIRY are used to specify a tolerance on total indicator runout, or TIR, which measures the amount of "wedge" in a lens.

The min and max values are twice the amount of "sag" in lens units measured at the maximum radial aperture of the surface where the maximum radial aperture is defined by the semi-diameter of the surface. The change in sag as a function of the x or y normalized coordinate for a TIRX or TIRY value is given by:

$$\Delta Z_x = \frac{TIRX}{2} \rho_x; \Delta Z_y = \frac{TIRY}{2} \rho_y.$$

For example, if the TIRX tolerance is 0.10 mm, then the change in sag at the maximum +x aperture of the lens is 0.05 mm, and the deviation at the maximum -x aperture is -.05 mm, for a "total" TIR of 0.10 mm. A similar discussion applies to TIRY. The min and max values are used to model the tilt of the surface in each direction. The tilt angle that is actually placed on the surface is given by

$$\theta_x = \tan^{-1}\left(\frac{\Delta Z_y}{S}\right); \theta_y = \tan^{-1}\left(\frac{\Delta Z_x}{S}\right),$$

where S is the semi-diameter of the surface. Note that a sag along Y implies a rotation about the X axis, and a sag along X implies a rotation about the Y axis.

The analysis of TIRX and TIRY uses the Irregular surface type. See the discussion “Tolerancing with the Irregular surface type” on page 555.

TIRR: Tolerance on surface irregularity



For a more detailed treatment of irregularity, see the discussion on the TEXI and TEZI operands which follows.

TIRR is used to analyze irregularity of a Standard surface. The Int1 value indicates the number of the surface, and this surface must be a Standard surface type. Analysis of irregularity on surface types other than Standard is not directly supported.

Modeling irregularity is somewhat more problematic than other types of tolerances. This is primarily because irregularity by nature is random, and not deterministic such as a change in radius. Therefore, some assumptions about the nature of the irregularity need to be made in order to perform the analysis. The assumption ZEMAX makes when using TIRR is that the irregularity is half spherical aberration, and half astigmatism. This is less restrictive model than assuming 100% astigmatism, because astigmatism cannot be compensated by focus, and is therefore a more serious defect in the lens.

The min and max values are the irregularity in units of fringes measured at the maximum radial aperture of the surface where the maximum radial aperture is defined by the semi-diameter of the surface. The TWAV operand is used to define the test wavelength.

ZEMAX assumes fringes are measured in a double pass Newton's rings type test. For example, a TIRR of "W" fringes would yield a change in sag of the surface of

$$\Delta z = \frac{\lambda_t W}{4} (\rho^4 + \rho_y'^2),$$

where λ_t is the test wavelength (defined by the TWAV operand), ρ is the normalized radial coordinate, and ρ_y' is the normalized and rotated radial coordinate as defined in “Irregular” on page 317. The change in wavefront optical path is related to the change in sag and the index of refraction of the two media the surface separates:

$$\Delta OP = \Delta z(n_2 - n_1).$$

The analysis of TIRR uses the Irregular surface type.

When computing the Monte Carlo analysis, the angle of the astigmatism is chosen randomly between 0 and 360 degrees. This allows simulation of randomly oriented astigmatic error, which is less severe and more realistic than placing all the astigmatism along the y axis of each element.

See the discussion “Tolerancing with the Irregular surface type” on page 555 for more information.

TEXI: Tolerance on surface irregularity using the Fringe Zernike model



See also the discussion for TEZI on page 551. TEZI is a superior alternative to TEXI.

TEXI is used to analyze random irregular deviations of small amplitude on a surface that is either a Standard, Even Aspheric, or Zernike Fringe Sag surface. Analysis of irregularity on surface types other than Standard, Even Aspheric, or Zernike Fringe Sag is not directly supported. The Int1 value indicates the number of the surface, Int2 defines the maximum Fringe Zernike term (must be between 3 and 37), and Int3 defines the minimum Fringe Zernike term (must be between 2 and the maximum term).

TEXI uses the Zernike Fringe Sag surface (see “Zernike Fringe Sag” on page 341) to model the irregularity rather than using the third order aberration formulas used by TIRR. When using TEXI, the min and max tolerance

values are interpreted to be the *approximate* magnitude of the zero to peak error of the surface in double-pass fringes at the test wavelength. The zero to peak is only a very rough measure of the irregularity. Whether the zero to peak and peak to valley are the same depends upon the particular Zernike term used. The min tolerance value is automatically set to the negative of the max value; this is done to yield both positive and negative coefficients on the Zernike Fringe Sag surface. The vagueness of defining irregularity by the zero to peak measure is the prime disadvantage to TEXI, and the reason why the newer TEZI operand is the preferred alternative.

ZEMAX computes the coefficients on the individual terms of the Zernike polynomial that defines the deviation in the shape of the surface using the following formula:

$$c = \frac{\lambda f}{2\sqrt{n}},$$

where f is the number of double pass fringes, n is the number of Zernike terms used (note n is given by $\text{Int}2 - \text{Int}3 + 1$), and λ is the test wavelength. The coefficients are scaled by one over root n to account for the fact that a random collection of Zernike terms will generally sum in an RSS sense; so the PTV error is not linear with the number of terms. Since it is convenient to specify the approximate overall PTV, the terms on each Zernike are computed according to the formula above. Note there is a "c" value for both the min and the max tolerance in fringes.

For the sensitivity analysis, the surface is converted to a Zernike Fringe Sag surface and all the coefficients of the Zernike polynomial are set to the "c" value from the above equation. Note that since the maximum surface sag deformation is the same for all Zernike Fringe polynomial terms at the edge of the aperture, a "c" value of 0.001 when using 20 Zernike terms will yield a maximum sag deviation of $20c$.

For the Monte Carlo analysis, the surface is converted as for the sensitivity analysis, but each polynomial term is assigned a coefficient randomly chosen which lies between the min and max "c" tolerance values. The random value is chosen using the statistical model selected for the operand; see the STAT command for a discussion.

Generally speaking, if lower order terms are used, the irregularity will be of low frequency, with fewer "bumps" across the surface. If higher order terms are used, there will be higher frequency irregularity, with more "bumps" across the surface. Note the TIRR irregularity operand models the lowest frequency form of irregularity, with just a quadratic and quartic deviation across the surface. TEXI can model much more irregular surfaces, and with 30 or more terms used, about 5-15 "bumps" will typically be seen over the surface.

Because the Zernike Fringe Sag surface sag expression contains portions of both the Standard and Even Aspheric surfaces, either of these surface types may be modeled by the Zernike Fringe Sag surface created with the TEXI operand. If the surface already is a Zernike Fringe Sag surface, then the deviations are added to the polynomial terms already there. If the surface is either a Standard or Even Aspheric surface, the normalization radius of the Zernike Fringe Sag surface is set to the semi-diameter of the surface. If the surface was already a Zernike Fringe Sag surface, then the min and max tolerances are assumed to be measured at the normalization radius already defined.

TEXI always ignores Zernike term 1, the piston term, and sets this value to zero.

TEZI: Tolerance on surface irregularity using the Standard Zernike model



See also the discussion for TEXI on page 550. TEZI is a superior alternative to TEXI.

TEZI is used to analyze random irregular deviations of small amplitude on a surface that is either a Standard, Even Aspheric, or Toroidal surface. Analysis of irregularity on surface types other than these is not directly supported. The $\text{Int}1$ value indicates the number of the surface, $\text{Int}2$ defines the maximum Standard Zernike term (must be between 3 and 231), and $\text{Int}3$ defines the minimum Standard Zernike term (must be between 2 and the maximum term).

TEZI uses the Zernike Standard Sag surface (see "Zernike Standard Sag" on page 343) to model the irregularity on Standard and Even Aspheric surfaces, while Toroidal surfaces use the Zernike terms already supported by the Toroidal surface. When using TEZI, the max tolerance value is the exact RMS error of the surface in lens units. The min tolerance value is automatically set to the negative of the max value; this is done

to yield both positive and negative coefficients for the Zernike terms. The resulting RMS is of course always a positive number whose magnitude is equal to the max tolerance value.

For the sensitivity analysis, the surface is converted to a Zernike Standard Sag or Toroidal surface and all the coefficients of the Zernike polynomial for terms greater than #1 (the "piston" term) are set to a value so that the square root of the sum of the squares of the coefficients yields the specified RMS value. All coefficients are set to the same value.

For the Monte Carlo analysis, the surface is converted as for the sensitivity analysis, but each polynomial term is assigned a coefficient randomly chosen between -1.0 and 1.0, and the resulting coefficients are then normalized to yield the exact RMS tolerance. The random value is chosen using the statistical model selected for the operand; see the STAT command for a discussion.

The number of terms is given by Int2 - Int3 + 1. Generally speaking, if lower order terms are used, the irregularity will be of low frequency, with fewer "bumps" across the surface. If higher order terms are used, there will be higher frequency irregularity, with more "bumps" across the surface. Note the TIRR irregularity operand models the lowest frequency form of irregularity, with just a quadratic and quartic deviation across the surface. TEZI can model much more irregular surfaces.

Because the Zernike Standard Sag surface sag expression contains portions of both the Standard and Even Aspheric surfaces, either of these surface types may be modeled by the Zernike Standard Sag surface created with the TEZI operand. If the surface is Toroidal, the Toroidal surface is retained since the Zernike terms are already supported with this surface type, however, the nominal value of all the Zernike terms must be zero if the nominal surface is Toroidal. The normalization radius for the Zernike terms is set to the semi-diameter of the surface.

TEZI always ignores Zernike term 1, the piston term, and sets this value to zero.

TPAI: Tolerance on the inverse of parameter data

TPAI is used to tolerance on the inverse of the parameter data. The Int1 number is the surface number. The Int2 number indicates the parameter number. See "SURFACE TYPES" on page 277 for tables listing which parameter numbers are used for each surface model. TPAI is very similar to TPAR, except the tolerance deviations are made on the inverse of the parameter value rather than the parameter directly. The toleranced parameter values are defined by

$$P' = \frac{1}{\left(\frac{1}{P}\right) + \Delta}$$

where P is the parameter value, P' is the perturbed parameter value, and Δ is the amount of tolerance perturbation. Min and max are the changes measured in the inverse of whatever units are used by the parameter.

TPAR: Tolerance on parameter data

TPAR is used to tolerance parameter data. The Int1 number is the surface number. The Int2 number indicates the parameter number. See "SURFACE TYPES" on page 277 for tables listing which parameter numbers are used for each surface model. TPAR is useful for determining tolerances on aspheric coefficients, and more. Min and max are the changes in whatever units are used by the parameter.

TEDV: Tolerance on extra data values

TEDV is the tolerance for extra data values. Extra data values are used by the extra data editor in ZEMAX-EE to define aspheric terms, diffraction coefficients, or other data. Int1 identifies the surface, Int2 identifies the extra data number, and min and max are the extreme errors in whatever units are used by the extra data value. See the Chapter "Surface Types" for tables listing which extra data numbers are used for each surface model.

TIND: Tolerance on index

TIND is the tolerance on index of refraction. Int1 identifies the surface, and min and max are the extreme errors in index of refraction. The index error is modeled as an "offset" which is independent of wavelength, unless either

the glass of the surface is a "model" glass defined by the index at d-light, Abbe number, and dPgF, or the glass of the surface is a catalog glass and all defined wavelengths are between 0.3 and 2.5 micrometers. In either of these latter cases, the TIND is interpreted to be a change in the d-light index, which will change the index at all wavelengths in a nonlinear fashion. The change in index is given by the change in the model glass index for the same change in the d-light index as a function of wavelength. For details on the model glass, see "Using model glasses" on page 601.

TABB: Tolerance on Abbe

TABB is the tolerance on the Abbe or Vd number. Int1 identifies the surface, and min and max are the extreme errors in the Abbe number. If either the glass of the surface is a "model" glass defined by the index at d-light, Abbe number, and dPgF, or the glass of the surface is a catalog glass and all defined wavelengths are between 0.3 and 2.5 micrometers, TABB is interpreted to be a change in the d-light Abbe number, which will change the index at all wavelengths in a nonlinear fashion. Otherwise, TABB is ignored.

The change in index is given by an estimate of the differential change in the index at any wavelength as a function of the d-light index, Abbe, and dPgF.

TCMU: Tolerance on coating multiplier

Coating multipliers are described in "Optimizing coatings with ZEMAX" on page 621 and "Surface coating tab" on page 86. The TCMU operand allows tolerancing of the coating multipliers, which in effect places a tolerance on the thickness of individual layers in the optical coating on a surface.

TCIO: Tolerance on coating index offset

Coating index offsets are described in "Optimizing coatings with ZEMAX" on page 621 and "Surface coating tab" on page 86. The TCIO operand allows tolerancing of the coating index offsets, which in effect places a tolerance on the real part of the material index of individual layers in the optical coating on a surface.

TCEO: Tolerance on coating extinction offset

Coating multipliers are described in "Optimizing coatings with ZEMAX" on page 621 and "Surface coating tab" on page 86. The TCEO operand allows tolerancing of the coating extinction offsets, which in effect places a tolerance on the imaginary part of the material index of individual layers in the optical coating on a surface.

TEDX, TEDY: Tolerance on element decenters

TEDX and TEDY are used to analyze decentration in X and Y respectively. The two surface numbers defined by Int1 and Int2 indicate the beginning and ending surfaces of a lens group. The min and max values are the decentring in lens units.

These tolerances require ZEMAX to insert a coordinate break surface before and after the group, and then decenter the entire group as a unit. For this reason, TEDX and TEDY generally should not be used on either side of an existing coordinate break. If you require this capability, see the discussion on TUDX and TUDY.

TEDX and TEDY may also be used to model decenter of surfaces, like TSDX and TSDY. TEDX and TEDY will work for surfaces of any type, including Standard surfaces and non-Standard surfaces, while TSDX and TSDY only work for Standard surfaces. To decenter a single surface using TEDX or TEDY,

set Int1 and Int2 to the same surface number.

To check to see if ZEMAX is actually doing what you want, see the SAVE command discussed in the section "Trouble shooting the tolerance results".

TEDX and TEDY operands may be nested. For example it is possible to analyze the decentration of one surface group defined by surfaces 5 through 20; while simultaneously analyzing decenters of elements defined by surfaces 5-8, 8-12, 14-20, etc. This capability can simulate the alignment errors of elements within an assembly, as well as the alignment error of the assembly as a whole. The nesting rules are fully described in the section "Nesting rules for Monte Carlo analysis" on page 569.

TETX, TETY, TETZ: Tolerance on element tilts

TETX, TETY, and TETZ are used to analyze tilts of either a surface or a lens group about the X, Y, or Z axes, respectively. The two surface numbers defined by Int1 and Int2 indicate the beginning and ending surfaces of a lens group. The min and max values are the tipping angles in degrees.

These tolerances require ZEMAX to insert a coordinate break surface before and after the lens group, and to use a dummy surface at the end of the group to return to the front vertex. The entire group may then be pivoted about a point as a unit. For this reason, TETX/Y/Z should not be used when the range of surfaces defined includes a coordinate break which is related to a following coordinate break outside the surface range controlled by the TETX/Y/Z by pickup solves. The hazard is if the two resulting tilt ranges overlap, then the position of the elements may not be what is intended. If you require this capability, see the discussion on TUTX, TUTY, and TUTZ.

TETX and TETY may also be used to model tilt of single surfaces, sometimes called "wedge", like TSTX and TSTY. TETX and TETY will work for surfaces of any type, including Standard surfaces and non-Standard surfaces, while TSTX and TSTY only work for Standard surfaces. To tilt a single surface using TETX or TETY, set Int1 and Int2 to the same surface number.

TETX and TETY by default pivot about the front vertex of a lens group, however it is often advantageous to pivot about some other point. For example, well-designed lens mounts will pivot about the nodal point of a lens to maintain focus during alignment. This case is easily modeled by ZEMAX through the use of a dummy surface at the nodal point. Change the starting surf (Int1) to be at the nodal point dummy surface, and the pivoting will be about that point. The first surface may be located anywhere with respect to the rest of the lens group; and so the tilt may be about any point.

TETX, TETY, and TETZ operands may be nested. For example it is possible to analyze the tilt of one surface group defined by surfaces 5 through 20; while simultaneously analyzing tilts of elements defined by surfaces 5-8, 8-12, 14-20, etc. This capability can simulate the alignment errors of elements within an assembly, as well as the alignment error of the assembly as a whole. The nesting rules are fully described in the section "Nesting rules for Monte Carlo analysis" on page 569.

TOFF: Tolerance off (can be used for comments)

This operand is a place holder that has no affect on the tolerance analysis. The TOFF operand may be used to enter comments.

TUDX, TUDY, TUTX, TUTY, TUTZ: Tolerance on user defined tilts and decenters

These five tolerances, TUDX, TUDY, TUTX, TUTY, and TUTZ, are used for more general user defined tilts and decenters. The names are mnemonics for Tolerance User Decenter/Tilt X, Y, and Z. These are very similar to the TEDX, TEDY, TETX, and TETY tolerances. The difference is that ZEMAX does not automatically insert the required coordinate break surfaces to achieve the specified decenters and tilts in the tolerances. To use the TUDX, TUDY, TUTX, TUTY, and TUTZ commands, you must have already defined the surface specified by Int1 to be a coordinate break surface. Normally, but not always, there will be a second coordinate break later in the Lens Data Editor that has a pickup solve on the toleranced parameter. The pickup solve may be positive or negative, as the case warrants. This permits some complex pivoting and decentering about arbitrary points, with a certain finesse required. Min and max are in lens units for TUDX and TUDY, and in degrees for TUTX, TUTY, and TUTZ, just like the coordinate break surface.

TNPS, TNPA, TNMA: Tolerances on non-sequential data

TNPS defines tolerances on Non-Sequential Component (NSC) data. Three integer parameters are required: the NSC surface number (use 1 if the program mode is NSC), the object number, and a code. The code is 1, 2, or 3 for the object's X, Y, or Z position, respectively; the code is 4, 5, or 6 for the X-Tilt, Y-Tilt, or Z-Tilt, respectively.

TNPA defines tolerances on NSC parameter data. The surface and object integers are defined in the same way as for TNPS; the third integer is for the parameter number. See "NSC Objects" on page 481 for a detailed explanation of what parameter numbers are used by which object types.

TNMA defines tolerances on the index Nd and Abbe Vd of NSC material data. Three integer parameters are required: the NSC surface number (use 1 if the program mode is NSC), the object number, and a code. The code is 1 or 2 for the material Nd and Vd values, respectively. When using TNMA, the change in the material index is computed exactly as for the sequential glass offset solve, as described in "Glass: Offset" on page 478.

TMCO: Tolerance on multi-configuration data

TMCO specifies the tolerance on any numerical value in the Multi-Configuration Editor. The two integer arguments are the operand row number and configuration number.

Tolerancing with the Irregular surface type

The tolerance operands TSDX, TSDY, TSTX, TSTY, and TIRR all use the Irregular surface type to model the perturbation to the lens surface. For a detailed description of the Irregular surface type, see page 317.

The overwhelming advantages of using the Irregular surface type are speed, simplicity, and flexibility. Any Standard surface type can be converted to an Irregular surface type without the need for dummy surfaces or coordinate breaks. Furthermore, the aggregate effects of tilt, decenter, and irregularity may all be modeled simultaneously using the Monte Carlo analysis.

When ZEMAX computes the tolerance analysis using the operands TSDX, TSDY, TSTX, TSTY, or TIRR, the surface is first converted from a Standard to an Irregular surface type. This is why only Standard surface types are supported when using these operands.

Tolerance control operands

There are also a few tolerance control operands which can be entered in the Tolerance Data Editor. These operands are not tolerances, but are used to define compensators, to save intermediate results for further evaluation, to define statistical properties, and to define the test wavelength for fringe tolerances.

TOLERANCE CONTROL OPERANDS

Name	Int1	Int2	Description
CEDV	Surface #	Extra Data #	Sets an extra data value as a compensator.
CMCO	Operand #	Config #	Sets a multi-configuration operand value as a compensator. See the discussion on CMCO below for important information.
CNPA	Object #	Parameter #	Sets a non-sequential parameter as a compensator. Int3 is the surface number, use 1 for non-sequential mode.
CNPS	Object #	Code	Sets a non-sequential object position value as a compensator. Code is 1-6 for x, y, z position and x, y, z tilt, respectively. Int3 is the surface number, use 1 for non-sequential mode.
COMM	-	-	This operand is used to print the comment, or if no comment is defined, to print a blank line in the tolerance analysis output report.
COMP	Surface #	Code	Sets a compensator. Code is 0 for thickness, 1 for curvature (1 over the radius), 2 for conic.
CPAR	Surface #	Parameter #	Sets a parameter as a compensator. For a list of parameters that control surface attributes, see "SURFACE TYPES" on page 277.
SAVE	File Number	-	Saves the file used to evaluate the tolerance on the prior row in the editor. See the discussion below.
SEED	Seed #	-	Seeds the random number generator for Monte Carlo analysis. Use a value of 0 to choose a random seed.
STAT	Type	# of standard deviations	Sets the type of statistical distribution for randomly selected Monte Carlo parameter analysis. See the discussion below.
TWAV	-	-	This operand sets the test wavelength. The "Min" column is used to edit and display the test wavelength.

The tolerance control operands are described in detail below.

CEDV: Define extra data value compensator

CEDV is used to define the surface number and extra data number of the compensator to use for tolerance analysis. Int1 is used to specify the surface, and Int2 is used to define the extra data number. For example, to specify extra data value 17 of surface 9 is a compensator, use CEDV with Int1 = 9, Int2 = 17.

See "SURFACE TYPES" on page 277 for tables defining which extra data values are used by which surface types.

General comments about min and max values on compensators

The min and max values of compensator operands such as CEDV, CMCO, COMP, and CPAR indicate the maximum change allowed for the compensator. For example, if the nominal value of a compensator is 50, and the min and max values are -1.0 and 1.0, then the compensator will be restricted to lie between 49.0 and 51.0. This bounding of the compensators is ignored if "merit function" or "user script" is selected as the tolerance criterion, although the compensators will still be used.

There is no limit to the number of compensators that may be defined.

CMCO: Define multi-configuration operand compensator

CMCO is used to define the operand number and configuration number of the compensator to use for tolerance analysis. Int1 is used to specify the operand number, and Int2 is used to define the configuration number. For example, to specify multi-configuration operand 6 in configuration 4 is a compensator, use CMCO with Int1 = 6, Int2 = 4. The order of the CMCO operands matters. The CMCO operands must be defined in order of operand number then configuration. All CMCO operands for MCE operand 1 must precede CMCO operands for MCE operand 2, etc. Within each group of CMCO operands using the same MCE operand number, the CMCO operands must be listed in order of configuration number.

See "MULTI-CONFIGURATIONS" on page 579 for information on using multi-configuration operands.

See also "General comments about min and max values on compensators" on page 556.

CNPA: Define non-sequential parameter compensator

CNPA is used to define the object and parameter number of the compensator to use for tolerance analysis. Int1 is used to specify the object number, and Int2 is used to define the parameter number. Int3 is the surface number, use 1 for non-sequential mode.

See also "General comments about min and max values on compensators" on page 556.

CNPS: Define non-sequential position compensator

CNPS is used to define the object and position code of the compensator to use for tolerance analysis. Int1 is used to specify the object number, and Int2 is used to define the position code. Int3 is the surface number, use 1 for non-sequential mode. The position code is 1-6 for x, y, z position and x, y, z tilt, respectively.

See also "General comments about min and max values on compensators" on page 556.

COMP: Define compensator

COMP is used to define the surface number and type of compensator to use for tolerance analysis. Int1 is used to specify the surface, and Int2 is used to define the type. Int2 uses a "code" defined as follows:

Int2 = 0, compensator is thickness

Int2 = 1, compensator is curvature (1 over the radius)

Int 2 = 2, compensator is conic

See also "General comments about min and max values on compensators" on page 556.

CPAR: Define parameter compensator

CPAR is used to define the surface number and parameter number of the compensator to use for tolerance analysis. Int1 is used to specify the surface, and Int2 is used to define the parameter number. For example, to specify parameter 2 of surface 5 is a compensator, use CPAR with Int1 = 5, Int2 = 2.

See “SURFACE TYPES” on page 277 for tables defining which parameters are used by which surface types.

See also “General comments about min and max values on compensators” on page 556.

SAVE: Save sensitivity analysis lenses

The operation of the tolerance feature is not always transparent. To get a closer look at what the tolerance routine is actually doing, use the SAVE operand. SAVE can be inserted after any tolerance you would like to inspect in more detail. For example, suppose you had the tolerance operand TEDX in the Tolerance Data Editor. After reviewing the resulting sensitivities, the results do not appear to make sense. Edit the tolerance operands in the Tolerance Data Editor to add the SAVE command after the TEDX command. For the Int1 value, enter in any integer between 1 and 9999, inclusive. If the integer number is 0, no file is saved. The next time the tolerance analysis is run, ZEMAX will save the files used to compute the TEDX tolerance. The file names will be TSAV_MIN_XXXX.ZMX and TSAV_MAX_XXXX.ZMX for the min and max tolerance analysis, respectively, where XXXX is the integer number specified in the Int1 column.

The lens files will be saved in the same folder as the current lens. ZEMAX will save the lens used to evaluate the min and max values of the preceding tolerance in a ZMX format file. Run the tolerance file in the usual way, and then load one of the files generated. The file will contain the prescription data, coordinate breaks, pickups, variables (which are optimally adjusted compensators), and merit function used to determine the criterion data. This procedure will let you verify the configuration is in fact what you are trying to model, and will give you insight into how ZEMAX adjusts your lens prescription to determine the tolerances.

SEED: Seed the random number generator

By default, every Monte Carlo analysis will produce new, randomly generated systems. There are times when it is useful to have multiple Monte Carlo analysis runs produce the same perturbations every time. The SEED operand may be used for this purpose. If the SEED is any value other than zero, the random number generator will be seeded with the specified value. This value is automatically incremented for each generated system with a single Monte Carlo analysis so that each generated system is unique, yet random and reproducible. Using a value of 0 will select a random seed based upon the system clock, which will generally not be the same each time the Monte Carlo analysis is run.

The placement of the SEED operand matters. All tolerance operands listed under a SEED operand will be affected by the seed value provided. Multiple SEED operands may be placed in the tolerance data editor, so that different operands may use the same seed every time, while others may use random seeds.

STAT: Define statistics

STAT is used to define the statistics "on the fly" during the Monte Carlo analysis. The STAT command takes two integer arguments. Int1 specifies the statistics type: 0 for normal distribution, 1 for uniform distribution, 2 for parabolic distribution, and 3 for user defined distribution statistics. The Int2 value is used only by the normal distribution, and it defines "n", the number of standard deviations between the mean and the tolerance extremes.

The statistics types are defined in detail in “Monte Carlo analysis” on page 566.

TWAV: Test wavelength

This operand sets the test wavelength. When setting the default tolerances, ZEMAX adds a TWAV operand with a value of 0.6328 micrometers (HeNe) for the test wavelength. If no TWAV operand is defined, then ZEMAX defaults to the primary wavelength. More than one TWAV may be placed in the operand list; each operand defines the test wavelength for the operands that follow. Only operands whose min and max values are measured in fringes are affected by this setting. The "Min" column in the Tolerance Data Editor is used to edit and display the test wavelength.

Defining default tolerances

The default tolerances may be defined by selecting Tools, Default Tolerances from the Tolerance Data Editor menu bar. The Tolerance Data Editor may be activated by selecting Editors, Tolerance Data from the main menu.

The default tolerances dialog box consists of several sections grouped by tolerance type:

Surface tolerances

Radius: If this box is checked, then the default radius tolerances will be included. The default tolerances may

be specified by a fixed distance in lens units or by fringes of power at the test wavelength (defined by the TWAV operand). This tolerance is only placed on surfaces which have optical power, which excludes dummy surfaces which have the same index on both sides. If the surface is plano, then the default tolerance is specified as a variation in fringes, even if another option is selected.

Thickness: If checked, a thickness tolerance is specified on each vertex separation. It is assumed that all variations in thickness affect only that surface and any surfaces in contact with that element; therefore, the first air space after the thickness is used as an adjuster. See the detailed discussion in "TTHI: Tolerance on thickness" on page 548.

Decenter X/Y: If checked, decenter tolerances are added to each individual lens surface. Tolerances are defined as a fixed decenter amount in lens units. ZEMAX uses TSDX and TSDY for Standard surface decenters, and TEDX and TEDY for non-Standard surfaces.

Tilt (TIR) X/Y: If checked, a tilt or "total indicator runout" tolerance in either lens units or degrees is added to each lens surface. ZEMAX uses TSTX and TSTY for Standard surface tilts in degrees, TIRX and TIRY for Standard surface total indicator runout (tilt or wedge) in lens units, and TETX and TETY for non-Standard surface tilts in degrees. If lens units are selected, ZEMAX automatically converts the dimension to degrees of tilt if required for the TSTX/TSTY or TETX/TETY operands. The conversion from total indicator runout in lens units to tilt in degrees is given by

$$\theta_x = \frac{\Delta_y}{2S}, \text{ where}$$

S is the semi-diameter of the surface and Δ is the total indicator runout, with a similar expression for θ_y . This expression assumes the angles are relatively small. Note that a length of total indicator runout in the x direction corresponds to a tilt about the y axis, and a length of total indicator runout in the y direction corresponds to a tilt about the x axis.

S + A Irreg: If checked, a spherical and astigmatism irregularity is specified on each Standard surface type. For details, see "TIRR: Tolerance on surface irregularity" on page 550.

Zernike Irregularity: If checked, a Zernike irregularity is specified on each Standard surface type. For details, see "TEZI: Tolerance on surface irregularity using the Standard Zernike model" on page 551. By default, ZEMAX uses Zernike terms 2 through 45.

Element tolerances

Decenter X/Y: If checked, decenter tolerances are added to each lens group. Tolerances may be defined as a fixed decenter amount in lens units.

Tilt X/Y: If checked, a tilt tolerance in degrees is added to each lens group and surface. It is important to note that lens groups are by default tipped about the vertex of the first surface in the group. See the section "TETX, TETY, TETZ: Tolerance on element tilts" on page 554 for information on tipping about some other point.

Index tolerances

Index: TIND is used to model changes in the index of refraction. The units are change in relative index.

Abbe: TABB is used to model changes in the Abbe number. The units are percentage of the catalog Abbe value.

Other controls

Start At Row: This control indicates where in the Tolerance Data Editor the default tolerances should be placed. If the row number is greater than 1, then the new default tolerances will be appended starting at the specified row number. If the row number is -1, then the tolerances will be appended to the end of the current list.

Test Wavelength: The wavelength, in micrometers, used to measured fringes of power or irregularity.

Use Focus Comp: If checked, then a default compensator of the back focus (the thickness prior to the image surface) will be defined. Using at least one compensator can greatly relax certain tolerances, however whether or not compensators should be used depends upon the specifics of the design. Other compensators may be defined. See the Section "Defining compensators" for more information.

Start / Stop At Surface: These controls define the range of surfaces to apply the default tolerances to.

There are also six buttons:

OK: Accept these settings and generate the default tolerances.

Cancel: Close the dialog box without changing the default tolerances.

Save: Save these settings for future use.

Load: Restore the previously saved settings.

Reset: Restore the settings to the default values.

Help: Invoke the help system.

By default, the Monte Carlo analysis that ZEMAX performs draws random values from a Gaussian "normal" distribution. Once the default tolerances are defined, they are saved with the lens file automatically. If additional surfaces are inserted in the Lens Data Editor, the tolerance surfaces are renumbered automatically.

Defining compensators

Many different types of compensators may be defined; thicknesses (most commonly used), curvature, conic constants, any parameter or extra data value, of any surface or surfaces. Multi-configuration operands may also be defined as compensators. The parameter values are useful for using tilts and decenters of particular components for compensation. The surface to be tilted must already be defined as a coordinate break (or perhaps a tilted surface), with any appropriate pickup solves as required.

The image surface focus is specified as a compensator by default. You may add or delete compensators to customize the analysis to your particular situation, and may use as many compensators as you like. Generally, using more compensators will loosen tolerances, and complicate the actual alignment of the system.

All compensators are defined using the COMP, CPAR, CEDV, and CMCO tolerance operands. Defining compensators is described in the section "Tolerance control operands".

When using a script to define the tolerance procedure, the compensators may be altered and different sets of compensators may be defined at different stages of the simulated alignment procedure. For details on this option, see "Using Tolerance Scripts" on page 569.

Performing the tolerance analysis

Once all the tolerance operands and compensators have been defined, the tolerance analysis may be performed. To perform a tolerance analysis, select Tolerancing from the Tools menu in the main program menu bar.

The dialog box which appears has several controls which are described below.

Set-Up tab

Mode: Four modes are supported: Sensitivity, Inverse Limit, Inverse Increment, and Skip Sensitivity.

Sensitivity mode computes the change in the criterion for each of the extreme values of the tolerances.

Inverse Limit computes the value of each tolerance that will yield a criterion being equal to the value specified by the Limit parameter. The Limit parameter is only available if the mode is Inverse Limit. Inverse mode will change the min and max values of the tolerance operands. See "Inverse sensitivity analysis" on page 566.

Inverse Increment computes the value of each tolerance that will yield a *change* in the criterion equal to the value specified by the Increment parameter. The Increment parameter is only available if the mode is Inverse Increment. Inverse mode will change the min and max values of the tolerance operands. For a description of how change is defined, see "Change" below. See "Inverse sensitivity analysis" on page 566.

Skip Sensitivity will bypass the sensitivity analysis and proceed to the Monte Carlo analysis.

Polynomial: The Polynomial feature has three options: None, 3-Term, and 5-Term. If either 3-Term or 5-Term is selected, the tolerance sensitivity will be determined at 4 and 6 points, respectively, spanning the tolerance perturbation range and a polynomial fit of the resulting criterion as a function of the tolerance perturbation will

be computed and displayed. The 5-Term polynomial is:

$$P = A + B\delta + C\delta^2 + D\delta^3 + E\delta^4,$$

where δ is the tolerance perturbation value and P is the resulting criterion. The 3-Term polynomial omits the D and E terms. Because the polynomial fit requires evaluation of the criterion and compensators at multiple points, selecting this option increases the computation time. However, the polynomial values are cached, and once computed, can be used to greatly speed up subsequent tolerance runs with modified tolerance values. See Cache below for details.

Cache: The cache feature can be used to greatly speed up the sensitivity and inverse sensitivity tolerance analysis, but must be used with some care. The first time the tolerance analysis is run, all the tolerances, perturbed criterion, and polynomial fit data (if any) is saved in memory. This cached data may be used to quickly generate the tolerance analysis again. If "Recompute All" is selected, the cached values are ignored, and all the tolerances are computed in full, and the new values stored in the cache. If "Recompute Changed" is selected, then only those tolerances that have been modified are recomputed. Results for unmodified tolerances are extracted from the cache rather than computed, significantly speeding up the analysis. If "Use Polynomial" is selected, then the last polynomial fit computed for each tolerance is used to quickly estimate the perturbed values. ZEMAX will automatically invalidate the cached values if any of the tolerance settings, such as criterion or Sampling, are modified. However, if any significant changes are made to the lens file itself, such as modifying surface data, the cached tolerances may be incorrect but ZEMAX is unable to detect this. Therefore, the cache feature should only be used in a single session where no changes are being made to the lens. The intention of this feature is to significantly speed up a sequence of tolerance analysis where changes are being made to the values of individual tolerances, but no other changes are being made between tolerance runs. If there is ever a doubt as to the validity of the cached data, set "Cache" to "Recompute All" and repeat the tolerance analysis. When using the cache, compensator data is not displayed (although the data presented does consider the affects of compensation), and separation by fields and configurations is not supported.

Change: Change determines how the change in the criterion and the predicted performance are computed. If Linear Difference is selected, then the change due to a tolerance is computed as

$$\Delta = P - N,$$

where P is the perturbed criterion and N is the nominal criterion. If Change is set to RSS Difference, the change is computed as:

$$\Delta = S(P - N)\sqrt{|P^2 - N^2|},$$

where the function S(x) returns +1 if the $x \geq 0$ and -1 otherwise.

Force Ray Aiming On: If the lens being toleranced is already using ray aiming, then ray aiming will be used when evaluating the tolerances. If ray aiming is not already on, then ray aiming will only be used if this box is checked. Generally, using ray aiming yields more accurate results but slower computation speed. For preliminary or rough tolerance work, leave the switch at the default "off", but for final or precise work, set the switch "on".

CPU's: Selects the number of CPU's over which to spread the tolerance analysis task. More than 1 may be selected, even on a single CPU computer, in which case the single CPU will time share the multiple simultaneous tasks. The default is the number of processors detected by the operating system.

Separate Fields/Configs: If checked, the criterion will be computed and displayed for all field positions in all configurations individually. If unchecked, the criterion will be computed as an average over all field positions in all configurations. During inverse tolerancing, if Separate Fields/Configs is unchecked, the inverse analysis is done on the overall tolerance criterion, which is usually an average of performance over all fields in all configurations. The problem with using the average performance is that some fields or configurations may be significantly degraded by tolerance defects while other fields or configurations are not, and the average may not reveal the severity of the loss of performance at a few fields or configurations. If this option is checked, then ZEMAX computes the criterion at each field in each configuration individually, and verifies that each field meets either the Limit or Increment value. For Inverse Increment mode, ZEMAX computes the nominal

performance at every field position and reduces the tolerances until the criterion at every field is degraded no more than the increment value. See “Inverse sensitivity analysis” on page 566.

Criterion tab

Criterion: This control is used to specify what shall be used as the criterion for tolerancing. The options are:

RMS spot size (radius, x, or y): The best choice for systems which are not close to the diffraction limit; for example, systems with more than one wave of aberration. This is the fastest option. ZEMAX always uses a centroid reference for tolerancing.

RMS wavefront: The best choice for systems which are close to the diffraction limit; for example, systems with less than one wave of aberration. This is nearly as fast as RMS spot radius. ZEMAX always uses a centroid reference for tolerancing.

Merit Function: Uses whatever merit function has been defined for the lens. This is useful for user-defined tolerancing criterion. User-defined merit functions may also be required for systems with non-symmetric fields, or with significant surface apertures which remove rays. If the user-defined merit function is used, no boundary constraints on the compensators are automatically added to the merit function. If the merit function was generated by ZEMAX as one of the default merit functions, make sure the "Assume Axial Symmetry" option was checked OFF, see “Assume Axial Symmetry” on page 487 for a discussion. Separate Fields/Configs is not supported when using this criterion.

Geometric or Diffraction MTF (average, tangential, or sagittal): The best choice for systems which require an MTF specification. If average is selected, the average of the tangential and sagittal responses is used. Diffraction based MTF tolerancing can be problematic if the tolerances are loose, because the diffraction MTF may not be computable or meaningful if the OPD errors are too large. This is especially true if the spatial frequency is high enough and the performance poor enough for the MTF to go to zero at some frequency below the frequency being analyzed. MTF is the slowest of the default criterion. The frequency at which the MTF is computed is specified in the "MTF Frequency" control.

Boresight error: Boresight error is defined as the radial chief ray coordinate traced for the on axis field divided by the effective focal length. This definition yields a measure of the angular deviation of the image. ZEMAX models boresight error by using just one BSER operand (see the Optimization chapter for details on BSER). Any element or surface decenters or tilts will tend to deviate the chief ray and increase the values of the BSER operand. Boresight error is always computed at the primary wavelength in radians. Boresight error should only be used with radially symmetric systems. Note boresight error gives no indication of image quality; it is a measure of the deviation of the axis ray.

RMS Angular radial, x, or y aberrations. This tolerance criterion is best used for afocal systems (see “Afocal Image Space” on page 103). The angular aberrations are based upon the direction cosines of the output rays.

User Script: A user script is a macro-like command file which defines the procedure to be used for alignment and evaluation of the lens during tolerancing. For details on this option, see the section “Using Tolerance Scripts” on page 569. Separate Fields/Configs is not supported when using this criterion.

Limit: When using Inverse Limit mode, this control is active and is used to define the limit on the criterion for computing inverse tolerances. For example, suppose the Change is Linear Difference, Criterion is RMS Spot Radius, and the nominal RMS of a system is 0.035. If Limit is set to 0.050, then ZEMAX will compute the min and max value of each tolerance that degrades the performance to an RMS of 0.050. The Limit value must represent worse performance than the nominal system has. When using MTF as a merit, then Limit is the lower bound on MTF since lower numbers indicate worse performance. See “Inverse sensitivity analysis” on page 566. The nominal value for the currently selected criterion may be computed by pressing the "Check" button adjacent to the Limit edit window.

Increment: When using Inverse Increment mode, this control is active and is used to define the limit on the change in the criterion for computing inverse tolerances. For example, suppose the Change is Linear Difference, Criterion is RMS Spot Radius, and the nominal RMS of a system is 0.035. If Increment is set to 0.01, then ZEMAX will compute the min and max value of each tolerance that degrades the performance to an RMS of 0.045. The Increment value must be positive to represent a degradation in performance. When using MTF as a merit, the Increment is still positive, and ZEMAX automatically interprets this number as a

decrease in MTF from the nominal. See "Inverse sensitivity analysis" on page 566. The nominal value for the currently selected criterion may be computed by pressing the "Check" button adjacent to the Increment edit window.

Sampling: Sampling is used to set how many rays are traced when computing the tolerance criterion. Higher sampling traces more rays, and gives more accurate results. However, the execution time increases. If the selected criterion is RMS spot or RMS wavefront, then the sampling value is an integer that refers to the number of rays traced along a radial arm of the pupil in the Gaussian quadrature technique (see "Selecting the pupil integration method" on page 485 for a description of this technique). The number of arms is always twice the number of rays along each arm. If MTF is the selected criterion, then the sampling refers to the pupil grid size, with a sampling of 1 yielding a 32 x 32 grid, sampling of 2 yielding a 64 x 64 grid, etc. Usually, a sampling of 3 or 4 is sufficient for quality optical systems. Systems with high amounts of aberration require higher sampling than systems with low aberration. The most reliable method for determining the best sampling setting is to run the tolerance at a sampling of 3, then again for a sampling of 4. If the results change moderately, then use the higher setting. If they change substantially, check the next higher sampling setting. If the results change little, go back to the lower sampling. Setting the sampling higher than required increases computation time without increasing accuracy of the results.

MTF Frequency: If MTF is selected as the Merit, then this control is active and is used for defining the MTF frequency. MTF frequency is measured in MTF Units, see "MTF Units" on page 106.

Comp: This control determines how the compensators are evaluated. "Optimize All" will use the optimization capability of ZEMAX to determine the optimum values of all defined compensators. Although optimization is accurate, it is slow to execute. There are two different optimization algorithms available, DLS and OD. If "Optimize All (DLS)" is selected, ZEMAX will execute one cycle of the Orthogonal Decent algorithm then execute the Damped Least Squares algorithm. If "Optimize All (OD)" is selected, Then the Orthogonal Descent algorithm only is used. For more information, see "Performing an optimization" on page 526. If "Paraxial Focus" is selected, only the change in paraxial back focus error is considered as a compensator; all other compensators are ignored. Using Paraxial Focus is very useful for rough tolerancing, and is significantly faster than using "Optimize All". If "None" is selected, no compensation will be performed, and any defined compensators will be ignored.

Config: For multi-configuration lenses, indicates which configuration should be used for tolerancing. The selected configuration only will be considered, and the configuration number will be printed on the final report. If "All" is selected, then all configurations will be considered at once.

Fields: Generally speaking, the field definitions used for optimization and analysis are inadequate for tolerancing. for example, a rotationally symmetric lens may use field definitions of 0, 7, and 10 degrees. For tolerancing purposes, the lack of symmetry in the field definitions may cause inaccurate results when analyzing tilt or decenter tolerances. When constructing a merit function to use for tolerancing ZEMAX can use three different field settings:

Y-Symmetric: ZEMAX computes the maximum field coordinate, then defines new field points at +1.0, +0.7, 0.0, -0.7, and -1.0 times the maximum field coordinate, in the Y direction only. All X field values are set to zero. This is the default for rotationally symmetric lenses.

XY-Symmetric: Similar to Y-Symmetric, except there are 9 field points used. The 5 Y-Symmetric points are used, and -1.0, -0.7, +0.7, and +1.0 are added in the X axis direction only.

User Defined: Use whatever field definitions exist in the current lens file. This option is required when using vignetting factors, tolerancing multiple configuration lenses, or using tolerance scripts. It is also highly recommended when tolerancing non rotationally symmetric lenses or lenses with complex field weighting that user defined fields be used.

If user defined fields are used, no adjustment of the weights is performed. For the Y-Symmetric case, the center point has a weight of 2.0, all others have a weight of 1.0. For the XY-Symmetric case, the center point has a weight of 4.0, all others unity.

Cycles: This determines how rigorously ZEMAX will attempt to optimize compensator values. If set to Auto, then ZEMAX will call the optimizer in "Auto" mode, which will run the optimizer until the optimization of the compensators has converged. For rough tolerancing, a low number, such as 1, 2, or 3, may be used. If the compensators are difficult to optimize, a higher setting may increase accuracy. If too few optimization cycles are selected, then the tolerances will be pessimistic; the predicted performance will be worse than the actual performance. The "Auto" setting is the safest to use. Higher settings increase accuracy at the expense of run

time. This setting is only used if Comp is set to "Optimize All".

Script: The name of the script, if using the User Script criterion. User scripts must be text files ending in the extension TSC and be in the <data>\Tolerance folder (see "Folders" on page 68).

Status: This control is used by the tolerance algorithm to provide status messages during computation of the nominal criterion when "Check" is pressed.

Monte Carlo tab

Monte Carlo Runs: This control is used to specify how many Monte Carlo simulations should be performed. The default setting of 20 will generate 20 random lenses which meet the tolerances specified. See the section "Monte Carlo simulations" for more details. The number of Monte Carlo runs may be set to zero, which will omit the Monte Carlo analysis from the summary report.

Monte Carlo Save: This option is used to save a specific number of lens files generated during the Monte Carlo analysis. The value specifies the maximum number of lens files to be saved. For example, suppose 20 is selected. After the first Monte Carlo lens is generated, the lens file will be saved in the file "MC_T0001.ZMX". The second Monte Carlo lens file will be generated, then saved in "MC_T0002.ZMX", and so on. Only the first 20 Monte Carlo lenses will be saved (the last will be "MC_T0020.ZMX"). If fewer than 20 Monte Carlo runs were requested, then fewer than 20 lenses would be saved. Be sure you have no lens files with the names "MC_Txxxx.ZMX", as ZEMAX will overwrite these files without warning as the lenses are saved. The purpose of this feature is to allow further study on the lenses being generated by the Monte Carlo feature. Note that the file names given here assume that no File Prefix to the Monte Carlo file names was defined; see File Prefix for details.

Statistics: Choose either a Gaussian "normal" distribution, "uniform", or "parabolic" distribution. This setting is only used by the Monte Carlo analysis; see that discussion for details on the statistics mode as well as the "STAT" command which provides detailed control over the statistics model used.

File Prefix: If a prefix string is provided, then the string will be prepended to the names of the Monte Carlo files generated. For example, if the prefix string is "Fast_Doublet_" then the first Monte Carlo file saved will be file name "Fast_Doublet_MC_T0001.ZMX", and similar names for the remaining saved files.

Save Best and Worst Monte Carlo Files: If checked, the best and worst Monte Carlo lens files generated will be saved. The file names are "prefixMC_BEST" and "prefixMC_WORST" where "prefix" is the File Prefix string.

Overlay MC Graphics: If checked, each open analysis graphical window (such as a ray fan or MTF graph) will be updated and overlaid for each Monte Carlo generated lens. The resulting plots are useful for showing the total range of performance for the simulated lenses. Analysis graphs which do not automatically change scale and do not depend upon surface numbers are the most useful, such as MTF, MTF vs. Height, Encircled Energy, and other plots that allow a user defined fixed scale, such as ray fan plots. Plots which change scale dynamically or require a fixed range of surfaces, such as layout plots, will not work reliably and are not supported. Static, Text, and Editor windows are not updated. Overlaid graphic windows will be flagged as static after the tolerance analysis is complete. The time to compute each of the analysis graphics for each MC lens obviously will slow down the tolerance analysis.

Display tab

Show Descriptions: If checked, a full description of the meaning of each tolerance operator will be provided in the analysis report. If unchecked, only the tolerance operator abbreviations will be listed.

Show Compensators: By default, the compensator values are not printed out during the sensitivity analysis. If this box is checked, then each compensator value will be printed along with the change in the criterion for each tolerance.

Hide All But Worst: If checked, printing of all the sensitivity data will be turned off. This is useful for decreasing the size of the output report. The "hide" check box is normally used in conjunction with the "Show worst" control.

Show worst: The "Show worst" control can be set to sort and display only some number of tolerance operands; this permits limiting the printout to the most severe tolerances only.

Output To File: If a valid file name is provided, the tolerance analysis output window text will be saved to the specified file. The path is always the same as the current lens file. Only provide a file name without a path in this box, such as OUTPUT.TXT. Leave this control blank to not save the output text.

Other buttons

There are also six buttons on the bottom of the dialog box:

Save: Save the currently selected tolerance options in a user selected file name for future use. The file ends in the extension TOP and these options may be used by other features.

Load: Restores tolerance options from a previously saved file.

Reset: Restore the settings to the default values.

OK: Performs the tolerance analysis using the current options.

Cancel: Exits the dialog box without performing the tolerance analysis, restoring the previous settings.

Apply: Saves the changed options for the session, even if Cancel is subsequently selected.

Once all of the options have been selected, press OK to begin the tolerance analysis. Details on the method of calculation are provided in the next section.

How ZEMAX computes the tolerance analysis

ZEMAX begins the tolerance analysis by saving the lens in a temporary file which will be used to restore the lens after the tolerancing is complete. All of the changes made during tolerancing will ultimately be discarded, and the original lens file will be restored unmodified. The exception is during inverse sensitivity analysis, where the tolerance data min and max limits may be altered.

ZEMAX then removes all of the variables. Solves are left in place, however, solves may cause problems during tolerancing; see "Tolerancing with solves" on page 577 for more information.

The tolerance operators are read, and the compensating parameters defined by COMP and CPAR specified are set as variables. The compensators defined may be altered during the analysis if a tolerance script is being used. If ray aiming is on in the lens being toleranced, or if the "Force Ray Aiming On" switch is checked, then ray aiming will be used in evaluating the tolerances, otherwise, ray aiming is left off. Tolerances computed using ray aiming are more accurate, but execution speed is slower. See "Ray Aiming" on page 106 for details on ray aiming.

ZEMAX then uses the Criterion, Field, MTF Frequency, and Sampling settings on the tolerance dialog box to define an appropriate merit function for the tolerancing. Since this is done to the temporary file only, the original merit function defined for the lens is not disturbed.

Boundary constraints are added to the merit function to limit the compensators to the min and max boundaries specified using the COMP and CPAR commands. If using a user defined merit function, or if the "Comp" control is set to "Paraxial Focus", boundary constraints on the compensators are ignored.

If "Comp" is set to "Optimize All", ZEMAX then calls the optimization function to find the best values for the defined compensators. Otherwise, the back focal position is set to paraxial focus. The resulting lens file is then saved for subsequent use by the tolerancing algorithm.

The criterion of this lens is considered to be the "nominal". Note that the nominal criterion value generally will not be the same as the criterion value that is reported on the optimization or merit function editor windows, because ZEMAX constructs a new merit function just for use in tolerancing. ZEMAX then proceeds with either sensitivity, inverse, or Monte Carlo analysis as described in subsequent sections.

Evaluating compensators

ZEMAX has two methods for evaluating compensators selected by the "Comp" control: either "Optimize All" or "Paraxial Focus". If "Paraxial Focus" is selected, then several assumptions are made which greatly speeds up the tolerance evaluation. First, all defined compensators and compensator boundary constraints are ignored. Back focus is implemented as a compensator using a paraxial focus error solve. This means the focus is adjusted to maintain the amount of defocus in the nominal lens design, without attempting to reoptimize the back focus precisely. These assumptions greatly speed up the tolerance evaluation. If "Optimize All" is selected, then ZEMAX uses the slower but more accurate optimization algorithm to find the best value for all the compensators.

"Paraxial Focus" mode is fast, accurate, and should be used if back focus is the only compensator, and the system is well described by paraxial rays, such as systems with rotational symmetry. "Paraxial Focus" should not be used when the system is highly non-symmetric or there are multiple compensators. When in doubt, run both

modes and compare the results. Since "Paraxial Focus" does not precisely optimize the compensators, "Paraxial Focus" results are generally more pessimistic than "Optimize All" results.

Sensitivity analysis

For the sensitivity analysis, each tolerance is evaluated independently using the following algorithm:

The temporary lens is restored.

The parameter whose tolerance is being evaluated is adjusted to the extreme minimum value. For example, if the tolerance being evaluated is TRAD, and the nominal value is 100 mm, with a minimum tolerance of -0.1 mm, the radius is set to 99.9. If the tolerance is an element tilt or decenter tolerance, dummy coordinate breaks are inserted as required to model the perturbation. For surface tilts and decenters, such as TSDX, TSDY, TSTX, TSTY, TIRX, or TIRY, the Irregular surface type is used if the surface is initially of type Standard.

The compensators are adjusted.

The resulting criterion are printed on the report.

The procedure is repeated for the maximum tolerance.

This basic algorithm is repeated for each tolerance operand.

The value of the sensitivity analysis is that the tolerances which are too loose will, in general, have greater contributions to the increase in the criterion than other tolerances. This technique allows the designer to identify surfaces which are highly sensitive to certain errors, such as tilts or decenters. Different surfaces also will, in general, have very different sensitivities to the various errors. The sensitivity analysis aids in identification of which tolerances need to be tightened, and which might be loosened. It is also valuable for finding the optimum (and minimum) number of compensators, and the required range of adjustment. There are in fact many more applications for this feature; for example, designing lens mounts to maximize compensation leverage.



The sensitivity analysis aids in identification of which tolerances need to be tightened, and which might be loosened.

The amount of output can be overwhelming, especially for lenses with many elements and a corresponding large number of tolerances. Often, the sensitivity to the tolerances varies widely over all the possible tolerances. The "Show worst" control is extremely useful for summarizing the worst offenders, because it sorts the tolerances by contribution to the criterion and then prints them in descending order. The "Hide all but worst" control turns off the bulk of the printing if only the worst offenders are of interest.

The RSS estimated change

After all the individual tolerances are computed, ZEMAX then computes a variety of statistics, the most important of which is the estimated change in the criterion and associated estimated performance. ZEMAX uses a Root Sum Square (RSS) assumption for computing the estimated changes in the performance. For each tolerance, the change in performance from the nominal is squared and then averaged between the min and max tolerance values. Note that the change for each individual tolerance itself can be computed as a linear difference or as an RSS difference; see "Change" under "Performing the tolerance analysis" on page 559. The total change can then be computed from a sum over all tolerances of the averaged squared values as follows:

$$C^2 = \sum_i \frac{\Delta_{i,min}^2 + \Delta_{i,max}^2}{2}$$

The average of the min and max squared change is taken because the min and the max tolerance cannot both occur simultaneously, and so summing the squares would result in an overly pessimistic prediction. The final predicted, as-built performance is then computed. If the Change is computed using Linear Difference, the final criterion is defined as:

$$F = N + C,$$

where N is the nominal criterion. If the Change is computed using RSS Difference, the final performance is:

$$F = \sqrt{N^2 + C^2}.$$

Inverse sensitivity analysis

If an inverse sensitivity is being performed, the tolerances are computed the same way as for the sensitivity analysis. However, the calculation is performed iteratively inside a loop while adjustments are made to the min and max tolerances. For Inverse Limit mode, the adjustments are made until the resulting criterion is approximately equal to the Max Criterion. For Inverse Increment mode, the adjustments are made until the resulting change in criterion is approximately equal to the Increment value.

For example, if the mode is Inverse Limit, the Criterion is RMS spot radius, the nominal criterion is 0.035, and the Max Criterion is 0.050, ZEMAX will adjust the tolerances until the criterion is 0.050. If the mode is Inverse Increment, the Criterion is RMS spot radius, the nominal criterion is 0.035, and the Increment is 0.010, ZEMAX will adjust the tolerances until the merit is 0.045. For Inverse Limit, the Max Criterion must be greater than the nominal, except for MTF Criterion, in which case the Max Criterion must be less than the nominal, or an error message will be generated and the analysis terminated. For Inverse Increment, the Increment must always be positive.

For the inverse sensitivity iteration to converge to the desired tolerance, the criterion must smoothly and monotonically change with the tolerance value. For criterion such as RMS spot radius, this is normally the case. However, some criterion, most notably the MTF related tolerance criterion, may have multiple peaks and valleys as the tolerance value is changed. This is usually caused by the starting tolerance being too loose. If the inverse sensitivity cannot iterate to the desired criterion, reduce the starting tolerance limit and repeat the analysis.

If the option to "Separate Fields/Configs" is not selected, the criterion ZEMAX uses to perform the inverse analysis is the overall criterion, which is typically an average over all fields and configurations. The problem with using the average performance is that some fields or configurations may be significantly degraded by tolerance defects while other fields or configurations are not, and the average may not reveal the severity of the loss of performance at a few fields or configurations. If "Separate Fields/Configs" is checked, then ZEMAX computes the criterion at each field in each configuration individually, and verifies that each field meets either the Max Criterion or Increment value. For Inverse Increment mode, ZEMAX computes the nominal performance at every field position and reduces the tolerances until the criterion at every field is degraded no more than the increment value.

If the starting value of the tolerance yields performance better than the Max Criterion or Increment value would dictate, the tolerance is not adjusted. This means the tolerance will never be loosened; it can only be tightened during inverse sensitivity analysis. For example, if the nominal is 0.035, and the Max Criterion is 0.050, and the initial tolerance yields a criterion of 0.040, the tolerance will not increase. To compute the actual limit, the tolerance must first be loosened on the Tolerance Data Editor, and then inverse sensitivity repeated. This is done to prevent looser than necessary tolerances. Generally, tolerances that are looser than some reasonable value do not decrease manufacturing costs.

The estimated change in performance is computed in the same way as for the sensitivity analysis, using the newly adjusted tolerances. The inverse sensitivity analysis aids in tightening individual tolerances so no one defect contributes too much to performance degradation.

Note that the inverse sensitivity is computed for each tolerance individually. The overall degradation in performance estimate will still be given by the RSS of all the individual increases. Inverse tolerancing can be disabled for individual operands; which prevents the inverse tolerance algorithm from tightening tolerances even if the estimated performance exceeds the specified limits. The option to disable inverse tolerancing is on the "Operand Type" dialog box listed under the "Edit" menu of the Tolerance Data Editor.

Monte Carlo analysis

Unlike the sensitivity and inverse sensitivity analysis, the Monte Carlo analysis simulates the effect of all perturbations simultaneously.



The Monte Carlo analysis simulates the effect of all perturbations simultaneously.

For each Monte Carlo cycle, all of the parameters which have specified tolerances are randomly set using the defined range of the parameter and a statistical model of the distribution of that parameter over the specified range. By default, all parameters are assumed to follow the same normal distribution with a total width of four standard deviations between the extreme minimum and maximum allowed values. For example, a radius of 100.00 mm with a tolerance of +4.0/ - 0.0 mm will be assigned a random radius between 100.00 and 104.00 mm, with a normal distribution centered on 102.00 mm and a standard deviation of 1.0 mm.

This default model may be changed using the STAT command. Each tolerance operand may have a separate definition for the statistics, or operands with the same statistical distribution form may be grouped together. All tolerance operands which follow a STAT command use the statistical distribution defined by that STAT command. As many STAT commands as desired may be placed in the tolerance data editor.

The STAT command accepts two arguments, Int1 and Int2. Int1 should be set equal to 0 for normal, 1 for uniform, 2 for parabolic, and 3 for user defined statistics. For normal statistics only, the Int2 value should be set to the number of standard deviations between the mean and extreme values of the parameter.

The available statistical distributions are described below.

Normal statistical distribution

The default distribution is a modified Gaussian "normal" distribution of the form

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(\frac{-x^2}{2\sigma^2}\right), -n\sigma \leq x \leq n\sigma.$$

The modification is that the randomly selected value x (measured as an offset from the midpoint between the two extreme tolerances) is restricted to lie within "n" standard deviations of zero. The default value of "n" is two, however "n" may be changed using the Int2 argument of the STAT command defined earlier. This is done to ensure that no selected value will exceed the specified tolerances. The standard deviation is set to half the range of the tolerances divided by n. For example, if "n" is 2, and a thickness is nominally 100 mm with a tolerance of plus 3 or minus 1 mm, then the selected value will be drawn from a normal distribution with a 101 mm mean and a range of plus or minus 2 mm, with a standard deviation of 1.0. If "n" is 5, then the standard deviation will be 0.4. The larger "n" is, the more likely the selected value will be closer to the mean value of the tolerance extremes. The smaller "n" is, the more the normal looks like the uniform distribution.

Uniform statistical distribution

The uniform distribution is of the form

$$p(x) = \frac{1}{2\Delta}, -\Delta \leq x \leq \Delta.$$

The Δ value is one half of the difference between the max and min tolerance values. Note that the randomly selected value will lie somewhere between the specified extreme tolerances with uniform probability.

Parabolic statistical distribution

The parabolic distribution is of the form

$$p(x) = \left(\frac{3x^2}{2\Delta^3}\right), (-\Delta \leq x \leq \Delta),$$

where Δ is defined exactly the same as for the uniform distribution. The parabolic distribution yields selected values that are more likely to be at the extreme ends of the tolerance range, rather than near the middle, as for the normal distribution.

User defined statistical distribution

The user defined statistical distribution is defined by a text file with tabulated distribution data. A general probability function may be defined as

$$p(x_i) = T_i, 0.0 \leq x_i \leq 1.0,$$

where the T values are tabulated for some number of discrete X values. This general distribution may be numerically integrated, and from the integral of the tabulated values an estimated X may be randomly generated with the statistics matching that of the tabulated distribution.

The format of the file is two columns of data as follows:

```
X1 T1
X1 T2
X3 T3
...etc.
```

where the X values are monotonically increasing floating point numbers between 0.0 and 1.0, inclusively, and the T values are the probabilities of getting that X value. Note ZEMAX uses a probability distribution that covers the range from 0.0 to 1.0, and so the first X1 value defined MUST be equal to 0.0 (it may have any probability T1, including zero), and the last value defined MUST have an Xn value of 1.0. Up to 200 points may be used to define the distribution between X = 0.0 and X = 1.0; a warning will be issued if too many points are listed.

For each tolerance operand subsequently defined (until another STAT command is reached) the defined min and max values will determine the actual range of the random variable X. For example, if a value of 100.0 has a tolerance of -0.0 and +2.0, the probability distribution will extend over the range of 100.0 to 102.0.

Once the data is defined in a file, the file extension must be *.UDD and the file must be placed in the <data>\Tolerance folder (see "Folders" on page 68), and the file name (with extension) must be placed in the comments column of the tolerance data editor on the same line as the STAT command. The STAT type must be set to "3".



The file name of the user defined probability distribution must be placed in the comments column of the tolerance data editor for the STAT operand.

One possible distribution could be:

```
0.0 0.0
0.1 0.5
0.2 1.0
0.3 0.5
0.4 0.0
0.5 0.5
0.8 4.0
1.0 5.0
```

Note the X data values do not need to be evenly spaced; finer spacing may be used in regions of rapid change in probability. This distribution has two peaks, and the higher peak is highly skewed to the maximum side of the distribution.

The user defined statistical distribution is very flexible, and may be used to model any probability distribution, including skewed, multiply peaked, or measured statistical probability data. Multiple distributions may be defined and used in the same tolerance analysis.

Discussion of Monte Carlo analysis method

Note that going from normal to uniform to parabolic statistics yields successively a more pessimistic analysis, and thus more conservative tolerances.

For each cycle, the compensators are adjusted and then the criterion and compensator values are printed. After all of the Monte Carlo trials a statistical summary is provided.

The value of the Monte Carlo analysis is estimating the performance of the lens considering all of the tolerances simultaneously. Unlike the sensitivity analysis, which identifies the "worst offenders" in the system, the Monte Carlo analysis estimates the real-world performance of a system which meets the tolerances specified. The statistical summary provided may be highly useful for lens systems which are mass-produced. Lenses which are one of a kind, of course do not follow these statistics because of the inadequate sampling. However, the Monte Carlo analysis is still useful because it indicates the probability that a single lens will meet the required specification.

Nesting rules for Monte Carlo analysis

When performing the Monte Carlo analysis, all tolerances are considered simultaneously. It is possible to define element tilt and decenter operands that conflict or are ambiguous if certain rules are not followed.

Tolerances such as TEDX, TEDY, TETX, and TETY require ZEMAX to insert coordinate breaks before and after the surface group, and then tilt or decenter the group as a whole. The tilts and decenters performed by the first coordinate break must be "undone" by the second coordinate break. This can only be accomplished if the vertices of the first and second coordinate breaks are at the same location in 3D space. ZEMAX ensures this condition through the use of pickup and position solves.

This method will fail if the surface ranges specified by multiple tolerance operands overlap. For example, if there is a TETX on surfaces 3-8, and then a subsequent TETX on surfaces 5-12, the first coordinate break which tilts the 5-12 group will change the location of the surfaces 5-8, and the second coordinate break of the first group will be moved. In this case, ZEMAX cannot ensure the coordinate breaks will work as intended. In fact, overlapping coordinate breaks do not have a unique interpretation, and it is difficult to imagine an optical system where they even have a meaningful physical interpretation.

Tolerances may be nested however, because nesting implies an unambiguous order to the tilts and decenters. For example, a TETX on 5-12 followed by a TETX on 5-9 and a TETX on 10-12 is perfectly reasonable. This order would simulate the tilting of an assembly, composed of multiple elements, each of which may be tilted within the assembly itself.

The nesting rules are:

- 1) All element tilt and decenter tolerances must be nested.
- 2) The outermost pair of surfaces in each nested group must be first.

Here is an example of a valid set of operands:

```
TETX 5 12  
TETX 5 10  
TETX 11 12
```

Here is an invalid set of operands:

```
TETX 5 12  
TETX 9 15  
TETX 5 15
```

The second operand is invalid because it partially overlaps the first (this violates rule 1). The third operand is invalid because although nested with operand 1, it is the outermost range of the two (this violates rule 2). The second operand must be deleted or modified, but the third operand could be placed in front of the first operand to make a legal operand list:

```
TETX 5 15  
TETX 5 12
```

Note that an operand is considered nested even if it shares one or both surface limits with a previous operand, so that TETX 5-15 may be followed by another TETX 5-15 or by TETX 5-12 or TETX 13-15, but not TETX 4-13.

Using Tolerance Scripts



This feature is only available in the EE edition of ZEMAX.

Tolerance Script overview

Tolerance scripts are macro-like command files that define a procedure to follow for evaluating the performance of a lens during tolerancing. Scripts allow simulation of a complex alignment and evaluation procedure for a lens. With scripts, the following actions can be taken to evaluate a perturbed lens:

- Add or remove compensators.
- Load new merit functions.
- Optimize defined compensators using any merit function.
- Monitor and report any value ZEMAX can compute via a merit function (which includes essentially any value ZEMAX can compute since the merit function may call ZPL macros). Other data, such as Zernike coefficients, may also be computed.
- Write data out to text or binary files for later analysis.
- Save lenses at any stage of the analysis to a ZMX file.

Any number of these operations may be combined in the script; so multiple merit functions and compensator groups may be defined. The script is executed many times during the evaluation of the tolerance analysis, including:

- Once to compute the nominal data.
- Twice for each tolerance operand in the sensitivity analysis, (once for the min and once for the max tolerance).
- Multiple times for each tolerance operand in the inverse sensitivity analysis (inverse sensitivity may require iteration).
- Once for each generated Monte Carlo lens.

The tolerance script files must be text files ending in the extension TSC, and be placed in the <data>\Tolerance folder (see "Folders" on page 68).

When using tolerance scripts, the min and max boundary values of compensators are not enforced; see "General comments about min and max values on compensators" on page 556 for more information. If any compensators are defined in the script, then no compensators are allowed in the Tolerance Data Editor.

The Tolerance Script commands

The tolerance script commands are defined and described below.

!

Purpose:

Adds a comment to the script.

Syntax:

! A comment line!

Discussion:

The "!" symbol is used to define comments in the script that are ignored during execution of the script.

CEDV

Purpose:

Defines a new extra data value compensator.

Syntax:

CEDV surf param

Discussion:

CEDV defines a new extra data value compensator. The value "surf" is the surface number in the original lens file. ZEMAX automatically renumbers the surface if additional coordinate break or other dummy surfaces were inserted by the tolerance program. The value for "param" corresponds to the extra data number for the surface.

CLEARCOMP

Purpose:

Removes all current compensators.

Syntax:

```
CLEARCOMP
```

Discussion:

CLEARCOMP removes all current compensators. New compensators must be defined before calling OPTIMIZE or OPTIMIZE-OD.

CLOSEFILE

Purpose:

Closes the current output file.

Syntax:

```
CLOSEFILE
```

Discussion:

CLOSEFILE closes the current output file, and no more data will be written to the output file until the next OPENFILE is executed. See OPENFILE below.

CMCO

Purpose:

Defines a new multi-configuration operand compensator.

Syntax:

```
CMCO operand config
```

Discussion:

CMCO defines a new multi-configuration operand compensator. The value "operand" is the operand (row) number in the original lens file. The value for "config" corresponds to the configuration number.

CNPA

Purpose:

Defines a new non-sequential parameter compensator.

Syntax:

```
CNPA object parameter surface
```

Discussion:

CNPA defines a new non-sequential parameter compensator. The value "object" is the object number in the original lens file. The value for "parameter" corresponds to the parameter number. Surface is the surface number of the NSC group, use 1 for non-sequential program mode.

CNPS

Purpose:

Defines a new non-sequential position compensator.

Syntax:

```
CNPS object code surface
```

Discussion:

CNPS defines a new non-sequential position compensator. The value "object" is the object number in the original lens file. The value for "code" is 1-6 for x, y, z position and x, y, z, tilt, respectively. Surface is the surface number of the NSC group, use 1 for non-sequential program mode.

COMP

Purpose:

Defines a new compensator.

Syntax:

```
COMP surf code
```

Discussion:

COMP defines a new compensator. The value "surf" is the surface number in the original lens file. ZEMAX automatically renumbers the surface if additional coordinate break or other dummy surfaces were inserted by the tolerance program. The value for "code" is 0 for thickness, 1 for curvature, 2 for conic.

CPAR**Purpose:**

Defines a new parameter compensator.

Syntax:

```
CPAR surf param
```

Discussion:

CPAR defines a new parameter compensator. The value "surf" is the surface number in the original lens file. ZEMAX automatically renumbers the surface if additional coordinate break or other dummy surfaces were inserted by the tolerance program. The value for "param" corresponds to the parameter number for the surface.

FORMAT**Purpose:**

Specifies the numerical precision format for subsequent REPORT statements.

Syntax:

```
FORMAT m.n [EXP]
```

Discussion:

The integers m and n are separated by a decimal point. The value m refers to the total number of characters to be printed, even though some of them may be blank. The value n refers to the number of places to display after the decimal point. Therefore, FORMAT 8.4 will cause subsequent REPORT statements to print 8 characters, with 4 numbers after the decimal point. FORMAT .5 will cause REPORT to show 5 decimal places, and as many total places as needed. FORMAT only affects numeric output from REPORT. If a number is too large to fit within the m decimal places, then the m portion of the FORMAT statement will be ignored. The optional keyword EXP after the m.n expression indicates exponential notation should be used. The default format is 16.8 EXP.

Example:

```
FORMAT 18.9 EXP
```

GETMERIT**Purpose:**

Evaluates the current merit function.

Syntax:

```
GETMERIT
```

Discussion:

GETMERIT should be used prior to any REPORT commands if the merit function has not yet been evaluated. LOADMERIT, OPTIMIZE, and OPTIMIZE-OD also update the current merit function.

HEADER**Purpose:**

Prints a text header to the beginning of an output file. The header can be used to identify to which tolerance operand or Monte Carlo file the subsequent data belongs to.

Syntax:

```
HEADER "text"
```

Discussion:

The HEADER command adds a text line to the output file when using OPENFILE in TXT mode. The HEADER command will output the user supplied text string, followed by either the tolerance operand number (during

sensitivity analysis) or the Monte Carlo file number being evaluated, or an indication that the nominal performance is being evaluated. For sensitivity analysis the min is done first, then the max.

LOADMERIT

Purpose:

Loads a new merit function to use as the tolerance criterion.

Syntax:

```
LOADMERIT filename.mf  
LOADMERIT "C:\SOME\PATH\SOME MERIT FILE NAME.MF"
```

Discussion:

The merit function file must reside in the <data>\MeritFunction folder (see “Folders” on page 68) if no path is specified. The file should be in the proper format as saved using the Tools, Save option of the Merit Function Editor, or may be a full ZMX file, in which case only the merit function portion of the file is read and loaded.

When the merit function is loaded, it replaces any existing merit function. Any operands in the merit function which reference surface numbers should be the surface numbers for the original, unaltered lens. ZEMAX automatically adjusts the surface numbers as required if additional coordinate break or other dummy surfaces were inserted by the tolerance program. The new merit function is then evaluated. The resulting merit function value is what will be returned to the tolerance analysis program as the criterion unless a subsequent GETMERIT, LOADMERIT, OPTIMIZE, or OPTIMIZE-OD command is executed.

NOMINAL

Purpose:

Allows some control over how the nominal performance is evaluated.

Syntax:

```
NOMINAL keyword
```

Discussion:

The NOMINAL script command allows some control over the evaluation of the script for the computation of the nominal performance only. The supported keywords are:

OPT_OFF: Turns off optimization of the compensators. Any subsequent OPTIMIZE or OPTIMIZE-OD commands are ignored.

OPT_ON: Turns on optimization of the compensators. Any subsequent OPTIMIZE or OPTIMIZE-OD commands are executed.

The NOMINAL command is only executed during computation of the nominal system performance, and not during the analysis of individual tolerances.

OPENFILE

Purpose:

Opens a text or binary file for subsequent output of data using REPORT.

Syntax:

```
OPENFILE "FILENAME" mode
```

Discussion:

FILENAME may be any valid file name, with the full path and extension included, such as "C:\DATA\My File.DAT". The mode parameter should be BIN for a binary file, or TXT for a text file. OPENFILE will open a new data file the first time the script is run for a tolerance analysis run. Each subsequent call to the script will append data to the same file.

Data is written to the file whenever the REPORT OR HEADER commands are executed. However, the text string is not included in REPORT commands; just the numerical data the REPORT command prints to the tolerance output file. If mode is TXT, the file will be text with each REPORT value written on a separate line. If mode is BIN, each value will be written out as a 64-bit double precision number.

Because data is written out each time the script is called, the amount of data in the file depends upon how many tolerance operands and Monte Carlo runs are being performed. The first call to the script is to compute the

nominal criterion, then the sensitivity analysis, then the Monte Carlo analysis. OPENFILE should not be used during an inverse sensitivity run, because there are multiple calls to execute the script and no practical way to discern which data belongs to which operand.

The file is automatically closed after the script is finished executing. Multiple OPENFILE commands may be used within the same file to write different data to different files. See also CLOSEFILE.

OPTIMIZE

Purpose:

Optimizes the system using the current merit function and compensators.

Syntax:

```
OPTIMIZE n
```

Discussion:

OPTIMIZE calls the damped least squared optimizer and executes "n" cycles of optimization. If n is zero or is omitted the optimizer runs in automatic mode, terminating when convergence is detected.

OPTIMIZE-OD

Purpose:

Optimizes the system using the current merit function and compensators.

Syntax:

```
OPTIMIZE-OD n
```

Discussion:

OPTIMIZE-OD calls the orthogonal descent optimizer and executes "n" cycles of optimization. If n is zero or is omitted the optimizer runs in automatic mode, terminating when convergence is detected.

PERTURB

Purpose:

Randomly perturbs a parameter in the lens.

Syntax:

```
PERTURB type int1 int2 int3 stat nstd min max
```

Discussion:

PERTURB is used to randomly change a numerical parameter in the lens. The integer parameter type must be 0 for surface data, 1 for parameter data, 2 for extra data, 3 for multi-configuration data, 4 for non-sequential component position data, and 5 for non-sequential parameter data. The arguments to PERTURB change meaning depending upon the value of type as described below.

Type = 0, surface data: int1 is used for the surface number, int2 is 0 for thickness, 1 for radius, 2 for conic, and 3 for semi-diameter. The value of int3 must be set to zero.

Type = 1, parameter data: int1 is the surface number, int2 is the parameter number, int3 must be set to zero.

Type = 2, extra data: int1 is the surface number, int2 is the parameter number, int3 must be set to zero.

Type = 3, multi-configuration data: int1 is the configuration number, int2 is the operand number, int3 must be set to zero.

Type = 4, NSC position data: int1 is the surface number, int2 is the object number, int3 must be 1 for x, 2 for y, 3 for z, 4 for tilt x, 5 for tilt y, or 6 for tilt z.

Type = 5, NSC parameter data: int1 is the surface number, int2 is the object number, int3 is the parameter number.

The stat integer indicates the statistical model to use: 0 for normal (Gaussian), 1 for uniform, and 2 for parabolic. Only these 3 distributions are currently supported. The nstd is a floating point value that indicates the number of standard deviations from the center of the distribution to either edge. A small value for nstd (less than 1.0) will yield a nearly uniform distribution. Large values of nstd will yield a sharply peaked distribution, with most

of the randomly generated values being clustered close to the center of the range. Only the normal distribution uses the nstd value, however, it must be provided and should be set to zero if the other stat types are used.

The min and max values indicate the range over which the random values are selected.

For example, to randomly perturb the thickness of surface 6 by -0 to +1 with a normal distribution covering plus/minus 3 standard deviations, the syntax would be:

```
PERTURB 0 6 0 0 0 3.0 0.0 1.0
```

The resulting random value will be between 0.0 and 1.0, and will likely be close to 0.5 because the number of standard deviations is moderately large. This random value will then be added to the current value for the thickness of surface 6. Note that PERTURB works on all values, whether they are compensators or not. To perturb values controlled by multi-configuration operands, use type 3 to perturb the operand directly rather than the parameter under control.

The PERTURB command is ignored while evaluating the nominal criterion.

QUIET

Purpose:

Controls whether data sent to a file using REPORT is also sent to the tolerance analysis output window..

Syntax:

```
QUIET ON  
QUIET OFF
```

Discussion:

The default setting for QUIET is OFF. If QUIET is OFF, then the output of all REPORT commands is sent to the tolerance analysis window. If QUIET is ON, then the output of REPORT commands will not be sent to the tolerance analysis window. Any output will still be sent to any files opened with OUTPUTFILE, however.

REPORT

Purpose:

Send numerical or text data to an output file.

Syntax:

```
REPORT "text" operand  
REPORT "text" SZERNIKE term field wave sampling maxorder  
REPORT "text" -1  
REPORT_NB (see discussion under syntax 4 below)
```

Discussion:

There are different ways to use the REPORT or REPORT_NB commands:

Syntax 1: REPORT will print any user defined text to the tolerance output window, along with the value of any operand in the currently loaded merit function. The value for operand is an integer corresponding to the operand number (row) of the value to print (note that any value may be computed in a merit function, and if it is not needed for optimization, it may be weighted to zero and is still available for reporting). If operand is zero, then the value of the entire merit function will be printed out. The numeric format is set by the FORMAT command.

Syntax 2: REPORT, when followed by a text string and then the keyword SZERNIKE, can be used to compute and report Standard Zernike coefficients (See "Zernike Standard Coefficients" on page 199). If the term number is zero, then the other values for field, wave, sampling, and maxorder are used to define the Zernikes to be computed. The computed values are then saved in a buffer, and no output is generated. If term is between 1 and the most recent maxorder value used in a previous SZERNIKE call, then that Zernike term is reported. A sample session follows:

```
REPORT "Compute Only!" SZERNIKE 0 1 1 1 4  
REPORT "Zernike 1 = " SZERNIKE 1  
REPORT "Zernike 2 = " SZERNIKE 2  
REPORT "Zernike 3 = " SZERNIKE 3  
REPORT "Zernike 4 = " SZERNIKE 4
```

Note that only when the term is zero are the Zernikes actually computed; the following calls retrieve the computed Zernikes one at a time.

Syntax 3: REPORT, when followed by a text string in quotes and then the integer -1, can be used to print any text string (or blank line if the string is empty).

Syntax 4: If the command REPORT_NB is used instead of REPORT, then no line break is printed at the end of the line, so multiple data may be displayed on one line. The last report issued should use REPORT and not REPORTNB to properly terminate the line. The remaining arguments are supported as described for syntax 1, 2, and 3.

See also "HEADER" on page 572 and "QUIET" on page 575, and "OPENFILE" on page 573.

SAVE

Purpose:

Saves the current lens file.

Syntax:

SAVE n

Discussion:

SAVE will save the current lens file in a ZMX file; with the name TSAVnnnn.ZMX where nnnn is the four digit integer representation of the number "n". For example, if "n" is 6, the file will be saved as TSAV0006.ZMX.

UPDATE

Purpose:

Updates the current lens.

Syntax:

UPDATE

Discussion:

Updates the current lens. Use after PERTURB commands.

Tolerance Script example

As an example, suppose an optical system is assembled from multiple lens elements. As part of the alignment and evaluation process the system is adjusted as follows:

Element number 2 is decentered until an axial test beam is centered on the image.

Element number 4 is then shifted along the axis until the proper magnification is achieved.

Finally, the back focus is adjusted to maximize the on-axis MTF only.

The distortion is then measured and logged.

The MTF at 5 field points is evaluated.

Assume that element number 2 is decentered using parameters 1 and 2 of surface 3, which is an existing coordinate break, the position of element 4 is thickness 10, and the back focus is thickness 15. Furthermore, assume a merit function which centers the axis ray on the image surface is defined in CENTER.MF, the magnification enforcing merit function is MAGNIFY.MF, the MTF merit function is MTF.MF, and the evaluation merit function is EVALUATION.MF. The corresponding script might look like this:

```
! clear any existing compensators for a clean start
CLEARCOMP
! load the centering merit function
LOADMERIT CENTER.MF
! define the two compensators to decenter element 2
CPAR 3 1
CPAR 3 2
! optimize 4 cycles
OPTIMIZE 4
! clear the decenters, load the magnification merit, and adjust thickness 10
CLEARCOMP
LOADMERIT MAGNIFY.MF
COMP 10 0
OPTIMIZE 4
```



```

! now load the MTF merit function, and adjust back focus
CLEARCOMP
LOADMERIT MTF.MF
COMP 15 0
OPTIMIZE 4
! load the evaluation merit function, report the distortion and 5 MTF values
! these should be the first 6 operands in EVALUATION.MF.
CLEARCOMP
LOADMERIT EVALUATION.MF
REPORT "Distortion      = " 1
REPORT "MTF at field 1 = " 2
REPORT "MTF at field 2 = " 3
REPORT "MTF at field 3 = " 4
REPORT "MTF at field 4 = " 5
REPORT "MTF at field 5 = " 6

```

At the end of the script, the entire merit function value resulting from the last LOADMERIT or OPTIMIZE command is returned as the criterion which ZEMAX reports and uses.

Tolerancing multi-configuration (zoom) lenses

Tolerance analysis can be performed on all or any specific configuration of a multiple configuration lens. Select the desired configuration from the "Configuration #" drop down box. When using inverse tolerances, successive tolerance analysis in each configuration individually will yield the tightest tolerances which apply to all configurations collectively.

In some cases, ZEMAX cannot perform tolerances correctly on surfaces controlled by multi-configuration operands. For example, consider an element decenter defined by the tolerance operand TEDX 3 4, while the radii and thicknesses of these surfaces are also being controlled by thermal operands in the multi-configuration editor. In this case ZEMAX will issue an error message. The solution is to separate the two effects by adding dummy surfaces before and after the element. The new dummy surfaces (3 and 6) are added before and after the real lens surfaces (now numbered 4 and 5). Now the decenter can be specified using TEDX 3 6, and the thermal controls may still be applied to surfaces 4 and 5.

Tolerancing with solves

Generally speaking, delete all ray-based solves before tolerancing a lens. Ray based solves, such as marginal ray height solves, no longer have meaning when the optical system contains tilted or decentered elements. Even if the system is initially rotationally symmetric, most tolerance commands such as TETX and TETY force the system to be non-rotationally symmetric. Curvature solves, such as the F/# solve, would dynamically alter the curvature of a surface during tolerancing, which does not reflect the behavior of the real system.

There are occasions where a pickup solve is required, and thus should not be deleted. For example, if the lens is used in double-pass, and a tilt on one element implies a tilt on a subsequent element, then a pickup solve can be used to pickup the tilts from the first element. For this case use the TUTX and TUTY tolerance operands.

Trouble shooting the tolerance results

See the discussion "SAVE" on page 576.

If any of the calculated tolerance data show a value of "Infinity", this means the criterion could not be evaluated for the specified tolerance. Usually the criterion cannot be evaluated because of total internal reflection of some of the ray targets. The statistical data which follows the sensitivity analysis is usually meaningless if any of the tolerances have a criterion value of infinity. One solution is to decrease by a factor of two or more the tolerance value, and then repeat the analysis.

Optimizing for tolerance sensitivity

See "Optimizing tolerance sensitivity" on page 530.

Pitfalls when tolerancing

One possible error is non-physical propagation of rays when using tilt tolerances such as TETX and TETY. If two elements are separated by a very small air space or dummy surface, the default tolerances will include tilting of each element independently. If the element spacing is small, it is possible for one of the elements to be tipped

to the point where it "collides" with the other element. This could not happen in practice, however, for small amounts of tilt the tolerance values are still a good indication of performance.

If the nominal criterion the tolerance algorithm reports is different from the expected value, check to see if ray aiming is "off" in the lens file (see "Ray Aiming" on page 106 for details about ray aiming). Try using ray aiming "on" and re-optimizing. Generally, if there is a big difference between criterion when ray aiming is turned on and off, then it should be left on.

Summary

The tolerance routine is quite flexible and powerful. ZEMAX uses no approximations, extrapolations, or estimations to compute the tolerances. For this reason, it gives useful results for both conventional and complex systems. It is very important to appreciate that tolerancing is a complex procedure, and the algorithms used by ZEMAX to manipulate the lens data are not infallible. Therefore, it is the designer's responsibility to verify that the program is computing reasonable results.

Introduction

ZEMAX supports a very general capability for defining, analyzing, and optimizing optical systems which are used in multiple configurations. Multi-configurations are used to design zoom lenses, or to optimize lenses tested and used at different wavelengths, or optical systems used in various configurations, to name a few. Like other ZEMAX features, multi-configuration is well integrated. However, like tolerancing, it requires a little more care and practice to become proficient.

ZEMAX uses a substitution procedure for defining multiple configurations. The configurations are differentiated by different values for the same parameter. For example, in a zoom lens, the spacings between various elements may take on more than one value. Each set of values used together forms one configuration.

The first step

By far the most important step is to define one configuration using the normal ZEMAX mode first. It is a good idea to start with the most complex configuration first. If all of the configurations have the same number of elements, pick any one of them. Once you have the basic configuration defined, it is time to define new configurations which are variations of the first. The first configuration does not need to be optimized yet, you can optimize across configurations later.

Select Editors, Multi-Configuration from the main menu bar. The spreadsheet which appears is the multi-configuration editor (MCE). Using the menu bar on the MCE, both configurations (columns) and operands (rows) can be inserted and deleted. The insert and delete keys will also add or remove new operand rows. The data entered on the MCE will be saved automatically whenever you save your lens file.

Summary of multi-configuration operands

To change the operand type, double click on the type column. A dialog box will appear where the type and number of the multi-configuration operand can be changed. The operands are also summarized in the following table.

SUMMARY OF MULTI-CONFIGURATION OPERANDS

Type	Numbers 1,2,3	Description
AFOC	Ignored	Afocal Image Space mode. See "Afocal Image Space" on page 103.
APDF	Ignored	System apodization factor. See also APDT.
APDT	Ignored	System apodization type. Use 0 for none, 1 for Gaussian, 2 for cosine cubed. See also APDF.
APDX	Surface #	Surface aperture X- decenter. The surface must have a defined aperture (NOT semi-diameter).
APDY	Surface #	Surface aperture Y- decenter. The surface must have a defined aperture (NOT semi-diameter).
APER	Ignored	System aperture value. If the system aperture type is float by stop size this is the semi-diameter of the stop surface. See also SATP.
APMN	Surface #	Surface aperture minimum value. The surface must have a defined aperture (NOT semi-diameter). This same operand also works to control the first parameter of all surface aperture types, such as the X-Half Width on rectangular and elliptical apertures.
APMX	Surface #	Surface aperture maximum value. The surface must have a defined aperture (NOT semi-diameter). This same operand also works to control the second parameter of all surface aperture types, such as the Y-Half Width on rectangular and elliptical apertures.

Type	Numbers 1,2,3	Description
APTP	Surface #	Surface aperture type. The integer values indicating the aperture type are 0-10 for none, circular aperture, circular obscuration, spider, rectangular aperture, rectangular obscuration, elliptical aperture, elliptical obscuration, user aperture, user obscuration, and floating aperture; respectively.
CADX	Surface #	Surface Tilt/Decenter after surface decenter x.
CADY	Surface #	Surface Tilt/Decenter after surface decenter y.
CATX	Surface #	Surface Tilt/Decenter after surface tilt x.
CATY	Surface #	Surface Tilt/Decenter after surface tilt y.
CATZ	Surface #	Surface Tilt/Decenter after surface tilt z.
CAOR	Surface #	Surface Tilt/Decenter after surface order. Use 0 for Decenter then Tilt, or 1 for Tilt then Decenter.
CBDX	Surface #	Surface Tilt/Decenter before surface decenter x.
CBDY	Surface #	Surface Tilt/Decenter before surface decenter y.
CBTX	Surface #	Surface Tilt/Decenter before surface tilt x.
CBTY	Surface #	Surface Tilt/Decenter before surface tilt y.
CBTZ	Surface #	Surface Tilt/Decenter before surface tilt z.
CBOR	Surface #	Surface Tilt/Decenter before surface order. Use 0 for Decenter then Tilt, or 1 for Tilt then Decenter.
CONN	Surface #	Conic constant.
COTN	Surface #	The name of the coating, if any, to be applied to the surface.
CROR	Surface #	Coordinate Return Orientation. Use 0 for none, 1 for Orientation only, 2 for Orientation XY, and 3 for Orientation XYZ. See "Using the Coordinate Return" on page 85.
CRSR	Surface #	Coordinate Return Surface. See "Using the Coordinate Return" on page 85.
CRVT	Surface #	Curvature of surface.
CSP1	Surface #	Curvature solve parameter 1.
CSP2	Surface #	Curvature solve parameter 2.
CWGT	Ignored	The overall weight for the configuration. This number only has meaning relative to the weights in other configurations.
EDVA	Surface, Extra Data Number	The EDVA operand is used to assign multiple values to the extra data values. This operand requires 2 numerical arguments: the surface number and the extra data value number.
FLTP	Ignored	Field type. Use 0 for angle in degrees, 1 for object height, 2 for paraxial image height, 3 for real image height.
FLWT	Field #	Field weight.
FVAN	Field #	Vignetting factor VAN.
FVCX	Field #	Vignetting factor VCX.

Type	Numbers 1,2,3	Description
FVCY	Field #	Vignetting factor VCY.
FVDX	Field #	Vignetting factor VDX.
FVDY	Field #	Vignetting factor VDY.
GCRS	Ignored	The global coordinate reference surface.
GLSS	Surface #	Glass.
GPEX, GPEY	obsolete	obsolete
GPJX	Ignored	Global Jones polarization vector component Jx.
GPJY	Ignored	Global Jones polarization vector component Jy.
GPIU	Ignored	Global polarization state "is unpolarized", 1 if polarization state is unpolarized, otherwise state is polarized.
GPPX	Ignored	Global polarization state phase x.
GPPY	Ignored	Global polarization state phase y.
HOLD	Ignored	Holds data in the multi-configuration buffer, but has no other effect. Useful for temporarily turning off one operand without losing the associated data.
IGNR	Surface #	Ignore This Surface status. Use 0 to consider the surface, and 1 to ignore the surface. See "Ignore This Surface" on page 78.
LTTL	Ignored	Lens title. The string length is limited to 32 characters.
MABB	Surface #	Model glass Abbe.
MCOM	Surface #	Surface comment.
MDPG	Surface #	Model glass dPgF.
MIND	Surface #	Model glass index.
MOFF	Ignored	An unused operand, may be used for entering comments.
NCOM	Surface, Object	Modifies the comment for non-sequential objects in the NSC Editor. The string value is limited to 32 characters.
NCOT	Surface, Object, Face #	Modifies the coating on each face for non-sequential objects in the NSC Editor.
NGLS	Surface, Object	The material type for non-sequential objects in the NSC Editor.
NPAR	Surface, Object, Parameter	Modifies the parameter columns for non-sequential objects in the NSC Editor.
NPOS	Surface, Object, Position	Modifies the x, y, z, tilt x, tilt y, and tilt z position values for non-sequential objects in the NSC Editor. The position flag is an integer between 1 and 6 for x, y, z, tilt x, tilt y, and tilt z, respectively.

Type	Numbers 1,2,3	Description
NPRO	Surface, Object, Property	Modifies various properties of NSC objects. Property is an integer value indicating what data is controlled: 1 - Inside of object number 2 - Reference object number 3 - Do Not Draw Object (0 = no, 1 = yes) 4 - Rays Ignore Object (0 = no, 1 = yes) 5 - Use Pixel Interpolation (0 = no, 1 = yes) 201-212 - User defined gradient index parameters 301-312 - User defined diffraction parameters 401-416 - User defined bulk scatter parameters 500 - Media is birefringent. Use 0 for false and 1 for true, 501 - Birefringent mode. use 0-3 for ordinary and extraordinary rays, ordinary rays only, extraordinary rays only, and waveplate mode, respectively. 502 - Birefringent Reflections. Use 0 for refracted and reflected rays, 1 for refracted rays only, and 2 for reflected rays only. 503-505 - Birefringent crystal axis orientation x, y, and z.
PAR1	Surface #	Parameter 1. Obsolete, use PRAM instead.
PAR2	Surface #	Parameter 2. Obsolete, use PRAM instead.
PAR3	Surface #	Parameter 3. Obsolete, use PRAM instead.
PAR4	Surface #	Parameter 4. Obsolete, use PRAM instead.
PAR5	Surface #	Parameter 5. Obsolete, use PRAM instead.
PAR6	Surface #	Parameter 6. Obsolete, use PRAM instead.
PAR7	Surface #	Parameter 7. Obsolete, use PRAM instead.
PAR8	Surface #	Parameter 8. Obsolete, use PRAM instead.
PRAM	Surface, Parameter	Parameter value. This operand controls any of the parameters. See "Parameter data" on page 277.
PRES	Ignored	Air pressure in atmospheres. Zero means vacuum, 1 means normal air pressure. See "Environment" on page 108.
PRWV	Ignored	Primary wavelength number.
PSCX	Ignored	X Pupil Compress. Used for ray aiming, see "Ray Aiming" on page 106.
PSCY	Ignored	Y Pupil Compress. Used for ray aiming, see "Ray Aiming" on page 106.
PSHX	Ignored	X Pupil Shift. Used for ray aiming, see "Ray Aiming" on page 106.
PSHY	Ignored	Y Pupil Shift. Used for ray aiming, see "Ray Aiming" on page 106.
PSHZ	Ignored	Z Pupil Shift. Used for ray aiming, see "Ray Aiming" on page 106.
PSP1	Surface #	Parameter solve parameter 1 (the pickup surface). This operand requires 2 numerical arguments: the surface number and the parameter number.
PSP2	Surface #	Parameter solve parameter 2 (the scale factor). This operand requires 2 numerical arguments: the surface number and the parameter number.

Type	Numbers 1,2,3	Description
PSP3	Surface #	Parameter solve parameter 3 (the offset). This operand requires 2 numerical arguments: the surface number and the parameter number.
PUCN	Ignored	Used for picking up a range of values from a previous configuration. If a positive integer configuration number is provided, then all values below the PUCN operand will be picked up from the configuration number specified. If the configuration value is negative, then a negative pickup will be used. If the configuration number is zero, then the values below the PUCN operand will not have pickup solves applied. Note two PUCN operands can be used to define the beginning and end of a range of values to be picked up. All specified configuration numbers must be less than the configuration the PUCN data is provided for.
RAAM	Ignored	Ray aiming. Use 0 for off, 1 for paraxial, and 2 for real.
SATP	Ignored	System aperture type. Use 0 for Entrance Pupil Diameter, 1 for Image Space F/#, 2 for Object Space NA, 3 for Float By Stop Size, 4 for Paraxial Working F/#, 5 Object Cone Angle. See also APER.
SDIA	Surface #	Semi-diameter.
STPS	Ignored	Stop surface number. The stop can be moved to any valid surface number (excluding the object and image surfaces) by specifying an integer argument for each configuration.
TCEX	Surface #	Thermal coefficient of expansion.
TELE	Ignored	Telecentric in object space, 0 for no, 1 for yes.
TEMP	Ignored	Temperature in degrees Celsius. See "Environment" on page 108.
THIC	Surface #	Thickness of surface.
TSP1	Surface #	Thickness solve parameter 1. See "SUMMARY OF SOLVES" on page 473.
TSP2	Surface #	Thickness solve parameter 2. See "SUMMARY OF SOLVES" on page 473.
TSP3	Surface #	Thickness solve parameter 3. See "SUMMARY OF SOLVES" on page 473.
UDAF	Surface #	User defined aperture file. Surface must use either a user defined aperture or user defined obscuration aperture type. See "Surface properties aperture tab" on page 79.
WAVE	Wave #	Wavelength.
WLWT	Wave #	Wavelength weight.
XFIE	Field #	X-field value.
YFIE	Field #	Y-field value.

Comment about operands that define character strings

When using operands that define string values, such as LTLT, NCOM, and others, the string length is limited to 32 characters.

Defining the number of configurations

The number of configurations (also called the number of zoom positions) is changed by inserting or deleting configurations using the Edit menu option.

Defining each configuration

To define a multi-configuration operand, double click on the name of the operand you wish to change (the names are listed in the left most column). For example, suppose you want to enter in a multiple value for the thickness of surface 5. Assume you want three different configurations. Insert 2 new configurations using the Edit menu. Double click on the left most column in row 1. Select "THIC" from the drop-down list on the dialog box for the Operand Type. Select "5" for the surface number, then click on OK. Now in each configuration column, enter the desired thickness in that configuration.

Adding and removing elements

One application for the multi-configuration feature is to design systems that can be operated with or without certain elements being present. This can easily be handled by using the "Ignore This Surface" feature described in "Ignore This Surface" on page 78. The ignore surface setting can be controlled by the multi-configuration operand IGNR to ignore different surfaces in different configurations.

Changing configurations

To perform any analysis, proceed as usual. The program will use the current configuration for all calculations and graphics. To change configurations, double click on the column header at the top of the configuration you wish to change to on the MCE. The shortcut keys Ctrl-A and Shift-Ctrl-A also work.

Optimization with multi-configurations

ZEMAX will optimize multi-configuration data as readily as the conventional spreadsheet data. To make a multi-configuration parameter variable, place the cursor on the parameter, and press Ctrl-Z. This is a toggle; pressing Ctrl-Z again would eliminate the variability. When optimization is invoked, the optimization will automatically count the new variable. As many multi configuration variables as you like may be defined.

To optimize across configurations, select "Default Merit Function" from the Tools menu of the merit function editor. ZEMAX will build an appropriate merit function for you. Optimization across configurations is performed using the CONF operand. This special operand changes the current configuration during the evaluation of the merit function. This means all operands defined after CONF will be relevant to the new configuration. CONF may be used multiple times in the operand sequence to evaluate various parameters.

Boundary constraints and other user entered optimization operands in multiple-configuration merit functions are evaluated, and therefore enforced, only in the configuration in which they are defined. For example, if a CONF 1 operand is followed by various operands such as EFFL or REAY, these will only be evaluated in configuration 1. To enforce these same operands in configuration 2, the same operands need to be repeated under the CONF 2 operand.

The advantage to this system is that the entered operands, or their respective targets or weights, may be different in each configuration. The disadvantage is the need to copy the operands that apply in more than one configuration to each configuration.

Suggestions for organizing multiple-configuration merit functions

There are two common different ways to organize a multiple-configuration merit function. The first way is to add the user defined operands within each CONF group of a default merit function:

```
CONF 1
User operands for configuration 1...
Default operands for configuration 1...
CONF 2:
User operands for configuration 2...
Default operands for configuration 2...
CONF 3:
...etc.
```

The other method is to build the default merit function, then add all the user operands at the top of the merit function to keep them all in one place:


```

CONF 1
User operands for configuration 1...
CONF 2
User operands for configuration 2...
CONF 3
User operands for configuration 3...
etc...
DMFS
CONF 1
Default operands for configuration 1...
CONF 2
Default operands for configuration 3...
CONF 3
Default operands for configuration 3...
etc...

```

Both merit functions will accomplish the same job, but the first method executes faster because the overhead in changing between configurations is reduced. The second method is easier to edit and maintain. Note the use of the DMFS operand after the user defined operands. This operand serves as a marker, so that when the default merit function is reconstructed, it will be appended after the DMFS operand, and the user entered operands are not lost.

Note that if you change the field angles, heights, or weights, or the wavelength values or weights on the multi-configuration screen you should rebuild the default merit function. When the default merit function is constructed, it uses the data for each configuration to determine the rays traced, and the appropriate weighting.



Note that if you change the field angles, heights, or weights, or the wavelength values or weights on the multi-configuration editor you should rebuild the default merit function.

Solves for multi-configuration data

The following types of solve status are supported on values placed in the multi-configuration editor.

Fixed

If the solve status is fixed, then the value in the MCE does not change. Fixed is not really a solve but is included here for completeness.

Variable (numerical value operands only)

If the status is variable, then the optimization algorithms may adjust the value during optimization.

Substitute (glass operands only)

If the glass status is set to substitute, then the global optimization algorithms may choose another glass. Optionally, a specific glass catalog to use for substitution may be specified. The name of the optional catalog is limited to 40 characters. For more information on glass substitution, see “Using glass substitution” on page 542.

Pickup

Suppose there are three configurations, and in two of the configurations the values for one of the rows (say a THIC or GLSS) need to be the same value. A pick-up solve on one of the configurations for that operand can be used to ensure they are always the same.

To set a pickup solve, double click on the row and configuration on which the solve is to be placed. A dialog box will appear which allows definition of the solve type, the target configuration and operand, and a scale factor and offset. The new cell value is defined by: $\text{new_value} = \text{target_value} * \text{scale} + \text{offset}$. Note the target_value may be from any other cell in the MCE, as long as the target configuration and operand number is less than or equal to the current configuration and operand number, respectively.

Thermal pickup


ZEMAX-EE also supports a thermal pickup solve which incorporates thermal effects; see “THERMAL ANALYSIS” on page 603 for details.

Introduction

The index of refraction for refractive surfaces and solids in ZEMAX may be defined several different ways. Most of this chapter is related to dispersion data as defined in the glass catalogs. For alternate means of defining index data, see “Alternate methods of defining dispersion data” on page 600. There are several standard catalogs supplied with ZEMAX, and custom catalogs may be created.

ZEMAX computes indices of refraction from formulas and coefficients entered into the glass catalogs. When you specify a glass name such as "BK7" in the glass column of the LDE, ZEMAX looks for the name in each of the currently loaded glass catalogs. If the glass is found, ZEMAX uses the coefficients for that glass, and then using the formula for the glass selected in the catalog, computes the indices at each of the defined wavelengths.

It is important to note that all ZEMAX glass catalogs assume that the index computed by the dispersion formulas is the relative index of refraction computed as a function of the relative wavelength. Relative means relative to air at 1.0 atmosphere and a reference temperature defined in the glass catalog. For details on how ZEMAX considers temperature, pressure, and the index of air see “Index of refraction computation” on page 603.

 ***ZEMAX assumes that the index computed by the dispersion formulas is the relative index of refraction computed as a function of the relative wavelength.***

This method may seem more complex than directly entering in the indices of refraction, but the advantages are numerous. For one, the formulas are generally more accurate than user-entered data. Catalog data is more convenient as well, requiring the user to only supply the glass name. This is a particular advantage during the glass selection phase of the design. Additionally, any wavelength may be used, even if no explicit index data at that wavelength is available. The primary disadvantage is that the coefficients must be calculated, although this data is either readily available in catalogs or easily calculated. If you have the index data for some material not in the catalog, or if you feel your data is better than the catalog data, ZEMAX will compute the coefficients for you automatically; see “Fitting index data” on page 592.

 ***For important comments about obsolete glasses and differences between glasses with identical names, see “Obsolete catalog data” on page 599.***

Specifying which glass catalogs to use

This chapter describes the loading, editing, and managing of the glass catalogs. To specify that a particular catalog be used for a particular lens, use the Glass Catalog tab of the System, General dialog box. Multiple simultaneous catalogs may be used by listing them with spaces between the names. For example, to use the "schott" and the "hoya" catalogs, enter "schott hoyo". To use the "ohara", "schott", and "infrared" catalogs, enter "ohara schott infrared". Because ZEMAX uses spaces as delimiters between glass catalog names, spaces are not allowed within the names of the glass catalog files.

All the listed catalogs must be in the Glass Folder folder (see “Folders” on page 68 to edit this path). The reason that glass catalog names are specified on this screen is that the catalog choice is stored with each lens separately. If you now save the lens, and reload it at a later time, the correct glass catalogs, and only those catalogs, will automatically be loaded.

Editing and reviewing glass catalogs

To edit or review data in an existing glass catalog, select Tools, Glass Catalogs. Select the catalog name from those listed in the drop-down list on the dialog box. Once the catalog is selected, you may insert, cut, copy, paste, or modify data in the catalog, as is described in the following sections. You can save the newly modified catalog to either the same name or a new name. When editing the glass catalogs supplied with ZEMAX, be sure to save the modified data to a new name using the "Save Catalog As" button. This is important because future releases of ZEMAX may include an updated catalog which will be installed over the existing catalog, and any changes that had been made to the existing catalog will be lost.

To edit a large amount of data in the glass catalog, to create an extensive custom catalog, or to convert data in another format to the format used by ZEMAX, it may be easier to edit the glass catalog file directly. For information on the file format see “The AGF and BGF file formats” on page 599.



Be extremely careful when editing the glass catalogs; erroneous ray trace data will result if the catalog data is incorrectly modified.

Description of catalog data

The glass catalog dialog box displays data about each individual glass. The data fields are described in the following table.

GLASS CATALOG DATA FIELDS

Item	Description
Catalog	Used to specify which of the .AGF format catalogs to display. The comment string to the right of the catalog name may be used to describe the catalog.
Glass	Used to specify which glass within the catalog to display data for.
Rename	In the event that a glass needs to be given a different name, the name of the currently selected glass can be modified in this field.
Formula	The dispersion of each glass is described by a formula. This control allows selection of which formula is used. If this setting is changed, then the dispersion data becomes invalid unless the appropriate coefficients are also entered. See “The glass dispersion formulas” on page 590.
Status	The status indicates the general availability of the glass. The available settings are Standard, Preferred, Obsolete, Special, and Melt. The status values are generally specified by the manufacturer. Standard glasses are generally available for purchase. Preferred glasses are usually frequently melted glasses, and more likely to be available upon demand. Obsolete glasses are no longer manufactured, but may be available. Special is a general category used to indicate a glass that does not fall into one of the other categories. Melt is a flag used by ZEMAX to indicate glasses that have been created in the catalog by the melt fit feature, see “Fitting melt data” on page 593.
Index Nd	The index at d-light, or 0.587 micrometers. ZEMAX does not use this number when computing the index of refraction. It is displayed solely for reference. The entry may be meaningless for some materials which do not transmit well in the region around 0.587 micrometers.
Abbe Vd	The Abbe value at d-light. This number is not used by ZEMAX when computing the index of refraction. It is displayed solely for reference.
Ignore Thermal Expansion	This switch allows accurate thermal modeling of non-solid materials, such as gasses and liquids, by allowing direct specification of the thermal coefficient of expansion in the lens data editor rather than in the glass catalog. Only the edge effects are considered; the radius of curvature and other thermal pickup solves will all use the adjacent material TCE, rather than the gas or liquid TCE.
Exclude Substitution	If checked, then this glass will not be selected during global optimization, conversion from model to real glasses, or be considered by the RGLA optimization operand.
Meta Material (Negative Index)	If checked, then the dispersion formula will be used to determine the absolute value of the index of refraction, but the sign will be reversed to yield a negative index of refraction. So-called Meta Materials have a negative index of refraction.
K1, L1... A0, A1..., A, B, C, etc.	The first eight rows in the center column of the dialog box display the dispersion coefficient data. The names of these coefficients changes depending upon the glass formula.

Item	Description
TCE	The TCE value is the thermal coefficient of expansion. This is a dimensionless parameter. The value displayed or entered needs to be multiplied by 1e-6 to yield the actual values. ZEMAX currently uses the TCE value to model linear thermal expansion independent of the temperature range being used. The catalogs values supplied by the glass vendor are usually the TCE for the temperature range from -30 to +70 degrees celsius, but ZEMAX will use the TCE value for any temperature range.
Temp	The reference temperature in degrees Celsius. This is the temperature of the air to which the relative wavelength and index data is referenced. This is usually the temperature of the glass at which the data was measured, but not always. See also "Defining temperature and pressure" on page 603.
D0, D1, D2, E0, E1, Ltk	These are the thermal coefficients used by the thermal analysis model. See "Index of refraction computation" on page 603.
p	p is the material density in grams per cubic centimeter.
dPgF	The value is the deviation of the relative partial dispersion from the normal line.
Min Wavelength, Max Wavelength	The minimum and maximum wavelengths in micrometers over which the dispersion formula returns valid index data.
Melt Frequency	This integer value indicates the relative frequency with which the glass is melted by the manufacturer. The convention used is 1 means the glass is melted very frequently, 2 means less frequently, etc. up to a value of 5 to indicate a very infrequently melted glass.
Comment	This is an optional comment specific to the individual glass.
Relative Cost	This number is intended to indicate the approximate relative cost of the glass as compared to BK7. For example, a value of 3.5 would indicate the glass costs about 3.5 times as much as BK7 per pound. The number is only intended as a rough guide. The exact glass cost may vary with the quality, amount, form, and availability of the glass being purchased. Contact the glass vendor for more information.
CR, FR, SR, AR, PR	These are the glass codes that indicate how resistant the glass is to various environmental effects. The codes correspond to Climate Resistance (CR), Stain Resistance (FR), Acid Resistance (SR), Alkali Resistance (AR), and Phosphate Resistance (PR). Generally, the lower the number the more durable the glass is. If the value -1 or the symbol "?" is entered or displayed, no data is available. For a complete description of these codes and the test conditions used to measure them, contact the glass vendor.

Creating a new catalog

To create a new glass catalog, choose "Save Catalog As" and specify a new name. After the new catalog is created, any unwanted data may be removed using the "Cut Glass" button.



Note that ZEMAX does not allow any spaces in the names of glasses or glass catalogs.

Copying or moving glass catalog files

To copy or move a glass catalog file, copy or move only the file ending with the AGF extension. When ZEMAX reads the catalog with the AGF extension, a file with the same name will be created with a BGF extension. For example, ZEMAX will read in the file SCHOTT.AGF and create a file named SCHOTT.BGF. The BGF version is a binary format version of the same data stored as text in the AGF file. ZEMAX will subsequently read the BGF version of the glass catalog because the binary version can be read much faster than the text version of the file. ZEMAX will automatically recreate the BGF version if the date stamp of the AGF file changes, or if the BGF file is deleted or missing, or if the ZEMAX version number changes.

The BGF file is of no value without the AGF file; so if a glass catalog must be copied, moved, or sent via email to another folder or location, copy the AGF file and not the BGF file. ZEMAX will automatically recreate the BGF whenever it is needed.

The glass dispersion formulas

The coefficients in the catalog are used in any one of several polynomial formulas that ZEMAX recognizes. There is also a dispersion formula described by the six-digit MIL number, but those indices are calculated directly from the MIL number entered on the spreadsheet. The MIL number formula is not part of the glass catalog, and so it will not appear. See “Using MIL number glasses” on page 600 for a discussion of the MIL number glass formula. You may add new glasses to the currently loaded catalog if you have index data given in the form of one of the following equations. In all of the equations λ is in micrometers.

The Schott formula

The Schott constants of dispersion formula is

$$n^2 = a_0 + a_1\lambda^2 + a_2\lambda^{-2} + a_3\lambda^{-4} + a_4\lambda^{-6} + a_5\lambda^{-8}.$$

The required coefficients are available in most manufacturers glass catalogs. Schott no longer uses this formula, but it is widely used by other glass manufacturers. See also “The Extended formula” on page 592.

The Sellmeier 1 formula

The Sellmeier 1 formula is

$$n^2 - 1 = \frac{K_1\lambda^2}{\lambda^2 - L_1} + \frac{K_2\lambda^2}{\lambda^2 - L_2} + \frac{K_3\lambda^2}{\lambda^2 - L_3}.$$

Coefficients for all three terms may be entered to describe the material, although fewer terms may be used. See also the Sellmeier 3 and Sellmeier 5 formulas.

The Sellmeier 2 formula

The Sellmeier 2 formula is

$$n^2 - 1 = A + \frac{B_1\lambda^2}{\lambda^2 - \lambda_1^2} + \frac{B_2}{\lambda^2 - \lambda_2^2}.$$

Only two terms are used, there is no wavelength dependence in the numerator of the second term, and there is a constant term.

The Sellmeier 3 formula

The Sellmeier 3 formula is just like the Sellmeier 1 formula, with one additional term added:

$$n^2 - 1 = \frac{K_1\lambda^2}{\lambda^2 - L_1} + \frac{K_2\lambda^2}{\lambda^2 - L_2} + \frac{K_3\lambda^2}{\lambda^2 - L_3} + \frac{K_4\lambda^2}{\lambda^2 - L_4}.$$

The Sellmeier 4 formula

The Sellmeier 4 formula is:

$$n^2 = A + \frac{B\lambda^2}{\lambda^2 - C} + \frac{D\lambda^2}{\lambda^2 - E}.$$

The Sellmeier 5 formula

The Sellmeier 5 formula is just like the Sellmeier 3 formula, with one additional term added:

$$n^2 - 1 = \frac{K_1\lambda^2}{\lambda^2 - L_1} + \frac{K_2\lambda^2}{\lambda^2 - L_2} + \frac{K_3\lambda^2}{\lambda^2 - L_3} + \frac{K_4\lambda^2}{\lambda^2 - L_4} + \frac{K_5\lambda^2}{\lambda^2 - L_5}.$$

The Herzberger formula

The Herzberger expression is

$$n = A + BL + CL^2 + D\lambda^2 + E\lambda^4 + F\lambda^6,$$

$$L = \frac{1}{\lambda^2 - 0.028}.$$

The Herzberger formula is used mainly in the infrared spectrum.

The Conrady formula

The Conrady formula is

$$n = n_0 + \frac{A}{\lambda} + \frac{B}{\lambda^{3.5}}.$$

The Conrady formula is extremely useful for fitting to sparse data. For example, if you have only three index-wavelength pairs of data, fitting to the six-term Schott formula would yield meaningless data at the intermediate wavelengths.

The Handbook of Optics 1 formula

There are two similar formulas from the Handbook of Optics. The "Handbook 1" formula is:

$$n^2 = A + \frac{B}{(\lambda^2 - C)} - D\lambda^2,$$

The Handbook of Optics 2 formula

The "Handbook 2" formula is:

$$n^2 = A + \frac{B\lambda^2}{(\lambda^2 - C)} - D\lambda^2.$$

The Extended formula

The Extended constants of dispersion formula is

$$n^2 = a_0 + a_1\lambda^2 + a_2\lambda^{-2} + a_3\lambda^{-4} + a_4\lambda^{-6} + a_5\lambda^{-8} + a_6\lambda^{-10} + a_7\lambda^{-12}.$$

This is similar to the Schott formula, with two additional terms added.

The Extended 2 formula

The Extended 2 constants of dispersion formula is

$$n^2 = a_0 + a_1\lambda^2 + a_2\lambda^{-2} + a_3\lambda^{-4} + a_4\lambda^{-6} + a_5\lambda^{-8} + a_6\lambda^4 + a_7\lambda^6.$$

This is similar to the Extended formula, with a variation in powers of the wavelength in the last two terms.

The Extended 3 formula

The Extended 3 constants of dispersion formula is

$$n^2 = a_0 + a_1\lambda^2 + a_2\lambda^4 + a_3\lambda^{-2} + a_4\lambda^{-4} + a_5\lambda^{-6} + a_6\lambda^{-8} + a_7\lambda^{-10} + a_8\lambda^{-12}.$$

This is similar to the Extended formula, with another variation in powers of the wavelength.

General comments on using dispersion formulas

It is important to note that some publications use equations similar, but not identical to any of these expressions. It is often possible to rearrange the expressions to get them into the required form, and then recompute the required coefficients.

It is also a good idea to check the coefficients against a tabular listing of the index data available in many handbooks and publications. Use the dispersion plot or listing feature, or the prescription data report, which lists index data at each surface. If there are discrepancies, check the data you have entered carefully, and verify that the correct units and formula are being used.

Fitting index data



See also the discussion "Fitting melt data" below.

It is often the case that the materials you are designing with are in the catalog already. If they are not, you can enter the coefficients for the formulas described previously. As an alternative, ZEMAX will compute either the Schott, Herzberger, Conrady, or Sellmeier 1 dispersion formula coefficients for you. With the glass catalog dialog box displayed, click on "Fit Index Data" and the Fit Index Data dialog box appears. For a more comprehensive tool, see also "Glass Fitting" on page 245.

On the left side of this screen is a two-column spreadsheet editor. Using the mouse, enter in the wavelength (in micrometers) and index data you have. The more data you enter, the more accurate will be the fit.

If you have more data than will fit in the spreadsheet, use the data most closely representing the wavelength region of interest. At least three points are required to get a good fit if you are using the Conrady formula, six or more (and preferably twelve to fifteen) for the Schott, Herzberger, or Sellmeier 1 formulas. Select the formula to use by selecting the formula name from the drop-down list. You may want to try each of the models in turn, to see which gives the lowest residual.

The RMS error is the RMS fit error between the given data, and the index data generated using the resulting fit coefficients. The max error is the largest error between the fit and any one data point. Both numbers can be compared to the magnitude of the index of refraction; which is of course dimensionless. Because the Sellmeier 1 formula has non-linear coefficients, the fitting is iterative, and this formula takes much more computer time to fit the data than the other formulas do.

Now move the cursor over to the "Name" field, and enter the material name for the catalog. Select "Fit" and ZEMAX will compute the optimal coefficients. The residual RMS error, and the maximum single-point error, are listed on the bottom of the display. To enter this data in the currently loaded catalog, select "Add to catalog". ZEMAX will issue a message verifying that the glass was saved.



When a glass is added to the catalog, the transmission data, if any, needs to be added as described in the next section. Otherwise, the default internal transmission of 1.0 is used at all wavelengths.

The index and wavelength data may also be saved to a text file for later use, and loaded for fitting again by selecting the appropriate buttons. The text file may also be edited outside of ZEMAX, and then loaded for fitting.

Fitting melt data



See also the discussion "Fitting index data" above.

It is important to understand that the index of refraction values computed by ZEMAX, or listed in the catalog of the glass manufacturer, are average values for the index of refraction over a large number of "melts" or batches of the glass. A specific piece of glass from one melt will deviate from the catalog or nominal values slightly. The deviation is typically small, but the difference between the nominal and actual index values may be important for some systems.

Usually, when quality optical glass is shipped from the manufacturer, a data sheet will accompany the glass which indicates the index of refraction for the supplied glass at a few wavelengths, either as an offset from the nominal catalog values or as the measured index directly. Typically 3-5 wavelength-index data points are provided. This data is called "melt" data because it is specific to the batch of glass melted at one time.

The Melt Data tool available from the glass catalog dialog box is a handy utility for converting the limited melt data provided into a usable new glass type in the glass catalog.

There are a maximum of 8 wavelength-index points allowed for melt data. If you have more than 8 points, use the "Fit Index" tool described in the previous section. The minimum number of points allowed is 3, however, at least 4 and preferably 5 points should be used to get a good melt fit. The wavelength range defined by the melt index data should be as broad as possible, and should at a minimum cover the intended range of wavelengths to be used for ray tracing through the melt glass. In all cases, the fitted data should be inspected carefully for accuracy before it is used.

The Melt Data tool supports the following controls:

Glass: The name of the nominal glass in the selected glass catalog.

Melt Name: The name of the new glass to be created. The default is the nominal glass name with "_MELT" appended on. The name length may not exceed 20 characters.

Fit only these wavelengths: If checked, the melt index fit will only be done over the wavelength range defined by the minimum and maximum wavelengths of the provided melt data points. This allows much more precise fitting of the glass data, however, the melt glass cannot be used outside of this wavelength range. If unchecked, ZEMAX will attempt to extrapolate the data (see "Discussion of melt fitting method:" on page 594) to create a melt fit that is valid over the entire wavelength spectrum of the original glass data. This is the key difference between fitting index data (see "Fitting index data" on page 592) and fitting melt data.

Formula: The dispersion formula to use for the new melt glass. Choose either Schott, Herzberger, Conrady, or Sellmeier 1. The default value is the Schott formula unless the nominal glass uses one of these formulas; in which case the same formula as the nominal glass will be used.

Use: This box turns "on" and "off" each row of data.

Wavelength: The wavelength in micrometers for the index values to be entered.

Nominal: The index of refraction at the defined wavelengths using the nominal glass dispersion.

Actual: The actual measured index from the melt data. Note if the actual value is edited, the "delta" value is automatically adjusted to keep the data consistent.

Delta: The difference between the actual and nominal index of refraction. Note if the delta value is edited, the "actual" value is automatically adjusted to keep the data consistent.

Fit/Insert: Choosing this button starts the fitting process as described below.

Cancel: Aborts the melt fitting process.

After the fitting is finished, the new melt glass will be inserted in the catalog, the catalog will be saved, and a report summarizing the fit will be presented.

Discussion of melt fitting method:

The problem with fitting melt data is the generally low number of points available; typically 3-5. Most fitting routines need at least 8 points for good accuracy. So, the problem is to extrapolate from a few points the variation in index over a large enough number of points to fit the resulting dispersion accurately. The accuracy of the resulting melt fit at the defined data points depends largely upon the extent of the wavelength range. Greater accuracy is achieved if "Fit only these wavelengths" (see discussion above) is checked, at the expense of a reduced spectral range over which the fit is valid.

ZEMAX does melt fitting using the following algorithm:

First, a fit of the actual dispersion data is computed using the Conrady formula. The Conrady formula is used because it is stable and reasonable when as few as three points are defined.

Then a Conrady fit of the nominal data is computed using only the defined wavelength points.

A large number of index points covering either the entire usable wavelength range (if "Fit only these wavelengths" is unchecked) or the wavelength range defined by the melt points (if "Fit only these wavelengths" is checked) of the nominal glass is generated. To each nominal index value, an offset is added which is the difference between the two Conrady fits that were generated using only the melt data wavelengths.

Finally, the resulting data is fit using the selected formula (not necessarily the Conrady). This is the final fit inserted into the catalog.

After the melt fitting is finished, a report summarizing the melt fitting process is presented. Check this report carefully before using the new melt glass!

ZEMAX automatically copies over all transmission, density, cost factor, and other data from the nominal glass to the melt glass.



Check the generated melt fitting report for index accuracy carefully before using the new melt glass!

Defining Transmission Data

Selecting the "Transmission" button invokes the transmission data editor within the glass catalog. Transmission refers to the intensity transmittance of light that depends upon the thickness of the glass as well as the wavelength. ZEMAX models the transmitted intensity using Beer's law:

$$t = e^{-\alpha\tau},$$

where α is the absorption coefficient and τ is the path length through the glass. The parameter α generally depends upon wavelength and has units of inverse length. See the Chapter "Polarization Analysis" for information on polarization ray tracing and transmission.

The transmission is defined in the glass catalog by a series of 3 numbers: the wavelength in micrometers, the intensity transmittance, and the reference thickness in mm. For example, the transmission of a glass at 0.35 micrometers may be 0.65 for a thickness of 25 mm. Multiple data points may be defined in the transmission data editor. Internally, ZEMAX converts the data to "per mm", and interpolates between defined wavelengths. If ray tracing is being performed at a wavelength outside of the defined wavelengths, then the data for the closest wavelength is used; otherwise, ZEMAX performs a linear interpolation.

To save the transmission data to a text file, use the "Save To File" button. The file format is three columns of data: wavelength in micrometers, intensity transmittance, and reference thickness in mm. Data in this same format may be read from a file and placed in the catalog using the "Load From File" button. To save the modified data directly in the original glass catalog, use the "Save Catalog" button.

Not all of the glasses listed in the supplied catalogs have valid transmission data, especially for infrared materials and other non-commercial glass types. If transmission data is supplied by the manufacturer, it is usually included. If no reliable data is available, or if the data has been omitted, the default internal transmission of 1.0 is used at all wavelengths.

The Transmission Data editor ignores all data after the first entry with a wavelength (λ) of zero. Only the data entered before the first zero wavelength is considered.

Modeling gases and liquids

Once a material is defined in the glass catalog, ZEMAX uses the TCE specified for that material to determine the thermal expansion of the radius, center thickness, and other data for the surfaces using the material. However, if the material is not a solid, but is instead a gas or a liquid, then generally the thermal expansion is not governed by the material properties, but is instead determined by the edge thickness of the mounting material.

In this special case, ZEMAX needs to use the TCE supplied in the Lens Data Editor to define the mounting material properties, rather than the TCE supplied in the glass catalog. This can be accomplished by setting the "Ignore Thermal Expansion" switch for the material in the glass catalog.

Finding a glass quickly

The fastest way to view the data for any glass is to click once on the name of the glass in the Lens Data Editor, then select Tools, Glass Catalogs (or select the GLA button if displayed). The correct catalog and glass will be displayed.

Glass catalog sources

The data for the glass catalogs included with ZEMAX are generally provided by the manufacturers. The data has been converted to ZEMAX format. The following vendors have supplied data for the ZEMAX glass catalogs:

GLASS VENDORS

Catalog	Vendor	Website
ARCHER	Archer OpTx	www.archeroptx.com
CORNING	Corning France Optical Division	www.corning.com
CDGM	Chengdu Guangming	www.cdgmgd.com
HEARAEUS	Heraeus Quarzglas GmbH & Co. KG	www.heraeus.de
HIKARI	Hikari Glass USA, Inc.	www.hikariglass.com
HOYA	Hoya Optics, Inc.	www.hoya.co.jp
IRPHOTONICS	IRphotonics Inc.	www.irphotonics.com
LIGHTPATH	LightPath Technologies	www.lightpath.com
LZOS	Russian optical glass fabricator	www.lzos.ru/indexe.htm
NHG	Hubei New Huaguang Information Materials Co., Ltd	www.hbnhg.com
OHARA	Ohara Corporation	www.oharacorp.com
PILKINGTON	Pilkington Special Glass Limited	www.pilkington.com
RPO	Rochester Precision Optics	www.rpoptics.com
SCHOTT	Schott Glass Technologies	www.schott.com
SUMITA	Sumita	www.sumita-opt.co.jp

Catalog	Vendor	Website
TOPAS	Polyplastics Co., Ltd.	www.polyplastics.com
UMICORE	Umicore Electro-Optic Materials	www.optics.unicore.com
ZEON	Zeon Chemicals L.P.	www.zeonchemicals.com

Although the data contained in the glass catalogs is generally reliable, it is always possible for errors to be made during translation or editing of the data. *It is absolutely crucial that all index data be verified for accuracy by the end user! This is especially true where fabrication of the optics is being considered.*



The catalog data may be in error, and it needs to be verified before it can be trusted.

The data contained in other catalogs provided with ZEMAX, such as the BIREFRINGENT, INFRARED, and MISC catalogs have been compiled from published sources as described in the following table. Some materials, such as calcite, are defined in more than one catalog. Materials may be included in the catalog which are not referenced in the table, or the source is specified as "unknown". These materials should be especially scrutinized before trusting the data accuracy.

BIREFRINGENT CATALOG DATA SOURCES

Material	Source
ADP	Handbook of Optics Vol. II
AG3ASS3	Handbook of Optics Vol. II
AGGAS2	Handbook of Optics Vol. II
AGGASE2	Handbook of Optics Vol. II
AL2O3	Handbook of Optics Vol. II
ALN	Handbook of Optics Vol. II
BATIO3	Handbook of Optics Vol. II
BBO	Handbook of Optics Vol. II
BEO	Handbook of Optics Vol. II
CALCITE	Handbook of Optics Vol. II
CAMOO4	Handbook of Optics Vol. II
CAWO4	Handbook of Optics Vol. II
CDGEAS2	unknown
CDS	Handbook of Optics Vol. II
CDSE	Handbook of Optics Vol. II
CUGAS2	Handbook of Optics Vol. II
KDP	Handbook of Optics Vol. II
LAF3	Handbook of Optics Vol. II
LINBO3	Handbook of Optics Vol. II
LIO3	Handbook of Optics Vol. II

Material	Source
LIYF4	unknown
MGF2	M.J. Dodge, "Refractive Properties of Magnesium Flouride," Applied Optics, Vol. 23 , No. 11: p1980 (1985)
PBMOO4	Handbook of Optics Vol. II
QUARTZ	Handbook of Optics Vol. II
SRMOO4	Handbook of Optics Vol. II
TAS	Handbook of Optics Vol. II
TE	Handbook of Optics Vol. II
TEO2	Handbook of Optics Vol. II
YV04	Handbook of Optics Vol. I
ZNGEP2	Handbook of Optics Vol. II
ZNO	Handbook of Optics Vol. II

INFRARED CATALOG DATA SOURCES

Material	Source
ALN	Handbook of Optics Vol. II
ALON	Handbook of Optics Vol. II
AMTIR1	Amtir Spec Sheet
AMTIR3	Amtir Spec Sheet
BAF2	Handbook of Optics Vol. II
BEO	Handbook of Optics Vol. II
CAF2	Handbook of Optics Vol. II
CALCITE	Handbook of Optics Vol. II
CDSE	Handbook of Optics Vol. II
CDTE	Handbook of Optics Vol. II
CLEARTRAN	Spec Sheet (Rohm & Haas)
CLEARTRAN_OLD	Spec Sheet (Morton)
CSBR	Handbook of Optics Vol. II
F_SILICA	The Infrared & Electro-Optical Systems Handbook V. III
GAAS	Amtir Spec Sheet
GEO2	Handbook of Optics Vol. II
GE_LONG	JOSA, 47, 244 (57): 48, 579 (58)
GE_OLD	JOSA, 47, 244 (57): 48, 579 (58)
GERMANIUM	Handbook of Optical Constants of Solids

Material	Source
IRG2	Schott Spec Sheet
IRG3	Schott Spec Sheet
IRGN6	Schott Spec Sheet
IRG7	Schott Spec Sheet
IRG9	Schott Spec Sheet
IRG11	Schott Spec Sheet
IRG15	Schott Spec Sheet
IRG100	Schott Spec Sheet
KBR	ISP Optics data sheet (www.ispoptics.com)
KCL	ISP Optics data sheet (www.ispoptics.com)
KRS5	Handbook of Optics Vol. II
LIF	Handbook of Optics Vol. II
MGF2	M.J. Dodge, "Refractive Properties of Magnesium Fluoride," Applied Optics, Vol. 23 , No. 11: p1980 (1985)
MGO	Handbook of Optics Vol. II
NACL	The Infrared & Electro-Optical Systems Handbook V. III
PBF2	Handbook of Optics Vol. II
SAPPHIRE	The Infrared & Electro-Optical Systems Handbook V. III
SILICON	Handbook of Optics Vol. II
SRF2	Handbook of Optics Vol. II
SRTIO3	Handbook of Optics Vol. II
ZBLA	Handbook of Optics Vol. II
ZBLAN	Handbook of Optics Vol. II
ZNGEP2	Handbook of Optics Vol. II
ZNSE	Handbook of Optics Vol. II
ZNS	The Infrared & Electro-Optical Systems Handbook, V.III

MISC CATALOG DATA SOURCES

Material	Source
ACRYLIC	Handbook of Optics Vol. II
BASF5	Laikin, Lens Design
BASF55	Laikin, Lens Design
CAF2	Handbook of Optics Vol. II
CDS	Handbook of Optics Vol. II

Material	Source
COC	Hoechst Celanese Spec sheet
CR39	The Photonics Design and Applications Handbook, Optical Plastics: Properties and Tolerances, H. D. Wolpert, pp. H-300 - H-307, (1991).
KDP	Handbook of Optics Vol. II
LAF3	Handbook of Optics Vol. II
LIYF3	Handbook of Optics Vol. II
PMMA	Handbook of Optics Vol. II
POLYCARB	Handbook of Optics Vol. II
POLYSTYR	Handbook of Optics Vol. II
PYREX	Laikin, Lens Design
QUARTZ	Handbook of Optics Vol. II
SAN	Handbook of Optics Vol. II
SEAWATER	Laikin, Lens Design
SILICA	Handbook of Optics Vol. II
TEO2	Handbook of Optics Vol. II
TYPEA	Laikin, Lens Design
VACUUM	Laikin, Lens Design
WATER	Handbook of Optics Vol. I

Obsolete catalog data

Changes in environmental protection laws have required the discontinuation of the manufacture of many optical glasses between the years 1990 - 2000. Roughly 2 out of 3 glasses which were listed in earlier catalogs of the various glass manufacturers have been discontinued. In some cases, glasses were reformulated to comply with new environmental restrictions. These new glasses may have identical or similar names as the old glasses, depending upon the glass vendor. The reformulated glasses may have different index data than the glasses they replace.

Obsolete glasses still exist in optical shops and may be used for new designs, if a supply can be located. Also, many older designs using the old glasses need to be modeled and ray traced, especially when designing new optics to work in harmony with the existing optical system. For these reasons, data for obsolete glasses are still provided with the catalogs distributed with ZEMAX. Obsolete glasses are indicated by the "obsolete" status displayed in the glass catalog dialog box.

Because some new glasses may have the same name as old glasses, although the exact composition may have changed, optical engineers need to be especially vigilant about checking the index data predicted by the software against the melt sheets of the glass which will actually be used.

Never blindly trust the accuracy of the index computed using catalog coefficients. There are numerous sources of potential error, such as the measurement of the original sample, the fitting of the data, the typing in of the data in the manufacturers catalog, and then retyping into ZEMAX, and finally, the ZEMAX code itself.

The AGF and BGF file formats

ZEMAX stores glass catalog data in two file formats, called the ANSI Glass Format (AGF) and the Binary Glass Format (BGF). The glass catalog data supplied with ZEMAX is in the AGF format, and AGF files may be used when glass catalog data needs to be modified or created by the glass manufacturer. When ZEMAX runs, AGF

files are automatically converted, if required, to BGF files. The BGF files are only used by ZEMAX to speed up the loading of the glass catalogs, and should never be created, edited, or distributed. ZEMAX will create or update a BGF whenever required. The AGF file is the "master" file used to define glass catalog data

The AGF file consists of a header line followed by a series of records, one for each glass. The header line format is:

```
CC <Catalog Comment>
```

The mnemonic CC stands for Catalog Comment, and any comment string up to 140 characters may be placed here to define a comment for the catalog as a whole. Each subsequent record consists of a two letter mnemonic followed by one or more data items separated by spaces. For a description of these data see "Description of catalog data" on page 588. The format of a single record is:

```
NM <glass name> <dispersion formula #> <MIL#> <N(d)> <V(d)> <Exclude Sub> <status>
<melt freq>
GC <Individual Glass Comment>
ED <TCE (-30 to 70)> <TCE (100 to 300)> <density> <dPgF> <Ignore Thermal Exp>
CD <dispersion coefficients 1 - 10>
TD <D0> <D1> <D2> <E0> <E1> <Ltk> <Temp>
OD <rel cost> <CR> <FR> <SR> <AR> <PR>
LD <min lambda> <max lambda>
IT <lambda> <transmission> <thickness>
IT <lambda> <transmission> <thickness>
... multiple IT lines may follow
```

Comments on the individual fields follow.

NM: Glass name and other data. The dispersion formula number is 1 for Schott, 2 for Sellmeier 1, 3 for Herzberger, 4 for Sellmeier 2, 5 for Conrady, 6 for Sellmeier 3, 7 for Handbook of Optics 1, 8 for Handbook of Optics 2, 9 for Sellmeier 4, 10 for Extended, 11 for Sellmeier 5, and 12 for Extended 2. The MIL# is provided for back compatibility and is not used, but a placeholder value must be provided. The nd and vd values are also provided for reference but are not used. The "exclude sub" flag is 0 for no and 1 for yes. Status is 0 for Standard, 1 for Preferred, 2 for Obsolete, 3 for Special, and 4 for Melt. Melt Freq is an integer between 1 and 5 to indicate the relative frequency of melting by the manufacturer.

GC: Glass comment. Any text string up to 140 characters long may be placed here.

ED: Extra Data. The values are the TCE in the -30 to 70 degree Celsius range, the TCE in the 100 to 300 degree Celsius range (this value is currently not used), the density, dPgF, and the Ignore Thermal Expansion flag, which is 0 for no and 1 for yes.

CD: Coefficient Data. Up to 10 coefficients may be provided. The meaning of these coefficients depends upon the dispersion formula used.

TD: Thermal Data. These are the D0, D1, D2, E0, E1, Ltk, and reference temperature values.

OD: Other Data. The relative cost, CR, FR, SR, AR, and PR values. For these values, -1 should be specified if the data is not available.

LD: Lambda Data. The minimum and maximum wavelength in micrometers over which the dispersion formula is valid.

IT: Internal Transmittance. These lines are the internal transmittance. Each line defines a wavelength, transmission, and thickness in millimeters. The wavelengths must be in ascending order. Up to 100 points may be defined for each glass record.

Alternate methods of defining dispersion data

There are alternate methods to define dispersion data that do not involve the glass catalogs and dispersion formulas. These methods are described in the following sections.

Using MIL number glasses

MIL number glasses are those described by a six-digit number, such as 517640 for BK7. The first three digits in the MIL number is the d-light index minus one, without the decimal place. The last three digits is the Abbe V-number times 10. To use a MIL number glass, type the 6 digit number in directly for the glass name.

ZEMAX uses a formula for computing the index at each defined wavelength based upon the index and Abbe number defined by the MIL number. The formula is based upon a least-squares fit of coefficient data of many typical glasses. Typically, the index data calculated is accurate to roughly 0.001. At wavelengths in the deep UV or infrared the index value becomes less reliable. MIL number glasses are generally inferior substitutes for the constants of dispersion (or other) models for the glass, however they are useful if no other data is available.



Note MIL number glasses are an approximation, although usually a very good approximation in the visible range. Outside the visible wavelength range, such as in the ultraviolet or infrared, the MIL number glass is not accurate and should not be used.

It is important to note that the indices calculated from the six-digit MIL number are not the same as those calculated from the glass catalog, even if the MIL number you are using corresponds to a glass in the catalog. Index data is calculated directly from the MIL number entered on the main screen; not from the glass catalog data, even if a glass with that MIL number is in the catalog.

Because any glass with a six digit name is assumed to be a MIL number glass defined by ZEMAX's internal equations, six digit numbers may not be used for glass names defined in the glass catalogs. MIL number glasses are not used by the glass substitution feature (see "Using glass substitution" on page 542).

Using table glasses

Table glasses are defined by data stored in a text file ending in the extension ZTG (for ZEMAX Table Glass). The file must be placed in the Glass Folder folder (see "Folders" on page 68 to edit this path). The file format is:

```
! Comment line
DENSITY grams_per_cc
wavelength_1 index_1 transmission_1 thickness_1
wavelength_2 index_2 transmission_2 thickness_2
wavelength_3 index_3 transmission_3 thickness_3
etc..
```

Any blank line or a line starting with the "!" symbol is assumed to be a comment and is ignored.

The DENSITY data line defines the density of the glass in grams per cubic centimeter. If the DENSITY data line is omitted the density is assumed to be zero.

The wavelength must always be in units of micrometers and should be listed in ascending order. The index of refraction must be a positive value. The transmission is the internal transmittance of the glass and should be a positive value. The thickness is the thickness of glass in millimeters for the specified internal transmittance. The transmission and thickness values are used to determine the internal transmission per unit length. If both the transmission and thickness values are omitted the material is assumed to be 100% transmitting.

Up to 1200 lines of wavelength data may be defined in the file. At least 5 data points should be defined to provide a reasonable spline fit. If fewer than 5 points are provided, dummy points are added before and after the specified wavelength range to provide a smooth fit. The data should be checked for accuracy and reasonable fitting by using the dispersion plot data, see "Dispersion Diagram" on page 209.

To use the table data, enter the file name in the glass column for the desired surface, including the extension. For example, if the file name is "MYGLASS.ZTG" then "MYGLASS.ZTG" should be entered in for the glass name. Note the total name, including the extension, may not exceed 20 characters.

When computing the data for a specific wavelength, ZEMAX will use a cubic spline interpolation if the wavelength is between the shortest and longest wavelengths defined in the ZTG file. For wavelengths outside of the defined range, the data for the nearest data point is used; no extrapolation of the data is performed.

The advantage of a table glass is the index at the defined points is exactly the index specified in the table, unlike fitted functions which may deviate from the data being fitted at the data points.

Table glasses are not used by the glass substitution feature (see "Using glass substitution" on page 542).

Using model glasses

ZEMAX can idealize the dispersion of glass using the index at d-light (.5875618 micrometers), the Abbe number, and a term which describes the deviation of the partial dispersion from the "Normal Line". The index at

d-light is given the symbol N_d . The Abbe number (also called the V-number) is given the symbol V_d and is defined by

$$V_d = \frac{N_d - 1}{N_F - N_C},$$

where N_F and N_C are the indices of refraction at 0.4861327 and 0.6562725 micrometers, respectively. The partial dispersion term is $\Delta P_{g,F}$.

ZEMAX uses a formula based upon the typical dispersion of standard glasses in the visible range to estimate the index at any defined wavelength within the visible range as a function of the N_d and V_d values. This formula is accurate to roughly 0.0001 for typical glasses.

The N_d , V_d , and $\Delta P_{g,F}$ values are specified on the glass solve dialog box, which can be reached from the Lens Data Editor.

See also “Using model glasses” on page 528 and “Optimizing glass selection” on page 541.



Note model glasses are an approximation, although usually a very good approximation in the visible range. Outside the visible wavelength range, such as in the ultraviolet or infrared, the model glass is not accurate and should not be used.

Introduction

Environmental factors such as temperature and pressure can affect the performance of optical systems. There are three primary factors to consider. First, the index of refraction of glass depends upon both temperature and wavelength; relative indices which are measured with respect to air also change with pressure. Second, glass expands and contracts with temperature, which can change the radius, thickness, or other dimensions of a lens. Third, the spacings between lenses changes due to the expansion and contraction of the mounting material.

The thermal analysis features of ZEMAX can account for all these effects. By accounting for thermal effects, ZEMAX can be used to analyze and optimize a design for any specific temperature or for a range of temperatures.

The index of refraction data given by the dispersion formulas are referenced to a specific temperature and pressure, which is typically 20 or 25 degrees Celsius (depending upon the manufacturer) and 1 atmosphere. Also, index of refraction data is by convention referenced to that of air, which means thereafter the air has an index of unity. The index referenced to air is called the relative index of refraction. When the index is referenced to vacuum (which truly has an index of unity) then the index is called the absolute index of refraction. The difference between these two references for any glass is a function of wavelength, temperature, and pressure.

Defining temperature and pressure

Two user provided values define the environment: the system temperature in degrees Celsius and the air pressure measured in atmospheres. These values are set on the dialog box which can be found under the Environment tab of the System, General dialog box. By default, any defined temperature and pressure applies to all surfaces in the optical system. However, optical systems which require multiple temperatures to be defined in the same system are also supported. This is required for optical systems that have some lenses in a different pressure (such as a vacuum) or in a region that is cooler or warmer than the system environment.

Defining wavelengths

Wavelength data are always measured in micrometers referenced to "air" at the system temperature and pressure. The default system temperature is 20 degrees Celsius, and the default air pressure is 1.0 atmospheres. If the system temperature and/or pressure is modified, or under the control of multi-configuration operands, care must be taken to adjust the wavelengths to the new air temperature and pressure.

Wavelength data is entered on the "Wavelength Data" dialog box; see the "Wavelengths" on page 116 for details.



Wavelength data are always measured in micrometers referenced to "air" at the system temperature and pressure.

Index of refraction computation

ZEMAX always uses relative, not absolute index data. One way to think about relative index is that the index of "air" (indicated by a blank entry in the glass column) is always defined to be exactly 1.0 at all wavelengths at the system temperature and pressure. Air and glass spaces at other temperatures and pressures are normalized relative to 1.0.

There are several steps required for ZEMAX to compute the relative index data. The basic steps for computing the index for each glass type at each wavelength are:

- Scale the wavelength to air at the reference temperature of the glass and a pressure of 1.0 atmospheres.
- Compute the relative index of the glass at the reference temperature from the dispersion formula.
- Compute the index of air at the reference temperature of the glass.
- Compute the absolute index of the glass (relative to vacuum) at the reference temperature of the glass.
- Compute the change in absolute index of refraction of the glass at the surface temperature.
- Compute the index of air at the system temperature and pressure.

-Compute the index of the glass relative to the air at the system temperature and pressure.

The end result, which is the index of the glass at the surface temperature and pressure relative to air at the system temperature and pressure, is what is used by ZEMAX for ray tracing. The index of a surface made of air is treated in the same way. This means that air surfaces at different temperatures and pressures from the system temperature and pressure will have slightly lower or higher index. For example, if the system pressure is 1.0 and the pressure of one air surface is set to 0.0, the index of this one surface will be approximately 0.99973. If the system pressure is 0.0 and the air space has a pressure of 1.0, the index for that surface will be approximately 1.00027. Remember, air at the system temperature and pressure is defined to be 1.0, all other indices are relative.



Air at the system temperature and pressure is defined to be 1.0, all other indices are relative.

Note also that ZEMAX can easily model systems used in a vacuum by changing the air pressure to zero. If only some surfaces are at vacuum, this can be set up using the TEMP and PRES commands on the multi-configuration editor.

For the index of refraction of the glass, ZEMAX uses the dispersion formulas and data stored in the glass catalog. For details, see “Description of catalog data” on page 588. For the index of air, ZEMAX uses the following formula:

$$n_{air} = 1 + \frac{(n_{ref} - 1)P}{1.0 + (T - 15) \cdot (3.4785 \times 10^{-3})}$$

where

$$n_{ref} = 1 + \left[6432.8 + \frac{2949810\lambda^2}{146\lambda^2 - 1} + \frac{25540\lambda^2}{41\lambda^2 - 1} \right] 1.0 \times 10^{-8},$$

T is the temperature in Celsius, P is the relative air pressure (dimensionless) and λ is measured in micrometers. This formula for the index of air is from F. Kohlrausch, *Praktische Physik*, 1968, Vol 1, page 408.

The change in absolute index of the glass with temperature is given by the following expression:

$$\Delta n_{abs} = \frac{n^2 - 1}{2n} \left[D_0 \Delta T + D_1 \Delta T^2 + D_2 \Delta T^3 + \frac{E_0 \Delta T + E_1 \Delta T^2}{\lambda^2 - \lambda_{tk}^2} \right],$$

where n is the relative index at the reference temperature of the glass, ΔT is the change in temperature relative to the reference temperature of the glass (ΔT is positive if the temperature is greater than the reference temperature of the glass), and the six other constants are provided by the glass manufacturer to describe the thermal behavior of the glass. This model for the change in the glass index was developed by Schott Glass Technologies, Inc.

The six constants must all be provided in the glass catalog for the computation to be valid. The default values of zero for all six constants yields zero for the change in index; therefore if no thermal data has been added to the catalog, no thermal effects are considered. ZEMAX cannot compute thermal index variation for any glass type without these six constants provided. However, some approximations are available if the six coefficients are not available. See the section “Adding thermal index variation data” on page 607 for details.

Environmental effects on index for gradient index, MIL number, and model glasses

ZEMAX only partially considers environmental effects on the index of refraction of gradient index surfaces, MIL number glasses, or model glasses which are described only by index and Abbe number. The relative index of refraction computed using these methods is adjusted to account for the change in the index of surrounding air at

the system temperature and pressure. However, the change in the absolute index of these glasses with temperature and pressure is not considered, because no "dn/dt" data is defined for these glass types. However, it is possible to manually define the variations in properties across multiple environments using the multi-configuration editor.

Defining multiple temperature and pressure values

To analyze or optimize an entire lens at a specific temperature and pressure, all that is required is to define the relevant data on the environment tab of the general data dialog box described in "Environment" on page 108. All the radii and thickness data are then assumed to be measured at that temperature, and ZEMAX computes the index data accordingly. However, the real power of the thermal analysis feature comes into play when a lens must be analyzed at or optimized for multiple environments, such as a broad temperature range or varying altitudes (or both). There are several new issues these lens systems introduce:

A means of specifying the nominal temperature at which radii and thickness are measured must be provided.

The change in index, radii, and thickness must be accounted for as the environment changes.

The thermal effects on mounting material must be considered.

Some surfaces may be at one temperature or pressure, while other surfaces are at a different temperature or pressure.

ZEMAX is up to the task in all particulars. The basic approach to setting up a multiple environment lens is:

Define the lens at some nominal temperature and pressure. This should be the environment under which the lens will be fabricated. All radii and thickness data will be specified only at this temperature and pressure.

Now define additional configurations using the multi-configuration feature (see the chapter Multi-configurations). In each additional configuration, the temperature and pressure will be specified, and special solves called "thermal pickup solves" will be used to adjust the radii and thickness data for each configuration. The multi-configuration operands that control temperature and pressure are TEMP and PRES.

TCE stands for thermal coefficient of expansion. When a glass element changes temperature, the linear change in size is given by the expression

$$L' = L(1 + \alpha\Delta T),$$

where L is a linear dimension, α is the TCE, and ΔT is the change in temperature. As the material expands, the radius of curvature also expands. Therefore, both the thickness and the radii of a glass surface change linearly with temperature. The assumption of linearity is only an approximation, but a reasonably good one for most materials and temperature ranges.

The TCE coefficient is defined along with the glass dispersion data in the glass catalog. See the chapter "Using Glass Catalogs" for details on the TCE data.

Defining which parameters consider thermal effects

On the multi-configuration editor, there is a special solve called a thermal pickup. This pickup solve is used to compute a new value for a multi-configuration parameter based up on the temperature and pressure of the new configuration as compared to the "reference" configuration. Thermal pickup solves only affect data for certain types of multi-configuration values, as described below.

Radius of curvature (CRVT) values

If the multi-configuration operand is a CRVT, then the curvature of that surface for that configuration is computed from the curvature in the nominal configuration, the difference in temperature between the two configurations, and the TCE for that material.

If the glass type for a surface is specified in the catalog (such as BK7 or F2) then the TCE from the catalog is used. If the glass type is air (a blank entry) or "MIRROR" then there are two possibilities: if the previous surface glass type is a catalog glass, then the TCE for that glass is used, otherwise, the TCE value entered in the TCE column will be used. These rules have one important ramification: if an element is a cemented doublet, the TCE for the first surface will be that of the first glass, while the TCE for the second and third surfaces will be that of the

second glass. ZEMAX ignores the "stress" induced in the common surface. This assumption may not be accurate for a large temperature range. Mirrors surrounded by air use the TCE from the TCE column, mirrors preceded by glass are assumed to be made of that same glass, and therefore use the TCE for the preceding glass.

Thickness (THIC) values

If the operand is a THIC, then there are two possibilities. If the surface is composed of a catalog glass, then the TCE for that glass is used. Otherwise, the TCE specified in the TCE column of the surface is used. The TCE column is used for entering user-defined TCE data for the material used to manufacture the mount. There is one other important point: if the surface is not composed of a catalog glass, then the thermal expansion is computed along a length of material which extends from the edge of the surface to the edge of the next surface. Since the material expands along the length of the edges, not at the center thickness, this is a more accurate computation.

For example, suppose two lenses were separated by aluminum with a center thickness spacing of 80 mm. If the sag of the first surface (the rear of the first lens) lens was -5 mm, and the sag of the second surface (the front face of the second lens) was 8 mm, the total edge thickness is then 93 mm. If the TCE for aluminum is 23.50 E-6, then under a temperature change of +20 degrees the edge thickness would change from 93 to 93.0437 mm. Ignoring the change in diameter of the spacer and the change in the sag of the two lens faces (ZEMAX actually accounts for both of these effects, but for simplicity they are ignored here), the center thickness would then change to 80.0437. Note that this is a different amount of expansion than if the computation had only considered the center thickness.

Because the expansion along the edge accounts for the change in radii of the adjacent surfaces, even a TCE of 0.0 will result in a change of thickness if the radii change. To turn off thermal expansion of a thickness, do not use a TCE of 0.0, rather, delete the thermal pickup solve entirely.

Parameter values

The thermal pickup solve properties for parameter values depends upon the parameter number and surface type. If the parameter is interpreted as having units of length, then the appropriate scaling is performed just like the radius of curvature described earlier. If the units are in powers of length, such as length squared or inverse length, then appropriate scaling is also performed. Otherwise, the thermal pickup solve ignores thermal effects and picks up the value from the nominal configuration.

Extra data values

The thermal pickup solve properties for extra data values depends upon the surface type. Generally, only the "normalization radius" needs to be scaled. This convenience is due to the fact that thermal expansion can be considered a scaling of length. Since most extra data surfaces use dimensionless coefficients, only the normalization radius needs to be scaled. This works effectively for binary optic and polynomial aspheric surfaces equally well. For extra data surface types which do not use a normalization radius, the thermal pickup solve ignores thermal effects and picks up the value from the nominal configuration.

Special surface types, such as polynomial aspheres, binary optics, holograms, and other surfaces that use parameter or extra data need to have operands and thermal pickup solves added manually to the multi-configuration editor; the automatic thermal setup tool does not add these operands automatically.

All other values

All other values are straight pickup solves, and the value will be identical to that in the nominal configuration. Thermal effects are ignored.


Defining multiple environments within a single configuration


Sometimes it is required to have different parts of the optical system at different temperatures and pressures. Note this is different from an entire system being at different environments across multiple configurations.

Groups of surfaces can be assigned their own temperature and pressure using the TEMP and PRES multi-configuration operands, even if there is only 1 configuration defined. The key is that each TEMP and PRES operand defines the environment for all operands that follow in the multi-configuration editor. The last TEMP and PRES operands listed in the editor define the "global" environment, which will apply to all data not listed in the multi-configuration editor.

For example, suppose an optical system model requires surfaces 1-5 to be at a temperature of 20 degrees C, while surfaces 6-10 are at 50 degrees C. The first operand listed should be TEMP (this same discussion applies to PRES), defining the initial environment of 50 degrees C. All the curvatures, thicknesses, semi-diameters, glasses, and other values for the surfaces 6-10 should be listed after the TEMP operand. Then the list should end with another TEMP, this one defining the 20 degree "global" environment. The resulting system will be evaluated at the respective temperatures (and/or pressures) for each surface.

It is extremely important to understand the two basic rules:

-  **1. All multi-configuration editor operands that follow a TEMP or PRES operand are evaluated at that temperature or pressure.**

-  **2. The last TEMP and PRES listed in the multi-configuration editor define the temperature and pressure for all other lens data, on the multi-configuration editor or not.**

By far the most important step in setting up a complicated multiple environment lens is to check the set up carefully. Two excellent tools for doing this are the index of refraction data and multi-configuration data table on the prescription report. This table lists the temperature and pressure for each glass type, and the thermal pickup relationships.

Checking the thermal pickup solves on each parameter is also a good idea; the data should be at least partially verified by hand to verify the correct temperature range and expansion are being used.

Automatic thermal setup

A tool for automatically setting up a lens for thermal analysis is described in "Tools" on page 90.

Adding TCE data

There are two kinds of TCE data. For surfaces which use a glass type named in one of the catalogs (such as BK7 from the Schott catalog), ZEMAX uses the TCE data specified in the catalog. See the chapter "Using Glass Catalogs" for the description of α , the thermal coefficient of expansion value.

If the surface does not use a catalog glass, then the TCE value is extracted directly from the TCE column in the Lens Data Editor. The TCE column is the last column in the spreadsheet, to the right of the parameter columns.

TCE is measured in units of 1E-06 per degree C. Therefore, a value of 23.50 E-06 per degree C would be entered as 23.5. ZEMAX automatically considers the 1E-6 factor when computing thermal effects.

Modeling gases and liquids

See "Modeling gases and liquids" on page 595 for important information on thermal modeling of non-solid materials.

Adding thermal index variation data

The variation of index of refraction with temperature, air pressure, and wavelength is modeled for any glass using the polynomial expression given earlier. The expression requires six coefficients to define the temperature and wavelength dependence on the change in absolute index. It is often the case that for materials added by the user, the six coefficients describing the variation are not available. However, most glass catalogs contain at least a single linear approximation of the ratio of index change with temperature change. This value is called dn/dt . If the only data available is a single dn/dt , then an approximation to the general expression can be made assuming all the coefficients except D_0 are zero:

$$\Delta n_{abs} = \frac{n^2 - 1}{2n} [D_0 \Delta T],$$

which implies that a reasonable approximation to D_0 is given by

$$D_0 = \frac{2n}{n^2 - 1} \frac{dn}{dt}.$$

The D_0 needs to be calculated and entered into the glass catalog. The relative index at the center wavelength measured at the reference temperature for the glass is an adequate value for the index n . Extreme care should be taken to ensure this is an adequate approximation by subsequently checking the computed index values at various wavelengths and temperatures. Note that the dn/dt data should be the absolute, not the relative dn/dt .



Extreme caution and suspicion is warranted when using only a single dn/dt value.

Using a single dn/dt value to estimate the D_0 term is only a rough approximation. The actual index variation with temperature is not likely to be linear over any extended wavelength or temperature range. Therefore, extreme caution and suspicion is warranted when using only a single dn/dt value.

Optimizing athermal lenses

To optimize an athermal lens, first define the multiple configurations required to simulate the lens at each temperature using the methods described in the previous sections. Then define the variables for optimization only in the nominal configuration. For example, suppose the nominal configuration is number 1, and configurations 2, 3, and 4 are defined using thermal pickup solves for every curvature and thickness. Make only the curvatures and thicknesses in the nominal configuration variable.

It is also possible to optimize the TCE of the spacing material between lens groups. To do so, set a variable on the value in the TCE column of the Lens Data Editor.

Limitations of thermal analysis

There are several limitations to the accuracy of the thermal analysis capabilities of ZEMAX. First of all, the TCE data should always be checked for accuracy over the temperature range of use. The index data coefficients should also be verified with the manufacturer of the glass being used.

The thermal analysis does not necessarily work correctly with tilted, decentered, or otherwise unconventional optical systems. The difficulty arises when computing edge thickness displacements on components that are not symmetric; for example, between two lenses tilted with respect to one another.

The thermal index and TCE data provided for Schott glasses comes from Schott, and they state the data is accurate over the temperature range from -40 to +80 degrees Celsius, and over the wavelength range 0.435 to 0.644 micrometers. The data may be extrapolated out to 1.06 micrometers with reduced accuracy. Data for other glasses is provided as is and the range of accuracy is unknown.

Because of the complexity of thermal effects modeling, none of the data should be trusted in critical applications and all computations, index values, and TCE data should be verified independently of ZEMAX. This is true even when working with Schott glasses within the ranges specified above.

Introduction



This feature is only available in the EE edition of ZEMAX.

Ray tracing programs generally treat rays as purely geometric entities, which have only a position, orientation, and phase. For example, a ray is completely described at a surface by the ray intercept coordinates, the direction cosines which define the angles the ray makes with respect to the local coordinate axes, and the phase, which determines the optical path length or difference along the ray.

At the boundary between two media, such as glass and air, refraction occurs according to Snell's law. Usually, the effects at the interface which do not affect beam direction are ignored. These effects include amplitude and phase variations of the electric field which depend upon the angle of incidence, the incident polarization, and the properties of the two media and any optical coatings at the interface.

Polarization analysis is an extension to conventional ray tracing which considers the effects that optical coatings and reflection and absorption losses have on the propagation of light through an optical system.

Review of polarization concepts

The ZEMAX User's Guide is not intended to be a tutorial on polarization theory. The subject is extensive and is treated well elsewhere. See for example "Polarization" by Jean M. Bennett, and "Optical Properties of Films and Coatings" by John A. Dobrowski, both in *The Handbook of Optics Volume I*, McGraw Hill, 1995.

The electric field vector

The amplitude and polarization state of the electric field is described by a vector:

$$E = \begin{bmatrix} E_x \\ E_y \\ E_z \end{bmatrix},$$

where E_x , E_y , and E_z are complex valued. The electric field vector must be orthogonal to the propagation vector of the ray. At a boundary between two media, the transmittance, reflectance, and phase of the electric field is different for the S and P components of the field. The S component of the field is the projection of E that lies along the axis orthogonal to the plane of incidence, while the P projection lies within the plane of incidence. The plane of incidence contains both the ray propagation vector and the surface normal vector at the intercept point. The electric field is then divided into E_s and E_p components, both of which are complex valued.

The electric field after the boundary is computed by the consideration of the S and P transmittance (if the surface is refractive) or reflectance (if the surface is reflective). ZEMAX computes the reflectance and transmittance coefficients, but only uses the one appropriate to propagation of the ray. For brevity, the rest of this discussion will use the terms transmittance and refraction to indicate the terms of interest. If the surface is a reflector (i.e. the glass type is MIRROR) then the reflectance coefficients will be used by the program automatically. The electric field after refraction is then

$$E_s' = E_s \tau_s, E_p' = E_p \tau_p,$$

where the transmittance coefficients, τ_s and τ_p , are complex valued. After the E_s and E_p projections are computed, they are recomposed into the E_x , E_y , and E_z electric field vector and propagation continues to the next surface. The computation of the transmittance coefficients depends upon the index of refraction of the incident media, the (possibly complex) index of refraction and thickness of each layer in the optical coating (if

any), and the (possibly complex) index of refraction of the substrate. The details of this computation are provided in *The Handbook of Optics Volume I*, McGraw Hill, 1995, and will not be repeated here.

Field vs. ray phase conventions

The thin film industry uses a different phase convention than is required for ray tracing. The thin film convention is to measure the phase shift along the normal vector as an imaginary plane wave propagates from the outermost coating layer to the substrate. This convention implies the phase shift is largest at normal incidence, and tapers off to zero by approximately the cosine of the angle of incidence for larger angles of incidence. ZEMAX calls the τ_s and τ_p values "Field" coefficients when using this reference.

For ray tracing, the optical phase advance or delay is measured along the ray. ZEMAX traces rays directly to the substrate, ignoring the coating thickness. This is because the thickness between surfaces specified in the Lens Data Editor is the thickness between the substrates, ignoring coatings. The coating is presumed to grow into the space preceding the surfaces. Correctly computing the phase of the ray requires back propagating the electric field to the point the coating begins, and adjusting the coating phase to be measured along the ray vector rather than the normal vector. ZEMAX calls the τ_s and τ_p values "Ray" coefficients when using this convention. A similar computation yields the factors for reflective surfaces, or for surfaces where the coating stack has automatically been reversed, and ZEMAX automatically chooses the correct factors in all cases.

The default in ZEMAX is to convert the Field coefficients to Ray coefficients, and this setting is recommended in all cases. However, the conversion may be disabled, see "Convert thin film phase to ray equivalent" on page 109.

The polarization ellipse

A convenient way of describing the polarization state of the electric field *for rays parallel to the local Z axis only* is to ignore the Ez component and consider only Ex and Ey. By plotting on a cartesian graph the endpoint of the vector (Ex, Ey) as time moves through one period, a figure is traced out which may be a straight line, a circle, or most generally an ellipse. This curve is defined by the major and minor axis lengths, and the angle the major axis makes with respect to the X axis on the cartesian graph. Note the minor axis may have a length of zero, in the case of linear polarization. For circular polarization the major and minor axis lengths are identical.

Whenever the polarization ellipse is used, it is important to remember that the Ez component is ignored; which limits the utility of the polarization ellipse in beams that are not nearly collimated along the local Z axis.

Definition of terms

There are a number of ways to define parameters of interest for describing polarized light beams. All of the expressions generally depend upon angle of incidence, wavelength, and polarization orientation. Here are the terms used frequently by ZEMAX and their definitions:

Intensity, I

The intensity is the sum of the squares of the individual components of the electric field, or equivalently, the dot product of the electric field and it's complex conjugate:

$$I = E \bullet \tilde{E},$$

where the symbol \tilde{E} denotes the complex conjugate of the vector E . Subscripts are applied to the symbol I to indicate which intensity is listed, for example, Ix would be the intensity in the x direction only.

Phase, P

The phase of any one component of the electric field can be computed with the arg function, which is the arctan(imaginary / real). The phase is given by

$$P = \arg(E) \quad .$$

Subscripts are applied to the symbol P to indicate which intensity is listed, for example, Ps would be the phase for the S component of the electric field.

Path length through medium, τ

For homogeneous medium, the path length is the distance in lens units along the ray from the previous surface to the current surface. Note the index is accounted for separately; this is not the "optical" path length. For gradient index medium, the effective path length (for purposes of computing internal absorption) is the actual integral of the optical path length divided by the base index of the medium (the index at $x=0, y=0, z=0$). The path length is only used for computing the internal absorption. The path length will be negative for virtual propagations.

Internal absorption per lens unit, α

This factor accounts for the internal or bulk absorption propagating through glass. The absorption is computed from the data provided in the glass catalog. See "Defining Transmission Data" on page 594.

Internal Transmittance, IT

The internal transmittance is computed using the path length τ and internal absorption α :

$$IT = e^{-\alpha\tau}.$$

Propagation Phase Factors, pc, ps

As the ray propagates from surface to surface, the electric field rotates according to

$$e^{i\theta}, \theta = (-2\pi\tau)/\lambda,$$

where τ is the path length and λ is measured in the medium. The propagation factors are written as the cosine and sine of the angle θ . The "E Field After" from the previous surface will be multiplied by these factors to yield the input electric field to the current surface. ZEMAX adopts the negative sign convention for the propagation phase factors to be consistent with the convention of negative imaginary index to indicate absorption in coating materials as discussed in "Defining coatings in ZEMAX" on page 613.

Amplitude Reflectance, ρ

The reflectance amplitude is the complex valued coefficient of reflectance for the electric field. These are computed for both S and P polarizations and using both field and ray phase conventions, as described above.

Amplitude Transmittance, τ

The transmittance amplitude is the complex valued coefficient of transmittance for the electric field. These are computed for both S and P polarizations and using both field and ray phase conventions, as described above.

Intensity Reflectance, Rs, Rp

Reflectance intensity is measured normal to the surface. It is given the symbol R, and is always real valued between 0 and 1. Reflectance can be computed from the reflectance amplitude as follows:

$$R = \rho^*\rho.$$

Intensity Transmittance, Ts, Tp

Transmittance intensity is measured normal to the surface. It is given the symbol T, and is always real valued between 0 and 1. Transmittance can be computed from the transmittance amplitude as follows:

$$T = \tau^*\tau.$$

Intensity Absorption, A_s, A_p

Absorption intensity is the intensity of the electric field which is neither transmitted nor reflected:

$$A = 1.0 - T - R.$$

Diattenuation, D

Diattenuation means "two attenuations", and is used to compare the loss of intensity of the S polarized light compared to the P polarized light. Diattenuation is defined as

$$D = \left| \frac{T_s - T_p}{T_s + T_p} \right|.$$

Coating Phase, P_s, P_p

The phase of the transmitted beam is generally different for the S and P polarizations. The phase is given by

$$P_s = \arg(E_s), P_p = \arg(E_p).$$

Retardance, S

For transmission, the retardance is the phase of the P polarization minus the phase of the S polarization:

$$S = \arg(E_p) - \arg(E_s).$$

For reflection, the retardance is the phase of the P polarization minus the phase of the S polarization plus π :

$$S = \arg(E_p) - \arg(E_s) + \pi.$$

The different conventions are used to agree with common practice in the thin films industry.

Phase Difference Between X and Y, P_{xy}

This quantity is the phase difference between the X and Y components of the electric field, or

$$P_{xy} = P_x - P_y.$$

Note this value ignores the Z component of the electric field.

Major and Minor semi axis of the Polarization Ellipse, EM, Em

An ellipse has a major axis and a minor axis; the major axis being the longer of the two. These magnitudes, along with the angle of the polarization ellipse (A_p) uniquely define the polarization state in the XY plane.

Angle of XY Polarization Ellipse, A_p

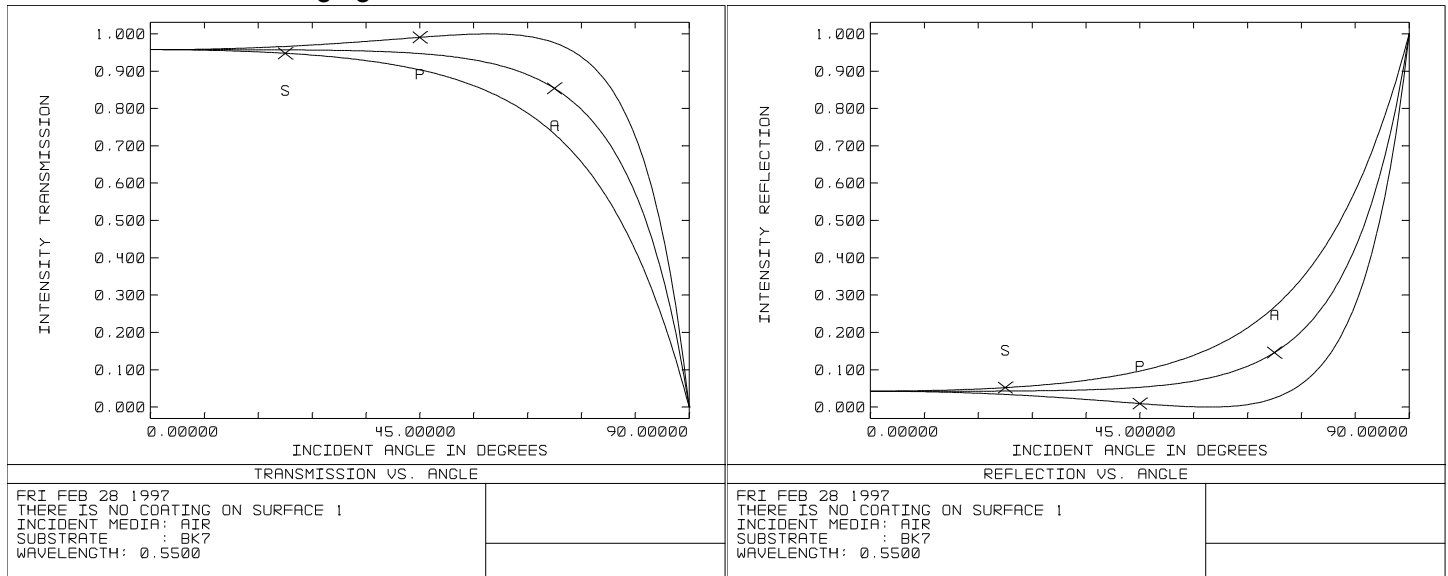
This parameter is the angle relative to the local +x axis that the major axis of the polarization ellipse makes. If the polarization ellipse is a circle, this angle is meaningless.

Properties of uncoated surfaces

At normal incidence only, uncoated surfaces between dielectric media have transmittance and reflectance amplitudes which are described by the Fresnel expressions:

$$R = \left[\frac{n_1 - n_2}{n_1 + n_2} \right]^2, \text{ and } T = \frac{4n_1 n_2}{[n_1 + n_2]^2}.$$

At angles other than normal incidence, the relations are more complex. For example, a surface between air and BK7 at a wavelength of 0.55 micrometers has reflectance and transmittance as a function of incident angle as shown in the following figures.



Note that the reflectance increases substantially with incident angle, until nearly all of the light is reflected rather than transmitted at grazing incidence angles. Note also that the S and P polarization states have different transmittance. This leads to a polarization dependent apodization of the aperture.

Defining coatings in ZEMAX

For dielectric materials, the index of refraction is purely real, and therefore the imaginary part of the index is zero. For metals, the index of refraction is complex. There are two sign conventions in common use for the imaginary part of the index of refraction. ZEMAX uses the following convention:

$$\eta = n + ik,$$

where n is the usual index of refraction and k is the extinction coefficient, a negative value for absorbing materials. The alternate sign convention is $\eta = n - ik$ where k is then a positive value for absorbing materials, but ZEMAX does not use this convention. For example, using the ZEMAX convention, the index of aluminum is approximately given by

$$\eta = 0.7 - 7.0i.$$

Note the extinction coefficient is negative using this convention for typical absorbing materials. This choice of index sign convention is related to the sign convention used by the propagation phase factors, see "Propagation Phase Factors, pc, ps" on page 611.

ZEMAX uses a text file format to define all coating data. The file may be of any valid file name. A sample file called COATING.DAT is supplied with ZEMAX, and COATING.DAT is the default coating file name for any new lens. Multiple coating files may be defined, and the name of the coating file use by any particular lens file is defined on the System, General, Files tab dialog box. It is highly recommended that any modifications to COATING.DAT be saved in a file with a different name, so that subsequent updates to ZEMAX will not overwrite the changes made to COATING.DAT.

In the coating file, The keywords MATE (for material), TAPR (for taper profile), COAT (for coating), TABLE (for coatings defined by a table of points) and IDEAL or IDEAL2 (for ideal coatings) are used to define different types of coating data. All of the material and taper definitions come first, then all of the coating definitions.

Coating file data syntax

The coating file consists of sections of data defining materials, tapers, and coatings. The data uses the following format and syntax. Each data section is described in detail in the text which follows.

```

! Any line starting with the ! symbol is a comment line
! For information on the MATE format see "The MATE data section" below.
MATE <material name>
wavelength real imaginary
wavelength real imaginary
....

```

```

MATE <next material name>
wavelength real imaginary
wavelength real imaginary
....

```

```

! For information on the TAPR format see "The TAPR data section" below.
TAPR <taper name>
DX decenterx
DY decentery
AN rotation_angle (degrees)
RT termnumber radialterm
CT termnumber cosineterm
PT termnumber polyterm
....

```

```

TAPR <next taper name>
DX decenterx
DY decentery
AN rotation_angle (degrees)
RT termnumber radialterm
CT termnumber cosineterm
PT termnumber polyterm
....

```

```

! For information on the COAT format see "The COAT data section" below.
COAT <coating name>
material thickness is_absolute loop_index tapername
material thickness is_absolute loop_index tapername
...

```

```

COAT <next coating name>
material thickness is_absolute loop_index tapername
material thickness is_absolute loop_index tapername
...

```

```

COAT I.transmission

```

```

ENCRYPTED filename

```

```

! For information on the IDEAL format see "The IDEAL data section" below.
IDEAL <name> <Transmitted intensity> <Reflected intensity>

```

```

! For information on the IDEAL2 format see "The IDEAL2 data section" below.
IDEAL2 <name> s_rr s_ri s_tr s_ti p_rr p_ri p_tr p_ti no_pi_flag

```

```

! For information on the TABLE format see "The TABLE data section" below.
TABLE <coating name>
ANGL <angle in degrees>
WAVE <wavelength 1 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp
WAVE <wavelength 2 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp
WAVE <wavelength 3 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp
...
ANGL <next angle in degrees>
WAVE <wavelength 1 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp
WAVE <wavelength 2 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp
WAVE <wavelength 3 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp
...

```

For details on each of these coating definition keywords, see the discussions which follow.

The MATE data section

Syntax:

```
MATE <material name>  
wavelength real imaginary  
wavelength real imaginary  
....
```

The material names may be any user defined name up to 20 characters in length, with no special characters permitted.

ZEMAX uses the material data in the following way:

The wavelength argument is always in micrometers. The wavelengths must be specified in ascending order.

The real value is the real index of refraction of the material at that wavelength.

The imaginary value is the extinction coefficient.

If a single wavelength value is provided for a material, then the real and imaginary parts of the index of refraction are used no matter what wavelength is being traced. Dispersion of the coating material is therefore ignored.

If two or more wavelengths are defined for a material, then for wavelengths shorter than the shortest defined wavelength, the shortest wavelength data is used. For wavelengths longer than the longest defined wavelength, the longest wavelength data is used. For wavelengths in between, linear interpolation is used.

The TAPR data section

Syntax:

```
TAPR <taper name>  
DX decenterx  
DY decentery  
AN rotation_angle (degrees)  
RT termnumber radialterm  
CT termnumber cosineterm  
PT termnumber polyterm  
....
```

Normally, ZEMAX assumes the coating thickness is uniform across the entire optical surface. Because of the way some optical coatings are applied, a thin film coating may have a thickness that varies over the optical surface. A tapered coating is defined by applying a taper function which computes a dimensionless coefficient to multiply the nominal coating thickness. Any layer of any coating may have a taper applied, and each layer may have a different taper. The form of the taper may be a radial polynomial, a cosine function, or a general polynomial in the x and y directions. The various models may all be used simultaneously, and the taper function may optionally be decentered and rotated.

The effective coating thickness is:

$$d_{eff} = d \times (f_r + f_c + f_p),$$

where d is the nominal coating thickness and f_r , f_c , and f_p are the radial, cosine, and polynomial taper factors, respectively. The taper factors are functions of the decentered and rotated coordinates (x' , y'), which are computed from the local surface coordinates using these expressions:

$$x' = (x - x_d)\cos\theta - (y - y_d)\sin\theta,$$

$$y' = (x - x_d)\sin\theta + (y - y_d)\cos\theta,$$

where θ is the taper function rotation angle specified in degrees. The radial coordinate is then defined as:

$$r = \sqrt{(x')^2 + (y')^2} .$$

The radial taper factor is defined as:

$$f_r = \sum_{i=0}^{22} \beta_i r^i ,$$

The cosine taper factor is defined as:

$$f_c = \sum_{i=0}^5 \delta_i (\cos \alpha)^i ,$$

where α is defined as:

$$\sin \alpha = \left| \frac{r}{R} \right| ,$$

where r is the radial coordinate as defined above. For coatings applied to sequential surfaces, R is the radius of curvature of the surface as defined by the "Radius" column in the Lens Data Editor. Cosine tapers should only be used on radially symmetric sequential surfaces for this reason. For coatings applied to non-sequential objects, R is always infinity, regardless of the shape of the object, and cosine tapers should not be used.

The polynomial taper factor is defined as:

$$f_p = \sum_{i=0}^{22} \gamma_i E_i ,$$

where E_i are the polynomial terms in x' and y' . The E_i terms are a power series:

$$E_0 = 1, E_1 = x', E_2 = y', E_3 = (x')^2, E_4 = x'y', E_5 = (y')^2, E_6 = (x')^3, \text{ etc.}$$

Note the cosine taper factor will yield different coatings on substrates that have a different radius of curvature. To define a coating whose thickness varies exactly as $\cos \theta$, the δ_0 cosine taper factor should be zero and the δ_1 cosine taper factor should be 1.0. When applying cosine taper factors to coatings on non-sequential component objects; the substrate radius is ignored and treated as a plane. The radial and polynomial terms are still used.

If the sum of the taper factors evaluates to a negative number, ZEMAX assumes it has a value of zero. The parameters $x_d, y_d, \theta, \gamma_i, \delta_i,$ and β_i are all defined in the taper data section of the coating file. Note that both even and odd powers of the radial coordinate are used, and that the zero order term is used. The coefficients generally have units and magnitudes that depend upon the lens units.



Care should be taken to check that the proper lens units are being used when defining or using tapered coatings.

The taper names may be any user defined name up to 20 characters in length, with no special characters permitted.

ZEMAX uses the taper data in the following way:

If the first keyword on any line following the TAPR keyword is DX or DY, then the data value after the DX or DY keyword on that same line is assumed to be the x_d or y_d in lens units, respectively.

If the first keyword on any line following the TAPR keyword is AN, then the data value after the AN keyword on that same line is assumed to be the taper rotation angle θ .

If the first keyword is RT, the second data value must be an integer between 0 and 20, inclusively. This integer determines the radial term number. The third data value is the coefficient on that term, β_i .

If the first keyword is CT, the second data value must be an integer between 0 and 5, inclusively. This integer determines the cosine term number. The third data value is the coefficient on that term, δ_i .

If the first keyword is PT, the second data value must be an integer between 0 and 20, inclusively. This integer determines the polynomial term number. The third data value is the coefficient on that term, γ_i .

Any terms ZEMAX uses are assumed to be zero unless otherwise defined.

If the entire layer of a coating is to be scaled by a different amount on each surface, or if the coating thickness needs to be optimized, see the discussion about coating multipliers in "Surface coating tab" on page 86 and "Optimizing coatings with ZEMAX" on page 621.

The COAT data section



The order of the coating layers is important! See the discussion in a following section.

Syntax:

```
COAT <coating name>
material thickness is_absolute loop_index tapername
material thickness is_absolute loop_index tapername
...
or
COAT I.transmission
```

The coating names may be any user defined name up to 20 characters in length, with no special characters permitted.

ZEMAX uses the coating data in the following way:

When the coating file is first read, ZEMAX verifies that each coating consists of materials that were defined in the material section. If referenced materials are not defined, an error is issued.

The coating thickness is measured either in units of primary wavelength thickness in that medium (relative definition), or in units of micrometers independent of any wavelength (absolute definition). If relative definition is used, the same coating used by two different lens files, with two different primary wavelengths, will have different actual coating thicknesses. The actual thickness of the coating is determined by:

$$d = \frac{\lambda_0}{n_0} T,$$

where λ_0 is the primary wavelength in micrometers, n_0 is the real part of the index of refraction of the coating at the primary wavelength, and T is the "optical thickness" of the coating specified in the coating file. For example, in a coating material whose real part of the index of refraction is 1.4, a quarter wave coating (T = 0.25) at a primary wavelength of 0.550 micrometers would be 0.0982 micrometers thick. Note that only the real part of the index of refraction is used for determining the thickness of the layer.

If the is_absolute integer is zero, the thickness is interpreted to be relative, otherwise, the coating thickness is absolute in micrometers. If the is_absolute, loop_index, and tapername are omitted, then is_absolute and loop_index are both assumed to be zero, and no taper is assumed.

The loop_index parameter determines how ZEMAX replicates groups of coating layers. See the next section for details.

Defining replicated groups of coating layers

Some coatings contain repeated groups of coating layers. For example, a coating definition may be defined as follows:

```
COAT 3GROUPS
MAT0 0.25 0 0
MAT1 0.25 0 0
MAT2 0.50 0 0
MAT3 0.25 0 0
MAT1 0.25 0 0
MAT2 0.50 0 0
MAT3 0.25 0 0
MAT1 0.25 0 0
MAT2 0.50 0 0
MAT3 0.25 0 0
MAT4 0.25 0 0
```

Note the sequence of layers with materials MAT1, MAT2, and MAT3 is repeated 3 times. This is a perfectly acceptable syntax, however, 11 text lines and 11 layers are required to define the coating. A shorthand syntax that replicates the layers is available using the loop_index parameter. The loop_index integer parameter is set to the number of times a group of layers is to be repeated. The loop_index must be set at BOTH the first AND the last layer (this is to indicate the range of layers, and ZEMAX also reads the coating data in both directions, as described in a following section). The coating above could be written as:

```
COAT 1GROUP
MAT0 0.25 0 0
MAT1 0.25 0 3
MAT2 0.50 0 0
MAT3 0.25 0 3
MAT4 0.25 0 0
```

Note the "3" parameter appears on both the first and third layers listed. This syntax not only saves typing, it reduces the chance of careless typing errors, eases editing, and conserves the total number of layers; which is limited for each coating definition (see the following section on coating data limits). The coating above only requires 5 "layers" to define, although there are physically 11 layers modeled.

There are some important points regarding the use of the loop_index:

There are no limits on the number of different groups that may be defined, or on how many layers each group contains, as long as the total number of lines under the COAT heading does not exceed the maximum number of allowed layers, defined below.

A loop_index of zero is the same as a loop_index of 1; that is, zero does not mean do not include a group at all!

Loops may not be nested, nor may they overlap.

Defining simple ideal coatings

The COAT keyword also may be used to define simple ideal coatings. The syntax is

I.transmission

where transmission is numerical data. For example, the command

```
COAT I.75
```

will create a coating named I.75 which transmits 75% of the ray energy. Ideal coatings defined in this way transmit the specified fraction of light, and reflect the rest, independent of the ray wavelength, angle of incidence, incident material, or substrate material. Note this type of ideal coating has no absorption, and the transmission must be less than 1.0 and greater than 0.0. If the ray total internal reflects, the reflected intensity factor is assumed to be 1.0. For a more general ideal coating, see the IDEAL and IDEAL2 coating definitions which follow.

The ENCRYPTED data section

Syntax:

ENCRYPTED filename

Encrypted coating files are created when a coating is exported to a ZEMAX Encrypted Coating file format as described in “Export Encrypted Coating” on page 249. The resulting ZEC file must be placed in the coating file folder. The filename should have no path or extension; the extension ZEC will automatically be appended.

Encrypted coatings and material data are not listed in the coating listing or prescription reports.

The IDEAL data section

Syntax:

```
IDEAL <name> T R TIR
```

Ideal coatings are defined only by three intensity coefficients: transmission, reflection, and total internal reflection. Ideal coating names may be any user defined name up to 20 characters in length, with no special characters permitted. After the ideal coating name, the T, R, and TIR values are listed. If T+R exceeds 1.0, the values will be scaled so that T+R = 1.0. The absorption coefficient is computed automatically via $A = 1.0 - R - T$, to conserve energy. To make a coating that absorbs 100% of the energy, set both T and R to zero. ZEMAX assumes the phase of transmission is zero. For mirror substrates, the phase of reflection is π . For other substrates, the phase of reflection is zero if the incident index is higher than the substrate index, otherwise the phase of reflection is π . The TIR intensity is used only if the ray total internal reflects. If the TIR value is omitted, 1.0 is assumed.

The IDEAL2 data section

Syntax:

```
IDEAL2 <name> s_rr s_ri s_tr s_ti p_rr p_ri p_tr p_ti no_pi_flag TIR
```

The Ideal2 keyword provides an alternate means of defining ideal coatings. Ideal2 coatings are defined by the real and imaginary amplitude transmission and reflection coefficients for S and P polarized light separately. Ideal2 coating names may be any user defined name up to 20 characters in length, with no special characters permitted. After the coating name, the reflection real, reflection imaginary, transmission real, and transmission imaginary values are listed for S polarized light, followed by these same four values for P polarized light. To make a coating that absorbs 100% of the energy, set both transmission and reflection values to zero. The phase of reflection and transmission is defined by the complex coefficients. If the substrate is not a mirror, and the incident index is higher than the substrate index, ZEMAX changes the sign of the reflected real and imaginary coefficients to yield an additional phase change upon reflection of π , unless the no_pi_flag is set to any value other than zero; in which case no changes are made to the phase of the coating. The TIR intensity is used only if the ray total internal reflects. If the TIR value is omitted, 1.0 is assumed.

The TABLE data section

Syntax:

```
TABLE <coating name>  
ANGL <angle in degrees>  
WAVE <wavelength 1 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp  
WAVE <wavelength 2 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp  
WAVE <wavelength 3 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp  
...  
ANGL <next angle in degrees>  
WAVE <wavelength 1 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp  
WAVE <wavelength 2 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp  
WAVE <wavelength 3 in micrometers> Rs Rp Ts Tp Ars Arp Ats Atp  
...
```

Table coatings are similar to IDEAL coatings, except the transmission and reflection may be a function of incident angle and wavelength and may be specified separately for S and P polarizations. For multiple discrete angles of incidence and wavelengths, four intensity values are defined: Rs, Rp, Ts, and Tp; for intensity reflection and transmission and the s and p subscripts indicate the polarization state. There are also four angles (in degrees) defined: Ars, Arp, Ats, and Atp; these values define the phase angle of the R and T values so ZEMAX can compute the real and imaginary parts of the transmission and reflection. The real part is proportional to the cosine of the angle, and the imaginary part is proportional to the sine of the angle. If the angles are omitted or zero the reflection or transmission is real and no phase change is introduced by the coating. The absorption is computed from $A =$

1 - R - T for each polarization state. If the substrate is not a mirror and the incident index is higher than the substrate index, the phase of reflection is changed by a factor of π .

The primary purpose of the table model is to allow accurate ray tracing where the layer by layer prescription of the coating is not known or is unavailable because the coating vendor will not provide the design data.

The syntax consists of the keyword TABLE followed by the name of the coating. After the TABLE keyword, blocks of data are defined for each incident angle. After each angle definition, R and T values are defined for each wavelength in the table. An example syntax would be:

```
TABLE MYTABLECOAT
ANGL 0
WAVE 0.35 0.014 0.014 0.986 0.986 0.0 0.0 0.0 0.0
WAVE 0.70 0.013 0.013 0.987 0.987 0.0 0.0 0.0 0.0
ANGL 5
WAVE 0.35 0.015 0.015 0.985 0.985 0.0 0.0 0.0 0.0
WAVE 0.70 0.012 0.012 0.988 0.988 0.0 0.0 0.0 0.0
ANGL 45
WAVE 0.35 0.065 0.070 0.935 0.930 0.0 0.0 0.0 0.0
WAVE 0.70 0.062 0.065 0.938 0.935 0.0 0.0 0.0 0.0
ANGL 80
WAVE 0.35 0.600 0.600 0.400 0.400 0.0 0.0 0.0 0.0
WAVE 0.70 0.600 0.600 0.400 0.400 0.0 0.0 0.0 0.0
```

There are several very important points to note:

Angles of incidence less than the smallest angle defined will use the smallest angle data. Angles of incidence greater than the largest angle defined will use the largest angle data. Other values are linearly interpolated in cosine of the angle space. Note ZEMAX interpolates but does not extrapolate. Angles must be listed in ascending order. If data for an angle of incidence of 0.0 degrees is defined, the S and P data should be identical, since S and P are indistinguishable at normal incidence.

Angles of incidence are assumed to be in the incident media, and for coatings placed on refractive surfaces, that the incident media has a lower index of refraction than the substrate. For coatings placed on surfaces where the rays go from high index to low index (such as glass to air), the refracted angle in the low index media is used as the angle of incidence. ZEMAX assumes the phase of transmission is zero if no phase angles are defined. The phase data is angles measured in degrees if the data is provided. If the substrate is not a mirror and the incident index is higher than the substrate index, the phase of reflection is changed by a factor of π .

Wavelengths in between wavelength data points will linearly interpolate in wavelength space.

All wavelengths listed must be identical for each angle. The same ascending order values for the wavelengths is required at each angle for accurate interpolation. An error message will result if the wavelengths are not properly defined.

Generally much more than the brief set of data listed here is required for adequate accuracy. There are no limits on the number of angle and wavelength data points associated with each table coating.

Adding comments to the coating file

Any text following the "!" symbol is considered a comment.

Limits on the amount of coating data

The number of different coating names, materials, table coatings, and other data are in general only by available memory. There are some limits to data that may be stored in any one catalog as follows:

400 layers per coating (see the section "Defining replicated groups of coating layers" on page 618 to work around this limit if required).

120 dispersion data points per material.

60 different tapers.

One way to work around these limits is to use a different coating file for each lens, if required. These limits only apply to a single coating file; multiple coating files may be used for different lenses.

Editing the coating file

There are a few ways of accessing the data in the coating file. First, the file can be edited outside of ZEMAX using any text editor. Second, there is a "Edit Coating File" option under the "Tools" menu on the main menu bar. This invokes the Windows NOTEPAD editor. Also, there is a "Coating Listing" option under the "Tools" menu, which lists the coating and material definition data. If the coating file is edited, the file must be reloaded in order to update the new coating data; there is a menu option on the Tools menu for reloading the catalog.

Optimizing coatings with ZEMAX

ZEMAX is not intended to be an optical coating design program. Optimization of coating designs generally requires some automation of the coating material selection and number of layers.

However, ZEMAX does have the ability to optimize and tolerance coating layer thicknesses, indices, and extinction coefficients, once the number of layers and material selection is made. ZEMAX uses a dimensionless "coating multiplier" which linearly scales the thickness of any layer of a coating, and dimensionless offsets of the index and extinction coefficients. Optimization may then be made on the usual coating performance data using the CODA operand. To optimize the layer thicknesses or offsets, first define a coating with the correct number of layers and materials for each layer. The initial coating thicknesses should be set to either a known good starting value, or 1 wave (in this latter case the coating multipliers may be interpreted to be in units of waves). Do not use zero for a coating layer thickness as this thickness cannot be linearly scaled. Coating multipliers may not exceed the value 10.0. The initial material index and extinction should be set to 1 and 0, respectively, or to a known good starting value.

Coating multipliers and index offsets may then be defined on the surface coating tab described in "Surface coating tab" on page 86. This tab supports setting the multipliers and index offsets to a variable status, which allows for optimization of the relative coating layer thickness and index of refraction of the coating materials. To constrain the coating multiplier values, see "CMGT" on page 494 and the related operands CMLT and CMVA. To constrain the index and extinction offsets, see "CIGT" on page 494 and the related operands CILT, CIVA, CEGT, CELT, and CEVA.

Default materials and coatings supplied with ZEMAX

The default COATING.DAT file supplied with ZEMAX contains several materials and a few common optical coatings. The following table describes some of the included materials and coatings.

DEFAULT MATERIALS AND COATINGS

Materials	Description
AIR	Unity index, used for including air gaps in coatings
AL2O3	Aluminum Oxide, index 1.59
ALUM	Aluminum, index 0.7-7.0i
ALUM2	An alternate definition for aluminum
CEF3	Cerous Flouride, index 1.63
LA2O3	Lanthanum Oxide, index 1.95
MGF2	Magnesium Fluoride, index 1.38
N15	Imaginary material of index 1.50
THF4	Thorium Fluoride, index 1.52
ZNS	Zinc Sulphide, index 2.35
ZRO2	Zirconium Oxide, index 2.1
Coatings	Description
AR	General anti reflection, defined as a quarter wave of MGF2.

Materials	Description
GAP	A small gap of air used to show evanescent propagation
HEAR1	High performance anti reflection coating
HEAR2	High performance anti reflection coating
METAL	A thin layer of aluminum used to make beamsplitters
NULL	A zero thickness coating used primarily for debugging
WAR	Anti reflection "W" coat, defined as a half wave of LA2O3 followed by a quarter wave of MGF2.

These materials and coatings are meant to be used as examples, and may not be applicable in any particular situation. Always check with coating manufacturer or designer for details on coating material specifics. For more detailed information on how these materials and coatings are defined, use the Coating Listing report described in the chapter "Reports menu".

Specifying coatings on surfaces

Once a coating has been defined, it can be applied to a surface by specifying the coating name in the coating column; located at the extreme right side of the lens data editor. ZEMAX interprets the coating definition using one of four rules:

If the surface specified is a boundary going from air to glass, the coating layer order is interpreted exactly as specified in the coating file. The incident media is air, then the outermost layer is listed first (at the top) of the coating definition, then the next layer, etc., with the substrate being the glass type on the surface. The definition of the coating should not include the substrate index or material definition. The term glass here means the glass type is not "MIRROR", and not blank (which is treated as unity index for air). Gradient index lenses are considered glass.

If the surface specified is a boundary between air and air, or glass and glass, the coating is also interpreted exactly as for air to glass, with the appropriate calculations done for the incident and substrate media.

If the surface specified is a boundary going from glass to air, then the order of the layers is automatically reversed, so that the coating is the same as if it had been applied going from air to glass. Therefore, a coating defined by ALHHS, would be interpreted as SHHLA if the boundary goes from glass to air.

If the surface type is a mirror, then the coating definition must include the substrate index. The last layer in the coating definition is then assumed to be a semi-infinite thickness of substrate material.

ZEMAX may also model certain limiting cases, such as frustrated TIR; see "Modeling frustrated total internal reflection" on page 625.

What ZEMAX does if no coating is specified

If the coating column is left blank, then the following is assumed:

If the surface is a mirror, then the surface is assumed to be coated with a thick layer of aluminum, with an index of refraction 0.7 - 7.0i. The aluminum layer is assumed to be thick enough that no light propagates past the layer.

If the surface is a dielectric interface, bare, uncoated surfaces are assumed.

Defining the initial polarization

The individual analysis settings, such as those for the polarization ray trace (see "Polarization Ray Trace" on page 216) allow for specification of the input polarization. The polarization is defined using a Jones vector:

$$J = \begin{bmatrix} J_x \\ J_y \end{bmatrix},$$

where J_x and J_y have both a magnitude and a phase, and the symbol J is used instead of E to distinguish the 2D Jones vector from the 3D electric field vector E . ZEMAX normalizes the specified J_x and J_y values to have unity magnitude, and then scales the intensity as appropriate if any pupil apodization has been specified. The J_x and J_y values are therefore measured in terms of relative electric field amplitude.

Let the ray vector be K , which has X , Y , and Z direction cosines (l , m , n). For rays traveling parallel to the Z axis, or $K = (0, 0, 1)$, the electric field in the Z direction is identically zero, and the conversion from the Jones vector to the Electric field is trivial:

$$E_x = J_x, E_y = J_y, E_z = 0.$$

For a more general ray, the conversion of the Jones vector (J_x, J_y) to the 3D electric field values (E_x, E_y, E_z) is ambiguous. It is not possible to interpret the J_x and J_y values as meaning that for any ray, the J_x value should be applied in such a way to leave the E_y zero, and for the J_y to be applied so that the E_x is zero. The reason is that the E resulting for the Jones vectors ($J_x = 1, J_y = 0$) and ($J_x = 0, J_y = 1$) must be orthogonal to both K and to each other.

ZEMAX allows user selection of three different methods to perform the conversion from J to E . In each method, the vector K refers to the ray vector, the J_x value is the field along a vector S , and the J_y value is the field along a vector P . Note that K , S , and P must all be unit vectors and orthogonal to each other. The three methods are:

X Axis Reference: The P vector is determined from K cross X , and $S = P$ cross K . This method is the default.

Y Axis Reference: The S vector is determined from Y cross K , and $P = K$ cross S .

Z Axis Reference: The S vector is determined from K cross Z , and $P = K$ cross S .

When the object is at infinity, the method selected will change the polarization orientation of S and P for different fields, but all rays from the same field will have the same polarization since all rays are parallel to each other. For finite conjugates, especially when the object space numerical aperture is large, the S and P vector orientations will vary for different rays in the pupil. No matter which method is selected, the transmission results for unpolarized light will be unaffected because any two orthogonal rays may be traced to compute transmission. For systems where a particular polarization is desired, care should be taken to verify that the conversion from J to E yields the intended polarized ray. To review the detailed conversion results, see "Polarization Ray Trace" on page 216.

Defining polarizing components

Any boundary between two media can polarize a beam. However, ZEMAX supports an idealized model for a general polarizing device. The model is implemented as a special "Jones Matrix" surface type for sequential ray tracing, and a "Jones Matrix" object type for non-sequential ray tracing. The Jones matrix modifies a Jones vector (which describes the electric field) according to

$$\begin{bmatrix} E'_x \\ E'_y \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} E_x \\ E_y \end{bmatrix},$$

where A , B , C , D , E_x , and E_y are all complex numbers. In the lens data and in the non-sequential components editor, ZEMAX provides cells for defining A real, A imag, etc.

Note that what happens to the z component of the electric field is not defined by the Jones matrix surface, which implicitly assumes the ray is at normal incidence. The E component of the electric field must satisfy the boundary condition:

$$\hat{k} \cdot \vec{E} = 0,$$

where \vec{E} is the electric field after the Jones matrix and \vec{k} is the ray vector. ZEMAX will adjust either the Ez component or the Ex and Ey components, such that the magnitude of \vec{E} does not increase. The adjustment may however require a reduction in the magnitude of \vec{E} , and thus an associated loss of transmitted energy.

The Jones matrix can be used to define a wide variety of polarizing components. For example, see the sample Jones matrices in the following table. The format of the numbers is (a, b) where the a value is the real part and b is the imaginary part.

SAMPLE JONES MATRICES

Component	Matrix
Null matrix	$\begin{bmatrix} (1, 0) & (0, 0) \\ (0, 0) & (1, 0) \end{bmatrix}$
X analyzer	$\begin{bmatrix} (1, 0) & (0, 0) \\ (0, 0) & (0, 0) \end{bmatrix}$
Y analyzer	$\begin{bmatrix} (0, 0) & (0, 0) \\ (0, 0) & (1, 0) \end{bmatrix}$
Quarter-Wave plate in X direction	$\begin{bmatrix} (0, 1) & (0, 0) \\ (0, 0) & (1, 0) \end{bmatrix}$
Quarter-Wave plate in Y direction	$\begin{bmatrix} (1, 0) & (0, 0) \\ (0, 0) & (0, 1) \end{bmatrix}$
Half-Wave plate in X direction	$\begin{bmatrix} (-1, 0) & (0, 0) \\ (0, 0) & (1, 0) \end{bmatrix}$
Half-Wave plate in Y direction	$\begin{bmatrix} (1, 0) & (0, 0) \\ (0, 0) & (-1, 0) \end{bmatrix}$
Intensity filter, 25% transmission	$\begin{bmatrix} (0.5, 0) & (0, 0) \\ (0, 0) & (0.5, 0) \end{bmatrix}$

What ZEMAX can compute using polarization analysis

ZEMAX can generate plots of R, T, A, D, P, and S as a function of either wavelength for a given incidence angle or as a function of angle given the wavelength. See the "Analysis" chapter for details.

Using the "Polarization Ray Trace" feature, ZEMAX can compute and tabulate the detailed calculations used to trace the polarization state through the system for any given ray. Other features directly compute polarization dependent results; see the "Analysis Menu" chapter for details.

Bulk absorption and transmission

ZEMAX also can accurately compute transmission of individual rays, or compute an average over the pupil. Transmission calculations are based upon surface effects, such as reflection losses, as well as bulk transmission through glass according to Beer's law:

$$t = e^{-\alpha\tau},$$

where α is the absorption coefficient and τ is the path length through the glass. The parameter α generally depends upon wavelength and has units of inverse length. See "Defining Transmission Data" on page 594 for details on defining transmission data. Polarization effects are not considered in all ZEMAX computations.

Modeling birefringent materials


Birefringent materials in sequential optical systems are modeled using the birefringent in/out surface types, described in "Birefringent In and Birefringent Out" on page 287. For non-sequential systems, see "Birefringence tab" on page 560.

Modeling frustrated total internal reflection

Frustrated TIR (FTIR) occurs when a ray of light traveling through glass strikes an interface at an angle exceeding the critical angle. If another dielectric medium, such as another piece of glass, is placed close, but not touching the interface, some light will transmit through the thin gap and propagate, even though the refraction at the boundary cannot satisfy Snell's law. Both the reflected and the transmitted beams will be affected, depending upon the thickness of the gap. In the limit of the gap having zero thickness, the light will continue through as if there were no boundary. In the limit of a large gap, more than a fraction of a wavelength, then virtually all of the light is perfectly reflected.

Although this appears to be a special case, the very general model incorporated into the polarization routines of ZEMAX handle FTIR as any other coating problem. The trick is to define a coating that reflects the situation. One of the default coatings is GAP, defined to be 0.1 waves of AIR, with index unity. Using the tilted surface with a Y-Tangent of 1.0, and placing the coating GAP on the boundary yields the following system, with correctly modeled polarization properties. Note the air gap may be changed by editing the coating file as described in "Defining coatings in ZEMAX" on page 613.

With the GAP coating in place, ZEMAX can now compute either transmitted or reflected polarization properties.



Transmission vs. angle for a FTIR beamsplitter.

A FTIR beamsplitter in transmission. The tilted surface between the two glass blocks has a GAP coating.

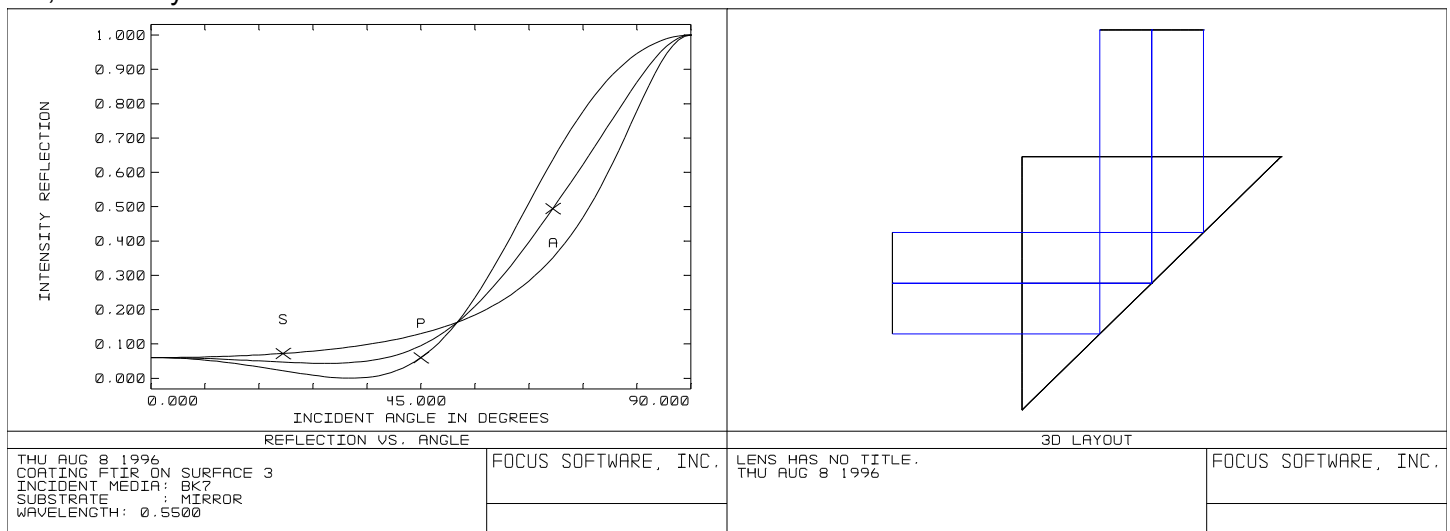
However, the ray tracing for other purposes (such as computing system transmission, ray fans, optimization) will only proceed along the transmitted path. To model the reflected path, give the internal tilted surface a glass type of MIRROR, add a 90 degree coordinate break after the tilted surface, and change the sign of the remaining thicknesses as usual.

The coating in this case now needs to be modified to indicate the presence of the second dielectric block, even though no ray tracing will be conducted in the transmitted path; just the reflected path. A new coating needs to be defined in the coating file (called FTIR in this example) that looks like:

```
COAT FTIR
AIR .1
N15 1
```

The "coating", when applied to a mirror, treats the mirror as a gap of air followed by piece of glass. Note that detailed modeling of the index and dispersion of the substrate glass is certainly possible by definition of a glass in the coating file; here N15 was used for simplicity. The thickness of the last material, N15, which defines the substrate, is irrelevant because ZEMAX always assumes the substrate is semi-infinite when applied to a MIRROR. The resulting lens figure and reflection plot is shown. Both the reflection and transmission plots are of course identical for either system model.

Finally, note that conventional TIR can also be modeled by increasing the gap thickness. Distances much greater than a fraction of a wavelength are essentially equivalent to an infinite gap. However, very large thicknesses for evanescent wave propagation should not be defined in the coating file because of "underflow" numerical errors which arise because the transmission is exceedingly low. One wave, or an optical thickness of 1.0, is usually sufficient.



Reflection vs. angle for a FTIR beamsplitter.

A FTIR beamsplitter in reflection. The tilted "mirror" surface has a FTIR coating.

Limitations of polarization analysis

There are a few limitations to the polarization analysis capabilities of ZEMAX. The data in the coating file should always be carefully checked for accuracy whenever fabrication decisions are being made. Also, the reflectivity and transmission curves generated by ZEMAX should always be presented to the coating manufacturer as a check to ensure the coatings will behave as intended.

The polarization ray tracing algorithm generally does not return correct results for certain types of surfaces, such as paraxial surfaces, diffraction gratings, binary optics, or Fresnel surfaces. Generally, only surfaces which have refraction or reflection described by Snell's law have valid polarization data computed. In particular, the polarization algorithm assumes that rays remain in the plane of incidence after refraction or reflection. This is true for conventional refractive optics (it is a condition of Snell's law) but is not true in general for diffractive optics and "imaginary" surfaces such as paraxial lenses.

The polarization effects of propagation through gradient index media are not considered, although the surface effects (reflection, transmission coefficients) are accounted for. Because the rays take curved paths through gradient media, the polarization vector is rotated in a manner ZEMAX does not account for.

Introduction



This feature is only available in the EE edition of ZEMAX. Physical optics propagation is complex, and the user is urged to read and understand this chapter before using the feature.

Geometrical optics is the modeling of optical systems by tracing rays. Rays are imaginary lines which represent normals to the surfaces of constant phase, called the wavefront. Either rays or wavefronts can be used to represent a beam. However, rays and wavefronts are propagated differently. Rays propagate along straight lines without interfering with one another, wavefronts propagate while coherently interfering with themselves. For this reason, the ray model and the wavefront model yield different representations of the beam as it propagates through free space or through optical components. The ray method is fast, flexible, and extremely useful for modeling almost any optical system. However, rays are not well suited to modeling certain important effects, primarily diffraction.

ZEMAX does have some ray based diffraction computations, such as the diffraction MTF or PSF. These diffraction computations make a simplifying approximation: that all the important diffraction effects occur going from the exit pupil to the image. This is sometimes called the "single step" approximation. Rays are used to propagate the beam from the object, through all the optics and intervening spaces, all the way to the exit pupil in image space. The ray distribution in the exit pupil, with transmitted amplitude and accumulated OPD used to compute the phase, is used to form a complex amplitude wavefront. Then, in a single step, a diffraction computation is used to propagate this complex amplitude wavefront to the region near focus.

Geometrical optics and the single step approximation work quite well for the majority of traditional optical designs, where the beam is not near focus anywhere except the final image. However, the model breaks down for several important cases:

When the beam comes to an intermediate focus, especially near optics that truncate the beam (rays by themselves do not predict the correct distribution near focus).

When the diffraction effects far from focus are of interest (rays remain uniform in amplitude and phase, wavefronts develop amplitude and phase structure).

When the propagation length is long and the beam is nearly collimated (collimated rays will remain collimated over any distance, real beams diffract and spread).

Physical optics is the modeling of optical systems by propagating wavefronts. The beam is represented by an array of discretely sampled points, analogous to the discrete sampling using rays for a geometric optics analysis. The entire array is then propagated through the free space between optical surfaces. At each optical surface, a transfer function is computed which transfers the beam from one side of the optical surface to the other. The physical optics model allows very detailed study of arbitrary coherent optical beams, including:

Gaussian or higher order multi-mode laser beams of any form (beams are user definable).

Beams may be propagated along any arbitrary field position (skew beams).

Amplitude, phase, and intensity may be computed at any surface in the optical system.

Effects of finite lens apertures, including spatial filtering, may be modeled.

Accurate computation of propagation through any optical component ZEMAX can model via ray tracing.

The physical optics model is generally more accurate at predicting the detailed amplitude and phase structure of the beam away from focus than conventional ray tracing. However there are some disadvantages to the physical optics propagation analysis:

Physical optics is generally slower than geometrical optics.

Because the entire beam array must be stored in computer memory at once, the required RAM may be quite large for large sampling arrays.

The sampling limits the amount of aberration in the beam that can be accurately modeled. For highly aberrated systems, geometrical optics should be used.

The following sections will summarize the physical optics propagation algorithms and the information needed to properly use this advanced ZEMAX feature.

Support for multiple processors

The physical optics propagation routines in ZEMAX are designed to work on multiple CPU computers. Computers with two or more CPU's will execute the physical optics propagation feature faster than single CPU computers.

Diffraction propagation

The theory and methodology of diffraction propagation is well understood and widely detailed in available literature. The methods used in ZEMAX are based upon these references:

Goodman, Joseph W., *Introduction to Fourier Optics*, McGraw-Hill, New York (1968).

Lawrence, George N. "Optical Modeling", in *Applied Optics and Optical Engineering*, Volume 11, R. R. Shannon and J. C. Wyant, eds., Academic, New York (1992).

Only the material relevant to using the physical optics propagation feature in ZEMAX will be summarized here.

Representation of the electric field

The electric field may be represented in three dimensions as

$$\vec{E}(x, y, z) = E_x \hat{x} + E_y \hat{y} + E_z \hat{z},$$

where the E values are all complex and \hat{x} , \hat{y} , and \hat{z} are the cartesian unit vectors. The coordinate system used by ZEMAX is that the beam propagates primarily down a local Z axis. The Z axis used to represent the beam is aligned with a reference chief ray in each optical space, and therefore this Z axis is not generally the same as the Z axis defined by the Lens Data Editor which is used to position optics. Because the beam is propagating along the local Z direction, the first approximation made is to neglect the E_z component. Since the electric field must always be normal to the ray propagation direction, E_z can be reconstructed from other data when required, as will be described later.

By keeping track of the electric field components along both the X and Y axes, effects due to polarization may be studied, such as transmission and reflection losses, polarization aberrations, and of course the polarization state of the beam. If polarization effects are not required, the Y component of the field may be ignored, speeding the computations.

The Fresnel number

A very useful concept in physical optics modeling is the Fresnel number. Strictly speaking, the definition of the Fresnel number only applies to unaberrated rotationally symmetric beams with a finite extent. However, the concept is still useful in cases that do not meet these criteria. The Fresnel number depends upon the diameter of the beam, the radius of curvature of the wavefront phase, and the distance to an observation point where the complex amplitude of the field is desired. Conceptually the Fresnel number is the number of annular "Fresnel zones" from the center of the beam to the edge. Fresnel zones are the radial zones where the phase as seen from the observation point changes by π .

A perfectly collimated beam will have a Fresnel number given by

$$F_n = \frac{2}{\lambda} [\sqrt{Z^2 + A^2} - Z],$$

which for Z greater than A reduces to approximately

$$F_n = \frac{A^2}{\lambda Z},$$

where A is the radial size of the beam and Z is the distance from the beam to the observation point. The Fresnel number becomes small as Z grows large.

For beams that are not collimated, the concept is the same. A converging beam will have a very small Fresnel number if the observation point is near focus. A perfectly spherical beam converging to focus will have a Fresnel number of zero, since there are no zones where the observed phase reaches π . As the observation point moves from the focal region, the Fresnel number increases.

Near and far field

If the Fresnel number is small, less than roughly 1, then the beam at the observation point is said to be in the "far field" relative to the current beam. For Fresnel numbers larger than 1, the beam at the observation point is said to be in the "near field" relative to the current beam.

It is important to consider the terms near and far as being relative to the propagation from the present location of the beam to the observation point at which the Fresnel number is computed, rather than having any rigid relationship to the beam position alone. For example, a beam in the exit pupil of an optical system is typically called the near field because the far field is at focus. However, a short propagation from focus to a slightly out of focus observation point is likely a near field propagation if the defocus is small.

The decision as to whether a propagation is in the near or far field will determine the choice of diffraction propagation methods.

Angular spectrum propagation

It is well known that plane waves remain plane waves while propagating in homogeneous medium. A plane wave is represented by

$$e^{i\vec{k} \cdot \hat{z}},$$

where \vec{k} is the wave vector, with magnitude $(2\pi)/\lambda$, and \hat{z} is the local z direction. The vector \vec{k} points along the normal to the wavefront in the direction of propagation. This normal vector has direction cosines α , β , and γ , where

$$\alpha^2 + \beta^2 + \gamma^2 = 1.$$

The plane wave can then be written as

$$e^{\frac{i2\pi}{\lambda}(\alpha x + \beta y + \gamma z)}$$

Now recall the definition of the Fourier transform and the inverse Fourier transform:

$$FF[A] \equiv \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} [A] e^{-i2\pi(x\xi + y\eta)} dx dy ,$$

$$FF^{-1}[A] \equiv \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} [A] e^{i2\pi(x\xi + y\eta)} d\xi d\eta .$$

Let $G = FF[E]$, that is, G is the Fourier spectrum representation of the electric field E . By definition then

$$E(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} G(\xi, \eta) e^{i2\pi(x\xi + y\eta)} d\xi d\eta .$$

The electric field can therefore be interpreted as being the integral of a collection of plane waves propagating with direction cosines

$$\alpha = \lambda\xi, \beta = \lambda\eta, \text{ and } \gamma = \sqrt{1 - (\lambda\xi)^2 - (\lambda\eta)^2} .$$

Eliminating γ , and **making the approximation** that the plane wave propagates at a small angle with respect to the Z axis (for more information on these approximations see "Algorithm assumptions" on page 648), the plane wave equation can be rewritten

$$e^{\frac{i2\pi z}{\lambda} \sqrt{1 - \alpha^2 - \beta^2}} \approx e^{\frac{i2\pi z}{\lambda}} e^{-\frac{i\pi z(\alpha^2 + \beta^2)}{\lambda}} .$$

The $e^{\frac{i2\pi z}{\lambda}}$ term is just a phase propagation term that is normally neglected. The term which depends upon α and β is the transfer function for a plane wave in free space. Defining $\rho^2 = \xi^2 + \eta^2$ the plane wave transfer function can then be rewritten as

$$G(\xi, \eta, z) = G(\xi, \eta, 0) e^{-i\pi\lambda z \rho^2} .$$

To propagate an electric field from one plane to another, the field needs to be Fourier transformed, the plane wave propagator applied, and then the resulting distribution inverse Fourier transformed. These operations may be summarized by defining the plane to plane (PTP) operator:

$$E(x, y, z_2) = PTP[E(x, y, z_1), (z_2 - z_1)] , \text{ where}$$

$$PTP(E, \Delta z) \equiv FF^{-1}[T(\Delta z)FF[E]] , \text{ and}$$

$$T(\Delta z) = e^{-i\pi\lambda\Delta z\rho^2} .$$

Note that the transfer function $T(\Delta z)$ has unity amplitude but a complex phase. This phase varies slowly from point to point in the frequency domain representation G if $\lambda\Delta z\rho^2$ is small. But if $\lambda\Delta z\rho^2$ grows large, the phase variations become increasingly rapid. If the phase changes by more than about $\pi/2$ between adjacent points in the finite array, the phase becomes ambiguous, and a phenomenon known as *aliasing* occurs. For this reason, the angular spectrum method works very well if the propagation distances are fairly short or if the beam is nearly collimated. Although the diffraction theory is accurate for any propagation distance, when the beam is represented by a finite sized array of discrete points, the phase of the beam cannot be accurately represented if the phase of the angular spectrum propagator changes too rapidly between points.

When using the angular spectrum propagator, the phase of the electric field is measured relative to a plane. Positive phase indicates the wavefront is advanced along the local +z axis relative to the plane, regardless of the direction of propagation.

The angular spectrum propagator is useful when the Fresnel number is large. This includes the important case of propagating a beam a short distance. However, the angular spectrum propagator also works well for propagating a large distance when the divergence of the beam (and thus ρ) is small. A good rule of thumb to use is that if the beam does not change size significantly, the angular spectrum propagator may be used. To propagate beams with small Fresnel numbers, where the beam will change size significantly, requires a separate theoretical and numerical method.

Fresnel diffraction

For small Fresnel numbers, the appropriate theory is Fresnel diffraction. For a complete discussion and derivation, see the Goodman reference given in the introduction. The key assumptions in Fresnel theory require that the field being computed is not too close to the initial field, namely, if $z_2 - z_1 = \Delta z$, then Δz is large compared to the region over which the field at z_2 is to be determined. Another way of saying this is that the beam cannot diverge too quickly; very fast $F/\#$ beams cannot be accurately modeled with Fresnel diffraction theory (for more information on these approximations see “Algorithm assumptions” on page 648).

In the Fresnel region the electric field distribution is given by

$$E(x_2, y_2, z_2) = \left[\frac{e^{ikz}}{i\lambda\Delta z} \right] q(r_2, \Delta z) \int_{-\infty-\infty}^{\infty} \int_{-\infty-\infty}^{\infty} E(x_1, y_1, z_1) q(r_1, \Delta z) e^{-\frac{i2\pi}{\lambda\Delta z}(x_1x_2 + y_1y_2)} dx_1 dy_1, \text{ where}$$

$$q(r, \Delta z) = e^{(i\pi r^2)/(\lambda\Delta z)}.$$

Each of the terms in the above expression has a useful physical interpretation. The leading term indicates that as the beam propagates, the phase changes along the z axis, just like the plane wave described earlier. The amplitude also decreases linearly with distance, or the intensity (E^*E) falls quadratically. The expression for $q(r, \Delta z)$, called the quadratic phase factor, indicates that the phase is referenced to a sphere of radius Δz (strictly speaking it is a parabola, but we have already assumed in the Fresnel development that $r_2 \ll \Delta z$). This is a very useful property; all that is required in the representation of the electric field is the phase difference relative to the reference sphere. This significantly reduces the number of sample points needed to accurately define the phase of the beam.

When using the Fresnel propagator, the phase of the electric field is measured relative to a reference sphere with a radius equal to that of the distance from the beam waist. This is not the same radius as the phase radius of curvature of the Gaussian beam. Positive phase indicates the wavefront is advanced along the local +z axis relative to the reference sphere, regardless of the direction of propagation.

Another important property of $q(r, \Delta z)$ is that as Δz gets larger, $q(r, \Delta z)$, varies *more slowly* in phase. This is the opposite of the $T(\Delta z)$ operator, which varies rapidly in phase as Δz gets larger. Accordingly, Fresnel diffraction is useful when the Fresnel number is small.

Selecting the correct propagator

ZEMAX will automatically choose the angular spectrum propagator when the Fresnel number is large, and the Fresnel propagator when the Fresnel number is small. However, there are times when the angular spectrum propagator is a better choice than the Fresnel, and ZEMAX supports a surface specific option to choose the angular spectrum propagator rather than the default choice.

Fraunhofer diffraction

Consider the Fresnel diffraction expression. If Δz is very large, then $q(r, \Delta z)$ may be neglected. This yields the Fraunhofer diffraction expression, which is

$$E(x_2, y_2, z_2) = \left[\frac{e^{ikz}}{i\lambda\Delta z} \right] \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E(x_1, y_1, z_1) e^{-\frac{i2\pi}{\lambda\Delta z}(x_1x_2 + y_1y_2)} dx_1 dy_1, \text{ or}$$

$$E(x_2, y_2, z_2) = FF(E(x_1, y_1, z_1)),$$

where phase and amplitude factors are omitted. The far field distribution is just a scaled version of the Fourier transform of the near field distribution. Fraunhofer diffraction is only valid if the Fresnel number is nearly zero.

The ray based diffraction features in ZEMAX, such as the diffraction MTF and PSF assume Fraunhofer diffraction. This is why ZEMAX cannot compute the ray based diffraction MTF or PSF if the beam is too much out of focus. The Fraunhofer assumption is never used by the physical optics propagation algorithm in ZEMAX, it is presented here for completeness.

The pilot beam

Consider a Gaussian beam with waist ω_0 . The Rayleigh range is given by

$$z_r = \frac{\pi\omega_0^2}{\lambda}.$$

The phase radius of curvature of the beam is a function of the distance from the beam waist, z :

$$R(z) = z + \frac{z_r^2}{z}.$$

Note that radius is infinite at $z = 0$, reaches a minimum of $2z_r$ at $z = z_r$, and asymptotically approaches infinity as $z \rightarrow \infty$. The phase of the Gaussian beam along the axis is defined by the Gouy shift, given by

$$\theta(z) = \tan^{-1}\left(\frac{z}{z_r}\right).$$

For example, the axial phase is $\pi/4$ at a distance of plus one Rayleigh range. The beam size is also a function of the distance from the waist:

$$\omega(z) = \omega_0 \left[1 + \left(\frac{z}{z_r} \right)^2 \right]^{1/2}.$$

Note for large distances the beam size expands linearly. The divergence angle of the beam is given by

$$\theta = \tan^{-1} \frac{\omega_0}{z_r}.$$

Now consider the problem of numerically representing this beam by a discrete sampling of points. If a constant spacing between points is used, the beam will expand beyond the edges of the array if the propagation proceeds too far from the waist. Therefore, far from the waist, a linearly expanding coordinate system where the point spacing is proportional to z is best. However, near the waist, the beam size does not decrease to zero, but remains reasonably constant. In this domain, a constant sampling is most convenient.

The compromise sampling system is to use a constant spacing near the waist, and a linearly scaled spacing far from the waist. The transition point is somewhat arbitrary, however for optical design applications a value of twice the Rayleigh range has proved useful, and ZEMAX uses this threshold. For this reason, in the following discussion about propagation in and out of the Rayleigh range, the actual threshold point may be different than the Rayleigh distance. The concept however remains the same.

The diffraction theory developed earlier did not assume any particular shape or form to the electric field being propagated. The algorithms are for the most part independent of the field distribution. However, the problem of sampling remains. It is also impractical to compute (or even define) Fresnel numbers for arbitrary, aberrated beams with irregular or non-existent apertures.

For this reason, a *pilot beam* is used to assist the physical optics propagation algorithm in determining which propagation algorithm to select. The pilot beam is an ideal Gaussian beam, with a waist, beam size, phase radius, and relative z position. The initial parameters may be generated by fitting the Gaussian beam equations to the initial distribution. The pilot beam is then propagated from surface to surface. At each surface, new beam parameters, such as the new waist, phase radius, or position are computed. The properties of the pilot beam are then used to determine if the actual distribution is inside or outside the Rayleigh range, and what propagation algorithms are appropriate.

After passing through an aperture that significantly truncates the beam, such as a pinhole aperture, it may be required to recompute the pilot beam parameters, as described later in this chapter.

Sign conventions for phase data

It is very important to understand the phase sign conventions when constructing beams using DLL files, beam files, or when performing fiber coupling calculations where phase matching is critical.

As shown in the previous sections, the angular spectrum propagator works best when the beam is nearly collimated, while the Fresnel theory works best when the beam is diverging. When using the angular spectrum propagator, the phase of the electric field is referenced to a plane.

Once the beam propagates past the Rayleigh range, the Fresnel propagator is used, and the phase of the electric field is referenced to a sphere whose radius is the distance from the beam waist to the current position of the pilot beam.

The sign of the phase is positive if the wavefront is to the "right" of the reference surface, with "right" being towards the positive local Z axis direction. For a beam propagated to the +z side just inside the Rayleigh range, the phase slope is negative because the wavefront is left of the reference plane. Just outside of the Rayleigh range, the phase slope is positive because the wavefront is now to the right of the reference sphere.

Therefore, the slope of the phase of the electric field will "flip" from negative to positive when crossing from inside to outside the Rayleigh range on the +z side.

For example, consider a Gaussian beam. At the waist, the phase is zero everywhere on the reference plane. If the beam is propagated to just inside the Rayleigh range on the positive side of the waist, the phase at the center of the beam will change to $\pi/4$ radians, this is the Gouy shift given by $\varphi = \tan^{-1}(z/z_r)$. The phase radius of curvature of the beam will become twice the Rayleigh range distance. The wavefront phase will be increasingly negative relative to the center of the wavefront as the radial aperture is increased, because the curved wavefront phase is measured relative to a plane.

If the beam now propagates a small distance so the pilot beam is now just on the far side of the Rayleigh range, the phase will remain $\pi/4$ radians at the center but will now be referenced to a sphere whose radius is the distance to the waist, which is the Rayleigh range for this case. Because the radius of the reference sphere is smaller than the curvature of the wavefront, the slope of the phase of the beam will now "flip" to be increasingly positive with radial aperture.


Finally, as the beam propagates further past the Rayleigh range, the phase relative to the reference sphere will tend toward a constant value of $\pi/2$ radians (the limiting value of the Gouy shift).

The two representations of the beam are equivalent, however care must be taken to account for the proper phase reference surface when constructing or comparing different beams.

Propagating in and out of the Rayleigh range

The naming conventions used in this section are from the Lawrence reference cited earlier in this chapter.

It has already been shown that short propagations are well modeled using the PTP operator. This operator has the property that the sample spacing remains constant.



It is most convenient to use the Fresnel propagator when the beam is already at the waist of the pilot beam, and the field is desired far from the waist relative to a reference sphere. For this reason, the Fresnel propagator is redefined as the waist to sphere (WTS) operator:

$$E(x_2, y_2, z_2) = WTS[E(x_1, y_1, 0), z_2] \quad , \text{ where}$$

$$WTS(E, \Delta z) \equiv \left[\frac{1}{i\lambda z_2} \right] FF^s[q(r_1, z_2)E], \quad s = \frac{z_2}{|z_2|},$$

$$\Delta x_2 = \frac{\lambda|z_2|}{n_x \Delta x_1}, \quad \text{and} \quad \Delta y_2 = \frac{\lambda|z_2|}{n_y \Delta y_1},$$

where n_x and n_y are the number of points in the x and y directions of the array. The last two expressions yield the new linearly scaled sample spacings after application of the operator. Reversing the order of operations results in the sphere to waist (STW) operator:

$$STW(E, \Delta z) \equiv q(r_1, z_2) FF^s[i\lambda z_2 E],$$

with a similar change in the sample spacing. There are four possible general propagation cases to consider:

II: propagation from inside to inside the Rayleigh range.

IO: propagation from inside to outside the Rayleigh range.

OI: propagation from outside to inside the Rayleigh range.

OO: propagation from outside to outside the Rayleigh range.

All these cases can be handled with the appropriate combination of the PTP, WTS, and STW operators:

$$II(z_1, z_2) = PTP(z_2 - z_1)$$

$$IO(z_1, z_2) = WTS(z_2 - z_0)PTP(z_0 - z_1)$$

$$OI(z_1, z_2) = PTP(z_2 - z_0)STW(z_0 - z_1)$$

$$OO(z_1, z_2) = WTS(z_2 - z_0)STW(z_0 - z_1)$$

where z_0 is the pilot beam waist position, z_1 is the starting beam position, and z_2 is the end beam position.

Separation of X and Y propagation

For rotationally symmetric optics, a single rotationally symmetric pilot beam and propagation method is appropriate. When the beam is astigmatic or anamorphic, and/or the optical components are cylindrical or toroidal, then a more efficient model for the beam is to separate the X and Y propagation. Separation of X and Y direction propagation is accomplished by choosing the appropriate option, see "General Tab:" on page 226.

If the propagation in X and Y is separated, then there will be two pilot beams; one each for the X and Y direction. The phase references will also be distinct in the X and Y directions. In place of a spherical reference surface, a "bispherical" surface is used instead. The sag of a bispherical surface is given by:

$$z = \frac{c_x x^2 + c_y y^2}{1 + \sqrt{1 - c_x^2 x^2 - c_y^2 y^2}},$$

where

$$c_x = \frac{1}{R_x}, c_y = \frac{1}{R_y}.$$

The pilot beam, including beam phase radius, waist, position, Rayleigh Range, divergence, are all distinct in the X and Y directions. There is a small amount of additional computational overhead in propagating the X and Y portions of the beam independently, and so this option should only be used when required.

Comments about point spacing and sampling

Although the total number of array points n_x and n_y , remain constant, the array size and point spacings Δx and Δy will change as the beam propagates. If the array width is very large at the beam waist relative to the waist size, then there are relatively few points across the beam waist. This will result in a smaller array size far from the waist, with a relatively large number of points across the beam size. Conversely, if the array size is small at the waist, the array size will grow large compared to the beam far from the waist, leaving few sample points to represent the beam. This inverse relationship is a necessary but frequently inconvenient product of the Fourier transform theory used to model the diffraction. The exact equations describing the change in point spacing are given in the previous section.

There is clearly a tradeoff between good sampling of the beam near the waist and good sampling far from the waist. It can be shown that to achieve approximately uniform sampling relative to the beam size at both the waist and far from the beam waist the array size at the beam waist in the X and Y directions should be

$$X = \omega_{0x} \sqrt{\pi n_x}, Y = \omega_{0y} \sqrt{\pi n_y}.$$

The physical optics analysis feature settings include an "Auto" button which will use this formula to set a suggested initial array width. See "Physical Optics Propagation" on page 225.

Propagation through arbitrary optical surfaces

The methods described thus far are good for propagating through homogeneous space. However, the prime interest is in propagating the beam through optical components such as lenses and apertures.

It is not practical to directly perform a diffraction propagation through an arbitrary surface shape. The difficulty lies in the representation of the complex amplitude at discrete points in a plane or spherical phase referenced array. When the beam is incident on a curved surface, different parts of the beam intercept the surface at different points along the local z axis.

To avoid this problem, the properties of the pilot beam are used to generate a set of rays that represent the wavefront incident upon the surface. This set of rays, called the *probing rays*, is then traced through the optical surface using conventional geometric optics. The path length of each ray, and the positional and angular aberrations generated, are then used to reconstruct the pilot beam and complex amplitude after the surface. The probing ray set is used to generate the transfer through the surface.

The probing ray set can be used to determine:

The effective power of the surface with respect to the pilot beam.

The direction of the beam after leaving the surface.

The polarization phase and amplitude transmission of a surface.

The vignetting of the surface with any supported surface apertures (or ray errors). Note that surfaces without hard apertures defined will pass the entire beam without vignetting, independent of how the system aperture is set.

The phase aberrations.

The anamorphic and non-linear stretching and compression (distortion) of the beam.

The probing ray method can be applied to propagation through a single surface, or multiple surfaces at once. This is a very desirable property, because geometrical optics may be used to propagate through whole optical components that would be difficult to model with physical optics propagation. These include highly tilted surfaces and gradient index lenses, to name a few. Propagating through multiple surfaces at once using rays also speeds up the analysis. Some special surfaces, such as the ABCD matrix surface, and some types of Fresnel surfaces, do not allow POP analysis at all because there is no way to compute the effective phase of these surface types. A warning is issued by ZEMAX if the POP analysis cannot proceed due to the presence of these special surface types.

Propagating through non-sequential surfaces

Non-sequential component (NSC) surfaces pose a difficult challenge for physical optics propagation. A NSC surface may represent a large number of discrete components. The incoming beam may not hit all or any of these components, or perhaps a portion of the beam will hit one component and split off from the rest of the beam, or any number of other possibilities. The Fresnel propagation method used by POP is not suitable for computing the effects of propagation through arbitrary NSC surfaces.

There are a limited number of cases where POP can be used to propagate a beam through a NSC surface with reasonable confidence in the results. These include:

- NSC surfaces with a single lens or mirror.
- Prisms with an effectively sequential path and apertures large compared to the beam.

Even in these cases, POP can only use ray tracing to model the effective phase the NSC surface imparts to the beam. Diffraction propagation is not possible. For this reason, only relatively slow beams propagating through thin NSC surfaces should be used. Always check the POP results carefully before having confidence in results of propagation through NSC surfaces.

Accounting for polarization

If polarization is used, ZEMAX will use polarization ray tracing to determine the properties of the probing ray set and corresponding transfer function. Polarization ray tracing permits the modeling of the effects of optical coatings on the phase and amplitude of the transmitted or reflected beam.

ZEMAX assumes the Ex and Ey portions of the field account for all the energy in the beam at any given array point. However, the Ez component is required for polarization ray tracing. ZEMAX recreates the Ez as needed by applying the condition that E must be perpendicular to the propagation vector k. The vector k is computed from the pilot beam properties and the phase errors present on the beam. The Ex and Ey fields are renormalized to account for the correct beam intensity. After polarization propagation, the Ex and Ey components are renormalized to again hold the energy present in the Ez component, which is then discarded.

Memory requirements

Geometric optics ray tracing generally requires very little RAM, because each ray may be traced from object to image, independently of all other rays. Physical optics is far more demanding, because the entire beam must be stored at once. ZEMAX uses 8 byte "double precision" floating point numbers to store the beam data. There are actually 16 bytes per array position because the electric field amplitude is a complex number. Here are some guidelines for how much RAM is required for various parts of the propagation algorithm:

Task	RAM Requirement
Storing an unpolarized beam	16 bytes per grid point
Storing a polarized beam	32 bytes per grid point
Computing the surface transfer probing ray set	56 bytes per grid point Some special surface types such as paraxial and birefringent surfaces require additional data and may require up to 90 bytes per grid point

For example, propagating a 2048 x 2048 unpolarized beam through normal refractive surfaces will require about 288 Mb. ZEMAX currently supports a minimum array size of 32 x 32 and a maximum array size of 8192 x 8192. ZEMAX requires the sampling to be an integral power of 2; for example, 32, 64, 128, 256, etc.

If not enough RAM is available, the local hard disk drive may be used for temporary storage. If the option "Use Disk Storage To Save Memory" is checked, ZEMAX will store the surface transfer array (at 56 bytes per grid point) in a temporary file. Although this saves a substantial amount of RAM, the analysis is much slower. Typically, using the disk storage option will slow the analysis by a factor of 10 to 20. If the computer does not have sufficient RAM, it may automatically "page" to the hard disk. This will allow the computation to proceed, but at a rate hundreds or even thousands of times slower than if sufficient RAM were available.

32-bit versions of Windows do not allow ZEMAX to allocate more than 3.0 Gb of RAM, and in many cases, the actual limit is much lower depending upon how much actual RAM the computer has. This limit is imposed by Windows, and not ZEMAX. Both ZEMAX and Windows are available as 64-bit versions, and these versions allow use of memory well beyond 3.0 Gb. For complex POP calculations, use of 64-bit Windows and ZEMAX is strongly recommended.

Defining the initial beam

The menu option Analysis, Physical Optics, Physical Optics Propagation invokes the Physical Optics Propagation feature. The settings box for this feature allows user selection of many options, including those needed to define the initial beam, sampling, surface ranges, and field positions. See "Physical Optics Propagation" on page 225 for a discussion of the available options and settings.

The X- and Y-Sampling determine the number of points used to sample the beam. Larger values are generally more accurate at the expense of longer computation times and larger RAM requirements.

The X- and Y-Width are measured in lens units. The larger the width, the more empty space around the portion of the beam with non-zero intensity there will be. This space is called the *guard band*, and it is vital that sufficient empty space exist around the beam. The empty space allows room for the beam to expand as aberrations are introduced to the beam. If portions of the beam get too close to the edge of the array, they will "alias" and reflect back into the beam, decreasing the accuracy of the computed results.

The initial beam may be any of these types: Gaussian Waist, Gaussian Angle, Gaussian Size+Angle, Top Hat, File, DLL, or Multimode. See the following sections for a detailed discussion of each beam type.

The beam may then be aligned along the chief ray for any defined field position, in the optical space preceding any surface. The starting location for the beam may also be offset from the starting surface by the "Surf To Beam" setting (see "Physical Optics Propagation" on page 225). Propagation then proceeds to the ending surface. Definitions of beam waist, Rayleigh range, and divergence angle may be found in "Propagating the embedded beam" on page 223.

Gaussian Waist

The Gaussian Waist beam is an optionally truncated and decentered Hermite-Gaussian of arbitrary order defined by:

$$E(x, y) = E_0 T(x, y) G_l \left[\frac{\sqrt{2}(x - dx)}{\omega_{0x}} \right] G_m \left[\frac{\sqrt{2}(y - dy)}{\omega_{0y}} \right], \text{ where}$$

$$G_i(u) = H_i(u) \exp\left(\frac{-u^2}{2}\right),$$

and $H_i(u)$ is the Hermite polynomial function of order i . Note the order may be specified in the x and y directions independently. The x order is defined by the integer l , and the y order is defined by the integer m . If both l and m are zero, the simple "Gaussian" TEM(0,0) beam is generated. Higher order modes may be generated by modifying the order values; for example, to generate TEM(1,2) set $l = 1$ and $m = 2$. For a discussion of Hermite-Gaussian beams see Saleh, B. E. A., and Teich, M. C., *Fundamentals of Photonics*, John Wiley & Sons, New

York (1991). If the order number is set higher than 30, the order zero will be used to prevent excessively long computation time.

The dx and dy values are used to decenter the beam. The transmittance function $T(x, y)$ is used to (optionally) truncate the beam to a finite aperture. The transmittance function is defined as

$$T(x, y) = 1, \text{ if } \left(\frac{x-dx}{A_x}\right)^2 + \left(\frac{y-dy}{A_y}\right)^2 \leq 1, \text{ and}$$

$$T(x, y) = 0 \text{ otherwise.}$$

A_x and A_y are the truncating aperture values. If A_x or A_y are zero, no truncating aperture is used. A smoothing function is used near the edge of the truncating aperture to reduce pixel related errors. The smoothing function weights the pixel amplitude by the area of the pixel inside the truncating aperture. The truncating aperture is useful for modeling receiver fiber modes, where the truncating aperture is typically 15% greater than the core size.

The value for E_0 is chosen to yield the peak irradiance in power per unit area or the total beam power as defined in the settings box.

The beam is defined at the waist and will generally diverge to a larger beam size as it propagates away from the waist. See Gaussian Angle and Gaussian Size+Angle below.

Gaussian Angle

The Gaussian Angle beam is an optionally decentered TEM(0, 0) mode whose waist size is determined by the far-field divergence half-angle in degrees, measured in air. ZEMAX uses the specified divergence angle to compute the beam waist. The beam waist is given by:

$$\omega_0 = \frac{\lambda}{\pi \tan \theta}.$$

Note that the angle is the half-angle of divergence in degrees, and the wavelength is not scaled by the index of the starting medium. If the x and y angle values are different, then an elliptical beam will be generated. The beam is defined at the waist.

Gaussian Size+Angle

The Gaussian Size+Angle beam is an optionally decentered TEM(0, 0) mode defined by the beam size (not waist) and the far-field divergence half-angle in degrees, measured in air. ZEMAX uses the specified beam size and divergence angle to compute the beam waist, position, and phase. The beam waist is given by:

$$\omega_0 = \frac{\lambda}{\pi \tan \theta}.$$

Note that the angle is the half-angle of divergence in degrees, and the wavelength is not scaled by the index of the starting medium. If the resulting beam waist is larger than the specified beam size (a physical impossibility), the waist value will be used for the beam size. The z position of the beam relative to the waist is then computed using:

$$z = z_r \sqrt{\left(\frac{\omega}{\omega_0}\right)^2 - 1},$$

where z_r is the Rayleigh range. If the x and y direction values for the various data are different, then an elliptical beam with a possibly toroidal phase will be generated.

Top Hat

The top hat beam is an optionally decentered uniform amplitude beam defined as

$$E(x, y) = E_0 \text{ if } \left(\frac{x-dx}{w_x}\right)^2 + \left(\frac{y-dy}{w_y}\right)^2 \leq 1, \text{ and}$$

$$E(x, y) = 0 \text{ if } \left(\frac{x-dx}{w_x}\right)^2 + \left(\frac{y-dy}{w_y}\right)^2 > 1,$$

where E_0 is chosen to yield the peak irradiance in power per unit area or the total beam power as defined in the settings box.

File

A beam may be defined in a user defined table properly formatted into a file. The table of values must be placed in either a binary or text format file and read from disk. The file must end in the extension ZBF (for ZEMAX Beam File). The binary format is identical to the format written out by ZEMAX if the "Save Output Beam To:" option is selected. The Ex and Ey values are defined such that the Ex*Ex + Ey*Ey is in units of watts. If the units flag indicates the beam units are different from the current lens units, the beam is automatically scaled to the current lens units when read into ZEMAX.



All ZBF Beam Files must be placed in the <pop> folder (see "Folders" on page 68).

Beams in ZEMAX are always centered on the chief ray for the selected field and wavelength. Therefore, the data in the beam file should be positioned relative to the chief ray that will be used to align the beam. The center point in the beam file is at the coordinate (nx/2+1, ny/2+1). ZEMAX requires the values of nx and ny to be an integral power of 2; for example, 32, 64, 128, 256, etc. The minimum sampling is 32 and the current maximum sampling is 8192. Fiber coupling data is ignored when reading beams, and will be zero if fiber coupling is not computed on output. Note the total fiber coupling is the product of the receiver and system efficiency. The first data point is at the -x, -y corner, and the data proceeds across the x rows first. The Rayleigh distance is ignored on input and is automatically recomputed by ZEMAX. The wavelength value stored in the ZBF file is scaled by the index of the media the beam is currently in.

ZEMAX Beam File (ZBF) binary format

The ZBF binary file format is defined as follows. All integers are 4 bytes, all doubles are 8 bytes.

1 integer: The format version number, currently 1.
1 integer: The number of x samples (nx).
1 integer: The number of y samples (ny).
1 integer: The "is polarized" flag; 0 for unpolarized, 1 for polarized.
1 integer: Units, 0 for mm, 1 for cm, 2 for in, 3 for meters.
4 integers: Currently unused, may be any value.
1 double: The x direction spacing between points.
1 double: The y direction spacing between points.
1 double: The z position relative to the pilot beam waist, x direction.
1 double: The Rayleigh distance for the pilot beam, x direction.
1 double: The waist in lens units of the pilot beam, x direction.
1 double: The z position relative to the pilot beam waist, y direction.
1 double: The Rayleigh distance for the pilot beam, y direction.
1 double: The waist in lens units of the pilot beam, y direction.
1 double: The wavelength in lens units of the beam in the current medium.
1 double: The index of refraction in the current medium.
1 double: The receiver efficiency. Zero if fiber coupling is not computed.
1 double: The system efficiency. Zero if fiber coupling is not computed.
8 doubles: Currently unused, may be any value.
2*nx*ny double: Ex values.
If polarized, 2*nx*ny Ey values follow the Ex values.

ZEMAX Beam File (ZBF) text format

The ZBF text file format is defined as follows. The first line must be the single character "A", followed by the other data values specified.

A: indicates a text file.
version: The format version number, currently 1.
nx: The number of x samples.
ny: The number of y samples.
ispol: The "is polarized" flag; 0 for unpolarized, 1 for polarized.
units: 0 for mm, 1 for cm, 2 for in, 3 for meters.
unused 1: Currently unused, may be any value.
unused 2: Currently unused, may be any value.
unused 3: Currently unused, may be any value.
unused 4: Currently unused, may be any value.
dx: The x direction spacing between points.
dy: The y direction spacing between points.
zx: The z position relative to the pilot beam waist, x direction.
Rx: The Rayleigh distance for the pilot beam, x direction.
wx: The waist in lens units of the pilot beam, x direction.
zy: The z position relative to the pilot beam waist, y direction.
Ry: The Rayleigh distance for the pilot beam, y direction.
wy: The waist in lens units of the pilot beam, y direction.
lambda: The wavelength in lens units of the beam in the current medium.
index: The index of refraction in the current medium.
re: The receiver efficiency. Zero if fiber coupling is not computed.
se: The system efficiency. Zero if fiber coupling is not computed.
unused 1: Currently unused, may be any value.
unused 2: Currently unused, may be any value.
unused 3: Currently unused, may be any value.
unused 4: Currently unused, may be any value.
unused 5: Currently unused, may be any value.
unused 6: Currently unused, may be any value.
unused 7: Currently unused, may be any value.
unused 8: Currently unused, may be any value.
Ex real value for point 1
Ex imaginary value for point 1
Ex real value for point 2
Ex imaginary value for point 2
etc... for 2*n_x*n_y Ex values.
If polarized, followed by 2*n_x*n_y Ey values.

DLL

To define a beam using an external program, the algorithm which computes the initial complex electric field from user defined parameter data must be written and compiled into a Windows Dynamic Link Library, or DLL. Sample DLLs are provided with ZEMAX with source code. Studying an existing DLL is the best way to see how the DLL is used to define the beam amplitude and phase. New DLLs may be easily created with a suitable compiler.

Beam DLL parameters

Each DLL may use between zero and 8 user defined data values as parameters in the computation of the beam properties. These values are defined by the DLL and are only used by the DLL. The values may be entered or edited directly on the settings box when the DLL is selected as the beam type.

Creating a new Beam DLL

The DLL must include two functions:

UserBeamDefinition

UserParamNames

When initializing a beam using a DLL, ZEMAX passes to the UserBeamDefinition function the source parameters, the wavelength, and other data. The UserBeamDefinition function then is required to compute the electric field values. These values are returned to ZEMAX and are used to populate the beam array. The DLL

also returns values for the pilot beam, including the waist size, position, and Rayleigh range in the x and y directions. If all of these pilot beam values are zero, ZEMAX will automatically fit the pilot beam parameters to the electric field data.

The function UserParamNames is used to define the names of all used parameters. These names appear on the settings dialog box.

The best way to learn the use of Beam DLLs is to copy and study an existing DLL. The sample DLLs provided with ZEMAX include extensive documentation and comments on the data format; see any of the sample source code files for examples.



All Beam DLLs must be placed in the <program>\DLL\PhysicalOptics folder.

Multimode

A multimode beam consists of a sum of any number of other beams. The sum may be coherent or incoherent, and several options exist for scaling, randomizing, and altering the phase of the individual components of the beam. The multimode beam sum is defined in a user defined text file. The file must end in the extension ZMM (for ZEMAX Multi Mode).



All ZMM Files must be placed in the <pop> folder (see "Folders" on page 68).

The ZMM file uses a simple command syntax that defines the multimode beam. The "master" beam is initially set to zero amplitude over the entire beam. The individual component beams are generated in the sequence defined in the ZMM file. Each beam as defined in the ZMM file is then summed either coherently or incoherently with the master beam. The sum mode, either coherent or incoherent, is controlled by the COHERENT and INCOHERENT commands. The default sum mode is COHERENT. All of the available ZMM commands are described below.

!. Comment

Any data following the ! symbol is ignored. Comment lines may be used by starting the line with a ! symbol. Blank lines are also allowed and are ignored.

COHERENT

Sets the beam summing mode to coherent. The syntax is:

COHERENT

The coherent sum of any two beams is computed by summing the point-by-point real and imaginary parts of the two beams. The phase of the resulting beam thus depends upon the vector sum of the components. All subsequent beam sums will be coherent unless an INCOHERENT command is used. See also INCOHERENT. The default sum mode is COHERENT.

DLL

The DLL command defines a beam based upon a user-defined DLL, which is then summed to the master beam. The syntax is:

DLL weight DLL_NAME param1 param2 param3 ...

Note that the data values are separated by spaces. The weight value defines the dimensionless relative total power of this beam as compared to other beams in the same ZMM file. The weight value may be negative, see "Using random values" on page 643. The DLL_NAME is the name of the DLL file (with no extension). The DLL file must be placed in the correct folder, see "DLL" on page 641 for details. The other values are exactly as defined for the DLL parameters. No scaling of the parameters is supported, since the dimensionality and purpose of the DLL parameters cannot be known in advance.

GW

The GW command defines a Gaussian Waist beam which is then summed to the master beam. The syntax is:

GW weight waistx waisty decenterx decentery aperturx aperturey orderx ordery

Note that the data values are separated by spaces. The weight value defines the dimensionless relative total power of this beam as compared to other beams in the same ZMM file. The weight value may be negative, see “Using random values” on page 643. The other values are exactly as defined for the Gaussian Waist beam, see “Gaussian Waist” on page 638 for a description. The transverse data values, such as the waist and aperture sizes, may be scaled from the POP Beam Definition tab; see “Using the scale factor” on page 643 for details.

INCOHERENT

Sets the beam summing mode to incoherent. The syntax is:

```
INCOHERENT
```

The incoherent sum of any two beams is computed by summing the point-by-point intensity of the two beams, taking the square root to determine the amplitude, and setting the real part of the combined beam to this amplitude, and the imaginary part to zero. All phase information, including that of aberrations in the beam, is lost when performing incoherent sums. Therefore care should be taken when using this command. All subsequent beam sums will be incoherent unless a COHERENT command is used. See also COHERENT. The default sum mode is COHERENT.

PHASE

The PHASE command multiplies any subsequently defined beams by a complex factor determined by the defined phase angle. The syntax is:

```
PHASE angle
```

or

```
PHASE RANDOM
```

The angle is in units of degrees and may be any value. If the angle is set to the keyword RANDOM, the phase is randomly selected to be between -180 and 180 degrees. Note the random angle generated depends upon the random seed value used, see “Using random values” on page 643 for details. All subsequently defined beams are multiplied by the resulting complex factor before being added to the master beam.

TH

The TH command defines a Top Hat beam. The syntax is:

```
TH weight waistx waisty decenterx decentery
```

Note that the data values are separated by spaces. The weight value defines the dimensionless relative total power of this beam as compared to other beams in the same ZMM file. The weight value may be negative, see “Using random values” on page 643. The other values are exactly as defined for the Top Hat beam, see “Top Hat” on page 640 for a description. The transverse data values, such as the waist and aperture sizes, may be scaled from the POP Beam Definition tab; see “Using the scale factor” on page 643 for details.

Using random values

The weight value used by some of the beam commands may be negative, which means a random weight will be assigned to the beam. The random weight will be between 0.0 and the absolute value of the weight. For example, if the weight is -2.0, then a random value will be assigned to the weight between 0.0 and 2.0.

The PHASE RANDOM command uses the random number generator to compute a random phase.

The random values used will always be the same given the same initial "seed" value. The seed value for the random number generator is a user definable value and is Parameter 1 for the multimode beam type. Use a different seed value to generate different, random, but reproducible beams.

Using the scale factor

It is convenient to be able to scale a beam without changing the fundamental data as defined in the ZMM file. Parameter 2 of the Multimode beam type as defined on the POP Settings Beam Definition tab is called "scale". This value, if not zero, will be used to scale all transverse data for all beams in the file.

Surface specific settings

Each ZEMAX surface supports these settings relevant to Physical Optics Propagation:

Use Rays To Propagate To Next Surface: If checked, then the diffraction propagation algorithms will not be used to propagate the beam to the next surface. Instead, rays will be traced to the next surface and the

resulting ray transfer function will be used to propagate the beam and update the pilot beam parameters. Use of this feature is required for non-sequential, birefringent, and gradient index surfaces. Multiple sequential surfaces may all have this option checked, in which case the beam will advance through all the surfaces using only ray propagation. This feature may also be used to speed up the propagation algorithm by propagating through several surfaces at once, if the distance between the surfaces is small enough so that the diffraction propagation is unimportant. The default is unchecked. See "Considerations when using rays to propagate" on page 645 for important information about using this feature near surfaces where the rays come to a focus.

Re-Compute Pilot Beam Parameters: Some surfaces significantly alter the beam characteristics. A good example is a pinhole aperture. After passing the aperture, the beam waist size, divergence, and position may be significantly altered, even though technically the pinhole has no optical power (this is in fact the core difference between geometric and physical optics). After passing the pinhole, the pilot beam parameters need to be recomputed for subsequent accurate propagation. Checking this option will invoke an algorithm which finds the pilot beam parameters that best fit the actual beam.

Do Not Rescale Beam Size Using Ray Data: By default, ZEMAX will use the ray grid to determine any distortion, stretching, scaling or other change in the beam shape. A good example of this is the passing of the beam through a grating. The beam will become compressed along the direction of the diffraction. However, there are times when this computation fails. One such case is when the ray grid enters a caustic. The rays no longer accurately represent the beam, and should not be used to determine the beam shape. ZEMAX automatically skips this step if it detects the ray grid is in a caustic, however, there may be other cases where the algorithm should not be applied, and this option allows user selective disabling of the algorithm.

Use Angular Spectrum Propagator: If selected, the angular spectrum propagator will be used for propagation through the surface instead of the propagator ZEMAX would automatically select. This option should only be used if the beam size does not change dramatically over the propagation distance. The array width will remain constant using the angular spectrum propagator. The default is unchecked.

Draw "beamfile name" on shaded model: If selected, the ZBF of the displayed name will be drawn on shaded model layout plots at the location of the surface. The file name is the same as the name of the lens file, with a 4 digit suffix equal to the surface number. This is the same name convention that automatically saved beam files use. See "Display Tab:" on page 227 for information on saving beam files at every surface. If this option is selected, and the correct beam file name is found, then the surface itself will not be drawn, just the beam file at that surface. For this reason, this option works best on dummy surfaces. Drawing large beam files uses considerable amounts of memory and will slow down the drawing of shaded model displays.

Resample after refraction: This control allows the beam sampling and width to be changed after propagating through any surface. If selected, the new x and y sampling and beam width may be specified. Note that this control is ignored if rays are used to propagate through the surface. See also "Auto Resample" below.

Auto Resample: If checked, the beam will be automatically resampled. The algorithm used first recomputes the pilot beam parameters as described in "Re-Compute Pilot Beam Parameters" above. The X and Y-width is then set according to the equations described in "Comments about point spacing and sampling" on page 636. The number of points used in the X- and Y- sampling will not change.

Output Pilot Radius: The output pilot beam is used to reference the phase and other properties of the beam after refraction through the surface. This control defines the method of computation for the phase radius of curvature of the output pilot beam. "Best Fit" selects the optimum pilot beam based upon ray tracing through the surface and is the recommended setting. "Shorter" selects the shorter of the pilot X or Y phase radii, "Longer" selects the longer of the two, and "X" and "Y" makes the choice explicitly. The "Plane" option will select a pilot beam with a waist at the surface and infinite phase radius. The "User" option allows specification of the X and Y radius in lens units. Use the value of zero for infinite radius. If "Separate X, Y" is selected on the beam definition tab, then the best fit radii in the X and Y directions are automatically selected, unless "Plane" or "User" is chosen. If the radius is set to "User" and "Separate X, Y" is not selected; the X radius value is used for both X and Y directions.

Considerations when using rays to propagate

When using "Use Rays To Propagate To Next Surface", the POP algorithm must first convert the wavefront to a set of rays that represents the beam. The rays are then traced through all consecutive surfaces that have this option checked on. When the rays reach the first surface with this option checked off, the rays must then be converted back to a beam. If the rays form a focus on this surface, the conversion back to a beam cannot be done accurately, because the rays do not adequately represent the diffracted beam near focus.

The solution is to use an additional (usually dummy) surface to propagate the rays some distance away from focus (at least one Rayleigh range), and then end the "Use Rays" section of the propagation at the dummy surface. The rays can then be accurately converted to a wavefront and the propagation can proceed, even if a virtual propagation back to the focus is necessary.

Computing Fiber Coupling

The physical optics propagation algorithm may be used to compute fiber coupling efficiency. A ray based method is also supported, see "Fiber Coupling Efficiency" on page 206.

The fiber coupling receiver efficiency is defined as a normalized overlap integral between the fiber and beam complex amplitude:

$$T = \frac{\left| \iint F_r(x, y) W'(x, y) dx dy \right|^2}{\iint F_r(x, y) F_r'(x, y) dx dy \iint W(x, y) W'(x, y) dx dy},$$

where $F_r(x, y)$ is the function describing the receiving fiber complex amplitude, $W(x, y)$ is the function describing the complex amplitude of the beam coupling into the fiber, and the ' symbol represents complex conjugate. Note that these functions are all complex valued, so this is a coherent overlap integral. Maximum receiver efficiency ($T = 1.0$) is achieved when the mode of the beam perfectly matches the mode of the fiber in both amplitude and phase at all points. Any deviation in mode shape, or phase, will reduce the value of T to less than 1.0. Optical aberrations typically introduce phase deviations which reduce receiver efficiency.

Additional system losses may be caused by apertures which vignette the beam, losses due to reflection from air-glass boundaries, or bulk absorption. These loss factors reduce the system efficiency. The total coupling efficiency into the fiber is the system efficiency multiplied by the receiver efficiency.

Where the integral is computed

The fiber coupling receiver efficiency integral defined above is valid and may be computed anywhere in the optical system; including at the location of the fiber, or at some location near or far from the fiber. The critical requirement is that the fiber mode function describes the mode of the fiber at the point the integral is computed, which is always the "end" surface (see "Physical Optics" on page 222 for a list of analysis options, including specification of the end surface and the fiber mode). If the end surface represents the beam at the location of the receiving fiber, then the fiber mode parameters must correspond to the mode at the receiving fiber. If the end surface is 10 mm away from the fiber, then the fiber mode parameters must be selected to define the mode 10 mm away from the fiber.

Defining the fiber mode

The fiber mode may be a Gaussian or Top Hat function, or may be defined by a DLL or a data file. This allows very general and arbitrary fiber modes to be described, including multi-mode, aberrated, or arbitrary amplitude and phase fibers. The fiber mode may be defined using all the same options allowed for defining initial beams, see "Defining the initial beam" on page 638.

If the beam is not polarized, then the fiber mode is not polarized and only the X-direction of the E field of the fiber mode is used in the overlap integral. If the beam is polarized, then the fiber mode may be either polarized or unpolarized. If the setting "Ignore Polarization" on the Fiber Data Tab (see "Fiber Data Tab:" on page 228) is checked, then the fiber mode is unpolarized, and the X-direction E field is used to compute the coupling for both

the X- and Y-direction fields in the polarized beam. If "Ignore Polarization" is unchecked, then the coupling for the X- and Y-direction is done independently and then combined to yield the overall coupling.

Decenters and tilts

It is usually valuable to determine fiber coupling for decentered and tilted fibers.

Decenter X/Y values are in lens units and may be defined on the fiber data tab of the Physical Optics Propagation settings box ("Physical Optics Propagation" on page 225). These values decenter the fiber mode. The decenter values shift the center of the mode within the array defining the fiber mode at the end surface. For this reason, the decenter values should be small enough to keep the majority of the mode energy inside the array width/height at the end surface.

The tilt values are in degrees. Tilt of the fiber is modeled as adding a linear phase shift to the fiber mode. The phase shift is proportional to the fiber tilts and the mode width/height.

Note that if the end surface is at the receiving fiber location (typically near focus) then a tilt of the fiber introduces phase tilt. If the end surface is far from the fiber, then tilt of the fiber introduces a decenter of the mode. Therefore, great care must be used in selecting the tilt and decenter values; proper use of these values depends upon knowledge of where the overlap integral is being computed (at the end surface!) relative to the fiber location.

Choosing the location for the receiving fiber

The receiving fiber mode may be either centered on the incoming beam chief ray or at the end surface vertex. In this latter case, the fiber mode is decentered by the amount equal to the incoming chief ray's x and y coordinates on the end surface, and the fiber is tilted to be aligned with the local Z axis rather than the chief ray.

Quantitative beam analysis

There are many different ways to analyze a beam. ZEMAX can compute total power, peak irradiance, encircled energy, centroid locations, beam width, M-squared values, and various RMS values. Because of the large number of data that can be computed, some of the data values are only available via the POPD optimization operand; see "POPD" on page 515. The expressions in the following sections are explicit in the x direction coordinate, but similar expressions for y of course exist.

Beam coordinates and pilot beam properties

Most basic beam data is displayed on the POP feature output window, and is also available from the POPD operand.

Peak Irradiance and total power

The peak irradiance is the maximum power per unit area at any point in the beam. Irradiance is measured in Source Units per lens unit squared. See "Units" on page 104. The total power is the integral of the irradiance over the entire beam.

Centroid locations

The beam centroid is given by:

$$c_x = \frac{\int x I(x, y) dx dy}{\int I(x, y) dx dy},$$

where $I(x, y)$ is the irradiance of the beam.

Beam width and M-squared

The following method and theoretical basis for defining beam width and M-squared is from A. E. Siegman, and is applicable to all beams, Gaussian or not, coherent and incoherent.

The second moment of an arbitrary beam is defined as:

$$\sigma_x^2 = \frac{\int (x - c_x)^2 I(x, y) dx dy}{\int I(x, y) dx dy},$$

where c_x is the centroid coordinate as defined earlier. The beam radius is σ_x . The beam width is then defined as $W_x = 2\sigma_x$. Note the beam width will change as the beam propagates, and is thus a function of the beam z coordinate relative to the waist. The smallest value for the beam size is $W_x(0)$. The beam width at other locations is given by

$$W_x^2(z) = W_x^2(0) + M_x^4 \left(\frac{\lambda z}{\pi W_x(0)} \right)^2,$$

where M_x is a parameter characteristic of the beam. This expression allows definition of the desired M-squared factor as

$$M_x^2 \cong \left(\frac{\pi}{\lambda z} \right) W_x(0) W_x(z), \text{ for } z \rightarrow \infty.$$

Wavefront error and RMS beam deviations

The mean, RMS, or peak-to-valley (PTV) irradiance or phase variation may be computed for any beam by using the appropriate settings for the POPD operand.

Encircled energy

Encircled energy is the fraction of the total power or energy contained with a circle of a specified radius r , centered on the chief ray reference point, the beam centroid, or a surface vertex reference. This fraction is always between zero and 1 as is defined as:


$$f = \frac{\int_0^r I(r) dr}{\int_0^\infty I(r) dr}.$$

A variation on encircled energy is to compute the radial coordinate at which a specified fraction is achieved.

Suggestions for use

It frequently takes some experimentation with the sampling, width, and surface specific settings (see above) to get good, accurate results from this feature. Keep in mind the following suggestions:

Use an appropriate width for the initial beam. The best initial width to use is almost always that computed by the "Auto" button next to the width control, and the use of this control is highly recommended. Note that if the width is too large, then there will not be enough points across the beam to yield adequate sampling, if the width is too small, aliasing will occur. If the beam waist is smaller than the wavelength at any surface, including the initial beam, the POP results are probably not accurate, see "Algorithm assumptions" on page 648.

 **Always check the initial beam by setting the end surface to the start surface in the analysis window, and observe that the beam is properly sized for the selected array width.**

Use adequate sampling. As the guard beam increases in size relative to the non-zero amplitude portion of the beam, the number of points that have non-zero data will decrease. It may be required to increase the sampling to be sure enough points sample the beam.

If there are any cylindrical or toroidal optics in the system, or if the beam itself is astigmatic or anamorphic, try using "Separate X, Y" on the beam definition tab. Separating the X and Y propagation greatly reduces the sampling required for accurate results because a reference bispherical surface is used instead of a reference sphere, and the bisphere is capable of more closely fitting anamorphic beams.

If the results seem unbelievable, debug the propagation by propagating one surface at a time through the optical system, and study the results for plausibility. Typically this process will highlight at which surface the algorithm fails. It is frequently the case that certain surfaces, such as gradient index optics or highly tilted surfaces, are best handled by ray tracing. For these surfaces, choose the "Use Rays To Propagate To Next Surface" option.

Imagine a large diameter beam with a wavelength of 1.0 micrometers (0.001 mm), an array width of 64 mm, and a sampling of 64 x 64 points. The delta between points is 1.0 mm. Now imagine the beam going through a collimating lens with a focal length of 100.0 mm. The sampling spacing near focus is given by:

$$\Delta x_2 = \frac{\lambda |z_2|}{n_x \Delta x_1},$$

which yields a new sampling spacing of 0.0015625 and a total width of 0.1 mm. To decrease the sample spacing in the focal plane while keeping the array width constant, the product $n_x \Delta x_1$ must increase. Since the sample spacing is the array width divided by the number of points, *both* the number of points and the initial array width would have to be doubled. If the aberrations are so large that the beam size is on the order of the array width in the focal plane, then the array width must increase. Note the array width in the focal plane is given by:

$$n_x \Delta x_2 = \frac{\lambda |z_2|}{\Delta x_1}.$$

Increasing the array width can be accomplished by increasing the number of points while leaving the initial array width fixed.

 **Note that the beam sampling and width may be manually adjusted at any surface.**

See "Surface specific settings" on page 643 for information on resizing the beam at any surface.

Algorithm assumptions

There are numerous approximations made in the model and algorithms used for physical optics modeling. Some of these assumptions will limit the accuracy of the results in certain cases. The major assumptions are described below.

Both the angular spectrum and the Fresnel diffraction propagation algorithms are developed assuming the beam is not too fast. Neither theory accurately predicts the correct diffraction results if the F/# of the beam is too low. The errors in the propagation theory will reveal themselves as discontinuities in the data as the propagation algorithm crosses over the Rayleigh range boundary. The loss of accuracy cannot be precisely quantified (if ZEMAX could precisely quantify the error, we could apply that as a correction to have a more accurate theory). ZEMAX will report a warning in the propagation report if the Gaussian pilot beam waist becomes smaller than the wavelength. This is a good rule to use as a limit of the accuracy of POP in general. Gaussian beams whose waist is small compared to the wavelength, such as fibers, act essentially as point

sources radiating into a sphere. In this domain, the ray model is perfectly acceptable. In other words, **don't use POP for very fast beams!** The big differences between ray based and POP based results show up in very slow diffracting beams, not fast beams, so use ray based computations in ZEMAX to model very fast beams with confidence. Specifically, use the ray-based fiber coupling algorithm rather than POP for fast, accurate results in systems where the beam is fast and the diffraction effects from the edges of the lenses are not significant; this applies to virtually all fiber coupling systems.

Scalar diffraction theory applies. The vector nature of the electric field is ignored. For very fast beams, scalar diffraction theory is not valid and the results should be used with caution. Note that propagation of fast beams is generally well handled by geometric ray tracing, except near focus. Although there is no exact cut-off to the accuracy of scalar diffraction theory, beams faster than F/1 are probably not accurately modeled with scalar theory. This limitation is another reason not to use POP for very fast beams, as the previous paragraph explains.

Diffraction effects of surface structures are ignored. Gratings, binary optics, and diamond turned surfaces have a complex micro surface structure that can significantly effect the local electric field. This may or may not affect the propagation of the beam. Note that ray tracing, which also ignores the surface structure, has proved to be very useful and accurate for modeling these devices. Any surface adequately modeled by ray tracing will also work fine with physical optics propagation, since ray tracing is used to compute the surface transfer function.

The POP algorithm generally separates the electric field into two beam local orthogonal directions. The pilot beam properties, which are used to reference the phase of the electric field, are also computed along these two orthogonal directions. If the beam is incident upon an optical surface where the plane of incidence is not aligned with these projections, and if the optical surface has diffractive power that is also not aligned with these projections, the POP algorithm cannot accurately predict the correct electric field after diffraction, and the predicted beam will appear to rotate slightly. The error is usually only noticeable for beams that are incident at large skew angles to very steeply tilted diffraction gratings that have significant diffractive bending of the beam (that is, a diffractive order of 0 would describe a significantly different ray path than the diffractive order being used would). These types of systems are well modeled by conventional ray tracing and POP should not be used if the beam rotation is apparent.

When performing polarized beam analysis, the input and fiber mode beam (if any) must contain polarized electric field data. If a ZBF file or DLL is used to define the input or fiber mode beam, and no polarized data is contained in the file or created by the DLL, and a polarized beam analysis is requested, ZEMAX will create a polarized beam from the unpolarized data. The method used is to set the two orthogonally polarized electric fields equal in amplitude and phase, normalizing to conserve the total power in the beam. This allows polarized beam analysis, including transmission through coatings, interfaces, and bulk media.

Samples

Many sample files are included with ZEMAX to demonstrate some of the applications and proper use of the physical optics modeling features. The files are located in the <data>\Samples\Physical Optics folder. Some of these files are described below. Because of the large amount of computation involved in physical optics propagation, some of these files take some time to compute and display the results.

Free space propagation

The sample file "Basic Propagation" illustrates the case of a beam propagating through free space. The beam is a Gaussian defined to have a waist of 0.1 at surface 1. The wavelength is chosen to be π micrometers so that the Rayleigh range is 10.0 mm. At a distance of 1 Rayleigh range, the beam expands in size by $\sqrt{2}$, and the peak irradiance drops to 0.5. At a distance of 2 Rayleigh ranges, the peak irradiance drops to 0.2, and at 3 Rayleigh ranges the irradiance decreases to 0.1.

Note that beam may be virtually propagated backward using the normal ZEMAX sign convention of a negative thickness. Choosing the "End Surface" to be any of the surfaces in the LDE will show the beam intensity at that surface. Note the phase of the beam along the axis has the correct Gouy shift.

A pinhole aperture

The sample file "Pinhole Aperture" illustrates spatial filtering. The first lens forms an aberrated image, which can be seen by selecting the end surface to be 5. Surface 6 has a small circular aperture, which only allows the central portion of the beam to pass. The second lens recollimates the beam. The spatially filtered beam can be seen at the image of the stop on surface 10. There are several interesting things to note:

The transmitted irradiance drops to 0.12. The smaller the pinhole, the lower the transmitted irradiance.

Because the pinhole aperture changes the properties of the beam, the pilot beam parameters must be recomputed. Note the option to do so is selected on surface 6.

A lens array

The sample file "Lenslet Array" uses the user defined surface DLL "US_ARRAY.DLL" to create a 7 x 7 array of rectangular lenses. The focal distance is 100.0 mm, and the pilot beam will focus down to a waist size of about 10 micrometers. However, the lenslet array forms multiple focal spots. To see all of the spots requires high sampling so that all the spots can be seen without aliasing. Notice the decrease in intensity and the aberrations apparent in the spots towards the outside. It is also possible to see the rectangular structure in the diffracted image spots due to the rectangular apertures of the lenslets.

Talbot imaging

Talbot imaging refers to the property of a beam forming an image of itself as it propagates. The image reveals itself in both phase and amplitude modulations as the beam propagates. The sample file "Talbot Imaging" shows a user defined aperture on surface 1 that consists of 20 rectangular slits, like a grating. A propagation of 20 mm reveals the phase reversed Talbot image of the grating. An additional 20.0 mm shows the restored image of the grating. The image is not perfect because of the finite extent of the grating. Note the region near the center of the grating has a well formed image.

Fresnel lens

The sample file "Fresnel Zone Plate Lens" illustrates the focusing power of an aperture consisting of concentric rings. The spacing of the rings is selected to block light from every other Fresnel zone, so that at a distance of 200.0 mm the light diffracted from all the zones interferes constructively. The constructive interference creates a bright focused spot on axis; an effect not predicted by ray tracing. This file also illustrates the intensity of the beam at several planes between the Fresnel zone plate and focus; note the concentric ripples caused by interference from the edges of the zone plate.

Since ZEMAX does not have a multiple-zoned circular aperture, the zone plate was created by placing one annular obscuration on each of several surfaces; all located at the same plane. The option to "Use Rays To Propagate To Next Surface" is used to speed the computation through all the aperture surfaces.

Introduction

The ZEMAX Programming Language (ZPL) is a macro language specifically designed for use with ZEMAX. ZPL offers the power of user-extensibility. This means that if you need a particular calculation or graphical display which is not "built in", you can write your own ZPL macro to do the job.

ZPL is similar to the BASIC programming language, except not all BASIC constructs and keywords are supported, and capabilities and functions unique to ray tracing have been added. ZPL is easy to use, and this chapter will give you instructions and examples to get you started. For an alternate programming capability see "ZEMAX EXTENSIONS" on page 749.

ZPL is a powerful programming language. While it is easy to use, the user is responsible for error-checking, debugging code logic and for good programming practice. For this reason, technical support on macro writing is restricted to ensuring that all ZPL functions and keywords work as documented: we cannot advise on how to perform detailed calculations as part of technical support. If you need a ZEMAX macro, and do not possess the desire or ability to write it yourself, please feel free to contact ZEMAX Technical Support for a quote on developing a custom program to meet your requirements. We have considerable experience in developing these types of programs, and can generally write macros at very competitive rates on short notice.

Creating ZPL macros

To create a ZPL macro, it is probably easiest to start with an existing macro that performs a task similar to the one you want to achieve. If you are attempting to write your first ZPL macro, you may want to read the example sections at the end of this chapter. Some example ZPL macros can also be found in the <data>\Macros folder installed with the ZEMAX program.

Use any text editor to create the ZPL file (such as the NOTEPAD editor). The file may have any name but must end in the .ZPL extension. The file must be placed in the ZPL Folder, which by default is <data>\Macros. To change this folder see "Folders" on page 68.

There is a limit to the allowed complexity of any single line in a ZPL macro. If the "line too long" error occurs, try breaking the line into several smaller lines.

Running ZPL macros

To run a ZPL macro from the main menu, select Macros, Edit/Run ZPL Macros. The ZPL control dialog box will appear with the following options:

Active File: A drop-down list of macros available. All the listed macros are text files ending in the extension .ZPL. The files must be in the folder for ZPL Macros; see "Folders" on page 68. Macros may also be placed in any subfolder within the macros folder.

Close After Execution: If checked, the ZPL control dialog box will automatically close after the macro execution.

Quiet Mode: If checked, the default output text window will not be shown. This is useful for graphics macros that do not generate useful text.

Check Obsolete Syntax: If checked, ZEMAX will test the macro for use of obsolete syntax.

Status: During execution of the macro, ZEMAX will use this area to print a status message stating the line number of the macro being executed. The status message is updated every quarter second.

Terminate: The terminate button will stop execution of the macro currently running.

Cancel: The cancel button terminates the current macro if one is running. If no macro is running, cancel closes the ZPL control dialog box.

Edit: The edit button invokes the Windows NOTEPAD editor. The editor can be used to modify or rename the macro.

View: The view button will display the contents of the macro file in a text window which can be scrolled or printed. No editing is allowed in the view window.

Select the macro to run from the "Active File" list, and then click on Execute.

ZEMAX will begin running the macro. The entire macro will run, and any text output from PRINT commands or error messages will be displayed in a window. The display of this output window can be suppressed using the CLOSEWINDOW keyword.

Assigning ZPL macros to buttons

Frequently used macros may be assigned to buttons for convenient access. See “Buttons 1-16, Buttons 17-32, Buttons 33-48” on page 73 for details.

An overview of ZPL

A ZPL macro consists of a series of commands which are stored in a text file. The commands are either assignments, keywords, or comments. Assignments may be for either numeric or string (textual) data. Both assignments and keywords may accept expressions as arguments, although the syntax is slightly different as described below.

Assignments

The general syntax for an assignment is

```
variable = (expression)
```

The (expression) may consist of an explicit value, such as 5, or a variable name containing some preassigned value, or a complex mathematical expression involving functions, constants, and variables. In all cases, the expression on the right side of the equal sign is evaluated, and the result is assigned to the variable designated on the left.

The simplest form of an assignment is where the expression is a fixed value, such as:

```
x = 5
```

There are several important things to note in this command. First, no declaration of variables is required. This means "x", which is called a variable, did not need to exist before the value of 5 was assigned to it. If "x" had already been assigned a value, it would now be reassigned. Second, no special symbol is required to terminate a command, such as ";" in C. Because of this, each ZPL command must be on a line by itself.

Here are some examples of assignments with expressions:

```
x = SQRT(5)
y = SINE(x)
z = SQRT(x+5*(7-x))
```

The functions SQRT (square root) and SINE (sine) are built in to ZPL. There are many of these functions, all of which are defined in “Numeric functions” on page 656. Note that ZPL is not case-sensitive; SQRT() and sqrt() are the same function. This documentation will use the convention of capital letters for functions and keywords, and lower case for everything else.

There are also string assignments as described in “String variables” on page 655.

Keywords

The general syntax for a keyword is

```
KEYWORD argument1, argument2, argument3...
```

Some keywords have no arguments, others have many. Arguments may be either numeric expressions or string constants or string variables. Some keywords accept a mixture of numeric and string arguments.

One example of a keyword is PRINT. The PRINT keyword is followed by a list of items, separated by commas, to be printed. For example, the ZPL commands

```
x = 3
y = 4
z = SQRT(x * x + y * y)
PRINT "The hypotenuse is ",z
```

will print the following:

```
The hypotenuse is 5.0000
```

Note that ZPL enforces operator precedence. ZPL uses the following precedence from highest to lowest: Parentheses, functions (such as SQRT), logical operators (such as ==), multiplication and division, and then addition and subtraction.

There are many keywords, all are described in “Keywords” on page 670.

Comments

There are 3 ways to add comments to a ZPL macro: by starting a line with the REM keyword, by starting a line with the "!" symbol, or by placing the "#" symbol anywhere on the line as long as the # symbol is not inside a string. Blank lines are also allowed anywhere in the macro. Here are 3 examples of comments:

```
REM this is a remark
! This is also a remark
x = 5 # The # symbol allows comments on the same line as a valid command
```

Comments make macros easier to understand and modify, and have no affect on performance.

Creating graphics

There are a number of low-level functions for generating graphics, including GRAPHICS, GTEXT, GLEN- SNAME, and others. For details see “GRAPHICS” on page 688. An easier method of generating typical ZEMAX style line graphs and 2D graphs graphics is available using the PLOT and PLOT2D keywords; see “PLOT” on page 706 and “PLOT2D” on page 707.

Numeric variables

Variables provide temporary storage for numerical quantities whose exact value may not be known when the macro is written, but is defined when the macro is run. ZEMAX performs most of the work for you when you need a new variable. For example, the command

```
x = 5
```

will cause ZEMAX to allocate memory for the new variable "x" and keep track of the value associated with it. Once the variable is defined, it may be used in any subsequent expression. There are a few rules regarding ZPL variables:

Variable names must not contain any "special" characters that ZPL uses for logical operations or delimiting such as (,), =, +, -, *, /, !, >, <, ^, &, |, #, ", and the space character.

A variable cannot take on the same name as a keyword or function, such as THIC or RAYX. Since ZPL is not case-sensitive, you cannot use rayX or Thic to avoid this rule.

Each variable name is limited to 28 characters.

All ZPL numeric variables are stored internally as 64-bit double precision numbers.

Array variables

Array variables are single- or multi-dimensioned arrays of double precision or integer values. Unlike (scalar) numeric variables, array variables must be declared prior to their use. The declaration syntax is

```
DECLARE name, type, num_dimensions, dimension1 [, dimension 2 [, dimension 3 [, dimension 4] etc...]]
```

The name may be any legal variable name as described in the previous section. The type must be either DOUBLE or INTEGER; this value indicates the type of array variable. The integer value num_dimensions defines the number of dimensions of the array (not the size), and must be between 1 and 4, inclusive. The integers dimension1, dimension2, etc., define the size of the array in that dimension. Note that array variables start at index 1, and thus an array of size 10 has valid indices from 1 to 10.

Array variables may be defined anywhere inside the macro, they need not be declared at the beginning of the macro. To release the memory associated with an array variable, use the RELEASE keyword. The syntax is

```
RELEASE name
```

The RELEASE keyword is optional, as the memory associated with the declared variable is automatically released when the macro terminates. However the RELEASE keyword is useful for conserving memory if large arrays are only needed during a portion of the macro execution.

Array variables are assigned values using the following syntax:

```
name ( index1, index2, ... ) = value
```

The values stored in the array may be retrieved with the same basic syntax:

```
value = name ( index1, index2, ...)
```

The following sample code declares a two-dimensional array variable, assigns a value to each element, prints the values, and then releases the memory for the array:

```
DECLARE Z, DOUBLE, 2, 5, 5
FOR i, 1, 5, 1
  FOR j, 1, 5, 1
    Z(i, j) = i + j
  NEXT j
NEXT i

FORMAT 8.0
k = 0
FOR i, 1, 5, 1
  FOR j, 1, 5, 1
    PRINT k, i, j, Z(i,j)
    k = k + 1
  NEXT j
NEXT i

RELEASE Z
```

Numeric operations

ZPL macros support basic numeric operations such as add, subtract, multiply, and divide. The syntax for each is shown below.

```
x = y + z
x = y - z
x = y * z
x = y / z
```

All other operations are supported only through the use of numeric functions or numeric logical operators, both described in subsequent sections.

Numeric logical operators

Logical operators are used to construct complex commands which ultimately evaluate to one or zero. Most logical operations take the form (left_expression) (operator) (right_expression), similar to mathematical expressions such as 1 + 2. The exception is the not operator "!" which takes only a single argument, of the form ! (right_expression). The logical operators use the convention that zero is "false" and any non-zero value is "true". The not operator returns 1 (true) if the (right_expression) is 0 (false) and returns 0 (false) if (right_expression) is non-zero (true). One common use of the not operator is in IF commands such as:

```
IF !x THEN PRINT "x is zero."
```

The other logical operators can be also be used as part of the argument in IF commands. For example, an IF command may contain two conditions which must both be true for the THEN command to be executed:

```
IF ( x > 1 ) & ( y < 2 ) THEN PRINT "Both conditions are true."
```

These two conditions are related by an "and" expression denoted by &. Note the parentheses are used to force precedence. ZPL supports the logical operators described in the following table.

NUMERIC LOGICAL OPERATORS

Logical	Description
&	And, returns 1 only if both expressions are non-zero.
	Or, returns 1 if at least one expression is non-zero.
^	Xor, returns 1 if only one expression is non-zero.

Logical	Description
!	Not, returns 0 if (right_expression) is non-zero, else returns 1.
==	Equality, returns 1 if expressions are equal.
>	Greater than, returns 1 if left_expression is greater than right_expression.
<	Less than, returns 1 if left_expression is less than right_expression.
>=	Greater than or equal to, returns 1 if left_expression is greater than or equal to right_expression.
<=	Less than or equal to, returns 1 if left_expression is less than or equal to right_expression.
!=	Inequality, returns 1 if expressions are unequal.

String variables

ZPL supports string variables and basic string operations. String variables can hold a maximum of 260 characters. String variables do not need to be declared, but can be created at any time using a defining assignment command such as:

```
newstring$ = "Here is the new string"
```

Note that string variables are distinguished from numeric variables by the presence of the \$ character at the end of the string.

There are also string functions which can be used to extract text data, such as

```
title$ = $LENSNAME()
```

Note the function \$LENTITLE() starts with the \$ character. This identifies the function as returning a string result.

String operations

String variables can be concatenated using the + operator. The syntax is:

```
C$ = A$ + B$
```

Concatenation can also include constant strings:

```
total$ = "A$ is " + A$ + " and B$ is " + B$
```

The string functions can be used in a defining command such as

```
this$ = "Here is the lens title: " + $LENSNAME()
```

String variables are printed just like other strings:

```
PRINT "Here is A$: ", A$
```

Note that the PRINT function can only print single string variables; there is no support for the concatenation operand or string functions inside print commands. The correct procedure is to concatenate the strings into a new string and then print the new string:

```
A$ = B$ + C$
PRINT A$
```

Alternatively, the comma acts as a concatenation operand:

```
PRINT A$, B$, C$
```

String functions cannot be printed directly like this:

```
PRINT $LENSNAME() ! NOT CORRECT !!!
```

Instead, the correct procedure is to assign the function result to a variable and then print the variable:

```
Z$ = $LENSNAME()
PRINT Z$
```

String logical operators

String logical operators are very similar to the numeric logical operators defined in a previous section. The key difference is that the expressions being compared are strings rather than numbers. The supported string logical operators are defined in the following table.

STRING LOGICAL OPERATORS

Logical	Description
\$==\$	Equality, returns 1 if left_string and right_string are identical.
\$>	Greater than, returns 1 if left_string is greater than right_string.
\$<	Less than, returns 1 if left_string is less than right_string.
\$>=	Greater than or equal to, returns 1 if left_string is greater than or identical to right_string.
\$<=	Less than or equal to, returns 1 if left_string is less than or identical to right_string.
\$!=	Inequality, returns 1 if left_string and right_string are not identical.

For example, an IF command may compare strings as follows:

```
A$ = "TEST"  
B$ = "TEST"  
IF (A$ $==$ B$) THEN PRINT "Strings are identical."
```

Numeric functions

Numeric functions can be used on the right hand side of a numeric variable assignment, and in expressions which are arguments to keywords. These functions may require no arguments, one argument, or multiple arguments. All functions return a single value. Certain functions, such as PWAV() (primary wavelength), return a value which does not depend upon the argument, and therefore it is not required to provide one. The parentheses however, are still required.

In the following table, all ZPL functions are listed. If the syntax is given as FUNC(), then no arguments are required. FUNC(x) indicates one argument is required, FUNC(x,y) indicates two arguments, etc.

ZPL FUNCTION LISTING

Function	Argument	Return value
ABSO(x)	Numeric expression	The absolute value of the expression.
ACOS(x)	Numeric expression	Arc cosine in radians.
APMN(x)	Valid surface number	The minimum radius value. For spider apertures, this is the width of the arms. For rectangular and elliptical apertures, it is the x-half width of the aperture.
APMX(x)	Valid surface number	The maximum radius value. For spider apertures, this is the number of arms. For rectangular and elliptical apertures, it is the y-half width of the aperture.
APOI(px, py)	Normalized pupil coordinates	Ray intensity apodization factor.
APXD(x)	Valid surface number	The aperture x-decenter value.
APYD(x)	Valid surface number	The aperture y-decenter value.
APTP(x)	Valid surface number	An integer code describing the aperture type on the specified surface.

Function	Argument	Return value
ASIN(x)	Numeric expression	Arc sine in radians.
ASPR()	(null)	Aspect ratio of the current graphics device.
ATAN(x)	Numeric expression	Arc tangent in radians.
ATYP()	(null)	System aperture type code: 0 for EPD, 1 for F/#, 2 for NA, and 3 for float by stop size.
AVAL()	(null)	System aperture value.
CALD(i)	index	Returns a numeric value from the CALLMACRO buffer at the specified index. See "Calling a Macro from within a Macro" on page 747.
CONF()	(null)	Returns the current configuration number, which is between 1 and the number of configurations, inclusive.
CONI(x)	Valid surface number	Conic constant of the surface.
COSI(x)	Numeric expression in radians	Cosine of the expression.
CURV(x)	Valid surface number	Curvature of the surface.
EDGE(x)	Valid surface number	Edge thickness at the semi-diameter of that surface.
EOFF()	(null)	End of file flag. Returns 1 if the end of file has been reached, otherwise returns 0. Only valid after execution of READ keyword.
ETIM()	(null)	The elapsed time in seconds since last TIMER.
EXPE(x)	Numeric expression	e to the power of the expression.
EXPT(x)	Numeric expression	10 to the power of the expression.
FICL(vec#)	Vector #, between 1 and 4 inclusive.	The fiber coupling efficiency. See "Using the FICL() function" on page 668
FLDX(x)	Valid field number	X-angle or height of specified field.
FLDY(x)	Valid field number	Y-angle or height of specified field.
FTYP()	(null)	Returns 0 if the current field type is angles in degrees, 1 if object height, 2 if paraxial image height, and 3 if real image height. Heights are always in lens units (see UNIT).
FVAN(x)	Valid field number	Vignetting angle of specified field.
FVCX(x), FVCY(x)	Valid field number	Vignetting compression x or y of specified field.
FVDX(x), FVDY(x)	Valid field number	Vignetting decenter x or y of specified field.
FWGT(x)	Valid field number	Field weighting of specified field.
GABB(x)	Valid surface number	The glass catalog Abbe number of the glass for the specified surface.

Function	Argument	Return value
GAUS(x)	Standard deviation	Returns a random value with a Gaussian distribution, zero mean, and the specified standard deviation.
GETT(window, line, column)	The window number, line number, and column number of the numerical value to extract. Each Columns are delimited by spaces.	Extracts a numerical value from any open text window. This feature allows computations using any value ZEMAX can display in a window.
GIND(x)	Valid surface number	The glass catalog d-light index of the glass for the specified surface.
GLCA(x)	Valid surface number	Global vertex x vector of the specified surface.
GLCB(x)	Valid surface number	Global vertex y vector of the specified surface.
GLCC(x)	Valid surface number	Global vertex z vector of the specified surface.
GLCM(surf, item)	Surf must be a valid surface number, item is an integer between 1 and 12.	For item equal to 1-9, the return value is R11, R12, R13, R21, R22, R23, R31, R32, or R33. For item equal to 10-12, the return value is the x, y, or z component of the global offset vector. See "Global Coordinate Reference Surface" on page 111.
GLCX(x)	Valid surface number	Global vertex x-coordinate of the specified surface.
GLCY(x)	Valid surface number	Global vertex y-coordinate of the specified surface.
GLCZ(x)	Valid surface number	Global vertex z-coordinate of the specified surface.
GNUM(A\$)	Any string variable name	If the string A\$ is the name of a valid glass, such as BK7, then GNUM returns the number of the glass in the glass catalog. The glass number can subsequently be used by SETSURFACEPROPERTY to set the glass type on a surface. If A\$ does not correspond to any glass in the catalog, GNUM returns 0. GNUM returns -1 if the string is "MIRROR".
GPARG(x)	Valid surface number	The glass catalog partial dispersion coefficient of the glass for the specified surface.
GRIN(s, w, x, y, z)	Surface #, wavelength #, x, y, and z coordinates	Returns the index of refraction at the specified x, y, z position on surface s at wavelength number w. Works for gradient and non-gradient media.
GTEM(code)	For checkbox values, code is 0 for Use Glass Substitution Template, 1 for Exclude Glasses With Incomplete Data, 11 for Standard, 12 for Preferred, 13 for Obsolete, 14 for Special, and 21-26 for Use Relative Cost, CR, FR, SR, AR, and PR, respectively. For numerical values, code is 31-36 for Relative Cost, CR, FR, SR, AR, and PR, respectively.	Returns data from the Glass Substitution Template. Checkboxes return zero for false (unchecked) and any non-zero value for true (checked).

Function	Argument	Return value
IMAE(seed)	The seed value	Returns the geometric image analysis efficiency. See the discussion "Optimizing with the IMAE operand" on page 532. If seed is zero, the same random numbers will be used on every call to IMAE. If seed is not zero, then every call to IMAE will use unique random numbers.
INDX(surface)	Valid surface number	Index of refraction at the primary wavelength. See ISMS.
INTE(x)	Numeric expression	Returns the largest integer not greater than the argument.
ISMS(surface)	Valid surface number	If the surface is an odd mirror, or follows an odd mirror but is not a mirror, the return value will be one, otherwise the return value is 0.
LOGE(x)	Positive numeric expression	Log base e of the expression.
LOGT(x)	Positive numeric expression	Log base ten of the expression.
LOST(code)	Code is 1 to return lost energy due to errors, or 2 for lost energy due to thresholds.	The lost energy following the most recent NSTR trace.
MAGN(x,y)	x and y are any real numbers	Computes the square root of x squared plus y squared.
MAXF()	(null)	The maximum radial field angle in degrees, or radial object or image height in lens units. Which value is returned depends upon if the fields are defined by angle, object height, or image height.
MAXG()	(null)	The number of glasses currently loaded.

Function	Argument	Return value
MCON(row, config, data)	Row (operand number), Configuration number, and data value to extract from the Multi-Configuration Editor.	<p>Extracts data from any row and configuration of the Multi-Configuration Editor. This function is similar to MCOP with extended capabilities for extracting data.</p> <p>If the row and config number are both zero, MCON returns either the number of operands, the number of configurations, or the active configuration number for data = 0, 1, and 2, respectively.</p> <p>If the row number is between 1 and the number of multi-config operands, and the config number is zero, MCON returns the operand type, integer 1, integer 2, integer 3, and string flag for that specified row, for data = 0 through 4 respectively. The 3 integer values are used for various purposes for different operands, such as surface and wavelength numbers. The string flag is 1 if the operand data is a string value, such as a glass name, or 0 for numerical data.</p> <p>If the row number is between 1 and the number of multi-config operands, and the config number is valid, MCON returns either the numerical value or the string data for that operand.</p> <p>Note that all string data returned by MCON must be extracted with the \$buffer command after the call to MCON. For example, the following code sample will place the name of the operand on row 1 in a\$:</p> <pre>dummy = MCON(1, 0, 0) a\$ = \$buffer()</pre> <p>See also keyword SETMCOPERAND.</p>
MCOP(row, config)	Row (operand number) and configuration of the Multi-Configuration Editor.	Extracts data from any row and configuration of the Multi-Configuration Editor. A configuration number of zero indicates the current configuration. See also keyword SETMCOPERAND.
MFCN()	(null)	MFCN updates the lens, validates the merit function, updates the merit function, then returns the current merit function value. See OPER.
NCON()	(null)	Returns the number of configurations.
NFLD()	(null)	The number of defined fields.
NOBJ(surface)	Valid non-sequential surface number	The number of defined objects in the non-sequential surface.
NPAR(surf, object, param)	surf is the surface number, object is the object number, and param is the parameter number of the value to return.	Returns a value from the parameter columns of the non-sequential components editor. See "SETNSCPARAMETER" on page 720.
NPOS(surf, object, code)	surf is the surface number, object is the object number, and code is 0-6 for reference object, x, y, z, tilt x, tilt y, and tilt z position to return, respectively.	Returns a value from the position columns of the non-sequential components editor. See "SETNSCPOSITION" on page 721.

Function	Argument	Return value
NPRO(surf, object, code, face)	surf is the surface number, object is the object number, code values are as defined for the SETNSCPROPERTY keyword, and face is the face number.	Returns a numeric or string value from the property pages of objects defined in the non-sequential components editor. For information on the code values, see "SETNSCPROPERTY" on page 721. String values may be extracted using the \$buffer() function after calling this function with a code that returns string data. For example, to extract the comment column of object 5 from NSC surface 2 use these two commands: dummy = NPRO(2, 5, 1, 0) a\$ = \$buffer() The value dummy may be ignored. The string variable a\$ will contain the object comment.
NSDC(surf, object, pixel, data)	The surf value indicates the surface number of non-sequential group. The object number indicates the desired detector. Pixel indicates which pixel value is returned. Data is 0 for real, 1 for imaginary, 2 for the amplitude, and 3 for the coherent intensity. See also "NSTR" on page 702.	If the object number corresponds to a detector rectangle object, then the coherent intensity data from the specified pixel is returned. Coherent intensity consists of 4 numbers which are each summed over all incident rays: the real part, the imaginary part, the amplitude, and the coherent intensity. If the pixel number is zero, then the sum of the data for all pixels for that detector object is returned.
NSDD(surf, object, pixel, data)	The surf value indicates the surface number of non-sequential group. The object number indicates the desired detector. See the discussion to the right for details on the pixel and data arguments. See also "NSTR" on page 702.	If the object number is zero, then all detectors are cleared and the function returns zero. If the object number is less than zero, then the detector defined by the absolute value of the object number is cleared and the function returns zero. If the object number corresponds to a detector rectangle, surface, or volume object, then the incoherent intensity data from the specified pixel is returned. For a full discussion of the pixel and data arguments, see "NSDD" on page 509. This ZPL function supports the same pixel and data arguments as the NSDD optimization operand.
NSDE(surf, object, pixel, angle, data)	The surf value indicates the surface number of non-sequential group. The object number indicates the desired detector. See the discussion to the right for details on the pixel, angle, and data arguments. See also "NSTR" on page 702.	If the object number is zero, then all detectors are cleared and the function returns zero. If the object number is less than zero, then the detector defined by the absolute value of the object number is cleared and the function returns zero. If the object number corresponds to a detector color object, then the data from the specified pixel is returned. For a full discussion of the pixel, angle, and data arguments, see "NSDE" on page 510. This ZPL function supports the same pixel, angle, and data arguments as the NSDE optimization operand.

Function	Argument	Return value
NSDP(surf, object, pixel, data)	<p>The surf value indicates the surface number of non-sequential group.</p> <p>The object number indicates the desired detector.</p> <p>See the discussion to the right for details on the pixel and data arguments.</p> <p>See also "NSTR" on page 702.</p>	<p>If the object number is zero, then all detectors are cleared and the function returns zero. If the object number is less than zero, then the detector defined by the absolute value of the object number is cleared and the function returns zero.</p> <p>If the object number corresponds to a detector polar object, then the data from the specified pixel is returned.</p> <p>For a full discussion of the pixel and data arguments, see "NSDP" on page 510. This ZPL function supports the same pixel and data arguments as the NSDP optimization operand.</p>
NSUR()	(null)	The number of defined surfaces.
NWAV()	(null)	The number of defined wavelengths.
OBJC(A\$)	Any string variable name	Object with comment. Returns the first non-sequential object number where the comment matches the string A\$. The comparison is case insensitive. If no NSC object has the matching comment the function returns -1.
OCOD(A\$)	Any string variable or literal that is the name of a ZEMAX optimization operand.	The optimization operand code number used by the OPEV function.
ONUM(A\$)	Any string variable name	If the string A\$ is the name of a valid optimization operand, such as EFFL, then ONUM returns the id number of the operand. The operand id number can subsequently be used by SETOPERAND to set the operand type in the merit function editor. If A\$ does not correspond to any operand name, ONUM returns 0.
OPDC()	(null)	The optical path difference. Valid only after a RAYTRACE call. OPDC will not return valid data if the chief ray cannot be traced. RAYTRACEX does not support OPDC.
OPER(row, col)	Row (operand number) and column (data type) of the Merit Function Editor.	Extracts data from any row and column of the Merit Function Editor. The row is the same as the operand number; the column is 1 for type, 2 for int1, 3 for int2, 4-7 for data1-data4, 8 for target, 9 for weight, 10 for value, 11 for percent contribution, and 12-13 for data5-data6. OPER does not update the lens or the merit function, it returns the current data. See also OPEV, MFCN and keyword SETOPERAND.

Function	Argument	Return value
OPEV(code, int1, int2, data1, data2, data3, data4)	Code is the optimization operand code (see function OCOD), and int1-int2, and data1-data4 are the defining values for the operand. See "Optimization operands" on page 489.	Computes the same value as any optimization operand would, without the need to add the operand to the merit function. This is useful for computing numbers already available from the optimization operands. For example, to compute the EFL from the EFFL operand use this code: C = OCOD("EFFL") E = OPEV(C, 0, 1, 0, 0, 0, 0) See also OPER, OCOD, MFCN and keyword SETOPERAND. Some optimization operands, such as OPGT, only make sense when taken in the context of other operands. OPEV will only work for operands that do not depend upon the presence of prior operands. See also OPEW below.
OPEW(code, int1, int2, data1, data2, data3, data4, data5, data6)	Code is the optimization operand code (see function OCOD), and int1-int2, and data1-data6 are the defining values for the operand. See "Optimization operands" on page 489.	This function is very similar to OPEV, the key difference being that OPEW supports two additional arguments. Some optimization operands use up to 6 data values rather than 4.
OPTH(x)	Valid surface number	The optical path length along the ray to the specified surface. Unlike RAYT and RAYO, OPTH considers the phase added by diffractive surfaces such as gratings, holograms, and binary optics. Valid only after a RAYTRACE call. OPTH will not return valid data if the chief ray cannot be traced. RAYTRACEX does not support OPTH.
PARM(n,s)	Valid parameter number and surface number	Parameter "n" of surface "s".
PIXX(A\$)	A string variable containing the full path to a JPG or BMP graphic file.	The number of X-direction pixels in the graphic.
PIXY(A\$)	A string variable containing the full path to a JPG or BMP graphic file.	The number of Y-direction pixels in the graphic.
PMOD()	(null)	0 if paraxial mode is off, else 1.
POWR(x,y)	x and y are any numbers	Computes the absolute value of x to the power of y.
PVHX()	(null)	Data1 parameter from the ZPLM optimization operand.
PVHY()	(null)	Data2 parameter from the ZPLM optimization operand.
PVPX()	(null)	Data3 parameter from the ZPLM optimization operand.
PVPY()	(null)	Data4 parameter from the ZPLM optimization operand.
PWAV()	(null)	The primary wavelength number.

Function	Argument	Return value
RADI(x)	Valid surface number	Radius of curvature of surface. If the surface has an infinite radius, RADI returns 0.0. This possibility must be considered to avoid potential divide by zero errors.
RAGX(x)	Valid surface number	The global x coordinate of the ray intercept. Valid only after a RAYTRACE or RAYTRACEX call.
RAGY(x)	Valid surface number	The global y coordinate of the ray intercept. Valid only after a RAYTRACE or RAYTRACEX call.
RAGZ(x)	Valid surface number	The global z coordinate of the ray intercept. Valid only after a RAYTRACE or RAYTRACEX call.
RAND(x)	Positive numeric expression	Random floating point number uniformly distributed between 0 and the expression.
RANX(x)	Valid surface number	The X-direction cosine of the surface normal. Valid only after a RAYTRACE or RAYTRACEX call.
RANY(x)	Valid surface number	The Y-direction cosine of the surface normal. Valid only after a RAYTRACE or RAYTRACEX call.
RANZ(x)	Valid surface number	The Z-direction cosine of the surface normal. Valid only after a RAYTRACE or RAYTRACEX call.
RAYE()	(null)	The ray-trace error flag, 0 if no error. Valid only after a RAYTRACE or RAYTRACEX call. See the RAYTRACE keyword for details.
RAYL(x)	Valid surface number	The X-direction cosine of the ray following the surface. Valid only after a RAYTRACE or RAYTRACEX call.
RAYM(x)	Valid surface number	The Y-direction cosine of the ray following the surface. Valid only after a RAYTRACE or RAYTRACEX call.
RAYN(x)	Valid surface number	The Z-direction cosine of the ray following the surface. Valid only after a RAYTRACE or RAYTRACEX call.
RAYO(x)	Valid surface number	The ray optical path length from the previous surface to the specified surface. The optical path length is the path length times the index of refraction, either or both of which may be negative. For rays inside a non-sequential surface, RAYO returns the sum of the path length times the index of refraction of all objects the ray passed through. Valid only after a RAYTRACE or RAYTRACEX call. See also OPTH and RAYT.

Function	Argument	Return value
RAYT(x)	Valid surface number	The ray path length from the previous surface to the specified surface. The path length may be negative. Valid only after a RAYTRACE or RAYTRACEX call. See also OPTH and RAYO. For rays inside a non-sequential surface, RAYT returns the sum of the path lengths along all segments. See also RAYO.
RAYV()	(null)	0 if ray was not vignetted, else vignetted surface number. Valid only after a RAYTRACE or RAYTRACEX call.
RAYX(x)	Valid surface number	The X-coordinate of the ray intercept. Valid only after a RAYTRACE or RAYTRACEX call.
RAYY(x)	Valid surface number	The Y-coordinate of the ray intercept. Valid only after a RAYTRACE or RAYTRACEX call.
RAYZ(x)	Valid surface number	The Z-coordinate of the ray intercept. Valid only after a RAYTRACE or RAYTRACEX call.
RELI(f)	Relative illumination	RI for the specified field position.
SAGG(x,y,z)	x,y are the coordinates in lens units at the surface number z	Computes the sag in lens units at the specified point on the surface.
SCOM(A\$, B\$)	Any two string variable names	If the strings are equal, SCOM returns 0. If A\$ is less than B\$, then SCOM returns a value less than 0; otherwise, a value greater than 0.
SDIA(x)	Valid surface number	Semi-diameter of surface.
SIGN(x)	Numeric expression	Returns -1 if the argument is less than zero, 0 if the argument is zero, and +1 if the argument is positive.
SINE(x)	Numeric expression in radians	Sine of the expression.
SLEN(A\$)	Any string variable name	The number of characters in the string variable A\$
SOLM(operand, config, code)	Operand is the multi-configuration operand number. Config is the configuration number. Code is 0 for the status, 1 for the pickup configuration number, 2 for the pickup operand row, 3 for the pickup scale, 4 for the pickup offset, 5 for the current numerical or string value of the operand, 6 for the ZPL macro solve name, and 7 for the glass catalog name. Not all codes are used by all solve status types.	The return value is data about the solve type and data for the specified multi configuration operand and configuration. If the code is 0 for status, the returned value is 0 for fixed, 1 for variable, 2 for pick-up, 3 for thermal pick-up, 4 for ZPL Macro, and 5 for glass substitution solve. String values may be extracted using the \$buffer() function after calling this function with a code that returns string data. For more information on the solve data and the meaning of the solve parameters, see "SOLVES" on page 473.

Function	Argument	Return value
SOLV(surf, code, param)	Surf is the surface number. Code is 0 for curvature, 1 for thickness, 2 for glass, 3 for conic, 4 for semi-diameter, and 5 for TCE. For parameter data, the code is 100 plus the parameter number. For extra data, the code is 300 plus the extra data number. Param is an integer between 0 and 4, inclusive.	The return value is data about the solve type for the specified surface and data. If param is zero, then an integer corresponding to the solve type is returned. For param between 1 and 3, the data is the solve parameters. For param 4, the data is the pickup column number. For more information on the solve data and the meaning of the solve parameters, see "SOLVES" on page 473.
SOSO(code)	Code is 0 to return the surface number of the macro solve currently executing, or 1 to return the object number.	This function returns the surface number and/or object number of the macro currently being executed as a solve. If the macro is not being used as a solve, these functions return meaningless data.
SPRO(surf, code)	Surf is the surface number. The code values are as defined for the SETSURFACEPROPERTY keyword. Only integer code values, and not mnemonics are supported.	Returns a numeric or string value from the property pages or editors of surfaces defined in the lens data editor. For information on the code values, see "SETSURFACEPROPERTY, SURP" on page 725. String values may be extracted using the \$buffer() function after calling this function with a code that returns string data. For example, to extract the comment column of surface 4 use these two commands: dummy = SPRO(4, 1) a\$ = \$buffer() The value dummy may be ignored. The string variable a\$ will contain the surface comment. Properties where SETSURFACEPROPERTY requires the use of value2 are not supported by SPRO, use SPRX instead. These properties include parameter data, extra data, and coating layer multiplier and index offset values and status.
SPRX(surf, code, value2)	Surf is the surface number. The code and value2 values are as defined for the SETSURFACEPROPERTY keyword. Only integer code values, and not mnemonics are supported.	This function is similar to SPRO. The key difference is that SPRX supports one additional argument. Only parameters set by SETSURFACEPROPERTY that use the "value2" argument should use SPRX. See "SETSURFACEPROPERTY, SURP" on page 725.
SQRT(x)	Positive numeric expression	Square root of the expression.
STDD() - obsolete	This function is obsolete, use SRPO instead.	
SURC(A\$)	Any string variable name	Surface with comment. Returns the first surface number where the comment matches the string A\$. The comparison is case insensitive. If no surface has the matching comment the function returns -1.
SVAL(A\$)	Any string variable name	String value. Returns a floating point value of the string A\$.

Function	Argument	Return value
SYPR(code)	The code values are as defined for the SETSYSTEMPROPERTY keyword. Only integer code values, and not mnemonics are supported.	Returns a numeric or string value for the corresponding system data. For information on the code values, see "SETSYSTEMPROPERTY, SYSP" on page 728. String values may be extracted using the \$buffer() function after calling this function with a code that returns string data. For example, to extract the lens title use these two commands: dummy = SYPR(16) a\$ = \$buffer() The value dummy may be ignored. The string variable a\$ will contain the lens title. Properties where SETSYSTEMPROPERTY requires the use of value2 are not supported by SYPR.
TANG(x)	Numeric expression in radians	Tangent of the expression.
TMAS()	(null)	The total mass in grams of the lens from surface 1 to the image surface.
THIC(x)	Valid surface number	Thickness of the surface.
TOLV(op, col)	The integer op is the tolerance operand number, col is the column number in the Tolerance Data Editor. See discussion.	Returns a numeric or string value from the Tolerance Data Editor. If col is 1, 2, or 3, the return value is the integer 1, 2, or 3 value as used by that operand type. If col is 4 or 5, the min or max value of the tolerance is returned. If col is 6, the "Do Not Adjust During Inverse Tolerancing" flag is returned (1 for checked, 0 for unchecked). If col is 90, the nominal value of the tolerance is returned. If the col value is 0 or 99, the operand type name or operand comment is placed in the buffer string, respectively. For example, to extract the operand name of operand 7 use these two commands: dummy = TOLV(7, 0) a\$ = \$buffer() The value dummy may be ignored. The string variable a\$ will contain the operand name. See also the string functions \$TOLOPERAND and \$TOLCOMMENT which perform a similar function.
UNIT()	(null)	Returns 0, 1, 2, or 3, if the current unit type is millimeters, centimeters, inches, or meters, respectively.
VEC1(x)	Positive subscript value	Returns the value of the array variable at the specified subscript.
VEC2(x)	Positive subscript value	Returns the value of the array variable at the specified subscript.
VEC3(x)	Positive subscript value	Returns the value of the array variable at the specified subscript.
VEC4(x)	Positive subscript value	Returns the value of the array variable at the specified subscript.

Function	Argument	Return value
VERS()	(null)	Returns the version of the ZEMAX program. The version is of the form yymmdd. For example, the version for the June 15, 2005 release of ZEMAX will be 050615.
WAVL(x)	Valid wavelength number	Wavelength in micrometers.
WINL()	(null)	Returns the number of the analysis window most recently opened by the keyword OPENANALYSISWINDOW.
WINN()	(null)	Returns the number of open analysis windows.
WWGT(x)	Valid wavelength number	Wavelength weighting.
XMIN()	(null)	The minimum x-coordinate in the graphics window.
XMAX()	(null)	The maximum x-coordinate in the graphics window.
YMIN()	(null)	The minimum y-coordinate in the graphics window.
YMAX()	(null)	The maximum y-coordinate in the graphics window.

Using the FICL() function

The FICL function computes fiber coupling. Because the function takes so many arguments, the arguments are placed in a vector array and then the function is called. The vector array must be one of the four VEC arrays already defined. The values placed in the vector array are defined as follows:

- 0 = Sampling (1 for 32 x 32, 2 for 64 x 64, 3 for 128 x 128, etc..)
- 1 = Wavelength #
- 2 = Field #
- 3 = Ignore source flag (0 for false, 1 for true)
- 4 = Source NA x-direction
- 5 = Receiver NA x-direction
- 6 = Source x angle (deg)
- 7 = Source y angle (deg)
- 8 = Receiver tilt about x (deg)
- 9 = Receiver tilt about y (deg)
- 10 = Receiver decenter x
- 11 = Receiver decenter y
- 12 = Receiver decenter z
- 13 = Align source to chief ray flag (0 for false, 1 for true)
- 14 = Align receiver to chief ray flag (0 for false, 1 for true)
- 15 = Use polarization flag (0 for false, 1 for true)
- 16 = Source NA y-direction (if zero then NAy = NAx)
- 17 = Receiver NA y-direction (if zero then NAy = NAx)
- 18 = Use Huygens integral (0 for false, 1 for true)

The fiber coupling is computed by calling FICL(n) where n is the vector number containing the argument list.

String functions

The following table lists the available string functions.

ZPL STRING FUNCTIONS

Function	Description
\$BUFFER()	Returns the current string in the lens buffer. This function is used to extract string data from various ZPL keywords and functions.
\$CALLSTR(i)	Returns the string from the CALLMACRO string buffer at index i. See "Calling a Macro from within a Macro" on page 747.
\$COAT(i)	Returns the coating name for the ith surface.
\$COATINGPATH()	Returns the path name for coating files.
\$COMMENT(i)	Returns the comment string for the ith surface.
\$DATAPATH()	Returns the path name for data files.
\$DATE()	Returns the current date and time string. The formatting is specified by the Date/Time control settings described under "Miscellaneous" on page 70.
\$EXTENSIONPATH()	Returns the path name for ZEMAX extensions.
\$FILENAME()	Returns the current lens file name, without the path.
\$FILEPATH()	Returns the current lens file name, with the complete path.
\$GETSTRING(A\$, n)	Returns the nth sub-string for the string A\$ using spaces for delimiters. For example, if A\$ = "one two three", then \$GETSTRING(A\$, 2) returns "two".
\$GETSTRINGC(A\$, n)	Returns the nth sub-string for the string A\$ using commas for delimiters. For example, if A\$ = "one,two,three", then \$GETSTRING(A\$, 2) returns "two".
\$GLASS(i)	Returns the glass name of surface number i.
\$GLASSCATALOG(i)	Returns the name of the ith loaded glass catalog for the current lens. If i is less than 1, then the names of all the loaded catalogs separated by spaces are returned in a single string.
\$GLASSPATH()	Returns the path name for glass catalog files.
\$LEFTSTRING(A\$, n)	Returns the left most n characters in the string A\$. If A\$ has fewer than n characters, the remaining spaces will be padded with blanks. This allows formatting of strings with a fixed length.
\$LENSNAME()	Returns the lens title defined in the General System dialog box.
\$MACROPATH()	Returns the path name for macro files.

Function	Description
\$NOTE(line#)	<p>Returns the notes information defined in the General System dialog box. Because the notes may be very long, \$NOTE returns the characters from the notes in groups called lines. A line ends when a newline (carriage return) character is found, or when the total number of consecutive characters on the line reaches 100, whichever comes first. The line# indicates which line in the notes is to be returned. For example, \$NOTE(1) returns the first line in the notes, up to the first newline character. If there are more than 100 consecutive characters before the first newline is found, then line 1 will be the first 100 characters, and line 2 will be the remainder of the line up to the first newline, or the next 100 characters, whichever comes first.</p> <p>The maximum number of characters in the notes is currently 4000 (because of space limitations, it may not be possible to edit this many characters in the General dialog however). \$NOTE will return a null (empty) string if there are no defined characters in the notes for the specified line.</p> <p>The function SLEN can be used to determine the actual number of characters returned by \$NOTE using this syntax: A\$ = \$NOTE(1) N = SLEN(A\$) The value N will be the integer number of characters in the string A\$.</p>
\$OBJECTPATH()	Returns the path name for NSC object files.
\$PATHNAME()	Returns the path name only for the current lens file. This is useful for determining the folder where the lens file is stored.
\$PROGRAMPATH()	Returns the path name for program files.
\$QUOTE()	Returns the double quote character (").
\$RIGHTSTRING(A\$, n)	Returns the right most n characters in the string A\$. If A\$ has fewer than n characters, the remaining spaces will be padded with blanks. This allows formatting of strings with a fixed length.
\$STR(expression)	Returns a string formatted using the format defined by the FORMAT keyword. The numeric expression may be any equation, including combinations of constants, variables, and functions. See function SVAL(A\$) to convert strings to numbers.
\$TAB()	Returns the tab character (\t).
\$TEMPFILENAME()	Returns the name of a temporary file, with complete path, suitable for temporary storage of text or binary data. See keyword GETTEXTFILE.
\$TOLCOMMENT(operand)	Returns the comment for the specified tolerance operand.
\$TOLOPERAND(operand)	Returns the operand name for the specified tolerance operand.
\$UNITS()	Returns either MM, CM, IN, or M, depending upon the current lens units.

Keywords

Keywords provide the capability to direct program flow, to generate output, and to perform certain crucial tasks, such as ray tracing and modifying the lens prescription. Each keyword is described in detail in the following sections.

APMN, APMX, APTP, APXD, APYD

These commands are obsolete. See "SETSURFACEPROPERTY, SURP" on page 725.

ATYP, AVAL

These commands are obsolete. See "SETSYSTEMPROPERTY, SYSP" on page 728.

BEEP

Purpose:

Makes an audible beep sound.

Syntax:

```
BEEP
```

Discussion:

This command can be used to alert the user when a calculation is finished or input is required.

CALLMACRO

Purpose:

Calls another ZPL macro.

Syntax:

```
CALLMACRO filename.zpl
```

```
CALLMACRO name$
```

Discussion:

This command is used to call another ZPL macro. The filename should be the full macro name without the path. The file must exist in the folder for ZPL macro files. When the macro returns, any text output generated by the called macro will be copied into the calling macro's text output window. See "Calling a Macro from within a Macro" on page 747.

CALLSETDBL

Purpose:

Sets a double precision numeric value in the master macro's buffer.

Syntax:

```
CALLSETDBL index, value
```

Discussion:

This command is used to set a numeric value into the buffer maintained by the master macro. The index can be any value between 0 and 50, inclusive. The numeric value may be retrieved using the CALD function. See "Calling a Macro from within a Macro" on page 747.

CALLSETSTR

Purpose:

Sets a string value in the master macro's buffer.

Syntax:

```
CALLSETSTR index, text$
```

Discussion:

This command is used to set a string value into the buffer maintained by the master macro. The index can be any value between 0 and 50, inclusive. The string value may then be retrieved using the \$CALLSTRING function. See "Calling a Macro from within a Macro" on page 747.

COAT

This command is obsolete. See "SETSURFACEPROPERTY, SURP" on page 725.

CLOSE

Purpose:

Closes a file previously opened by the OPEN command.

Syntax:

```
CLOSE
```

Discussion:

See the description for OPEN.

CLOSEWINDOW

Purpose:

Suppresses the display of the default output window. This keyword can also be used to close any open analysis windows.

Syntax:

CLOSEWINDOW
CLOSEWINDOW n

Discussion:

CLOSEWINDOW (with no argument "n" provided) is used to run the ZPL macro in "quiet" mode. The text window normally displayed at the end of the macro execution will not be displayed if the CLOSEWINDOW keyword is included at any line in the macro. CLOSEWINDOW has no other effect on macro execution.

CLOSEWINDOW (with an integer argument "n" provided) will close analysis window number n.

COLOR

Purpose:

Sets the current pen color to use for graphics line and text functions.

Syntax:

COLOR n

Discussion:

The value n must evaluate to an integer corresponding to the desired color. The new pen color will be used for all subsequent line and text commands in graphics mode. Black is color 0, the other colors are as defined in "Colors 1-12, Colors 13-24" on page 73. There are 24 colors other than black available.

COMMAND

Purpose:

Executes a shell command.

Syntax:

COMMAND executable, arguments

Discussion:

The values executable and arguments are either string literals or string variables. The value of executable is the name (including folder path) to the executable to run. The value of arguments is optional, and includes any command line arguments to the executable.

Once the command is launched, the executable runs in its own thread. This command does not cause the ZPL macro or ZEMAX to "wait" until the executable completes its task. No error codes are returned, and ZEMAX has no way of knowing if the shell command was launched successfully.

COMMENT

This command is obsolete. See "SETSURFACEPROPERTY, SURP" on page 725.

CONI

This command is obsolete. See "SETSURFACEPROPERTY, SURP" on page 725.

CONVERTFILEFORMAT

Purpose:

Converts a text file from ANSI to Unicode format or from Unicode to ANSI format.

Syntax:

CONVERTFILEFORMAT filename, encoding

Discussion:

Filename is the name of the text file to convert. Encoding is 1 to convert the file from Unicode to ANSI, or 2 to convert ANSI to Unicode. Note converting a Unicode file into ANSI may cause the loss of data, if any of the characters have no ANSI equivalent.

COPYFILE

Purpose:

COPYFILE is used to make a copy of a file.

Syntax:

```
COPYFILE sourcefilename, newfilename
```

Discussion:

This keyword requires two file names, defined as literal string expressions in quotes or as string variables. The file sourcefilename is copied into the new file newfilename. If newfilename already exists it is overwritten without warning.

Example:

```
COPYFILE "C:\source.dat", "C:\copy.dat"
```

Related Keywords:

DELETEFILE

RENAMEFILE

CURV

This command is obsolete. See "SETSURFACEPROPERTY, SURP" on page 725.

DECLARE

See "Array variables" on page 653.

DEFAULTMERIT

Purpose:

Generates a default merit function.

Syntax:

```
DEFAULTMERIT type, data, reference, method, rings, arms, grid, delete, axial,  
lateral, start, xweight, oweight
```

Discussion:

This keyword generates a default merit function in the Merit Function Editor. Any existing default merit function will be deleted. For details see "Defining the default merit function" on page 483. The values are as follows:

type: use 0 for RMS, 1, for PTV.

data: use 0 for wavefront, 1 for spot radius, 2 for spot x, 3 for spot y, 4 for spot x + y.

reference: use 0 for centroid, 1 for chief, 2 for unreferenced.

method: use 1 for Gaussian quadrature, 2 for rectangular array.

rings: the number of annular rings (Gaussian quadrature only).

arms: the number of radial arms (Gaussian quadrature only). The number of arms must be even and no less than 6.

grid: the size of the grid. Use an integer, such as 8, for an 8 x 8 grid. n must be even and no less than 4.

delete: use 0 to not delete vignettted rays, 1 to delete vignettted rays.

axial: use -1 for automatic, which will use symmetry only if the system is axial symmetric. Use 1 to assume axial symmetry, 0 to not assume axial symmetry.

lateral: use 1 to ignore lateral color, 0 otherwise.

start: use -1 for automatic, which will add the default merit function after any existing DMFS operand. Otherwise use the operand number at which to add the default merit function. Any existing operands above the specified operand number will be retained.

xweight, oweight: the x direction weigh and overall weight for the merit function. Only the data "spot x + y" uses the xweight value.

DELETE

Purpose:

Deletes a surface from the spreadsheet.

Syntax:

```
DELETE n
```

Discussion:

The value n must evaluate to an integer expression. See also INSERT.

Example:

```
DELETE 5  
DELETE i+2*j
```

DELETECONFIG

Purpose:

DELETECONFIG deletes an existing configuration in the multi-configuration editor.

Syntax:

```
DELETECONFIG config
```

Discussion:

The value config must evaluate to an integer expression greater than 0 and less than or equal to the current number of configurations. See also INSERTCONFIG and DELETEMCO.

DELETEFILE

Purpose:

DELETEFILE is used to delete a file.

Syntax:

```
DELETEFILE filename
```

Discussion:

This keyword requires a file names, defined as literal string expression in quotes or as a string variable.

Example:

```
DELETEFILE XFILE$
```

Related Keywords:

COPYFILE
RENAMEFILE

DELETEMCO

Purpose:

DELETEMCO deletes an existing operand in the multi-configuration editor.

Syntax:

```
DELETEMCO row
```

Discussion:

The value row must evaluate to an integer expression greater than 0 and less than or equal to the current number of operands. See also INSERTMCO and DELETECONFIG.

DELETEMFO

Purpose:

DELETEMFO deletes an existing operand in the merit function editor.

Syntax:

```
DELETEMFO row  
DELETEMFO ALL
```

Discussion:

The value row must evaluate to an integer expression greater than 0 and less than or equal to the current number of operands. If the ALL argument is used all operands are deleted. See also INSERTMFO.

DELETEOBJECT

Purpose:

Deletes an existing NSC object.

Syntax:

```
DELETEOBJECT surf, object
```

Discussion:

The value surf must evaluate to an integer expression, and the specified surface must be a non-sequential component surface. Use a surf value of 1 if the program mode is non-sequential. The value object must evaluate to an integer that is between 1 and the current number of objects plus 1, inclusive. The specified object will be deleted, renumbering other objects as required. See also INSERTOBJECT and SETNSCPROPERTY.

Example:

```
DELETEOBJECT 1, 1
```

DELETETOL

Purpose:

DELETETOL deletes an existing operand in the tolerance data editor.

Syntax:

```
DELETETOL row
```

Discussion:

The value row must evaluate to an integer expression greater than 0 and less than or equal to the current number of operands. See also INSERTTOL and SETTOL.

EDVA

This command is obsolete. See “SETSURFACEPROPERTY, SURP” on page 725.

END

See GOSUB.

EXPORTBMP

Purpose:

Exports any graphic window as a BMP file.

Syntax:

```
EXPORTBMP winnum, filename, delay
```

Discussion:

The winnum value may be either an integer or an expression that evaluates to an integer. The integer winnum corresponds to the graphic window number that should be saved to a file. ZEMAX numbers windows sequentially as they are opened, starting with 1. Any closed windows are deleted from the window list, without renumbering the windows which remain. Any windows opened after another window has been closed will use the lowest window number available. The filename should be the full file name including the path, but with no extension. ZEMAX will automatically add the BMP extension. The optional delay parameter specifies a time delay in

milliseconds. For some complex graphics, a delay is required to allowed the graphic to be completely redrawn and the screen capture to complete. If the JPG files appear incomplete, try a delay value of 500 - 2500 milliseconds. See also EXPORTJPG and EXPORTWMF.

Example:

```
EXPORTBMP 1, "C:\TEMP\MYBMPFILE"  
EXPORTBMP k, A$
```

EXPORTCAD

Purpose:

Exports lens data as an IGES, STEP, SAT, or STL file for import into CAD programs.

Syntax:

```
EXPORTCAD filename
```

Discussion:

For a full description of CAD file export, see "Export IGES/SAT/STEP Solid" on page 257. The filename can be a literal string, such as "C:\DATA\FILE.IGS" or a string variable name, such as MYFILENAME\$. There are many other parameters required to define the export. These other data are placed in the vector 1 variable as follows:

VEC1(1): The file type. Use 0 for IGES, 1 for STEP, 2 for SAT, 3 for STL.

VEC1(2): The number of spline points to use (if required on certain entity types).

VEC1(3): The First surface to export. In NSC Mode, this is the first object to export.

VEC1(4): The last surface to export. In NSC Mode, this is the last object to export.

VEC1(5): The layer to place ray data on.

VEC1(6): The layer to place lens data on.

VEC1(7): Use 1 to export dummy surfaces, otherwise use 0.

VEC1(8): Use 1 to export surfaces as solids, otherwise use 0.

VEC1(9): Ray pattern. Use 0 for XY, 1 for X, 2 for Y, 3 for ring, 4 for list, 5 for none, 6 for grid, and 7 for solid beams.

VEC1(10): The number of rays.

VEC1(11): The wave number. Use 0 for all.

VEC1(12): The field number. Use 0 for all.

VEC1(13): Use 1 to delete vignettted rays, otherwise use 0.

VEC1(14): The dummy surface thickness in lens units.

VEC1(15): Use 1 to split rays from NSC sources, otherwise use 0.

VEC1(16): Use 1 to scatter rays from NSC sources, otherwise use 0.

VEC1(17): Use 1 to use polarization when tracing NSC rays, otherwise use 0. Polarization is automatically selected if splitting is specified.

VEC1(18): Use 0 for the current configuration, 1 - n for a specific configuration where n is the total number of configurations, n+1 to export "All By File", n+2 to export "All By Layer", and n+3 for "All At Once". For a more detailed explanation of the configuration setting, see "Export IGES/SAT/STEP Solid" on page 257.

VEC1(19): Tolerance setting. Use 0 for 1.0E-4, 2 for 1.0E-05, 3 for 1.0E-06, and 4 for 1.0E-07.

VEC1(20), VEC1(21): If the program mode is sequential, and the range of surfaces includes a non-sequential components surface, these values allow a range of objects to be exported. VEC1(20) is the first object to export, and VEC1(21) is the last object to export. If both values are zero or out of range all objects are exported.

Example:

```
VEC1 (1) = 0  
VEC1 (2) = 32  
VEC1 (3) = 1
```

```
etc...
VEC1(18) = 0
EXPORTCAD "C:\TEMP\MYCADFILE.IGS"
```

EXPORTJPG

Purpose:

Exports any graphic window as a JPG file.

Syntax:

```
EXPORTJPG winnum, filename, delay
```

Discussion:

See EXPORTBMP.

Example:

```
EXPORTJPG 3, "C:\TEMP\MYJPGFILE"
EXPORTJPG k, FILENAME$, 500
```

EXPORTWMF

Purpose:

Exports any graphic window as a Windows Metafile. Unlike EXPORTBMP and EXPORTJPG, EXPORTWMF requires the filename to have the desired extension included in the filename argument. The metafile format will be whatever the current selected metafile type is, see "Graphics" on page 69.

Syntax:

```
EXPORTWMF winnum, filename
```

Discussion:

See EXPORTBMP.

Example:

```
EXPORTWMF 3, "C:\TEMP\MYWMFFILE.EMF"
EXPORTWMF k, FILENAME$
```

FINDFILE

Purpose:

Used to find the names of files.

Syntax:

```
FINDFILE TEMPNAME$, FILTER$
```

Discussion:

This keyword requires two expressions, one to specify the string variable name to store the file name in, and another string variable which contains a "filter" string. The filter string usually specifies a path name and wildcards appropriate to the desired file type. See the example below.

FINDFILE is useful for listing all files of a certain type in a folder, or for analyzing large numbers of similar lens files. To reset FINDFILE back to the first file of any type, just call FINDFILE with a different filter, then call FINDFILE again with the original filter name. Each time FINDFILE is called with a new filter, it resets back to the first file that meets the filter specifications.

Example:

```
FILTER$ = "C:\ZEMAX\*.ZMX"
PRINT "Listing of all ZEMAX files in ", FILTER$
FINDFILE TEMPFILE$, FILTER$
LABEL 1
if (SLEN(TEMPFILE$))
    PRINT TEMPFILE$
    FINDFILE TEMPFILE$, FILTER$
    GOTO 1
ENDIF
PRINT "No more files."
```

FLDX, FLDY, FWGT, FVDX, FVDY, FVCX, FVCY, FVAN

These commands are obsolete. See "SETSYSTEMPROPERTY, SYSP" on page 728.

FOR, NEXT

Purpose:

The FOR and NEXT keywords define a program block which is executed a specific number of times in a loop.

Syntax:

```
FOR variable, start_value, stop_value, increment
    commands...
NEXT
```

Discussion:

The keyword FOR marks the beginning of a group of commands to be executed a multiple number of times. FOR requires a variable to be specified which acts as a counter (it need not be an integer), a starting value for the counter, a stop value, and an increment. The NEXT keyword marks the end of the group of commands. FOR-NEXT loops may be nested. The number of FOR and NEXT commands must be the same.

Upon reaching a FOR command, the expressions for the start, stop, and increment values are evaluated and saved. The stop and increment values are not evaluated again, even if the expressions defining the values consist of variables whose values change within the program block. Only the values valid at the beginning of the FOR loop are used.

If the start value and stop value are the same, the loop executes exactly once. If the start value is less than the stop value, then the loop continues until the counter variable is greater than the stop value. If the start value is greater than stop value, then the loop continues until the counter variable is less than the stop value.

Example:

```
FOR i, 1, 25, 1
    PRINT i
NEXT

j = 5
k = 0
FOR i, j, j + 5, 2
    k = i + j + k
NEXT
```

FORMAT

Purpose:

Specifies the numerical precision format for subsequent PRINT and \$STR commands.

Syntax:

```
FORMAT m.n
FORMAT m.n EXP
FORMAT m [INT]
FORMAT "C_format_string" [LIT]
```

Discussion:

The integers m and n are separated by a decimal point. The value m refers to the total number of characters to be printed, even though some of them may be blank. The value n refers to the number of places to display after the decimal point. Therefore, FORMAT 8.4 will cause subsequent PRINT commands to print 8 characters, with 4 numbers after the decimal point. FORMAT .5 will cause PRINT to show 5 decimal places, and as many total places as needed. FORMAT only affects numeric output from PRINT. If a number is too large to fit within the m decimal places, then the m portion of the FORMAT command will be ignored.

The optional keyword EXP after the m.n expression indicates exponential notation should be used.

The optional keyword INT indicates the value should be first converted to an integer and printed in integer format using the number of places specified by m.

The optional keyword LIT (for literal) indicates the value should be printed according to the "C" language format specifier. The C format specification can be found in any programming reference for the C language.

Example:

The macro:

```
X = 5
FORMAT 5.3
PRINT "FORMAT 5.3          :", X
FORMAT 12.2 EXP
PRINT "FORMAT 12.2 EXP    :", X
FORMAT 6 INT
PRINT "FORMAT 6 INT       :", X
FORMAT "%#06i" LIT
PRINT "FORMAT %#06i LIT   :", X
```

will produce this output:

```
FORMAT 5.3          :5.000
FORMAT 12.2 EXP     : 5.00E+000
FORMAT 6 INT        :      5
FORMAT %#06i LIT    :000005
```

ETYP

This command is obsolete. See "SETSYSTEMPROPERTY, SYSP" on page 728.

GCRS

This command is obsolete. See "SETSYSTEMPROPERTY, SYSP" on page 728.

GDATE

Purpose:

GDATE will place the current date under the lens name in the text box on user defined graphics screens.

Syntax:

```
GDATE
```

Discussion:

GDATE is primarily used for making graphics generated by ZPL macros look like other ZEMAX graphics.

Example:

See "GRAPHICS" on page 688.

GETEXTRADATA

Purpose:

Retrieves the extra data values from the extra data editor. The data is placed in one of the vector array variables (either VEC1, VEC2, VEC3, or VEC4).

Syntax:

```
GETEXTRADATA vector_expression, surface_expression
```

Discussion:

The data is stored in the specified VECn array variable. For example, if the command GETEXTRADATA 1, 5 is issued, the extra data for surface 5 will be placed in VEC1. The data is stored in the following format, where the first number in each line refers to the array position:

```
0:      The number of extra data values in the vector
1:      The first extra data value
n:      The nth extra data value
```

See "Extra data" on page 277 for descriptions of the extra data values.

GETGLASSDATA

Purpose:

Retrieves the data for any glass in the current catalogs. The data is placed in one of the vector array variables (either VEC1, VEC2, VEC3, or VEC4).

Syntax:

```
GETGLASSDATA vector_expression, glass_number
```

Discussion:

The data is stored in the specified VECn array variable. For example, if the command GETGLASSDATA 1, 32 is issued, the data for glass number 32 will be placed in VEC1. The glass number is returned by the GNUM function. The data is stored in the following format, where the first number in each line refers to the array position:

```
0:      The number of data values in the vector
1:      Formula number: The number indicates the formula as follows:
      1 = Schott, 2 = Sellmeier 1, 3 = Herzberger, 4 = Sellmeier 2
      5 = Conrady, 6 = Sellmeier 3, 7 = Handbook of Optics 1, 8 = Handbook of Optics 2
      9 = Sellmeier 4, 10 = Extended, 11 = Sellmeier 5, 12 = Extended 2
2:      Reference temperature in degrees c
3:      Nd
4:      Vd
5:      Thermal coefficient of expansion -30 to +70 c
6:      Thermal coefficient of expansion +20 to 300 c
7:      Density in g/cm^3
8:      Deviation from normal line P gf
9:      Lambda min
10:     Lambda max
11-16: Constants of dispersion (meaning depends upon formula)
17-22: Thermal constants of dispersion
23+:   Internal transmission coefficient (per mm) alpha, T = exp(alpha * path).
The alpha for wavelength 1 is stored in 23, wavelength 2 is in 24, etc.
```

Related Functions:

GNUM

GETLSF

Purpose:

Calculates the geometric edge and line response functions, similar to the “Geometric Line/Edge Spread” on page 160.

Syntax:

```
GETLSF wave, field, sampling, vector, maxradius, use_polarization
```

Discussion:

Wave is an integer corresponding to the wavelength number to use for the calculation. A value of zero indicates a polychromatic calculation. Field must be an integer between 1 and the maximum number of fields. The value indicates which field position to use. Sampling may be 1 (32 x 32), 2 (64 x 64), 3 (128 x 128), etc... up to 2048 x 2048. The vector argument must be an integer value between 1 and 4, and specifies which vector array the data should be placed in. The maxradius argument is the maximum radial coordinate of the edge and line spread functions; this is the half-width of the data range. Use 0 for a default width. If any of the arguments fall outside the valid ranges, then the nearest acceptable value is used instead.

The data is returned as an array of values in the specified vector. Vector position 0-3 will hold the number of points "N", the starting x coordinate (this is the negative of the half width of the data range), the delta coordinate, and the offset (defined below), respectively. The offset is the first position in the vector that holds the edge or line spread data. Starting at the offset, the first N value are the tangential LSF response. The next N values are the sagittal LSF response. The tangential and sagittal ERF values are in the next two groups of N data values.

If the current vector size is not large enough, ZEMAX will automatically increase the size of the vectors to hold the LSF data in the manner described in SETVECSIZE.

Example:

```
! Macro computes and prints the LSF and ERF for polychromatic light at field 1.
!
! Syntax is GETLSF wave, field, samp, vector, maxradius, usepol
```



```

!
GETLSF 0, 1, 3, 1, 0, 0
N_BINS = vec1(0)
STARTX = vec1(1)
DELTAX = vec1(2)
OFFSET = vec1(3)
FORMAT 15.0
PRINT "Number of Bins      = ", N_BINS
FORMAT 15.3 EXP
PRINT "Starting Coordinate = ", STARTX
PRINT "Delta Coordinate   = ", DELTAX
FORMAT 15.0
PRINT "Offset              = ", OFFSET
OFF1 = OFFSET
OFF2 = OFF1 + N_BINS
OFF3 = OFF2 + N_BINS
OFF4 = OFF3 + N_BINS
MAXI = N_BINS-1
FORMAT 16.3 EXP
PRINT
PRINT "          X          TLSF          SLSF          TERF          SERF"
PRINT
FOR i, 0, MAXI, 1
    PRINT STARTX + DELTAX*i,
    PRINT vec1(OFF1 + i),
    PRINT vec1(OFF2 + i),
    PRINT vec1(OFF3 + i),
    PRINT vec1(OFF4 + i)
NEXT i

```

GETMTF

Purpose:

Calculates tangential and sagittal MTF, real part, imaginary part, phase, or square wave response data for the currently loaded lens file, and places the data in one of the vector arrays (either VEC1, VEC2, VEC3, or VEC4).

Syntax:

GETMTF freq, wave, field, sampling, vector, type

Discussion:

The freq argument is the desired spatial frequency in MTF Units (see "MTF Units" on page 106). If the frequency is less than zero, or greater than the cutoff frequency, GETMTF returns zero. Wave is an integer corresponding to the wavelength number to use for the calculation. A value of zero indicates a polychromatic calculation. Field must be an integer between 1 and the maximum number of fields. The value indicates which field position to use. Sampling may be 1 (32 x 32), 2 (64 x 64), 3 (128 x 128), etc... up to 2048 x 2048. The vector argument must be an integer value between 1 and 4, and specifies which vector array the data should be placed in. The type argument refers to the data type: 1 for MTF, 2 for real part, 3 for imaginary part, 4 for phase in radians, 5 for square wave MTF. If any of the arguments fall outside the valid ranges, then the nearest acceptable value is used instead. This calculation uses the FFT MTF method (see "FFT MTF" on page 133).

The data is returned in one of the vector arrays with the following format: Vector position 0: tangential response; Vector position 1: sagittal response.

Example:

```

! This macro computes the T & S response at 30 lp/mm
! for the currently loaded lens, polychromatic,
! at the maximum defined field,
! and a 32x32 grid density (sampling = 1).
! Data will be placed in vector 1.
! This is all it takes to get the data:
GETMTF 30, 0, NFLD(), 1, 1, 1
PRINT "Tangential response:", vec1(0)
PRINT "Sagittal response  :", vec1(1)

```

GETPSF

Purpose:

Calculates the diffraction point spread function (PSF) using the FFT algorithm and places the data in one of the vector arrays (either VEC1, VEC2, VEC3, or VEC4).

Syntax:

GETPSF wave, field, sampling, vector, unnormalized, phaseflag, imagedelta

Discussion:

Wave is an integer corresponding to the wavelength number to use for the calculation. A value of zero indicates a polychromatic calculation. Field must be an integer between 1 and the maximum number of fields. The value indicates which field position to use. Sampling may be 1 (32 x 32), 2 (64 x 64), 3 (128 x 128), etc... up to 2048 x 2048. The vector argument must be an integer value between 1 and 4, and specifies which vector array the data should be placed in. The unnormalized flag is zero if the data should be normalized to a peak of 1.0, if the unnormalized value is 1, then the data is returned unnormalized. If phase flag is zero, the data returned is intensity, if 1, then the phase in degrees is returned. The imagedelta value is the spacing between PSF points in micrometers; use zero for the default spacing. The wavelength must be monochromatic to compute phase data. If any of the arguments fall outside the valid ranges, then the nearest acceptable value is used instead.

The data is returned in one of the vector arrays with the following format:

Vector position 0: the total number of PSF data points in the vector array. Usually, this number will be $4*n*n$ where n is the sampling size (32, 64, etc.). For example, if the sampling density is 2, the pupil sampling will be 64 x 64, and there will be 128 x 128 or 16,384 values in the array. This will require 8 bytes per number, or a total of 131 kb. A sampling density of 1024 will require at least 8 Mb just for the array; another 64 Mb or more to compute the PSF. Position 0 also returns other values as error codes. If position 0 is zero, then the computation was aborted. If -1, then the vector array is not large enough to hold all the data. Use SETVECSIZE to make the array bigger. If -2, then there is not enough system RAM to compute the PSF data. If -3, a general error occurred while computing the PSF.

Vector position 1 through $4*n*n$ holds the PSF data intensity. The first $2n$ values are the first row, going left to right from - x to $+x$, then each subsequent block of $2n$ values is another row, going from - y to $+y$. Vector position $4*n*n+1$ holds the spacing between data values in micrometers.

Example:

```
! This macro computes the PSF
! for the currently loaded lens, polychromatic,
! at the first field,
! and a 32x32 grid density (sampling = 1),
! data will be placed in vector 1,
! normalized to 1,
! no phase data,
! default image delta.
```

```
SETVECSIZE 4500
GETPSF 0, 1, 1, 1, 0, 0, 0
np = vec1(0)
IF (np == 0)
    PRINT "PSF Computation aborted."
    GOTO 1
ENDIF
IF (np == -1)
    PRINT "SETVECSIZE too small for PSF data."
    GOTO 1
ENDIF
IF (np == -2)
    PRINT "Not enough system RAM for PSF data."
    GOTO 1
ENDIF
PRINT "There are ", np, " data points, spaced ", vec1(np+1), " micrometers apart".
LABEL 1
```

GETSYSTEMDATA

Purpose:

Retrieves most system specific data, such as effective focal length, working F/#, apodization factors, and other data not associated with any particular surface. The data is placed in one of the vector array variables (either VEC1, VEC2, VEC3, or VEC4).

Syntax:

```
GETSYSTEMDATA vector_expression
```

Discussion:

The data is stored in the specified VECn array variable. For example, if the command GETSYSTEMDATA 1 is issued, the system data will be placed in VEC1. See also SETSYSTEMPROPERTY.

The data is stored in the following format, where the first number in each line refers to the array position:

```
0:      The number of system data values in the vector
1:      Aperture Value
2:      Apodization Factor
3:      Apodization Type (0:none, 1:gaussian, 2:tangent)
4:      Adjust Index Data To Environment setting (1 if true, 0 if false)
5:      Temperature in degrees c (valid only if Use Env Data true)
6:      Pressure in ATM (valid only if Use Env Data true)
7:      Effective Focal Length
8:      Image Space F/#
9:      Object Space Numerical Aperture
10:     Working F/#
11:     Entrance Pupil Diameter
12:     Entrance Pupil Position
13:     Exit Pupil Diameter
14:     Exit Pupil Position
15:     Paraxial Image Height
16:     Paraxial Magnification
17:     Angular Magnification
18:     Total Track
19:     Ray Aiming (0 for off, 1 for paraxial, and 2 for real)
20:     X Pupil Shift
21:     Y Pupil Shift
22:     Z Pupil Shift
23:     Stop Surface Number
24:     Global Coordinate Reference Surface Number
25:     Telecentric object space (0 for off, 1 for on)
26:     The number of configurations
27:     The number of multi-configuration operands
28:     The number of merit function operands
29:     The number of tolerance operands
30:     Afocal image space (0 for off, 1 for on)
31:     X Pupil Compress
32:     Y Pupil Compress
```

GETTEXTFILE

Purpose:

Creates a text file from any ZEMAX analysis window that supports text.

Syntax:

```
GETTEXTFILE textfilename, type, settingsfilename, flag
```

Discussion:

The textfilename must be in quotes, or be a string variable name, and include the full path, name, and extension for the file to be created. The string function \$TEMPFILENAME can be used to define a suitable temporary file name.

The type argument is a 3 character string code that indicates the type of analysis to be performed. The string codes are identical to those used for the button bar in ZEMAX. A list of string codes may be found on the "Buttons" tab of the File, Preferences dialog box. If no type is provided or recognized, a standard ray trace will be generated.

If a valid file name in quotes or a string variable name is used for the settingsfilename argument, ZEMAX will use or save the settings used to compute the text file, depending upon the value of the flag parameter.

If the flag value is 0, then the default settings will be used. If the lens file has it's own default settings, then those will be used; these are the settings stored in the "lensfilename.cfg" file. If no lens specific default settings exist, then the default settings for all ZEMAX files, stored in the file "ZEMAX.CFG" will be used, if any. If no previous settings have been saved for this or any other lens, then the default settings used are the "factory" defaults used by ZEMAX.

If the flag value is 1, then the settings provided in the settings file, if valid, will be used to generate the file. If the data in the settings file is in anyway invalid, then the default settings will be used to generate the file. The only valid settings files are those generated by ZEMAX, then renamed and saved to a new user defined file name. For example, on the Spot Diagram settings dialog, pressing "Save" will generate a "SPT.CFG" file with the saved settings. If this file is renamed to "MySPT.CFG", then the file, with a full path, may be used as the settingsfilename argument to GetTextFile.

If the flag value is 2, then the settings provided in the settings file, if valid, will be used and the settings box for the requested feature will be displayed. After the user makes any changes to the settings the file will then be generated using the new settings. The dialog boxes used to change the analysis settings use the data from the lens currently in the Lens Data Editor.

No matter what the flag value is, if a valid file name is provided for the settingsfilename, the settings used will be written to the settings file, overwriting any data in the file. To modify the settings defined within an existing settings file, use MODIFYSETTINGS.

Only text, and not graphic files, are supported by GETTEXTFILE.

See also OPEN, CLOSE, READ, READNEXT, READSTRING, and MODIFYSETTINGS.

Example:

```
! Macro sample to generate cardinal points in a file.
! Get a temporary file name
A$ = $TEMPFILENAME()
! Compute the data and place in the temp file
GETTEXTFILE A$, Car
! Open the new file and print it out
OPEN A$
LABEL 1
READSTRING B$
IF (EOFF()) THEN GOTO 2
PRINT B$
GOTO 1
LABEL 2
CLOSE
```

GETVARDATA

Purpose:

Retrieves the current number, type, and value of all optimization variables. The data is placed in one of the vector array variables (either VEC1, VEC2, VEC3, or VEC4).

Syntax:

```
GETVARDATA vector
```

Discussion:

The data is stored in the specified VECn array variable. For example, if the command GETVARDATA 1 is issued, the data will be stored in VEC1. The data is stored in the following format, where the first number in each line refers to the array position:

```
0:      n, the number of variables
1:      The type code for the first variable
```

2: Surface number for the first variable
 3: Parameter number for the first variable
 4: Object number for the first variable
 5: The value of the first variable
 5*q-4: The type code for the qth variable
 5*q-3: Surface number for the qth variable
 5*q-2: Parameter number for the qth variable
 5*q-1: Object number for the qth variable
 5*q: The value of the qth variable
 etc...

The integer q goes from 1 to n, where n is the number of variables. If n is zero, then no valid data is returned. The value of the number in array position zero, n, is always valid. The type codes for variables is as described in the following table. The surface number, parameter number, and object number may or may not have meaning depending upon the type of variable.

GETVARDATA TYPE AND ID CODES

Variable type	Type Code	Surface	Parameter	Object
Curvature	1	surface #	-	-
Thickness	2	surface #	-	-
Conic	3	surface #	-	-
Index Nd	4	surface #	-	-
Abbe Vd	5	surface #	-	-
Partial Dispersion $\Delta P_{g,F}$	6	surface #	-	-
TCE	7	surface #	-	-
Parameter Values	8	surface #	parameter #	-
Extra Data Values	9	surface #	extra data #	-
Multi-configuration Operand Values	10	oper #	config #	-
Non-Sequential Object Position X	11	surface #	-	object #
Non-Sequential Object Position Y	12	surface #	-	object #
Non-Sequential Object Position Z	13	surface #	-	object #
Non-Sequential Object Tilt X	14	surface #	-	object #
Non-Sequential Object Tilt Y	15	surface #	-	object #
Non-Sequential Object Tilt Z	16	surface #	-	object #
Non-Sequential Object Parameters	17	surface #	parameter #	object #

GETZERNIKE

Purpose:

Calculates Zernike Fringe, Standard, or Annular coefficients for the currently loaded lens file, and places them in one of the vector arrays (either VEC1, VEC2, VEC3, or VEC4).

Syntax:

GETZERNIKE maxorder, wave, field, sampling, vector, zerntype, epsilon, reference

Discussion:

The maxorder argument is any number between 1 and 37 for Fringe or between 1 and 231 for Standard or Annular coefficients (see the discussion of zerntype below), and corresponds to the highest Zernike term desired. Wave and field are the integer values for the wavelength and field number respectively. The value for sampling

determines the size of the grid used to fit the coefficients. Sampling may be 1 (32 x 32), 2 (64 x 64), etc.... up to 2048 x 2048. The vector argument must be an integer value between 1 and 4, and specifies which vector array the data should be placed in. The zerntype is 0 for "fringe" Zernike terms, 1 for "Standard" Zernike terms, and 2 for "Annular" Zernike terms. See "Zernike Fringe Coefficients" on page 196, "Zernike Standard Coefficients" on page 199, and "Zernike Annular Coefficients" on page 202 for descriptions of these different types of Zernike terms. For Annular Zernike Coefficients epsilon is the annular ratio; this value is ignored for other Zernike types. To reference the OPD to the chief ray, the reference value should be zero or omitted; use 1 to reference to the surface vertex. If any of the arguments fall outside the valid ranges, then the nearest acceptable value is used instead.

The data is returned in one of the vector arrays with the following format: Vector position 1: Peak to valley in waves; Vector position 2: RMS to the zero OPD line in waves (this value is not physically meaningful but is provided for reference); Vector position 3: RMS to the chief ray in waves; Vector position 4: RMS to the image centroid in waves (this is the most physically meaningful number related to image quality); Vector position 5: Variance in waves; Vector position 6: Strehl ratio; Vector position 7: RMS fit error in waves; Vector position 8: Maximum fit error (at any one point) in waves. The remaining vector positions contain the actual Zernike coefficient data. For example, Zernike term number 1 is in vector position 9, Zernike term 2 is in position 10, and so on.

Example:

```
! This macro computes the first 37 Zernike Fringe coefficients
! for the currently loaded lens, at wave 1, field 1
! and a 32x32 grid density (sampling = 1). The coefficients
! will be placed in vector 1. First get the data:
GETZERNIKE 37,1,1,1,1,0
! Now print it out:
FORMAT 16.6
PRINT "Peak to Valley      : ", vec1(1)
PRINT "RMS to chief       : ", vec1(3)
PRINT "RMS to centroid    : ", vec1(4)
PRINT "Variance           : ", vec1(5)
PRINT "Strehl ratio       : ", vec1(6)
PRINT "RMS Fit Error      : ", vec1(7)
PRINT "Maximum Fit Error  : ", vec1(8)
i = 1
label 1
    FORMAT 2.0
    PRINT "Zernike #", i, " = ",
    FORMAT 16.6
    PRINT vec1(8+i)
    i = i + 1
if (i < 38) THEN GOTO 1
PRINT "All Done!"
```

GLAS

This command is obsolete. See "SETSURFACEPROPERTY, SURP" on page 725.

GLASSTEMPLATE

Purpose:

GLASSTEMPLATE sets data on the Glass Substitution Template.

Syntax:

```
GLASSTEMPLATE code, data
```

Discussion:

For checkbox values, code is 0 for Use Glass Substitution Template, 1 for Exclude Glasses With Incomplete Data, 11 for Standard, 12 for Preferred, 13 for Obsolete, 14 for Special, and 21-26 for Use Relative Cost, CR, FR, SR, AR, and PR, respectively. The checkbox will be unchecked if the value of data is zero. Any non-zero value will check the checkbox.

For numerical values, code is 31-36 for Relative Cost, CR, FR, SR, AR, and PR, respectively.

Related Function:

GTEM

GLENSNAME

Purpose:

GLENSNAME will place the current lens name at the top left corner of the text box on user defined graphics screens.

Syntax:

GLENSNAME

Discussion:

GLENSNAME is primarily used for making your graphics look like other ZEMAX graphics.

Example:

See the section GRAPHICS.

GLOBALTOLOCAL

Purpose:

Converts coordinate breaks from global to local references.

Syntax:

GLOBALTOLOCAL surf1, surf2, direction

Discussion:

GLOBALTOLOCAL concatenates all adjacent coordinate breaks into single coordinate breaks. Surf1 and surf2 define the surface range. Direction is 0 for forward, 1 for reverse coordinate breaks. See "Global To Local" on page 256. See also "LOCALTOGLOBAL" on page 696.

GOSUB, SUB, RETURN, and END

Purpose:

These four keywords are used together to define and call subprograms within the ZPL macro file. Each keyword has a special purpose. GOSUB is used to direct the program flow to a defined subroutine. SUB is used to define the subroutine name, as well as indicate the beginning of the subroutine body. RETURN indicates the macro execution should continue at the point where the most recent GOSUB call was placed. END indicates that the macro should terminate immediately.

Syntax:

See the Example section for sample syntax.

Discussion:

There can be no more than 100 defined subroutines per ZPL macro file. Each subroutine must be terminated by a RETURN command. More than one return command may be placed in the body of a subroutine. If subroutines are defined, at least one END command must be used to indicate the end of the main macro body. The main macro body must be at the top of the file.

There can be no more than 100 "nesting levels" used in the ZPL macro. For example, if subroutine ABC calls subroutine XYZ, then the nesting level is 2. If subroutine XYZ then calls subroutine DEF, the nesting level is 3.

All variables in ZPL are global. Any variables used or defined within a subroutine exist within the main macro as well.

Example:

```
x = 1
y = 2
GOSUB add
print "the sum of ", x, " and ", y, " is ", z
END
```

```
SUB add
z = x + y
```

RETURN

GOTO

Purpose:

Normally, each macro line is executed in turn. GOTO allows execution to resume at an arbitrary point in the macro. GOTO is always used in conjunction with the LABEL command.

Syntax:

```
GOTO label_number  
GOTO text_label
```

Discussion:

There must be a LABEL command with the corresponding label_number or text_label somewhere in the macro, or an error will result.

Example:

```
LABEL 1  
x = RAND(10)  
if x <= 5 THEN GOTO 1  
PRINT " X is greater than 5 "
```

GRAPHICS

Purpose:

Creates a standard ZEMAX graphics frame with ruling lines for the plot title.



For an easy way of creating plots, see "PLOT" on page 706 and "PLOT2D" on page 707.

Syntax:

```
GRAPHICS  
GRAPHICS NOFRAME  
GRAPHICS OFF
```

Discussion:

If GRAPHICS is specified alone, then a standard ZEMAX graphics window will be created. If the optional argument NOFRAME is supplied, then the standard frame for the graph title will be suppressed. All subsequent graphics commands will be sent to this newly created window. GRAPHICS OFF will close any existing open graphics windows, and then display the closed window.

Example:

```
graphics  
xmx = xmax()  
xmn = xmin()  
ymx = ymax()  
ymn = ymin()  
xwidth = xmx-xmn  
ywidth = ymx-ymn  
xleft = xmn + ( .1 * xwidth )  
xright = xmn + ( .9 * xwidth )  
ytopp = ymn + ( .1 * ywidth )  
ybott = ymn + ( .7 * ywidth )  
line xleft,ytopp,xright,ytopp  
line xright,ytopp,xright,ybott  
line xright,ybott,xleft,ybott  
line xleft,ybott,xleft,ytopp  
gtitle " the rain in spain falls mainly on the plain"  
glensname  
gdate  
gtext xmx/2,ymx/2,0, " start this text in the center."  
gtextcent ymx*.05, " center this text near the top."  
gtext xmx*.05,ymx*.75,90, " place me vertically near left edge."  
gtext xmx*.15,ymx*.68,30, " orient me at 30 degrees."
```


graphics off

Related Keywords:

LINE, GITITLE, GLENSNAME, GDATE, GTEXTCENT, GTEXT, PIXEL, PLOT

GTEXT

Purpose:

GTEXT is used for labeling graphics plots with user defined text.

Syntax:

```
GTEXT x, y, angle, user_text  
GTEXT x, y, angle, A$
```

Discussion:

The coordinates x and y refer to the left edge of where the text string user_text will appear. "user_text" may be either a constant string in quotes or a string variable name. Angle specifies how the text is rotated with respect to the graphics frame, and defaults to 0 degrees (horizontal). See also SETTEXTSIZE.

Example:

See "GRAPHICS" on page 688.

GTEXTCENT

Purpose:

GTEXT is used for centering labels on graphics plots with user defined text.

Syntax:

```
GTEXTCENT y, user_text
```

Discussion:

The coordinate y refers to the vertical position of the text string user_text. See also SETTEXTSIZE.

Example:

See "GRAPHICS" on page 688.

GTITLE

Purpose:

GTITLE is similar to GTEXT except only the text needs to be specified, and the text will appear centered in the title bar on the graphics display. GTITLE is useful for making ZPL graphics functions look like standard ZEMAX graphics.

Syntax:

```
GTITLE user_title
```

Example:

See "GRAPHICS" on page 688.

HAMMER

Purpose:

Invokes the Hammer optimization algorithm to optimize the current lens with the current merit function.

Syntax:

```
HAMMER  
HAMMER number_of_cycles  
HAMMER number_of_cycles, algorithm
```

Discussion:

If no argument is provided, then the hammer optimization runs 1 cycle using Damped Least Squares. If an argument is provided, it must evaluate to an integer between 1 and 99, and the Hammer optimization algorithm will run the specified number of cycles. For the algorithm argument, use 0 for Damped Least Squares (the default) and 1 for Orthogonal Descent. For more information see "Performing an optimization" on page 526.

Related Functions:

MFCN

Example:

```
PRINT "Starting merit function:", MFCN()
HAMMER 3
PRINT "Ending merit function  :", MFCN()
```

IF-THEN-ELSE-ENDIF

Purpose:

IF provides conditional macro execution and branching.

Syntax:

```
IF (expression)
(commands)
ELSE
(commands)
ENDIF
```

or

```
IF (expression) THEN (command)
```

Discussion:

The IF-ELSE-ENDIF construction is used for conditional execution of either the group of commands following the IF command or the commands following the ELSE command, but not both. The value of expression is considered false if it is zero, otherwise it is considered true. The expression can be any valid ZPL expression, composed of functions, variables, operands, and constants. The IF command must be paired with an ENDIF, although the ELSE is optional. IF-ENDIF pairs may be nested to any level.

The IF-THEN construction is handy for conditional execution of a single instruction. If the THEN keyword is specified, the IF command is terminated, and there is no need for an ENDIF. The ELSE keyword is not supported in the IF-THEN construction.

Example:

```
x = 1
y = 2
if (x < y)
    PRINT "x is less than y"
ELSE
    if (x == y) THEN PRINT "x equals y"
    if (x > y) THEN PRINT "x is greater than y"
ENDIF
```

IMA

Purpose:

Computes the Geometric Image Analysis feature and saves the result to a BIM file. For a description of the BIM (Binary Image) format see "The BIM format" on page 168. For a description of the Geometric Image Analysis feature see "Geometric Image Analysis" on page 165.

Syntax:

```
IMA outfile, infile
```

Discussion:

This keyword requires the name of the output BIM file, and optionally, the name of the input IMA or BIM file. If the extension to the outfile is not provided, the extension BIM will be appended. The extension must be provided on the infile. The filenames must be enclosed in quotes if any blank or other special characters are used. The outfile will be placed in the <data>\<images> folder. The infile must also be placed in the <data>\<images> folder. No paths should be provided with the file names.

The settings for the Geometric Image Analysis feature will be those settings previously saved for the current lens. To make adjustments to the settings, open a Geometric Image Analysis window, choose the appropriate settings, then press "Save". Be certain to choose "Show As" as anything other than "spot diagram". All

subsequent calls to IMA will use the saved settings. The exceptions are the output file name, which is specified as the first argument after the IMA keyword, and the input source file, which is optionally specified as the second argument after the IMA keyword.

Example:

```
IMA "output.BIM", "LETTERF.IMA"
```

Related Keywords:

IMASHOW, IMASUM

IMAGECOMBINE

Purpose:

Combines two graphic files and writes the combined graphic to a new file.

Syntax:

```
IMAGECOMBINE source1$, source2$, destination$, mode$
```

Discussion:

This keyword opens the two source graphic files, combines them to a single image, and then writes the resulting image to the destination file. The files may be in JPG or BMP format or any combination of the two. Conversion between the formats is automatic and is based upon the extension provided. If the source images have the same number of rows, they can be combined left-to-right. If the source images have the same number of columns, they can be combined top-to-bottom. The combination used depends upon the mode\$ value. Use either "LEFTRIGHT" or "TOPBOTTOM" for the value of mode\$ to define how the images should be combined.

Example:

```
S1$ = "C:\TEMP\SOURCE1.JPG"  
S2$ = "C:\TEMP\SOURCE2.JPG"  
D$ = "C:\TEMP\DESTINATION.JPG"  
M$ = "LEFTRIGHT"  
IMAGECOMBINE S1$, S2$, D$, M$
```

Related Functions:

PIXX, PIXY

Related Keyword:

IMAGEEXTRACT

IMAGEEXTRACT

Purpose:

Extracts a rectangular region of a graphic file and writes the portion of the graphic to a new file.

Syntax:

```
IMAGEEXTRACT source$, destination$, startx, starty, width, height
```

Discussion:

This keyword opens the source graphic, and extracts a rectangular portion of the file, then writes the extracted portion to the destination file. The files may be in JPG or BMP format or any combination of the two. Conversion between the formats is automatic and is based upon the extension provided. The startx and starty values are the offsets from the left side and top side respectively, in pixels. The top row and left column start at number 1. The width and height values are the size in pixels of the rectangular region to extract.

Example:

```
S$ = "C:\TEMP\SOURCE.JPG"  
D$ = "C:\TEMP\DESTINATION.JPG"  
IMAGEEXTRACT S$, D$, 1, 1, 20, 20
```

Related Functions:

PIXX, PIXY

Related Keywords:

IMAGECOMBINE

IMASHOW

Purpose:

Displays an IMA or BIM file in a viewer window. For a description of the IMA and BIM file formats, see “The IMA format” on page 167 and “The BIM format” on page 168.

Syntax:

```
IMASHOW filename.ima
```

Discussion:

This keyword requires the name of the IMA or BIM file. The extension must be included. The filename may be enclosed in quotes if any blank or other special characters are used. The file must be located in the <data>\<images> folder. This command will open a new window to display the file.

Example:

```
IMASHOW "LETTERF.IMA"
```

Related Keywords:

IMA, IMASUM

IMASUM

Purpose:

Sums the intensity in a BIM file. Despite the name of this function, currently only BIM files, and not IMA files, may be summed. This limitation is due to the low number of grey scales that are supported in the IMA file format. For a description of the IMA and BIM file formats, see “The IMA format” on page 167 and “The BIM format” on page 168.

Syntax:

```
IMASUM filename1, filename2, outfile
```

Discussion:

This keyword requires the names of three BIM files (which need not be distinct). The BIM extension must be provided on all three file names. The filenames may be enclosed in quotes if any blank or other special characters are used. The files must be located in the <data>\<images> folder, which is where the output file will be. The two source files must have the same number of pixels and color channels, or an error message will result.

Example:

```
IMASUM "A.BIM", "B.BIM", "sum of A and B.BIM"
```

Related Keywords:

IMA, IMASHOW

IMPORTEXTRADATA

Purpose:

Imports data into the extra data editor from a file.

Syntax:

```
IMPORTEXTRADATA surface, filename
```

Discussion:

The surface may be any valid surface number that uses extra data values. The data to be imported should be a single column of numeric data in text format. See "Tools" under “Extra Data” on page 93, and “Importing grid data” on page 314.

INPUT

Purpose:

INPUT provides a means for prompting the user for numeric or text data when the macro is run.

Syntax:

```
INPUT "Prompt String" , variable  
INPUT variable
```

```
INPUT "Prompt String" , string_variable$
INPUT string_variable$
```

Discussion:

The variable may be any valid variable name. If the variable name is a string variable, then the input will be interpreted as a literal string; otherwise, as a numeric. The INPUT command will use a "?" prompt if no prompt string is supplied. The prompt is always displayed on the screen, and the input always is accepted from the keyboard only. The prompt string cannot exceed 200 characters in length.

Example:

```
INPUT "Enter value for x:", x
PRINT "X = ", x
INPUT "Enter a value for A$:", A$
PRINT A$
```

INSERT

Purpose:

INSERT inserts a new surface in the lens data editor.

Syntax:

```
INSERT surf
```

Discussion:

The value surf must evaluate to an integer expression. See also DELETE and SURFTYPE.

Example:

```
INSERT 5
INSERT i+2*j
```

INSERTCONFIG

Purpose:

INSERTCONFIG inserts a new configuration in the multi-configuration editor.

Syntax:

```
INSERTCONFIG config
```

Discussion:

The value config must evaluate to an integer expression greater than 0 and less than or equal to the current number of configurations plus 1. See also DELETECONFIG and INSERTMCO.

Example:

```
INSERTCONFIG 4
```

INSERTMCO

Purpose:

INSERTMCO inserts a new multi-configuration operand in the multi-configuration editor.

Syntax:

```
INSERTMCO row
```

Discussion:

The value row must evaluate to an integer expression greater than 0 and less than or equal to the current number of operands plus 1. See also DELETEMCO and INSERTCONFIG.

Example:

```
INSERTMCO 7
```

INSERTMFO

Purpose:

INSERTMFO inserts a new merit function operand in the merit function editor.

Syntax:

INSERTMFO row

Discussion:

The value row must evaluate to an integer expression greater than 0 and less than or equal to the current number of operands plus 1. See also DELETMFO and SETOPERAND.

Example:

INSERTMFO 23

INSERTOBJECT

Purpose:

Inserts a new NSC object.

Syntax:

INSERTOBJECT surf, object

Discussion:

The value surf must evaluate to an integer expression, and the specified surface must be a non-sequential component surface. Use a surf value of 1 if the program mode is non-sequential. The value object must evaluate to an integer that is between 1 and the current number of objects plus 1, inclusive. The new object will be inserted at the specified object number, renumbering other objects as required. See also DELETEOBJECT and SETNSCPROPERTY.

Example:

INSERTOBJECT 1, 1

INSERTTOL

Purpose:

INSERTTOL inserts a new operand in the tolerance data editor.

Syntax:

INSERTTOL row

Discussion:

The value row must evaluate to an integer expression greater than 0 and less than or equal to the current number of operands plus 1. See also DELETETOL and SETTOL.

Example:

INSERTTOL 3

LABEL

Purpose:

LABEL provides a destination for the GOTO command, see "GOTO" on page 688 for details.

Syntax:

LABEL label_number

LABEL text_label

Discussion:

The label_number must be an integer value greater than zero, such as 1 or 7. If a text label is used, the text label must not contain spaces or other special characters that are used as delimiters. LABEL has no effect on program flow by itself. There is a limit of 300 label commands in a macro.

Example:

LABEL 7

LABEL startover

LINE

Purpose:

LINE is the primitive line drawing function for graphical displays.

Syntax:

LINE oldx, oldy, newx, newy

Discussion:

LINE will evaluate the four expressions and draw a line connecting the points defined. The coordinates refer to the current graphics frame, and must be contained within the bounds defined by XMIN, YMIN, XMAX, and YMAX. Although only integer pixel values can actually be plotted, LINE will accept real values as arguments and round the coordinates to the nearest integer equivalents. LINE is only valid in graphics mode. The color of the line drawn is controlled by the current pen color, which is specified using the COLOR keyword.

Example:

See "GRAPHICS" on page 688.

LOADARCHIVE

Purpose:

Opens an existing ZEMAX archive (*.ZAR) file.

Syntax:

LOADARCHIVE filename, extractpath

Discussion:

The filename must include the name of the archive file including the file extension. If the filename does not include a complete path to the archive file, the default folder for lenses is assumed. If no extractpath is defined, the path is assumed to be the same as the filename. The lens and session files defined in the archive file are then opened. See "Restore From Archive File" on page 66.

LOADCATALOG

Purpose:

Reloads glass and coating catalogs for the currently loaded lens.

Syntax:

LOADCATALOG

Discussion:

When lenses are loaded, any associated glass catalogs and data files, including the COATING.DAT file, are automatically loaded if they are not already loaded. However, if these catalogs have been modified, perhaps by the ZPL macro itself; then the LOADCATALOG keyword may be used to force a reload of the catalogs. Use of this keyword is not required unless the COATING.DAT or glass AGF catalog files have been modified since the start of the current ZEMAX session.

LOADDETECTOR

Purpose:

Loads the data saved in a file to an NSC Detector Rectangle, Detector Color, Detector Polar, or Detector Volume object.

Syntax:

LOADDETECTOR surf, object, filename

Discussion:

This keyword requires numeric expressions that specify the surface and object, and a file name with or without a full path. Surf is the surface number of the non-sequential group; use 1 when using non-sequential mode. Object is the object number of the detector object. The filename may include the full path, if no path is provided the path of the current lens file is used. The extension should be DDR, DDC, DDP, or DDV for Detector Rectangle, Color, Polar, and Volume objects, respectively.

Related Keywords:

SAVEDETECTOR

LOADLENS

Purpose:

Loads a new lens file, replacing any lens file currently in memory.

Syntax:

```
LOADLENS filename, appendflag, session
```

Discussion:

LOADLENS will load a lens file. If the filename contains the complete path, such as C:\MYDIR\MYLENS.ZMX, then the specified file will be loaded. If the path is left off, then the default folder for lenses will be used (see "Folders" on page 68).

If the appendflag is zero or absent, then LOADLENS loads the file. If the appendflag is greater than zero, then the file is appended to the current lens starting at the surface specified by the value of the appendflag. If the session flag is non-zero, any associated session file will be loaded with the lens and all windows will be updated otherwise the lens session file is ignored.

Example:

```
LOADLENS "COOKE.ZMX"
```

Related Keywords:

SAVELENS

LOADMERIT

Purpose:

Loads a merit function file, replacing the merit function in the current lens.

Syntax:

```
LOADMERIT "filename"  
LOADMERIT file$
```

Discussion:

LOADMERIT will load a new merit function from a .MF file. If the filename contains the complete path, such as C:\MYDIR\MYLENS.MF, then the specified file will be loaded. If the path is left off, then the <data>\MeritFunction folder will be used (see "Folders" on page 68). This keyword can also load just the merit function from a ZMX file using the same syntax. See also SAVEMERIT.

LOADTOLERANCE

Purpose:

Loads a tolerance file, replacing the tolerances in the current lens.

Syntax:

```
LOADTOLERANCE "filename"  
LOADTOLERANCE file$
```

Discussion:

LOADTOLERANCE will load a new tolerance file. If the filename contains the complete path, such as C:\MYDIR\MYLENS.TOL, then the specified file will be loaded. If the path is left off, then the <data>\Tolerance folder will be used (see "Folders" on page 68). See also SAVETOLERANCE.

LOCALTOGLOBAL

Purpose:

Converts coordinate breaks from local to global references.

Syntax:

```
LOCALTOGLOBAL surf1, surf2, reference
```

Discussion:

LOCALTOGLOBAL modifies all coordinate breaks into groups of 3 coordinate breaks, providing a means to locate surfaces in global coordinates. Surf1 and surf2 define the surface range. Reference is the reference

surface number, which must precede s1. See “Local To Global” on page 256. See also “GLOBALTOLOCAL” on page 687.

LOCKWINDOW

Purpose:

Locks any one or all open windows.

Syntax:

```
LOCKWINDOW winnum
```

Discussion:

See “Graphic windows operations” on page 42. If the winnum is zero, then all open windows are locked. If the winnum argument is -1, then the currently executing window will lock at the end of the macro execution.

Example:

```
LOCKWINDOW 7
```

Related Keywords:

UNLOCKWINDOW

MAKEFACETLIST

Purpose:

Makes an ASCII listing of the vertices of a faceted description of any importable CAD file.

Syntax:

```
MAKEFACETLIST infile, outfile
```

Discussion:

The infile is a string variable or name with the full path to a CAD format file. The outfile is a string variable or name to the full path for the output file. The output file will be a text listing. There are 18 lines for each triangle in the faceted description of the object. The first 6 lines are data for vertex 1 of the first triangle in this format:

```
x coordinate  
y coordinate  
z coordinate  
x normal vector  
y normal vector  
z normal vector
```

The next 6 values are for vertex 2, and then 6 more values for vertex 3. The 18 values together define the first triangle. The pattern repeats for the next triangle; each represented by 18 values, until all the triangles are written.

MAKEFOLDER

Purpose:

Makes a folder for files.

Syntax:

```
MAKEFOLDER "C:\TEMP\TEST_FOLDER"  
MAKEFOLDER path$
```

Discussion:

The new folder will be created. If the folder already exists the keyword is ignored. If the path is invalid, an error message is issued and the macro execution will be terminated.

MODIFYSETTINGS

Purpose:

Modifies data within the settings files used by GETTEXTFILE.

Syntax:

```
MODIFYSETTINGS settingsfilename, type, value
```

Discussion:

The settingsfilename must be in quotes, or be a string variable name, and include the full path, name, and extension for the file to be modified.

The type argument is a text mnemonic that indicates which setting within the file is to be modified. The supported values for the type argument are listed in the table below. Because there are many different types of analysis windows, and each has many different settings available, the list of types does not include all possible settings. More type values may be added upon request; please contact technical support for assistance.

Type codes supported by MODIFYSETTINGS

Feature	Available type codes
2D Layout	LAY_RAYS: The number of rays.
Detector Viewer	DVW_SURFACE: The surface number. Use 1 for Non-Sequential mode. DVW_DETECTOR: The detector number. DVW_SHOW: The "show as" setting. The meaning depends upon the type of window displayed: For Graphics Windows: Use 0 for grey scale, 1 for inverted grey scale, 2 for false color, 3 for inverted false color, 4 for cross section row, and 5 for cross section column. For Text Windows: Use 0 for full pixel data, 1 for cross section row, and 2 for cross section column. DVW_ROWCOL: The row or column number for cross section plots. DVW_ZPLANE: The Z-Plane number for detector volumes. DVW_SCALE: The scale mode. Use 0 for linear, 1 for Log -5, 2 for Log -10, and 3 for Log - 15. DVW_SMOOTHING: The integer smoothing value. DVW_DATA: Use 0 for incoherent irradiance, 1 for coherent irradiance, 2 for coherent phase, 3 for radiant intensity, 4 for radiance (position space), and 5 for radiance (angle space). DVW_ZRD: The ray data base name, or null for none. DVW_FILTER: The filter string. DVW_MAXPLOT: The maximum plot scale. DVW_MINPLOT: The minimum plot scale.
Extended Diffraction Image Analysis	EXD_DISPLAYSIZE: The display size. EXD_FIELD: The field number. EXD_FILESIZE: The file size. EXD_WAVE: The wavelength number.
FFT Line/Edge Spread	LSF_COHERENT: Use 0 for incoherent, 1 for coherent LSF_TYPE: Use 0-9 for X-Linear, Y-Linear, X-Log, Y-Log, X-Phase, Y-Phase, X-Real, Y-Real, X-Imaginary, or Y-Imaginary, respectively. LSF_SAMP: The sampling, use 1 for 32 x 32, 2 for 64 x 32, etc. LSF_SPREAD: Use 0 for line, 1 for edge. LSF_WAVE: The wavelength number, use 0 for polychromatic (incoherent only) LSF_FIELD: The field number. LSF_POLARIZATION: Use 0 for unpolarized, 1 for polarized. LSF_PLOTSCALE: The plot scale.
FFT PSF	PSF_TYPE: Use 0-4 for Linear, Log, Phase, Real, or Imaginary, respectively. PSF_SAMP: The sampling, use 1 for 32 x 32, 2 for 64 x 32, etc. PSF_WAVE: The wavelength number, use 0 for polychromatic. PSF_FIELD: The field number. PSF_SURFACE: The surface number, use 0 for image. PSF_POLARIZATION: Use 0 for unpolarized, 1 for polarized. PSF_NORMALIZE: Use 0 for unnormalized, 1 for unity normalization. PSF_IMAGEDELTA: The image point spacing in micrometers.

Feature	Available type codes
FFT PSF Cross Section	PSF_TYPE: Use 0-9 for X-Linear, Y-Linear, X-Log, Y-Log, X-Phase, Y-Phase, X-Real, Y-Real, X-Imaginary, or Y-Imaginary, respectively. PSF_ROW: The row number (if doing an X scan) or column number (if doing a y scan). Use 0 for center. PSF_SAMP: The sampling, use 1 for 32 x 32, 2 for 64 x 32, etc. PSF_WAVE: The wavelength number, use 0 for polychromatic. PSF_FIELD: The field number. PSF_POLARIZATION: Use 0 for unpolarized, 1 for polarized. PSF_NORMALIZE: Use 0 for unnormalized, 1 for unity normalization. PSF_PLOTSCALE: The plot scale.
Footprint Diagram	FOO_RAYDENSITY: The ray density. Use 0 for ring, 1 for 10, 2 for 15, 3 for 20 etc. FOO_SURFACE: The surface number. FOO_FIELD: The field number. FOO_WAVELENGTH: The wavelength number. FOO_DELETEVIGNETTED: Delete vignetted, use 0 for no, 1 for yes.
Geometric Bitmap Image Analysis	GBM_FIELDSIZE: The field Y size. GBM_RAYS: The number of rays per source pixel. GBM_XPIX: The number of X pixels. GBM_YPIX: The number of Y pixels. GBM_XSIZ: The X pixel size. GBM_YSIZ: The Y pixel size. GBM_INPUT: The input file name GBM_OUTPUT: The output file name GBM_SURFACE: The surface number GBM_ROTATION: The rotation setting
Geometric Image Analysis	IMA_FIELD: The field size. IMA_IMAGESIZE: The image size. IMA_IMANAME: The image file name. IMA_KRAYS: The number of rays x 1000. IMA_NA: The numerical aperture. IMA_OUTNAME: The output file name. IMA_SURFACE: The surface number.
Huygens PSF	HPS_PUPILSAMP: The pupil sampling, use 1 for 32 x 32, 2 for 64 x 64, etc. HPS_IMAGESAMP: The image sampling, use 1 for 32 x 32, 2 for 64 x 64, etc. HPS_WAVE: The wavelength number, use 0 for polychromatic. HPS_FIELD: The field number. HPS_IMAGEDELTA: The image point spacing in micrometers. HPS_TYPE: The data type. Use 0-8 for Linear, Log -1, Log -2, Log -3, Log -4, Log -5, Real, Imaginary, or Phase, respectively.
Huygens PSF Cross Section	HPC_PUPILSAMP: The pupil sampling, use 1 for 32 x 32, 2 for 64 x 64, etc. HPC_IMAGESAMP: The image sampling, use 1 for 32 x 32, 2 for 64 x 64, etc. HPC_WAVE: The wavelength number, use 0 for polychromatic. HPC_FIELD: The field number. HPC_IMAGEDELTA: The image point spacing in micrometers. HPC_TYPE: The data type. Use 0-9 for X-Linear, X-Log, Y-Linear, Y-Log, X-Real, Y-Real, X-Imaginary, Y-Imaginary, X-Phase, or Y-Phase, respectively.

Feature	Available type codes
Illumination XY Scan	ILL_SOURCE: The source size. ILL_SMOOTH: The smoothing value to use. ILL_DETSIZE: The detector size. ILL_SURFACE: The surface number.
Image Simulation	ISM_INPUTFILE: The input file name. ISM_FIELDHEIGHT: The Y field height. ISM_OVERSAMPLING: Oversample value. Use 0 for none, 1 for 2X, 2 for 4x, etc. ISM_GUARDBAND: Guard band value. Use 0 for none, 1 for 2X, 2 for 4x, etc. ISM_FLIP: Flip Source. Use 0 for none, 1 for TB, 2 for LR, 3 for TB&LR. ISM_ROTATE: Rotate Source: Use 0 for none, 1 for 90, 2 for 180, 3 for 270. ISM_WAVE: Wavelength. Use 0 for RGB, 1 for 1+2+3, 2 for wave #1, 3 for wave #2, etc. ISM_FIELD: Field number. ISM_PSAMP: Pupil Sampling. Use 1 for 32x32, 2 for 64x64, etc. ISM_ISAMP: Image Sampling. Use 1 for 32x32, 2 for 64x64, etc. ISM_PSFY, ISM_PSFY: The number of PSF grid points. ISM_ABERRATIONS: Use 0 for none, 1 for geometric, 2 for diffraction. ISM_POLARIZATION: Use 0 for no, 1 for yes. ISM_SHOWAS: Use 0 for Simulated Image, 1 for Source Bitmap, and 2 for PSF Grid. ISM_REFERENCE: Use 0 for chief ray, 1 for vertex, 2 for primary chief ray. ISM_SUPPRESS: Use 0 for no, 1 for yes. ISM_PIXELSIZE: Use 0 for default or the size in lens units. ISM_XSIZE, ISM_YSIZE: Use 0 for default or the number of pixels. ISM_OUTPUTFILE: The output file name or empty string for no output file.
MTF - FFT	MTF_SAMP: The pupil sampling, use 1 for 32, 2 for 64, etc. MTF_WAVE: The wavelength number, use 0 for all. MTF_FIELD: The field number, use 0 for all. MTF_TYPE: Use 0 for modulation, 1 for real, 2 for imaginary, 3 for phase, 4 for square wave. MTF_SURF: The surface number, use 0 for image. MTF_MAXF: The maximum frequency, use 0 for default. MTF_SDLI: Show diffraction limit, 0 for no, 1 for yes. MTF_POLAR: Polarization, 0 for no, 1 for yes. MTF_DASH: Use dashes, 0 for no, 1 for yes.
NSC Object Viewer	SHA_ROT_X: The x rotation in degrees. SHA_ROT_Y: The y rotation in degrees. SHA_ROT_Z: The z rotation in degrees.
NSC Shaded Model	SHA_ROT_X: The x rotation in degrees. SHA_ROT_Y: The y rotation in degrees. SHA_ROT_Z: The z rotation in degrees.

Feature	Available type codes
Partially Coherent Image Analysis	PCI_FIELD: The field number. PCI_FILESIZE: The file size. PCI_WAVE: The wavelength number. PCI_RESAMPLE: The resample image setting, 0 for no 1 for yes. PCI_RSNX: The resample number x PCI_RSNY: The resample number y PCI_RSDCX: The resample decenter x PCI_RSDCY: The resample decenter y PCI_RSDLX: The resample delta x PCI_RSDLY: The resample delta y
Physical Optics Propagation - General Tab	POP_END: The end surface. POP_FIELD: The field number. POP_START: The starting surface. POP_WAVE: The wavelength number.
Physical Optics Propagation - Beam Definition Tab	POP_AUTO: Simulates the pressing of the "auto" button which chooses appropriate X and Y beam widths based upon the sampling and other settings. POP_BEAMTYPE: Selects the beam type. Use 0 for Gaussian Waist, 1 for Gaussian Angle, 2 for Gaussian Size + Angle, 3 for Top Hat, 4 for File, 5 for DLL and 6 for Multimode. POP_PARAMn: Sets beam parameter n, for example, use POP_PARAM3 to set parameter3. POP_PEAKIRRAD: Sets the normalization by peak irradiance. POP_POWER: Sets the normalization by total beam power. POP_SAMPX: The X direction sampling, use 1 for 32, 2 for 64, etc. POP_SAMPY: The Y direction sampling, use 1 for 32, 2 for 64, etc. POP_SOURCEFILE: The file name if the starting beam is defined by a ZBF file, DLL, or multimode file. POP_WIDEX: The X direction width. POP_WIDEY: The Y direction width.
Physical Optics Propagation - Fiber Data Tab	POP_COMPUTE: Use 1 to check the fiber coupling integral on, 0 to check it off. POP_FIBERFILE: The file name if the fiber mode is defined by a ZBF or DLL. POP_FIBERTYPE: Use the same values as POP_BEAMTYPE above, except for multimode which is not yet supported. POP_FPARAMn: Sets fiber parameter n, for example, use POP_FPARAM3 to set fiber parameter3. POP_IGNOREPOL: Use 1 to ignore polarization, 0 to consider polarization. POP_POSITION: Fiber position setting. Use 0 for chief ray, 1 for surface vertex. POP_TILTX: The X-Tilt. POP_TILTY: The Y-Tilt.
Shaded Model	SHA_ROTX: The x rotation in degrees. SHA_ROTY: The y rotation in degrees. SHA_ROTZ: The z rotation in degrees.
Spot Diagram	SPT_RAYS: The ray density.
Wavefront Map	WFM_SAMP: The sampling, use 1 for 32, 2 for 64, etc. WFM_FIELD: The field number. WFM_WAVE: The wavelength number. WFM_SUBSR: The sub aperture radius. WFM_SUBSX: The sub aperture X decenter. WFM_SUBSY: The sub aperture Y decenter.

The value parameter is the new data for the specified setting. The modified settings file is written back to the original settings file name.

See also GETTEXTFILE.

Example:

```
MODIFYSETTINGS "C:\MySPT.CFG", SPT_RAYS, 24
MODIFYSETTINGS "C:\MyPOP.CFG", POP_SOURCEFILE, "MyStartBeam.ZBF"
```

NEXT

See “FOR, NEXT” on page 678.

NSTR

Purpose:

Initiates a Non-sequential trace.

Syntax:

```
NSTR surf, source, split, scatter, usepolar, ignore_errors, random_seed, save,
savefilename, filter, zrd_format
```

Discussion:

Surf is an integer value that indicates the number of the Non-Sequential surface. If the program mode is set to Non-Sequential, use 1. Source refers to the object number of the desired source. If source is zero, all sources will be traced. If Split is non-zero, then splitting is on, otherwise, ray splitting is off. If Scatter is non-zero, then scattering is on, otherwise scattering is off. If Usepolar is non-zero then polarization will be used, otherwise polarization is off. If splitting is on polarization is automatically selected. If ignore_errors is non-zero, then errors will be ignored, otherwise ray errors will terminate the non-sequential trace and macro execution and an error will be reported.

If random_seed is zero, then the random number generator will be seeded with a random value, and every call to NSTR will produce different random rays. If random_seed is any integer other than zero, then the random number generator will be seeded with the specified value, and every call to NSTR using the same seed will produce identical rays. When using NSTR for optimization, it is recommended that a non-zero value be used for random_seed.

If save is omitted or is zero, the arguments savefilename, filter, and zrd_format need not be supplied. If save is not zero, the rays will be saved to a file. The saved data file will have the name specified by the savefilename. The file naming convention and destination folder are the same as described in “Saving ray data to a file” on page 582. The extension of savefilename should be provided, but no path should be specified. If save is not zero, then the optional filter name is either a string variable with the filter, or the literal filter in double quotes. For information on filter strings see “The filter string” on page 589. For ZRD files, the zrd_format can be 0, 1, or 2 for uncompressed full data, compressed basic data, or compressed full data, respectively. For more information on ZRD formats see “Ray database (ZRD) files” on page 583.

NSTR always calls UPDATE before tracing rays to make certain all objects are correctly loaded and updated.

Related Functions:

NSDD

Example:

```
NSTR 1, 0, 0, 0, 0, 1, 0, 1, "saverays.ZRD", "h2"
```

NUMFIELD

This command is obsolete. See “SETSYSTEMPROPERTY, SYSP” on page 728.

NUMWAVE

This command is obsolete. See “SETSYSTEMPROPERTY, SYSP” on page 728.

OPEN

Purpose:

Opens an existing text file for reading by the READ command.

Syntax:

```
OPEN "filename"  
OPEN A$
```

Discussion:

The filename provided must be a valid file within quotes, or a string variable name containing the file name. See the keywords READ and CLOSE. Always CLOSE a file after all the data has been read.

Related Functions:

EOFF

Related Keywords:

READ, READNEXT, CLOSE

Example:

```
PRINT "Reading the double-column file TEST.DAT!"  
OPEN "TEST.DAT"  
READ x1, y1  
READ x2, y2  
READ x3, y3  
CLOSE
```

OPENANALYSISWINDOW

Purpose:

Opens a new analysis window.

Syntax:

```
OPENANALYSISWINDOW type, settingsfilename
```

Discussion:

The type argument is a 3 character string code that indicates the type of analysis to be performed. The 3 letter string codes are identical to those used for the button bar in ZEMAX. A list of string codes may be found on the "Buttons" tab of the File, Preferences dialog box. If no string code is provided or recognized, no window is opened. The settingsfilename contains the name of the CFG settings file to use. If this argument is missing, the default settings will be used.

Related Functions:

WINL, WINN

Related Keywords:

CLOSEWINDOW, MODIFYSETTINGS

Example:

```
FORMAT 0.0  
PRINT "There are ", WINN(), " Windows currently open."  
OPENANALYSISWINDOW "mtf"  
PRINT "Just opened window ", WINL(), "."  
PAUSE TIME, 3000  
PRINT "There are ", WINN(), " Windows currently open."  
CLOSEWINDOW WINL()  
PAUSE TIME, 1000  
PRINT "There are ", WINN(), " Windows currently open."
```

OPTIMIZE

Purpose:

Invokes the optimization algorithm to optimize the current lens with the current merit function.

Syntax:

```
OPTIMIZE  
OPTIMIZE number_of_cycles  
OPTIMIZE number_of_cycles, algorithm
```

Discussion:

The expression for `number_of_cycles` must evaluate to an integer value between 1 and 99, and the optimization algorithm will run the specified number of cycles. If `number_of_cycles` evaluates to zero, then the optimization runs in "Automatic" mode, stopping when the algorithm detects the process has converged. For the algorithm argument, use 0 for Damped Least Squares (the default) and 1 for Orthogonal Descent. For more information see "Performing an optimization" on page 526.

To update the merit function without optimizing, use the MFCN function.

Related Functions:

MFCN

Example:

```
PRINT "Starting merit function:", MFCN()
OPTIMIZE
PRINT "Ending merit function  :", MFCN()
```

OPTRETURN

Purpose:

Used to return numerical values back to the optimization algorithm through the use of the ZPLM optimization operand.

Syntax:

```
OPTRETURN datafield, result
```

Discussion:

The `datafield` expression must evaluate to an integer between 0 and 50. The `datafield` refers to a position in an array where the value of the expression result may be stored. The sole purpose of OPTRETURN is to be able to optimize values computed within a ZPL macro.

The optimization operand ZPLM must be used in the merit function to call the ZPL macro and retrieve the value returned by OPTRETURN. See "Optimizing with ZPL macros" on page 533 for details.

Example:

```
x = sqrt(thic(3) + radi(5))
OPTRETURN j, x+5
```

OUTPUT

Purpose:

Specifies destination for text output. Output is either to the screen, or to a file.

Syntax:

```
OUTPUT SCREEN
OUTPUT filename
OUTPUT filename, APPEND
```

Discussion:

If OUTPUT SCREEN is specified alone, then all subsequently executed PRINT commands will be directed to the screen. If a valid filename is provided, then subsequent PRINT commands will output to the filename specified. To close the file created earlier, use OUTPUT SCREEN which will direct subsequent PRINT outputs to the screen. SHOWFILE will close the file and send it to the text viewer program for screen display. PRINTFILE will close the file and print it on the currently defined printer.

If the keyword APPEND follows the file name, then subsequent output will be appended to the file. Otherwise, the contents of the file will be overwritten.

Example:

```
OUTPUT "x.txt"
PRINT "This will not appear on the screen, but in the file x.txt."
OUTPUT SCREEN
PRINT "This will appear on the screen."
OUTPUT "x.txt", APPEND
```



```
PRINT "This will appear after the first line in the file x.txt."
```

Related Keywords:

CLOSE, OPEN, SHOWFILE, PRINTFILE

PARM

This command is obsolete. See "SETSURFACEPROPERTY, SURP" on page 725.

PARAXIAL

Purpose:

Used to control whether ray tracing is done with paraxial or real rays.

Syntax:

```
PARAXIAL ON  
PARAXIAL OFF
```

Discussion:

The current paraxial mode can be established by a call to the function PMOD, which returns 0 if paraxial mode is off, 1 if it is on. This feature is used for switching between real and paraxial ray traces. Certain calculations, such as measuring distortion and computing first-order properties such as effective focal length, require the tracing of paraxial rays.

Example:

```
mode = PMOD()  
IF mode THEN PRINT "Paraxial mode is on!"  
IF !mode THEN PRINT "Paraxial mode is not on!"  
PARAXIAL ON  
PRINT "Now paraxial mode is on!"  
PRINT "Restoring original mode..."  
if !mode THEN PARAXIAL OFF
```

PAUSE

Purpose:

Pauses macro execution, optionally while displaying a status message. The status message can be a string or numerical value. The macro continues once the user presses the "OK" button on the status dialog.

Syntax:

```
PAUSE  
PAUSE "Ready to continue..."  
PAUSE TIME, time  
PAUSE THREADS
```

Discussion:

This feature may be used for debugging, presenting results, or pausing the execution of the macro. If the keyword after the PAUSE command is TIME, then the macro execution will "sleep" for *at least* the number of milliseconds given by the value provided. For example, to pause the macro for 100 milliseconds, the syntax is

```
PAUSE TIME, 100
```

The macro execution will continue after the specified time interval has passed. No status message is displayed in this case. Note that the system timer has a resolution of about 10 milliseconds, so the actual delay introduced by PAUSE TIME will be between the requested time and the requested time plus (approximately) 10 milliseconds. The maximum allowed timer interval is 1,000,000 milliseconds (about 16.7 minutes).

If the keyword after the PAUSE command is THREADS, the macro execution will pause until all open windows complete whatever computation they are currently executing. This is a very useful feature when calling UPDATE ALL or LOADLENS using session files.

PIXEL

Purpose:

Turns on a single pixel on the current graphics screen.

Syntax:

PIXEL xcoord, ycoord

Discussion:

This feature is useful for making spot diagrams. See "GRAPHICS" on page 688.

PLOT

Purpose:

The PLOT keyword supports a number of arguments used to simplify the task of creating plots of numeric data.

Syntax:

```
PLOT NEW
PLOT TITLE, string
PLOT TITLEX, string
PLOT TITLEY, string
PLOT BANNER, string
PLOT WINASPECT, type
PLOT COMM1, string
PLOT COMM2, string
PLOT COMM3, string
PLOT COMM4, string
PLOT COMM5, string
PLOT COMM6, string
PLOT RANGEX, minx, maxx
PLOT RANGEY, miny, maxy
PLOT CHECK, x_increment, y_increment
PLOT TICK, x_increment, y_increment
PLOT FORMATX, format_string
PLOT FORMATY, format_string
PLOT DATA, x_array, y_array, number_of_points, color, style, options
PLOT LINE, x1, y1, x2, y2
PLOT LABEL, x, y, angle, size, string
PLOT GO
```

Discussion:

PLOT NEW is used to initialize a new plot. All data related to any previously generated plot is discarded. This should always be the first PLOT command used when making a new graphic.

PLOT TITLE, PLOT TITLEX, and PLOT TITLEY are used to define the main plot title, the x-axis title, and the y-axis title, respectively. The parameter "string" may be a literal string in quotes or may be a string variable.

PLOT BANNER is used to define the banner title. The parameter "string" may be a literal string in quotes or may be a string variable.

PLOT WINASPECT defines the aspect ratio format of the window that will display the window. The valid values for type are the integers 0 through 3, which yield aspect ratios of 4x3, 5x3, 3x4, and 3x5, respectively. If no WINASPECT command is issued a default aspect ratio is used.

PLOT COMMx is used to define up to 6 comments that appear on the bottom of the plot. The parameter "string" may be a literal string in quotes or may be a string variable.

PLOT RANGEX and PLOT RANGEY define the minimum and maximum values of the plot in the x- and y-directions respectively. If no RANGEX or RANGEY command is issued, a default range will be selected.

PLOT CHECK is used to define the size of the "X" marks used to mark data points in the PLOT DATA command. The units are dimensionless fractions relative to the width of the display. If no CHECK command is issued, a default size will be selected.

PLOT TICK is used to define the increments between tick marks on the x- and y-axis, respectively. If no TICK command is issued, a default increment will be selected.

PLOT FORMATX and PLOT FORMATY are used to define the format strings for numeric labels on the x- and y-axis, respectively. The parameter "format_string" may be a literal string in quotes or may be a string variable. The format_string must be a valid C-language format specifier. A common format string is "%M.nf" where M is the total number of spaces in the output string and n is the number of figures after the decimal point. For example,

the value for pi will print as "3.1415" if the format string is "%6.4f". To define an exponential format, the format string is of the form "%M.nE" which would yield a total of M characters and n is the number of figures after the decimal point. The value for pi would print as "3.14E+000" if the format_string was "%9.3E". If no FORMATX or FORMATY is specified a default format is used.

PLOT DATA is used to define a series of data points to be plotted. The variables x_array and y_array MUST be array variables of 1 dimension. For information on defining array variables see "Array variables" on page 653. The number_of_points defines how many points in the array are to be used. The color argument is 0 for black, or 1-16 to select different pen colors. To define pen colors, see "Colors 1-12, Colors 13-24" on page 73. The style argument defines the line style; the current supported values are 0 for a solid line, or 1-4 for various styles of dashed lines. The options value is 0 for no marks indicating the location of data points, 1 for both a line and data point marks, or 2 for no line and data point marks only. Multiple PLOT DATA commands may be issued, and each will create a separate line or curve to be plotted.

PLOT LINE makes a line between the points (x1, y1) and (x2, y2). The units of x and y here are normalized screen coordinates between 0.0 and 1.0. A maximum of 1000 lines may be defined.

PLOT LABEL prints any text string on the plot starting at the coordinates (x, y). The units of x and y here are normalized screen coordinates between 0.0 and 1.0. The angle is the angle from the +x direction in degrees. The size is a scaling constant greater than 0.0 that defines the size of the font. A size of 1.0 would print at the normal font size, while a value of 1.5 would print at 50% larger than normal size. The parameter "string" may be a literal string in quotes or may be a string variable. A maximum of 100 labels may be defined.

PLOT GO uses all settings and data from previous PLOT commands and generates the desired plot, and displays the plot in a window. After the plot is generated all the PLOT data is reset as though PLOT NEW had been executed.

Related Keywords:

PLOT2D

PLOT2D

Purpose:

The PLOT2D keyword supports a number of arguments used to simplify the task of creating plots of numeric data. This keyword makes it easy to generate surface, contour, grey scale, and false color plots.

Syntax:

```
PLOT2D NEW
PLOT2D TITLE, string
PLOT2D COMM1, string
PLOT2D COMM2, string
PLOT2D COMM3, string
PLOT2D COMM4, string
PLOT2D COMM5, string
PLOT2D RANGE, min, max
PLOT2D ASPECT, ratio
PLOT2D WINASPECT, type
PLOT2D DATA, arrayname
PLOT2D ACTIVECURSOR, left, right, bottom, top
PLOT2D DISPLAYTYPE, type
PLOT2D CONTOURINTERVAL, string
PLOT2D SURFACESCALE, scale
PLOT2D LOGPLOT, peak, decades
PLOT2D HIDEADDRESS, flag
PLOT2D CONFIG, configuration
PLOT2D GO
```

Discussion:

PLOT2D NEW is used to initialize a new plot. All data related to any previously generated plot2d is discarded. This should always be the first PLOT2D command used when making a new graphic.

PLOT2D TITLE is used to define the main plot title. The parameter "string" may be a literal string in quotes or may be a string variable.

PLOT2D COMMX is used to define up to 5 comments that appear on the bottom of the plot. The parameter "string" may be a literal string in quotes or may be a string variable.

PLOT2D RANGE define the minimum and maximum values of the plot. If no RANGE command is issued, a default range will be selected. The range should span the data values, or the data will be truncated to fit the range. The surface and contour plots ignore this command.

PLOT2D ASPECT defines the ratio of the x-width of the plot as compared to the y-width. The surface plots ignore this command.

PLOT2D WINASPECT defines the aspect ratio format of the window that will display the window. The valid values for type are the integers 0 through 3, which yield aspect ratios of 4x3, 5x3, 3x4, and 3x5, respectively. If no WINASPECT command is issued a default aspect ratio is used.

PLOT2D DATA is used to define the array variable containing the data to be plotted. The variable arrayname MUST be an array variables of 2 dimensions. For information on defining array variables see "Array variables" on page 653. The number of points in the x and y directions is defined by the array dimensions. Any number of points is allowed, however, the minimum number of points in either direction is 5. The data will be assuming the two dimensions form a rectangular grid with uniform spacing along each dimension, although an overall aspect ratio for the plot may be defined using PLOT2D ASPECT. Only one PLOT2D DATA command may be issued.

PLOT2D ACTIVECURSOR is used to define the left, right, bottom, and top boundaries of the plotted data for display on contour, grey scale, and false color plots.

PLOT2D DISPLAYTYPE is used to select the type of plot generated. The value for type should be an integer between 1 and 6, inclusive, for surface, contour, grey scale, inverse grey scale, false color, and inverse false color, respectively.

PLOT2D CONTOURINTERVAL is used to define the contour format string. Only contour plots use this command. For details on the contour interval format string, see "The Contour Format String" on page 151.

PLOT2D SURFACESCALE is used to define an overall vertical scaling of the surface plot. Only surface plots use this command. The default scaling value is 0.5 arbitrary units.

PLOT2D LOGPLOT is used to generate a log scale display. Only grey scale, inverse grey scale, false color, and inverse false color plots use this command. The peak value is an integer indicating the maximum power of 10 to plot. For example, is peak is 3, then the maximum value plotted will be 1E+03. The decades value is an integer indicating the number of log decades below the peak to plot. If peak is 3, and decades is 5, then the plot will span the range from 1.0E+03 to 1.0E-02. Note that the use of LOGPLOT does not actually take the logarithm of the input array data. The macro itself must take the log base 10 of the values in the array before calling PLOT2D GO.

PLOT2D HIDEADDRESS is used to hide or show the address data. If the flag value is zero, the address is shown, otherwise the address is not shown. See also "Address" on page 68.

PLOT2D CONFIG is used to define the configuration number shown if the address box is displayed and the address box displays information about the current configuration. If configuration is 0, the address box will show the current configuration number. If the configuration is an integer between 1 and the number of configurations, inclusive, the specific configuration number will be shown. If configuration is an integer less than 0, then "All" is shown for the configuration number.

PLOT2D GO uses all settings and data from previous PLOT2D commands and generates the desired plot, and displays the plot in a window. After the plot is generated all the PLOT2D data is reset as though PLOT2D NEW had been executed.

Related Keywords:

PLOT

POLDEFINE

Purpose:

Defines the input polarization state for subsequent POLTRACE calls.

Syntax:

POLDEFINE Jx, Jy, PhaX, PhaY

Discussion:

The POLDEFINE keyword is used to define the input polarization state for subsequent polarization ray tracing. POLDEFINE requires the Jx and Jy electric field magnitudes, as well as the X and Y phase angles in degrees. For more information on the meaning of Jx and Jy, see "Defining the initial polarization" on page 622. The input values are automatically normalized to have unity magnitude. The default values are 0, 1, 0, and 0, respectively. Once the polarization state is defined, it remains the same until changed.

Example:

```
POLDEFINE 2.0, 2.0, 45.0, -66.0
```

Related Keywords:

POLTRACE

POLTRACE

Purpose:

Calls the ZEMAX polarization ray tracing routines to trace a particular ray through the current system.

Syntax:

```
POLTRACE Hx, Hy, Px, Py, wavelength, vec, surf
```

Discussion:

The expressions Hx and Hy must evaluate to values between -1 and 1, and represent the normalized object coordinates. The pupil coordinates are specified by the expressions Px and Py, which also must be between -1 and 1. For more information about normalized coordinates, see the chapter "Conventions and Definitions" under "Normalized field and pupil coordinates". The wavelength expression must evaluate to an integer between 1 and the maximum number of defined wavelengths. The vec expression must evaluate to a number between 1 and 4, inclusive. The surf expression must evaluate to an integer between 1 and the number of surfaces, inclusive.

The input polarization state of the ray is defined by the POLDEFINE keyword.

Once the ray is traced, the polarization data for the ray is placed in the vector variable specified by the vec expression. For example, if the command "POLTRACE Hx, Hy, Px, Py, w, 2, n" is issued, the data will be stored in VEC2. The data is stored in the following format, where the first number in each line refers to the array position:

```
0:  n, the number of data entries in the vector
1:  The ray intensity after the surface
2:  E-Field X component, real
3:  E-Field Y component, real
4:  E-Field Z component, real
5:  E-Field X component, imaginary
6:  E-Field Y component, imaginary
7:  E-Field Z component, imaginary
8:  S-Polarization field amplitude reflection, real
9:  S-Polarization field amplitude reflection, imaginary
10: S-Polarization field amplitude transmission, real
11: S-Polarization field amplitude transmission, imaginary
12: P-Polarization field amplitude reflection, real
13: P-Polarization field amplitude reflection, imaginary
14: P-Polarization field amplitude transmission, real
15: P-Polarization field amplitude transmission, imaginary
16: E-Field X direction phase Px
17: E-Field Y direction phase Py
18: E-Field Z direction phase Pz
19: Major axis length of polarization ellipse
20: Minor axis length of polarization ellipse
21: Angle of polarization ellipse in radians
22: The surface number at which the ray was vignetted or zero if not vignetted
23: S-Polarization ray amplitude reflection, real
24: S-Polarization ray amplitude reflection, imaginary
25: S-Polarization ray amplitude transmission, real
26: S-Polarization ray amplitude transmission, imaginary
27: P-Polarization ray amplitude reflection, real
28: P-Polarization ray amplitude reflection, imaginary
```

29: P-Polarization ray amplitude transmission, real
30: P-Polarization ray amplitude transmission, imaginary

If the value in array position 0 is 0, then an error occurred and the polarization data is invalid. This may occur if the specified ray cannot be traced. See the RAYTRACE command to extract extended error information.

Example:

```
POLDEFINE 0, 1, 0, 0  
POLTRACE 0, 1, 0, 0, pwav(), 1, nsur()  
PRINT "Transmission of chief ray at primary wavelength is ", vecl(1)
```

Related Keywords:

POLDEFINE, RAYTRACE

POP

Purpose:

Computes the Physical Optics Propagation (POP) of a beam through the optical system and saves the surface by surface results to ZBF files. For a description of the POP feature see "Physical Optics Propagation" on page 225. For information on the ZBF file format, see "ZEMAX Beam File (ZBF) binary format" on page 640.

Syntax:

```
POP outfilename, lastsurface
```

Discussion:

This keyword requires the name of the output ZBF file, and an expression that evaluates to the last surface to propagate to. The filename must be enclosed in quotes if any blank or other special characters are used. The created ZBF files will be placed in the <pop> folder. No paths should be provided with the file names.

The settings for the POP feature will be those settings previously saved for the current lens. To make adjustments to the settings, open a POP window, choose the appropriate settings, then press "Save". All subsequent calls to POP within ZPL will use the saved settings. The exceptions are the output file name, which is specified as the first argument after the POP keyword, and the last surface number, which is optionally specified as the second argument after the POP keyword.

Example:

```
POP "pop_output.ZBF", 12
```

PRINT

Purpose:

Print is used to output constant text and variable data to either the screen or a file, depending upon the current status of the keyword OUTPUT.

Syntax:

```
PRINT  
PRINT X  
PRINT "The value of x is ", x  
PRINT " x = ", x, " x + y = ", x + y
```

Discussion:

PRINT alone will print a blank line. PRINT with a list of text arguments and expressions will print each text string (enclosed in double quotes) and the numeric value of each expression. PRINT uses the numerical output format specified by FORMAT. If the last item in the list is followed by a comma, PRINT will not end the line with a carriage return.

Example:

```
X = 3  
PRINT "X equals ", x
```

Related Keywords:

REWIND

PRINTFILE

Purpose:

Prints a text file.

Syntax:

```
PRINTFILE filename
```

Discussion:

The filename must be a valid file name. The file must be a text file (as would be created by OUTPUT and PRINT commands in ZPL) and must be in the current folder. PRINTFILE also closes the file if no CLOSE command has been executed.

Example:

```
OUTPUT "test.txt"  
PRINT "Print this to the printer."  
PRINTFILE "test.txt"
```

Related Keywords:

OPEN, OUTPUT, CLOSE, PRINT, PRINTFILE

PRINTWINDOW

Purpose:

Prints any open graphic or text window.

Syntax:

```
PRINTWINDOW winnum
```

Discussion:

The winnum value may be either an integer or an expression that evaluates to an integer. The integer winnum corresponds to the window number that should be printed. ZEMAX numbers windows sequentially as they are opened, starting with 1. Any closed windows are deleted from the window list, without renumbering the windows which remain. Any windows opened after another window has been closed will use the lowest window number available.

Example:

```
PRINTWINDOW 5
```

PWAV

Purpose:

Sets the primary wavelength.

Syntax:

```
PWAV n
```

Discussion:

The expression n is evaluated and the primary wavelength is set to the specified number. The UPDATE command must be issued before the new data takes effect.

Example:

```
PWAV 1
```

Related Functions:

WAVL, WWGT, PWAV

QUICKFOCUS

Purpose:

Adjusts the thickness of the surface prior to the image surface to minimize the specified criteria.

Syntax:

```
QUICKFOCUS mode, centroid
```

Discussion:

The expression for mode should evaluate to 0, 1, 2, or 3 for RMS spot radius, spot x, spot y, or wavefront OPD. The expression for centroid should evaluate to 0 or 1 to indicate the RMS should be referenced to the chief ray or image centroid, respectively. The "best" focus is chosen as a wavelength weighted average over all fields.

Example:

```
! Focus at best RMS wavefront to centroid
QUICKFOCUS 3, 1
```

RADI

This command is obsolete. See "SETSURFACEPROPERTY, SURP" on page 725.

RANDOMIZE

Purpose:

RANDOMIZE seeds the random number generator.

Syntax:

```
RANDOMIZE seed
```

Discussion:

If seed evaluates to zero, ZEMAX seeds the random number generator with a value based upon the CPU clock. Otherwise, the value provided is used to seed the random number generator. Using the same seed will reproduce the identical series of random numbers created by the RAND function.

Example:

```
RANDOMIZE
RANDOMIZE 250
```

Related Functions:

RAND

RAYTRACE

Purpose:

Calls the ZEMAX ray tracing routines to trace a particular ray through the current system.

Syntax:

```
RAYTRACE hx, hy, px, py, wavelength
```

Discussion:

The expressions hx and hy must evaluate to values between -1 and 1, and represent the normalized field coordinates. The pupil coordinates are specified by the expressions px and py, which also must be between -1 and 1. For more information about normalized coordinates see "Normalized field coordinates" on page 55. The wavelength expression is optional, defaulting to the primary wavelength, but if supplied must evaluate to an integer between 1 and the maximum number of defined wavelengths.

Once the ray is traced, the ray intercept coordinates and direction cosines may be determined using the ZPL functions RAYX, RAYY, RAYZ, RAYL, RAYM, and RAYN. If an error occurred during ray tracing, the function RAYE (for RAY Error) will return a value other than zero. If RAYE is negative, it indicates that total internal reflection occurred at the surface whose number is the absolute value of the value returned. If RAYE is greater than zero, then the ray missed the surface number returned. Checking RAYE is optional, however, the RAYX, RAYY, ... functions may return invalid data if RAYE is not zero. The functions RANX, RANY, and RANZ return the intercept surface normal direction cosines, and OPDC returns the optical path difference for the ray. The function RAYV returns the surface number at which the ray was vignetted, or it returns zero if the ray was not vignetted. Values returned for surfaces past the surface of vignetting may not be accurate.

Example:

```
PRINT "Tracing the marginal ray at primary wavelength!"
n = NSUR()
RAYTRACE 0, 0, 0, 1
y = RAYY(n)
```



```
PRINT "The ray intercept is ", y
PRINT "Tracing the chief ray at maximum wavelength!"
RAYTRACE 0,1,0,0,NWAV()
y = RAYY(n)
PRINT "The ray intercept is ", y
```

Related Keywords:

RAYTRACEX

RAYTRACEX

Purpose:

Calls the ZEMAX ray tracing routines to trace a particular ray from any starting surface through the current system.

Syntax:

```
RAYTRACEX x, y, z, l, m, n, surf, wavelength
```

Discussion:

The expressions x, y, z, l, m, and n define the input ray position and direction cosines in the local coordinates of the starting surface. The surface expression must evaluate to an integer between 0 and the number of surfaces minus one, inclusive. The wavelength expression is optional, defaulting to the primary wavelength, but if supplied must evaluate to an integer between 1 and the maximum number of defined wavelengths.

If the object has a thickness of infinity, and the surf parameter is zero, then the input coordinates are assumed to be relative to the first surface rather than the object surface; although the ray will still be defined in object space media. Otherwise, ZEMAX uses the specified coordinates without alteration.

Once the ray is traced, the ray intercept coordinates and direction cosines may be determined using the ZPL functions RAYX, RAYY, RAYZ, RAYL, RAYM, and RAYN. Note only data from surface AFTER the "surf" surface will be valid.

If an error occurred during ray tracing, the function RAYE (for RAY Error) will return a value other than zero. If RAYE is negative, it indicates that total internal reflection occurred at the surface whose number is the absolute value of the value returned. If RAYE is greater than zero, then the ray missed the surface number returned. Checking RAYE is optional, however, the RAYX, RAYY, ... functions may return invalid data if RAYE is not zero. The functions RANX, RANY, and RANZ return the intercept surface normal direction cosines, and RAYT returns the optical path length up to the surface for the ray. The function RAYV returns the surface number at which the ray was vignetted, or it returns zero if the ray was not vignetted. Values returned for surfaces past the surface of vignetting may not be accurate.

Example:

```
n = NSUR()
RAYTRACEX 0,1,0,0,0,1,0,NWAV()
y = RAYY(n)
PRINT "The ray intercept is ", y
```

Related Keywords:

RAYTRACE

READ

Purpose:

Reads data from an existing text file opened for reading by the OPEN command.

Syntax:

```
READ x
READ x, y
READ x,y,z,a,b,c,q
```

Discussion:

The file must be already open, see the keyword OPEN for details. Each READ command reads a single line from the file. The first valid data field from this line is placed in the variable first listed. The data from the second field is placed in the second variable listed, if any. Therefore, the number of variables listed in the read command

should match the number of columns in the text file. Numeric data in the file should be delimited by spaces. The data may be in free-form, and is internally promoted to double precision. A maximum of 2000 characters can be read in on any single line. The maximum number of variable arguments is 199; for reading longer lines with more arguments use READNEXT instead. The variables listed must be valid ZPL variable names. READNEXT does not support array variable names as arguments. The workaround is to read the data into a scalar variable and then set the array variable to the scalar variable value in a subsequent line like this:

```
READ x
data(i, j) = x
```

Always CLOSE a file after all the data has been read. See the function EOFF.

Example:

```
PRINT "Reading the double-column file TEST.DAT!"
OPEN "C:\DATA\TEST.DAT"
READ x1, y1
READ x2, y2
READ x3, y3
CLOSE
```

Related Functions:

EOFF

Related Keywords:

OPEN, CLOSE, READNEXT, READSKIP, READSTRING

READNEXT

Purpose:

Reads data from an existing text file opened for reading by the OPEN command.

Syntax:

```
READNEXT x
READNEXT x, y
READNEXT x, y, z, a, b, c, q
```

Discussion:

READNEXT is almost identical to READ. The key difference is READ will read the entire data line from the opened file, up to the newline character, while READNEXT reads only enough characters to fill the number of arguments.

For example, if a data file contains a line with this data:

```
3.0 4.0 5.0
```

The following two READNEXT commands will read the values 3.0, 4.0, and 5.0 for x, y, and z:

```
READNEXT x, y
READNEXT z
```

READNEXT is more useful than READ if the line is very long, or the number of arguments is large. READNEXT does not support array variable names as arguments. The workaround is to read the data into a scalar variable and then set the array variable to the scalar variable value in a subsequent line like this:

```
READNEXT x
data(i, j) = x
```

Example:

```
OPEN "C:\DATA\TEST.DAT"
READNEXT x1, x2
READNEXT x3
CLOSE
```

Related Keywords:

OPEN, CLOSE, READ, READSKIP, READSTRING

READSKIP

Purpose:

Reads any number of unwanted characters from an existing text file opened for reading by the OPEN command.

Syntax:

```
READSKIP n
```

Discussion:

The file must be already open, see the keyword OPEN for details. If n is positive, READSKIP reads off exactly n characters from the file. If n is negative, READSKIP reads to the end of the current line. The read characters are discarded. This function allows faster access to large files by effectively skipping over any number of unwanted characters. Always CLOSE a file after all the data has been read. See the function EOFF.

Example:

```
PRINT "Reading the contents of file TEST.DAT!"
OPEN TEST.DAT
READSKIP 59
READSTRING A$
PRINT A$
CLOSE
```

Related Keywords:

OPEN, CLOSE READ, READNEXT, READSTRING

READSTRING

Purpose:

Reads data from an existing text file opened for reading by the OPEN command.

Syntax:

```
READSTRING A$
```

Discussion:

The file must be already open, see the keyword OPEN for details. Each READ command reads a single line from the file. The entire line read is placed in the variable listed. A maximum of 359 characters can be read in and placed in a single string. The variable listed must be a valid ZPL string variable name, although it does not need to be previously referenced. Always CLOSE a file after all the data has been read. See the function EOFF.

Example:

```
PRINT "Reading the contents of file TEST.DAT!"
OPEN TEST.DAT
READSTRING A$
PRINT A$
CLOSE
```

Related Keywords:

OPEN, CLOSE READ, READNEXT, READSKIP

RELEASE

See "Array variables" on page 653.

RELOADOBJECTS

Purpose:

Reloads NSC objects into the NSC Editor.

Syntax:

```
RELOADOBJECTS surface, object
```

Discussion:

The expression for surface must evaluate to an integer surface number corresponding to the non-sequential component surface. For NSC mode, use 1. The expression for object must evaluate to an integer object number, or zero to reload all objects. This keyword may take a significant amount of time to execute, depending upon the number and type of objects defined.

Example:

```
RELOADOBJECTS 1, 0
```

REM, !, #

Purpose:

REM, !, and # are used to define remarks or comments.

Syntax:

```
REM (text)
! (text)
(valid zpl line) # (text)
```

Discussion:

The exclamation symbol may also be used to indicate a remark. Both the REM command and the "!" symbol are only recognized as remark indicators if they appear at the very beginning of the line. The # symbol may be used anywhere on the line, but will only be interpreted as the beginning of a comment if the # is not inside a string. Everything after the # symbol is ignored.

Example:

```
REM any text can be placed after the REM command.
! any text can also be placed
! after the exclamation symbol.
x = 5 # this syntax allows comments to be placed on the same line as a command.
```

REMOVEVARIABLES

Purpose:

Sets all currently defined variables to fixed status.

Syntax:

```
REMOVEVARIABLES
```

RENAMEFILE

Purpose:

RENAMEFILE is used to rename a file.

Syntax:

```
RENAMEFILE oldfilename, newfilename
```

Discussion:

This keyword requires two file names, defined as literal string expressions in quotes or as string variables. The file oldfilename is renamed newfilename.

Example:

```
RENAMEFILE AFILE$, BFILE$
```

Related Keywords:

COPYFILE
DELETEFILE

RETURN

See GOSUB.

REWIND

Purpose:

REWIND erases the last line printed by the PRINT command, up to the previous end of line. This allows printing a counter or other data over an existing line in the text output file.

Syntax:

```
REWIND
```

Example:

```
PRINT "First line"  
REWIND  
PRINT "New First line"
```

Related Keywords:

PRINT

SAVEARCHIVE

Purpose:

Saves the current lens to a ZEMAX archive (*.ZAR) file.

Syntax:

```
SAVEARCHIVE filename
```

Discussion:

The filename must include the name of the archive file including the file extension. If the filename does not include a complete path to the archive file, the default folder for lenses is assumed. See “Backup To Archive File” on page 65.

SAVEDETECTOR

Purpose:

Saves the data currently on an NSC Detector Rectangle, Detector Color, Detector Polar, or Detector Volume object to a file.

Syntax:

```
SAVEDETECTOR surf, object, filename
```

Discussion:

This keyword requires numeric expressions that specify the surface and object, and a file name with or without a full path. Surf is the surface number of the non-sequential group; use 1 when using non-sequential mode. Object is the object number of the detector object. The filename may include the full path, if no path is provided the path of the current lens file is used. The extension should be DDR, DDC, DDP, or DDV for Detector Rectangle, Color, Polar, and Volume objects, respectively.

Related Keywords:

LOADDETECTOR

SAVELENS

Purpose:

Saves the current lens file.

Syntax:

```
SAVELENS filename, session
```

Discussion:

SAVELENS will save the current lens file to the specified file name. The name of the current lens in memory will also be changed. If the file name is absent, then the lens data is stored in the current file name. If the session argument evaluates to anything other than zero, the session file will also be saved.

Example:

```
SAVELENS  
SAVELENS "NEWCOPY.ZMX"  
SAVELENS NEW$
```

Related Keywords:

LOADLENS

SAVEMERIT

Purpose:

Saves the current merit function to a file.

Syntax:

```
SAVEMERIT filename
```

Discussion:

SAVEMERIT will save the current merit function to a file. If the filename contains the complete path, such as C:\MYDIR\MYLENS.MF, then the specified path will be used. If the path is left off, then the <data>\MeritFunction folder will be used (see “Folders” on page 68). See also LOADMERIT.

SAVETOLERANCE

Purpose:

Saves the current tolerances to a file.

Syntax:

```
SAVETOLERANCE "filename"  
SAVETOLERANCE file$
```

Discussion:

If the filename contains the complete path, such as C:\MYDIR\MYLENS.TOL, then the specified file will be created. If the path is left off, then the <data>\Tolerance folder will be used (see “Folders” on page 68). See also LOADTOLERANCE.

SAVEWINDOW

Purpose:

Saves the text from any text window to a file.

Syntax:

```
SAVEWINDOW winnum, filename
```

Discussion:

The winnum value may be either an integer or an expression that evaluates to an integer. The integer winnum corresponds to the text window number that should be saved to a file. ZEMAX numbers windows sequentially as they are opened, starting with 1. Any closed windows are deleted from the window list, without renumbering the windows which remain. Any windows opened after another window has been closed will use the lowest window number available.

Example:

```
SAVEWINDOW 1, "C:\TEMP\TEXTFILE.TXT"  
SAVEWINDOW 3, A$
```

SCATTER

Purpose:

Used to control whether sequential surface scattering is done while tracing rays.

Syntax:

```
SCATTER ON  
SCATTER OFF
```

Discussion:

The default condition at the start of a macro is SCATTER OFF; and all rays will be traced deterministically. If SCATTER ON is executed, then sequential surface scattering will be enabled for all subsequent RAYTRACE commands.

SDIA

This command is obsolete. See “SETSURFACEPROPERTY, SURP” on page 725.

SETAIM

Purpose:

Sets the state of the ray aiming function.

Syntax:

```
SETAIM state
```

Discussion:

This keyword requires one numeric expression that must evaluate to either 0 or 1. The state is a code which is 0 for ray aiming off and 1 for ray aiming on.

Example:

```
SETAIM 1
```

Related Keywords:

SETAIMDATA

SETAIMDATA

Purpose:

Sets various data for the ray aiming function. Syntax:

```
SETAIMDATA code, value
```

Discussion:

The code values are used as follows:

Code	Property
1	Sets "Use Ray Aiming Cache" to true if value is 1, or false if value is 0.
2	Sets "Robust Ray Aiming" to true if value is 1, or false if value is 0.
3	Sets "Scale Pupil Shift Factors by Field" to true if value is 1, or false if value is 0.
4, 5, 6	Sets the value of the x, y, and z pupil shift, respectively.
7, 8	Sets the value of the x and y pupil compress, respectively.

Example:

```
SETAIMDATA 5, 0.34
```

Related Keywords:

SETAIM

SETAPODIZATION

This command is obsolete. See "SETSYSTEMPROPERTY, SYSP" on page 728.

SETCONFIG

Purpose:

Sets the current configuration for multi-configuration (zoom) systems.

Syntax:

```
SETCONFIG config
```

Discussion:

The expression config must evaluate to an integer between 1 and the maximum number of configurations.

Example:

```
SETCONFIG 4
```

Related Functions:

CONF, NCON

SETDETECTOR

Purpose:

Sets the coherent or incoherent detector data for any pixel on a detector rectangle or detector surface object.

Syntax:

```
SETDETECTOR surf, object, pixel, datatype, value
```

Discussion:

This keyword requires numeric expressions that specify the surface, object, pixel, and data type (all integers) and the new detector value. Surf is the surface number of the non-sequential group; use 1 when using non-sequential mode. Object is the object number of the detector. The object must be either a detector surface or a detector rectangle. Pixel is a number between 1 and the number of pixels the object supports. Datatype should be 0 for incoherent intensity, 1 for incoherent intensity in angle space, 2 for coherent real part, 3 for coherent imaginary part, and 4 for coherent amplitude. For more information on the use and meaning of these values see “Detector Rectangle object” on page 535. The units for incoherent intensity are source units (see “Source Units” on page 105). The units for coherent amplitude are the square root of source units.

Example:

```
SETDETECTOR 1, 5, 12, 0, 1.0056
```

Related Functions:

NSDD

SETMCOPERAND

Purpose:

Sets any row or configuration of the Multi-Configuration Editor to any numeric value.

Syntax:

```
SETMCOPERAND row, config, value, datatype
```

Discussion:

This keyword requires numeric expressions that evaluate to integers specifying the row and configuration of the Multi-Configuration Editor.

If the config number is 0, then the value is interpreted as follows:

datatype = 0, value is a string literal or variable that specifies the name of the operand.

datatype = 1, 2, or 3, value is the number 1, 2, or 3 value used as part of the multi-configuration operand definition. See the description of the multi-configuration operand numbers defined in “Summary of multi-configuration operands” on page 579.

If the config number corresponds to a defined configuration then the value is interpreted as follows:

datatype = 0, value is the value of the operand.

datatype = 1, value is the pickup offset of the operand.

datatype = 2, value is the pickup scale of the operand.

datatype = 3, value is the status of the operand, 0 for fixed, 1 for variable, 2 for pickup, 3 for thermal pickup.

datatype = 4, value is the pickup configuration number.

datatype = 5, value is the pickup row number.

Example:

```
SETMCOPERAND 3, 4, somevalue, 0  
SETMCOPERAND 1, 0, "THIC", 0
```

Related Functions:

MCOP

SETNSCPARAMETER

Purpose:

Sets the parameter values of any object in the NSC editor.

Syntax:

```
SETNSCPARAMETER surface, object, parameter, value
```

Discussion:

This keyword requires 3 numeric expressions that evaluate to integers specifying the non-sequential components surface number, the object number, and the parameter number. The fourth argument is the new value for the specified parameter.

Example:

```
SETNSCPARAMETER 4, 2, 15, newp15value
```

Related Functions:

NPOS, NPAR

Related Keywords:

INSERTOBJECT, SETNSCPOSITION, SETNSCPROPERTY

SETNSCPOSITION

Purpose:

Sets the x, y, z or tilt x, tilt y, tilt z position of any object in the NSC editor.

Syntax:

```
SETNSCPOSITION surface, object, code, value
```

Discussion:

This keyword requires 3 numeric expressions that evaluate to integers specifying the non-sequential components surface number, the object number, and a code. The code is 1 through 6 for x, y, z, tilt x, tilt y, tilt z, respectively. The fourth argument is the new value for the specified position.

Example:

```
SETNSCPOSITION 4, 2, 2, newyvalue
```

Related Functions:

NPOS, NPAR

Related Keywords:

INSERTOBJECT, SETNSCPARAMETER, SETNSCPROPERTY

SETNSCPROPERTY

Purpose:

Sets properties of NSC objects.

Syntax:

```
SETNSCPROPERTY surface, object, code, face, value
```

Discussion:

This keyword requires 4 numeric expressions that evaluate to integers specifying the non-sequential components surface number (use 1 if the program mode is non-sequential), the object number, a code which specifies what property of the object is being modified, and the face number (use 0 if not applicable to the property being set). The fifth argument is the new value for the specified property, and it may be either text in quotes, a string variable, or a numeric expression. The code is as follows:

Code	Property
The following codes set values on the NSC Editor.	
1	Sets the object comment.
2	Sets the reference object number.

Code	Property
3	Sets the "inside of" object number.
4	Sets the object material.
The following codes set values on the Type tab of the Object Properties dialog.	
0	Sets the object type. The value should be the name of the object, such as "NSC_SLEN" for the standard lens. The names for each object type are listed in the Prescription Report for each object type in the NSC editor. All NSC object names start with "NSC_".
13	Sets User Defined Aperture, use 1 for checked, 0 for unchecked.
14	Sets the User Defined Aperture file name.
15	Sets the "Use Global XYZ Rotation Order" checkbox, use 1 for checked, 0 for unchecked.
16	Sets the "Rays Ignore This Object" checkbox, use 1 for checked, 0 for unchecked.
17	Sets the "Object Is A Detector" checkbox, use 1 for checked, 0 for unchecked.
18	Sets the "Consider Objects" list. The argument should be a string listing the object numbers to consider delimited by spaces, such as "2 5 14".
19	Sets the "Ignore Objects" list. The argument should be a string listing the object numbers to ignore delimited by spaces, such as "1 3 7".
20	Sets the "Use Pixel Interpolation" checkbox, use 1 for checked, 0 for unchecked.
The following codes set values on the Coat/Scatter tab of the Object Properties dialog.	
5	Sets the coating name for the specified face.
6	Sets the profile name for the specified face.
7	Sets the scatter mode for the specified face, 0 = none, 1 = Lambertian, 2 = Gaussian, 3 = ABg, 4 = User Defined.
8	Sets the scatter fraction for the specified face.
9	Sets the number of scatter rays for the specified face.
10	Sets the Gaussian sigma for the specified face.
11	Sets the reflect ABg data name for the specified face.
12	Sets the transmit ABg data name for the specified face.
27	Sets the name of the user defined scattering DLL.
28	Sets the name of the user defined scattering data file.
21-26	Sets parameter values on the user defined scattering DLL.
29	Sets the "Face Is" property for the specified face. Use 0 for "Object Default", 1 for "Reflective", and 2 for "Absorbing".
The following codes set values on the Bulk Scattering tab of the Object Properties dialog.	
81	Sets the "Model" value on the bulk scattering tab. Use 0 for "No Bulk Scattering", 1 for "Angle Scattering", and 2 for "DLL Defined Scattering".
82	Sets the mean free path to use for bulk scattering.
83	Sets the angle to use for bulk scattering.

Code	Property
84	Sets the name of the DLL to use for bulk scattering.
85	Sets the parameter value to pass to the DLL, where the face value is used to specify which parameter is being defined. The first parameter is 1, the second is 2, etc.
86	Sets the wavelength shift string.
The following codes set values on the Diffraction tab of the Object Properties dialog.	
91	Sets the "Split" value on the diffraction tab. Use 0 for "Don't Split By Order", 1 for "Split By Table Below", and 2 for "Split By DLL Function".
92	Sets the name of the DLL to use for diffraction splitting.
93	Sets the Start Order value.
94	Sets the Stop Order value.
95	Sets the parameter values on the diffraction tab. These are the parameters passed to the diffraction splitting DLL as well as the order efficiency values used by the "split by table below" option. The face value is used to specify which parameter is being defined. The first parameter is 1, the second is 2, etc.
The following codes set values on the Sources tab of the Object Properties dialog.	
101	Sets the source object random polarization. Use 1 for checked, 0 for unchecked.
102	Sets the source object reverse rays option. Use 1 for checked, 0 for unchecked.
103	Sets the source object Jones X value.
104	Sets the source object Jones Y value.
105	Sets the source object Phase X value.
106	Sets the source object Phase Y value.
107	Sets the source object initial phase in degrees value.
108	Sets the source object coherence length value.
109	Sets the source object pre-propagation value.
110	Sets the source object sampling method; 0 for random, 1 for Sobol sampling.
111	Sets the source object bulk scatter method; 0 for many, 1 for once, 2 for never.
112	Sets the array mode; 0 for none, 1 for rectangular, 2 for circular, 3 for hexapolar, and 4 for hexagonal.
113	Sets the source color mode. For a complete list of the available modes, see "Defining the color and spectral content of sources" on page 563. The source color modes are numbered starting with 0 for the System Wavelengths, and then from 1 through the last model listed in the dialog box control.
114-116	Sets the number of spectrum steps, start wavelength, and end wavelength, respectively.
117	Sets the name of the spectrum file.
161-162	Sets the array mode integer arguments 1 and 2.
165-166	Sets the array mode double precision arguments 1 and 2.
181-183	Sets the source color mode arguments, for example, the XYZ values of the Tristimulus.

Code	Property
The following codes set values on the Grin tab of the Object Properties dialog.	
121	Sets the "Use DLL Defined Grin Media" checkbox. Use 1 for checked, 0 for unchecked.
122	Sets the Maximum Step Size value.
123	Sets the DLL name.
124	Sets the Grin DLL parameters. These are the parameters passed to the DLL. The face value is used to specify which parameter is being defined. The first parameter is 1, the second is 2, etc.
The following codes set values on the Draw tab of the Object Properties dialog.	
141	Sets the do not draw object checkbox. Use 1 for checked, 0 for unchecked.
142	Sets the object opacity. Use 0 for 100%, 1 for 90%, 2 for 80%, etc.
The following codes set values on the Scatter To tab of the Object Properties dialog.	
151	Sets the scatter to method. Use 0 for scatter to list, and 1 for importance sampling.
152	Sets the Importance Sampling target data. The argument should be a string listing the ray number, the object number, the size, and the limit value, all separated by spaces Here is a sample syntax to set the Importance Sampling data for ray 3, object 6, size 3.5, and limit 0.6: "3 6 3.5 0.6".
153	Sets the "Scatter To List" values. The argument should be a string listing the object numbers to scatter to delimited by spaces, such as "4 6 19".
The following codes set values on the Birefringence tab of the Object Properties dialog.	
171	Sets the Birefringent Media checkbox. Use 0 for unchecked, and 1 for checked.
172	Sets the Birefringent Media Mode. Use 0 for Trace ordinary and extraordinary rays, 1 for Trace only ordinary rays, 2 for Trace only extraordinary rays, and 3 for Waveplate mode.
173	Sets the Birefringent Media Reflections status. Use 0 for Trace reflected and refracted rays, 1 for Trace only refracted rays, and 2 for Trace only reflected rays.
174-176	Sets the Ax, Ay, and Az values.
177	Sets the Axis Length.
The following codes do not set values, but are included here to return values for the function NPRO.	
200	Used by function NPRO to determine the index of refraction of an object. The syntax is NPRO(surface, object, 200, wavenumber)
201-203	Used by function NPRO to determine the nd (201), vd (202), and dpgf (203) parameters of an object using a model glass. The syntax is NPRO(surface, object, 201, 0)

Other code values will be added in the future to set other properties of NSC objects as requested by users.

Example:

```
SETNSCPROPERTY 1, 2, 0, 0, "NSC_SLEN"
```

Related Functions:

NPOS, NPAR, NPRO

Related Keywords:

INSERTOBJECT, SETNSCPARAMETER, SETNSCPOSITION

SETOPERAND

Purpose:

Sets any row or column of the Merit Function Editor to any numeric value.

Syntax:

```
SETOPERAND row, col, value
```

Discussion:

This keyword requires numeric expressions that evaluate to integers specifying the row and column of the Merit Function Editor. The col integer is 1 for operand type, 2 for int1, 3 for int2, 4-7 for data1-data4, 8 for target, and 9 for weight. To set the comment string associated with an operand use col 10. To set the operand type, use a col value of 1 and the integer operand returned by the function ONUM. An alternate method for setting the operand type is to use a col value of 11 with the string name of the operand type, such as EFFL or DIST. If col is 10 or 11, value should be a string constant or variable. The col integer is 12 for data5 and 13 for data6. Note the value and percent contribution columns cannot be set but must be computed.

Example:

```
SETOPERAND 1, 8, tarvalue  
SETOPERAND 3, 11, "EFFL"  
SETOPERAND 5, 10, "Operand Number 5"
```

Related Functions:

MFCN, OPER, ONUM

SETSTDD

This command is obsolete. See "SETSURFACEPROPERTY, SURP" on page 725.

SETSURFACEPROPERTY, SURP

Purpose:

Sets properties of surfaces.

Syntax:

```
SETSURFACEPROPERTY surface, code, value1, value2  
SURP surface, code, value1, value2
```

Discussion:

Surface is an expression that evaluates to an integer specifying the surface number. The code may either be an expression that evaluates to an integer or a mnemonic which specifies what property of the surface is being modified. The third and fourth arguments are the new values for the specified property, and they may be either text in quotes, a string variable, or a numeric expression, depending upon the code. For most codes, the property value being modified is defined by the value1 argument. A few operands require both a value1 and a value2, as described in the table below.

If the property being modified is under control of the Multi-Configuration Editor, then the multi-configuration data for the current configuration only will also be modified to reflect the changed property.

To set a surface as the stop surface, see "STOPSURF" on page 737.

SURP is a shorthand for SETSURFACEPROPERTY and is functionally identical.

Code	Property
Basic Surface Data. See "Lens Data" on page 75.	
0 or TYPE	Surface type. The value should be the name of the object, such as "STANDARD" for the standard surface. The names for each surface type are listed in the Prescription Report in the Surface Data Summary for each surface type currently in the Lens Data Editor. To change the surface type to a user defined surface, first set the dll name using code 9 (SDLL), then set the new surface type to USERSURF. See also Code 17.

Code	Property
1 or COMM	Comment.
2 or CURV	Curvature (not radius) in inverse lens units. Use zero for an infinite radius.
3 or THIC	Thickness in lens units.
4 or GLAS	Glass name. See also Code 18.
5 or CONI	Conic constant.
6 or SDIA	Semi-diameter. If the value is zero or positive, the semi-diameter solve is set to "Fixed". If the value is negative, the semi-diameter solve is set to "Automatic" and the semi-diameter will be computed with the next UPDATE keyword.
7 or TCE	Thermal coefficient of expansion.
8 or COAT	Coating name. Use a blank string for value1 to remove the coating.
9 or SDLL	User defined surface DLL name.
10 or PARM	Parameter value. Value1 is the new value. Value2 is the parameter number.
11 or EDVA	Extra Data value. Value1 is the new value. Value2 is the extra data number.
12	Surface color, Use 0 for default.
13	Surface opacity.
14	Row color.
15	Surface cannot be hyperhemispheric. Use 1 to avoid surface being hyperhemispheric.
16	Ignore surface. Use 1 to ignore surface, 0 to not ignore surface.
17 or CODE	The integer code for the surface type. The integer code is an alternative to the surface name used by code 0. For full details see the documentation for code 0 above.
18 or GLAN	Glass number. See also Code 4.
Surface aperture data. See "Surface properties aperture tab" on page 79.	
20 or ATYP	Surface aperture type code.
21 or APP1	Surface aperture parameter 1.
22 or APP2	Surface aperture parameter 2.
23 or APDX	Surface aperture decenter x.
24 or APDY	Surface aperture decenter y.
25 or UDA	User Defined Aperture (UDA) file name.
26 or APPU	Surface aperture pick up from surface number. Use 0 for no pickup.
Physical Optics Propagation Settings. See "Surface specific settings" on page 643.	
30	Physical Optics setting "Use Rays To Propagate To Next Surface". Use 1 for true, 0 for false.
31	Physical Optics setting "Do Not Rescale Beam Size Using Ray Data". Use 1 for true, 0 for false.
32	Physical Optics setting "Use Angular Spectrum Propagator". Use 1 for true, 0 for false.
33	Physical Optics setting "Draw ZBF On Shaded Model". Use 1 for true, 0 for false.

Code	Property
34	Physical Optics setting "Recompute Pilot Beam Parameters". Use 1 for true, 0 for false.
35	Physical Optics setting "Resample After Refraction". Use 1 for true, 0 for false.
36	Physical Optics setting "Auto Resample". Use 1 for true, 0 for false.
37	Physical Optics setting "New X Sampling". Use 1 for 32, 2 for 64, etc..
38	Physical Optics setting "New Y Sampling". Use 1 for 32, 2 for 64, etc.
39	Physical Optics setting "New X-Width". New total x direction width of array.
40	Physical Optics setting "New Y-Width". New total y direction width of array.
41	Physical Optics setting "Output Pilot Radius". Use 0 for best fit, 1 for shorter, 2 for longer, 3 for x, 4 for y, 5 for plane, 6 for user.
42, 43	Physical Optics setting "X-Radius" and "Y-Radius", respectively.
Coating Settings. See "Surface coating tab" on page 86. See also code 8 above.	
50	Use Layer Multipliers and Index Offsets. Use 1 for true, 0 for false.
51	Layer Multiplier value. Value1 is the new value. Value2 is the layer number.
52	Layer Multiplier status. Value 1 is the status, use 0 for fixed, 1 for variable, or n+1 for pickup from layer n. Value2 is the layer number.
53	Layer Index Offset value. Value1 is the new value. Value2 is the layer number.
54	Layer Index Offset status. Value 1 is the status, use 0 for fixed, 1 for variable, or n+1 for pickup from layer n. Value2 is the layer number.
55	Layer Extinction Offset value. Value1 is the new value. Value2 is the layer number.
56	Layer Extinction Offset status. Value 1 is the status, use 0 for fixed, 1 for variable, or n+1 for pickup from layer n. Value2 is the layer number.
Surface Tilt and Decenter Data. See "Surface tilt/decenter tab" on page 84.	
60 or BOR	Before tilt and decenter order. Use 0 for dec/tilt, 1 for tilt/dec.
61 or BDX	Before decenter x.
62 or BDY	Before decenter y.
63 or BTX	Before tilt about x.
64 or BTY	Before tilt about y.
65 or BTZ	Before tilt about z.
66 or APU	After pick up status: 0 for explicit, 1/2 for pickup/reverse current surface, 3/4 for pickup/reverse current surface minus 1, 5/6 for pickup/reverse current surface minus 2, etc...
70 or AOR	After tilt and decenter order. Use 0 for dec/tilt, 1 for tilt/dec.
71 or ADX	After decenter x.
72 or ADY	After decenter y.
73 or ATX	After tilt about x.
74 or ATY	After tilt about y.

Code	Property
75 or ATZ	After tilt about z.
76	Coordinate Return status. Valid only on Coordinate Break surfaces. Use 0 for None, 1 for Orientation Only, 2 for Orientation XY, and 3 for Orientation XYZ.
77	Coordinate Return To Surface. Valid only on Coordinate Break surfaces.
Surface scatter data. See "Surface properties scattering tab" on page 84.	
80	Sets the scatter code; 0 for none, 1 for Lambertian, 2 for Gaussian, 3 for ABg, and 4 for DLL.
81	Sets the scatter fraction, should be between 0.0 and 1.0.
82	Sets the Gaussian scatter sigma.
83	Sets the ABg name.
84	Sets the name of the user defined scattering DLL. To set the parameters see Code 181.
85	Sets the name of the data file used by the user defined scattering DLL.
Surface draw data. See "Surface properties draw tab" on page 78.	
90	Sets the "Hide Rays To This Surface" checkbox status: 0 for off, 1 for on.
91	Sets the "Skip Rays To This Surface" checkbox status: 0 for off, 1 for on.
92	Sets the "Do Not Draw This Surface" checkbox status: 0 for off, 1 for on.
93	Sets the "Do Not Draw Edges From This Surface" checkbox status: 0 for off, 1 for on.
96	Sets the "Draw Edges As" status: 0 for squared, 1 for tapered, 2 for flat.
97	Sets "Mirror Substrate" status: 0 for none, 1 for flat, 2 for curved.
98	Sets the mirror substrate thickness value.
User defined surface scatter DLL parameters. See "Surface properties scattering tab" on page 84.	
181-186	Sets the user defined scatter DLL parameters 1-6.

Changes to surface properties usually require that a subsequent UPDATE (see "UPDATE" on page 738) be performed to update pupil positions, solves, and other data required for correct ray tracing. The UPDATE is not performed automatically because it is faster to execute all SETSURFACEPROPERTY keywords and then execute just one UPDATE. The function SPRO uses a very similar syntax and identical code values to "get" rather than "set" these same values.

Example:

```
! Set the glass type on surface 7 to BK7
SETSURFACEPROPERTY 7, GLAS, "BK7"
! Set the thickness of surface 2 to the thickness of surface 1
SETSURFACEPROPERTY 2, THIC, THIC(1)
! Set the value of parameter 4 on surface 11 to 7.3
SURP 11, PARM, 7.3, 4
```

Related Functions:

SPRO

Related Keywords:

SETSYSTEMPROPERTY, UPDATE

SETSYSTEMPROPERTY, SYSP

Purpose:

Sets properties of the system, such as system aperture, field, wavelength, and other data.

Syntax:

```
SETSYSTEMPROPERTY code, value1, value2
SYSP code, value1, value2
```

Discussion:

This keyword requires a numeric expression that evaluates to an integer code which specifies what property is being modified. The second and third arguments are the new values for the specified property, and they may be either text in quotes, a string variable, or a numeric expression, depending upon the code. For most codes, the property value being modified is defined by the value1 argument. A few operands require both a value1 and a value2, as described in the table below.

If the property being modified is under control of the Multi-Configuration Editor, then the multi-configuration data for the current configuration only will also be modified to reflect the changed property.

SYSP is a shorthand for SETSYSTEMPROPERTY and is functionally identical.

Code	Property
4	Adjust Index Data To Environment. Use 0 for off, 1 for on. See "Adjust Index Data To Environment" on page 108.
10	Aperture Type code. See "Aperture Type" on page 101 for code values and discussion.
11	Aperture Value. See "Aperture Value" on page 102.
12	Apodization Type code. Use 0 for uniform, 1 for Gaussian, 2 for cosine cubed. See "Apodization Type" on page 102.
13	Apodization Factor. See "Apodization Factor" on page 103.
14	Telecentric Object Space. Use 0 for off, 1 for on. See "Telecentric Object Space" on page 103.
15	Iterate Solves When Updating. Use 0 for off, 1 for on. See "Iterate Solves When Updating" on page 104.
16	Lens Title. See "Lens Title" on page 106.
17	Lens Notes. See "Notes" on page 106.
18	Afocal Image Space. Use 0 for off, 1 for on. See "Afocal Image Space" on page 103.
21	Global coordinate reference surface. See "Global Coordinate Reference Surface" on page 111.
23	Glass catalog list. Use a string or string variable with the glass catalog name, such as "SCHOTT". To specify multiple catalogs use a single string or string variable containing names separated by spaces, such as "SCHOTT HOYA OHARA".
24	System Temperature in degrees Celsius. See "Temperature in degrees C" on page 108.
25	System Pressure in atmospheres. See "Pressure in ATM" on page 108.
26	Reference OPD method. Use 0 for absolute, 1 for infinity, 2 for exit pupil, and 3 for absolute 2. See "Reference OPD" on page 110.
30	Lens Units code. Use 0 for mm, 1 for cm, 2 for inches, or 3 for Meters. Changing lens units does not scale or convert the lens data in any way, it only changes how the lens prescription data is interpreted. See "Lens Units" on page 104.
31	Source Units Prefix. Use 0 for Femto, 1 for Pico, 2 for Nano, 3 for Micro, 4 for Milli, 5 for None, 6 for Kilo, 7 for Mega, 8 for Giga, and 9 for Tera. See "Source Units" on page 105.
32	Source Units. Use 0 for Watts, 1 for Lumens, and 2 for Joules. See "Source Units" on page 105.

Code	Property
33	Analysis Units Prefix. Use 0 for Femto, 1 for Pico, 2 for Nano, 3 for Micro, 4 for Milli, 5 for None, 6 for Kilo, 7 for Mega, 8 for Giga, and 9 for Tera. See "Analysis Units" on page 105.
34	Analysis Units "per" Area. Use 0 for square mm, 1 for square cm, 2 for square inches, 3 for square Meters, and 4 for square feet.
35	MTF Units code. Use 0 for cycles per millimeter, or 1 for cycles per milliradian. See "MTF Units" on page 106.
40	Coating File name. See "Coating File" on page 109.
41	Scatter Profile name. See "Scatter Profile" on page 109.
42	ABg Data File name. See "ABg Data File" on page 109.
43	GRADIUM Profile name. See "GRADIUM Profile" on page 110.
50	NSC Maximum Intersections Per Ray. See "Maximum Intersections Per Ray" on page 113.
51	NSC Maximum Segments Per Ray. See "Maximum Segments Per Ray" on page 113.
52	NSC Maximum Nested/Touching Objects. See "Maximum Nested/Touching Objects" on page 114.
53	NSC Minimum Relative Ray Intensity. See "Minimum Relative Ray Intensity" on page 114.
54	NSC Minimum Absolute Ray Intensity. See "Minimum Absolute Ray Intensity" on page 114.
55	NSC Glue Distance In Lens Units. See "Glue Distance In Lens Units" on page 114.
56	NSC Missed Ray Draw Distance In Lens Units. See "Missed Ray Draw Distance in Lens Units" on page 114.
57	NSC Retrace Source Rays Upon File Open. Use 0 for no, 1 for yes. See "Retrace Source Rays Upon File Open" on page 115.
58	NSC Maximum Source File Rays In Memory. See "Maximum Source File Rays In Memory" on page 114.
59	Simple Ray Splitting. Use 0 for no, 1 for yes. See "Simple Ray Splitting" on page 114.
60	Polarization Jx. See "Jx, Jy, X-Phase, Y-Phase" on page 109.
61	Polarization Jy. See "Jx, Jy, X-Phase, Y-Phase" on page 109.
62	Polarization X-Phase. See "Jx, Jy, X-Phase, Y-Phase" on page 109.
63	Polarization Y-Phase. See "Jx, Jy, X-Phase, Y-Phase" on page 109.
64	Convert thin film phase to ray equivalent. Use 0 for no, 1 for yes. See "Convert thin film phase to ray equivalent" on page 109.
65	Unpolarized. Use 0 for no, 1 for yes. See "Unpolarized" on page 109.
66	Method. Use 0 for X-axis, 1 for Y-axis, and 2 for Z-axis. See "Method" on page 109.
70	Ray Aiming. Use 0 for off, 1 for on, 2 for aberrated. See "Ray Aiming" on page 106.
71, 72, 73	Ray aiming pupil shift x, y, and z. See "Pupil Shift, Pupil Compress" on page 107.
74	Use Ray Aiming Cache. Use 0 for no, 1 for yes. See "Use Ray Aiming Cache" on page 107.
75	Robust Ray Aiming. Use 0 for no, 1 for yes. See "Robust Ray Aiming (slow)" on page 107.

Code	Property
76	Scale Pupil Shift Factors By Field. Use 0 for no, 1 for yes. See "Pupil Shift, Pupil Compress" on page 107.
77, 78	Ray aiming pupil compress x, y. See "Pupil Shift, Pupil Compress" on page 107.
100	Field type code. See "Fields" on page 115.
101	Number of fields.
102, 103	The field number is value1, value2 is the field x, y coordinate
104	The field number is value1, value2 is the field weight
105, 106	The field number is value1, value2 is the field vignetting decenter x, decenter y
107, 108	The field number is value1, value2 is the field vignetting compression x, compression y
109	The field number is value1, value2 is the field vignetting angle
110	The field normalization method, value 1 is 0 for radial and 1 for rectangular
200	Primary wavelength number. See "Wavelengths" on page 116.
201	Number of wavelengths
202	The wavelength number is value1, value 2 is the wavelength in micrometers.
203	The wavelength number is value1, value 2 is the wavelength weight
901	The number of CPU's to use in multi-threaded computations, such as optimization. If the passed value is zero, the number of CPU's will be set to the default value. When testing this value using the function SYPR, this returns the total number of CPU's available as reported by the operating system.

Usually, changes to system properties do not become effective until after the UPDATE keyword is executed.

Example:

```
! Set the number of wavelengths to 3
SETSYSTEMPROPERTY 201, 3
! Set the number of fields to 4
SYSP 101, 4
```

Related Functions:

SYPR

Related Keywords:

SETSURFACEPROPERTY, UPDATE

SETTEXTSIZE

Purpose:

Changes the size of the characters drawn by the GTEXT command.

Syntax:

```
SETTEXTSIZE xsize, ysize
```

Discussion:

The arguments refer to the fraction of the graphic screen width that each character represents. For example, the default text size is 70 40. This means each character is 1/70 of the graphic screen width, and 1/40 of the screen height. An argument of zero restores the text size to the default.

Example:

```
! Make text twice default size
```

```
SETTEXTSIZE 35, 20
! Restore text size to default
SETTEXTSIZE
```

SETTITLE

This command is obsolete. See “SETSYSTEMPROPERTY, SYSP” on page 728.

SETTOL

Purpose:

Sets the data in the various columns for any tolerance operand in the tolerance data editor.

Syntax:

```
SETTOL row, column, data
```

Discussion:

The value row must evaluate to an integer expression greater than 0 and less than or equal to the current number of tolerance operands. To set the operand type, use a column value of 0. The data is either a literal string or a string variable that contains the name of the operand, such as "TTHI". To set the integer 1, 2, or 3 values, use a column number of 1, 2, or 3. To set the min or max tolerance values, use column 4 or 5, respectively. To set the "Do Not Adjust During Inverse Tolerancing" checkbox, use column 6. The data is 1 for checked, 0 for unchecked. To set the comment, use column 99. The data is either a literal string or a string variable that contains the comment. The nominal value cannot be set using SETTOL.

See also INSERTTOL and DELETETOL.

Example:

```
SETTOL 16, 0, "TRAD"
```

Related Functions:

TOLV, \$TOLOPERAND, \$TOLCOMMENT

SETUNITS

This command is obsolete. See “SETSYSTEMPROPERTY, SYSP” on page 728.

SETVAR

Purpose:

Changes the state of variables for optimization.

Syntax:

```
SETVAR surf, code, status, object
or
SETVAR config, M, status, operand
```

Discussion:

Surf must evaluate to an integer between 0 and the maximum number of surfaces. Config must evaluate to an integer between 1 and the number of configurations. Code must be one of the following strings:

```
R for radius of curvature
T for thickness
C for conic
G for glass
I for glass index
J for glass Abbe
K for glass dpgf
Pn for parameter n
D for thermal coefficient of expansion
En for extra data value n
M for multi-configuration data, see discussion below
Nn for non-sequential component position data, 1-6 for x, y, z, tx, ty, tz
On for non-sequential component parameter data, where n is the parameter number
```

If status evaluates to 0, then the variable status is removed. Otherwise, the value is made variable. If the code is Nn or On; the object number must be provided; otherwise, it should be omitted. If the code is M, then the syntax for this command is as shown under "syntax" above.

Examples:

```
SETVAR j+3, R, 1
SETVAR 5, P6, 0
SETVAR surfk+2, E06, status
SETVAR config, M, status, operand
SETVAR 1, O32, 1, 5
```

SETVECSIZE

Purpose:

Changes the maximum size of the VEC1, VEC2, VEC3, and VEC4 arrays.

Syntax:

```
SETVECSIZE n
```

Discussion:

The expression must evaluate to an integer argument between 1 and 18,000,000. All four vector variables are always the same size. All data in the vectors will be lost during the resize. The initial size of the vectors is 1000.

SETVIG

Purpose:

Sets the vignetting factors for the lens.

Syntax:

```
SETVIG
```

Discussion:

See "Vignetting factors" on page 61 for a description of vignetting factors.

SHOWBITMAP

Purpose:

Displays a JPG or BMP file in a viewer window.

Syntax:

```
SHOWBITMAP filename
```

Discussion:

This keyword requires the name of the JPG or BMP file. The extension must be included. The filename may be enclosed in quotes if any blank or other special characters are used. The file must be located in the <data><images> folder. This command will open a new window to display the file.

Example:

```
SHOWBITMAP "BARCHART.BMP"
```

SHOWFILE

Purpose:

Displays a text file to the screen using the ZEMAX file viewer.

Syntax:

```
SHOWFILE filename, saveflag
```

Discussion:

The filename must be a valid file name. The file must be a text file (as would be created by OUTPUT and PRINT commands in ZPL) and must be in the current folder. Once the file is displayed, it may be scrolled up and down and printed like any other text file. The ability to scroll and print the data is the primary advantage of using OUTPUT and SHOWFILE instead of PRINT commands. SHOWFILE also closes the file if no CLOSE command

has been executed. If the saveflag is zero or omitted, then the file is erased when the window is closed. if saveflag is any value other than zero, then the file remains even after the window is closed.

Example:

```
OUTPUT "test.txt"
PRINT "Print this to a file."
SHOWFILE "test.txt"
```

Related Keywords:

OPEN, OUTPUT, CLOSE, PRINT, PRINTFILE

SOLVERETURN

Purpose:

Returns a computed solve value back the the editor calling the solve.

Syntax:

```
SOLVERETURN x
```

Discussion:

This keyword is used only for supporting ZPL Macro solves. The macro computes a single value, and returns the value to the editor calling the solve using this keyword. If more than one SOLVERETURN keyword exists, only the last SOLVERETURN called is considered and prior calls are ignored.

If an error or invalid condition occurs in the macro, and the solve value cannot be computed, then the macro should return without calling SOLVERETURN. The absence of the SOLVERETURN call will indicate to ZEMAX that the solve cannot be computed and the optical system is in an error condition. This is particularly important during optimization.

See “Using ZPL Macro solves” on page 479.

SOLVETYPE

Purpose:

Changes the solve status on a given surface and value. Only some solve types are supported; contact ZEMAX technical support for information on setting other types of solves.

Syntax:

```
SOLVETYPE surf, CODE, arg1, arg2, arg3, arg4
```

Discussion:

Surf must evaluate to an integer argument between 0 and the maximum number of surfaces. The CODE must be a mnemonic as listed in the following table. The arg1 - arg4 expressions evaluate to the first - fourth solve parameters as specified in “SOLVES” on page 473. Note that for cross-column Pickup solves, the column numbers are defined in “Integer codes for column numbers” on page 480. For non-sequential pickup solves, the arguments are the first through fourth lines that follow the "Solve Type" selection on the NSC solve dialog box. Some codes do not use any or all of the arguments and the arguments may be omitted in these cases.

CODES FOR SOLVETYPE KEYWORD

Solve Type	CODE
Curvature: Fixed (turn solve off)	CF
Curvature: Variable	CV
Curvature: Marginal Ray	CM
Curvature: Chief Ray	CC
Curvature: Pickup	CP
Curvature: Marginal Ray Normal	CN
Curvature: Chief Ray Normal	CO

Solve Type	CODE
Curvature: Aplanatic	CA
Curvature: Element Power	CE
Curvature: Concentric With Surface	CQ
Curvature: Concentric With Radius	CR
Curvature: F/#	CG
Curvature: ZPL Macro	CZ
Thickness: Fixed (turn solve off)	TF
Thickness: Variable	TV
Thickness: Marginal Ray Height	TM
Thickness: Chief Ray Height	TC
Thickness: Edge Thickness	TE
Thickness: Pickup	TP
Thickness: Optical Path Difference	TO
Thickness: Position	TL
Thickness: Compensator	TX
Thickness: Center Of Curvature	TY
Thickness: Pupil Position	TU
Thickness: ZPL Macro	TZ
Glass: Fixed (turn solve off)	GF
Glass: Model	GM
Glass: Pickup	GP
Glass: Substitute	GS
Glass: Offset	GO
Semi-Diameter: Automatic	SA
Semi-Diameter: User Defined	SU
Semi-Diameter: Pickup	SP
Semi-Diameter: Maximum	SM
Semi-Diameter: ZPL Macro	SZ
Conic: Fixed (turn solve off)	KF
Conic: Pickup	KP
Conic: ZPL Macro	KZ
Parameter: Fixed (turn solve off). Replace "p" with the parameter number in the code, for example, PF_3 will turn off the solve on parameter 3.	PF_p
Parameter: Pickup. Replace "p" with the parameter number in the code, for example, PP_4 will set the solve on parameter 4.	PP_p

Solve Type	CODE
Parameter: Chief Ray. Replace "p" with the parameter number in the code, for example, PC_1 will set the solve on parameter 1.	PC_p
Parameter: ZPL Macro. Replace "p" with the parameter number in the code, for example, PZ_1 will set the solve on parameter 1.	PZ_p
Thermal Coefficient of Expansion: Fixed (turn solve off)	HF
Thermal Coefficient of Expansion: Pickup	HP
Extra Data Value: Fixed (turn solve off). Replace "e" with the extra data number in the code, for example, EF_3 will turn off the solve on extra data value 3.	EF_e
Extra Data Value: Pickup. Replace "e" with the extra data number in the code, for example, EP_4 will set the solve on extra data value 4.	EP_e
Extra Data Value: ZPL Macro. Replace "e" with the extra data number in the code, for example, EZ_4 will set the solve on extra data value 4.	EZ_p
Non-Sequential Component Pickup X, Y, Z, Tilt-X, Tilt-Y, Tilt-Z, Material. Replace "o" with the object number in the code, for example, NSC_PX_14 will set a pickup solve on object 14. NSC_PMAT_o is redundant with NSC_MATP_o.	NSC_PX_o, NSC_PY_o, NSC_PZ_o, NSC_PTX_o, NSC_PTY_o, NSC_PTZ_o, NSC_PMAT_o (see left)
Non-Sequential Component Material is fixed, model glass, pick up, or offset. Replace "o" with the object number in the code, for example, NSC_MATM_11 will set the material on object 11 to be a model glass. NSC_MATP_o is redundant with NSC_PMAT_o.	NSC_MATF_o, NSC_MATM_o, NSC_MATP_o, NSC_MOFF_o (see left)
Non-Sequential Component ZPL Macro solve on X, Y, Z, Tilt-X, Tilt-Y, Tilt-Z. Replace "o" with the object number in the code, for example, NSC_ZX_14 will set a macro solve on object 14.	NSC_ZX_o, NSC_ZY_o, NSC_ZZ_o, NSC_ZTX_o, NSC_ZTY_o, NSC_ZTZ_o (see left)
Non-Sequential Component Parameter Pickup. Replace "o" with the object number and "p" with the parameter number in the code, for example, NSC_PP_11_7 will set a pickup solve on object 11, parameter 7.	NSC_PP_o_p (see left)
Non-Sequential Component ZPL Macro solve. Replace "o" with the object number and "p" with the parameter number in the code, for example, NSC_ZP_11_7 will set a macro solve on object 11, parameter 7.	NSC_ZP_o_p (see left)

Example:

```
! The following line will add a glass pickup solve
! on surface 7, picking up from surface 5:
SOLVETYPE 7, GP, 5
! Add a thickness pickup with a scale factor of -1:
SOLVETYPE 7, TP, 5, -1

! Set a pickup solve on surface 1, NSC object 12 Z position,
! pick up from object 11, with a scale factor of 2, offset 3,
! from the parameter 7 column. Note the column number is argument 4.
! The column number is 0 for the same column, 1-6 for x, y, z, tilt x, tilt y, tilt z,
! respectively. The column number for the parameter columns
! is 6 + the desired parameter number.
! In summary, the syntax is:
! SOLVETYPE, surf, code, object, scale, offset, column
! where code has the object/parameter number imbedded as shown in the table above.
```


! The syntax for this example is:
SOLVETYPE 1, NSC_PZ_12, 11, 2, 3, 13

Related Functions:

SOLV

STOPSURF

Purpose:

STOPSURF sets the current stop surface location by number.

Syntax:

```
STOPSURF surf
```

Discussion:

The UPDATE command must be issued before the new data takes effect.

Example:

```
STOPSURF n+2
```

Related Keywords:

UPDATE

SUB

See GOSUB.

SURFTYPE

This command is obsolete. See “SETSURFACEPROPERTY, SURP” on page 725.

TELECENTRIC

This command is obsolete. See “SETSYSTEMPROPERTY, SYSP” on page 728.

TESTPLATEFIT

Purpose:

TESTPLATEFIT calls the test plate fitting routine. See “Test Plate Fitting” on page 242.

Syntax:

```
TESTPLATEFIT tpd_file, log_file, method, number_cycles
```

Discussion:

This keyword requires string expressions for the test plate data file, and the name of a file for the output log file. The method is an integer between 0 and 4, inclusive, for try all methods, best to worst, worst to best, long to short, and short to long, respectively. The integer number_cycles is 0 for automatic or the maximum number of optimization cycles to execute. Note the tpd_file name should NOT include the path, since all test plate files are in a fixed folder; while the path should be included for the log file.

This keyword may take a long time to execute. It is advisable to display the log file after completion of the fitting, or use some other means to indicate when the fitting is complete.

Example:

```
TESTPLATEFIT "optico.tpd", "c:\temp\logfile.dat", 0, 0  
SHOWFILE "c:\temp\logfile.dat"
```

THIC

This command is obsolete. See “SETSURFACEPROPERTY, SURP” on page 725.

TIMER

Purpose:

Resets the internal clock. This feature is used in conjunction with the ZPL function ETIM() for measuring the elapsed time since the last TIMER command.

Syntax:

TIMER

Discussion:

TIMER and ETIM() are used primarily for testing execution efficiency of the ZPL interpreter and various program architectures.

Example:

```
i = 0
TIMER
LABEL 1
x = RAND(1000)
i = i+1
if i < 10000 THEN GOTO 1
FORMAT .1
PRINT "Elapsed time:", ETIM(), "Seconds"
```

TOLERANCE

Purpose:

Runs a tolerance analysis and saves the tolerance report to a text file.

Syntax:

TOLERANCE top_file_name, out_file_name

Discussion:

TOLERANCE runs a tolerance analysis on the current lens based upon the settings in the tolerance options file defined by top_file_name. This file name should end in the extension TOP and be placed in the <data>\Configs folder (see “Folders” on page 68). TOP files are created by saving the tolerance options from the tolerance dialog, see “Other buttons” on page 564. The out_file_name will be a text file (any extension may be used) and the file will be placed in the same folder as the current lens file. Neither top_file_name nor out_file_name should specify a path.

Example:

```
UNLOCKWINDOW 2
```

UNLOCKWINDOW

Purpose:

Unlocks any one or all locked windows.

Syntax:

UNLOCKWINDOW winnum

Discussion:

See “Graphic windows operations” on page 42. If the winnum is zero, then all open windows are unlocked. If the winnum argument is -1, then the currently executing window will be unlocked at the end of the macro execution.

Example:

```
UNLOCKWINDOW 2
```

Related Keywords:

LOCKWINDOW

UPDATE

Purpose:

Updates pupil positions, index data, paraxial constants, semi-diameters, maximum field normalization values, solves, non-sequential objects, and other data required for both sequential and non-sequential ray tracing. The UPDATE keyword MUST be used before tracing or evaluating a system if the prescription data, such as surface

type, radii, thicknesses, system apertures, wavelengths, or non-sequential object types or parameters, have been altered since the last UPDATE.

If the UPDATE command is followed by the keyword "ALL", then all open windows, except ZPL macro windows, will be updated as well. If the UPDATE command is followed by an expression which evaluates to an integer corresponding to an open window, then the specified window is updated, as long as it is not the window of the macro currently executing. If UPDATE is followed by EDITORS, any open spreadsheet editors will be refreshed to show recent changes in the lens data. If UPDATE is followed by the keyword MACROS, then all open ZPL macro windows other than the calling window are updated.

Syntax:

```
UPDATE
UPDATE ALL
UPDATE n
UPDATE EDITORS
UPDATE MACROS
```

VEC1, VEC2, VEC3, VEC4

Purpose:

These keywords are used to set the vector variables VEC1, VEC2, VEC3, and VEC4. Each vector can store an array of double-precision floating point numbers.

Syntax:

```
VEC1 subscript, value
VEC2 subscript, value
VEC3 subscript, value
VEC4 subscript, value
```

Discussion:

VEC1..4 are used to store data in an array. The subscript value can be any expression which is rounded down to an integer. The resulting integer expression must be between 0 and the current maximum vector size, which is initially 1000. The ZPL functions VEC1..4 can be used to extract the data. To change the array size use the SETVECSIZE keyword.

Example:

```
i = 0
LABEL 1
i = i + 1
VEC1 i, i
IF i < 10 THEN GOTO 1
j = 0
LABEL 2
j = j + 1
VEC2 j, VEC1(j) * VEC1(j)
IF j < 10 THEN GOTO 2
i = 0
LABEL 3
i = i + 1
PRINT "x = ", VEC1(i), " x*x = ", VEC2(i)
IF i < 10 THEN GOTO 3
PRINT
PRINT "All done!"
```

WAVL, WWGT

These commands are obsolete. See "SETSYSTEMPROPERTY, SYSP" on page 728.

XDIFFIA

Purpose:

Computes the Extended Diffraction Image Analysis feature and saves the result to a ZBF file. For a description of the ZBF (ZEMAX Beam File) format see "ZEMAX Beam File (ZBF) binary format" on page 640. For a

description of the Extended Diffraction Image Analysis feature see “Extended Diffraction Image Analysis” on page 177.

Syntax:

```
XDIFFIA outfile, infile
```

Discussion:

This keyword requires the name of the output ZBF file, and optionally, the name of the input IMA or BIM file. If the extension to the outfile is not provided, the extension ZBF will be appended. The extension must be provided on the infile. The filenames must be enclosed in quotes if any blank or other special characters are used. The outfile will be placed in the <pop> folder. The infile must be placed in the <data>\<images> folder. No paths should be provided with the file names.

The settings for the Extended Diffraction Image Analysis feature will be those settings previously saved for the current lens. To make adjustments to the settings, open an Extended Diffraction Image Analysis window, choose the appropriate settings, then press "Save". All subsequent calls to XDIFFIA will use the saved settings. The exceptions are the output file name, which is specified as the first argument after the XDIFFIA keyword, and the input source file, which is optionally specified as the second argument after the XDIFFIA keyword.

Example:

```
XDIFFIA "output will be in this file.ZBF", "SOMEIMAFILE.IMA"
```

Related Functions:

ZBFCLR, ZBFSHOW, ZBFSUM, ZBFMULT

ZBFCLR

Purpose:

Clears the complex amplitude data in a ZBF file.

Syntax:

```
ZBFCLR filename
```

Discussion:

This keyword requires only the name of the ZBF file.

General comments about all ZBF related keywords

For a description of the ZBF (ZEMAX Beam File) format see “ZEMAX Beam File (ZBF) binary format” on page 640. ZEMAX requires the filename extension to be ZBF, and will append or replace the extension as required. The filename may be enclosed in quotes if any blank or other special characters are used. The file must be located in the <pop> folder. All ZBF output files will be placed in this same folder.

Example:

```
ZBFCLR "some beam file name.ZBF"  
ZBFCLR N$
```

ZBFMULT

Purpose:

Multiplies the complex amplitude data in a ZBF file by a complex factor.

Syntax:

```
ZBFMULT filename, Ax, Bx, Ay, By
```

Discussion:

This keyword requires the name of the ZBF file and the real (A) and imaginary (B) parts of the complex number to multiply every point in the ZBF file by. There are separate A and B for the x and y polarized light directions. The resulting data is written back to the same file name.

See “ZBFCLR” on page 740 for comments that apply to all ZBF related keywords.

Example:

```
ZBFMULT "some beam file name.ZBF", 0.0, 1.0, 0.0, -1.0
```

ZBFPROPERTIES

Purpose:

Opens the specified ZBF file and places various data about the beam in a vector variable.

Syntax:

```
ZBFPROPERTIES filename, vector
```

Discussion:

This keyword requires the name of the ZBF file and a value for the vector number to place the data in. The value of vector must be between 1 and 4, inclusive. After the ZBFPROPERTIES function executes, the following beam data will be placed in the specified vector: nx, ny, dx, dy, waist_x, waist_y, position_x, position_y, rayleigh_x, rayleigh_y, wavelength (in lens units), total power, peak irradiance (power per area), the is_polarized flag (0 for no, 1 for yes), and the media index; the values are placed in vector positions 1 through 15.

See "ZBFCLR" on page 740 for comments that apply to all ZBF related keywords.

Example:

```
ZBFPROPERTIES "TEST.ZBF", 1  
rayleighx = VEC1(9)
```

ZBFREAD

Purpose:

Opens the specified ZBF file and places the electric field and beam property data in two user-defined array variables.

Syntax:

```
ZBFREAD filename, beamname, propertyname
```

Discussion:

This keyword requires the name of the ZBF file and the name of two arrays defined by a previous call to DECLARE. The beamname must be a 3 dimensional array, of minimum size (nx, ny, 2) for an unpolarized beam and minimum size (nx, ny, 4) for a polarized beam. The propertyname array must be a one dimensional array of minimum size 14. After the ZBFREAD function executes, the following beam data will be placed in the specified propertyname array: nx, ny, dx, dy, waist_x, waist_y, position_x, position_y, rayleigh_x, rayleigh_y, wavelength (in lens units), total power, peak irradiance (power per area), the is_polarized flag (0 for no, 1 for yes), and the media index; the values are placed in array positions 1 through 15. The electric field data will be placed in the beamname array. The third dimension of the beamname array is 1 for Ex Real, 2 for Ex Imaginary, and if the beam is polarized, 3 for Ey Real, and 4 for Ey Imaginary.

See "ZBFCLR" on page 740 for comments that apply to all ZBF related keywords. See also "ZBFWRITE" on page 743.

Example:

```
! First get the beam size  
ZBFPROPERTIES "TEST1.ZBF", 1  
nx = vec1(1)  
ny = vec1(2)  
ip = vec1(14) ! The "is polarized" flag  
  
! Allocate enough memory to hold the beam  
IF (ip == 0) THEN DECLARE B, DOUBLE, 3, nx, ny, 2  
IF (ip == 1) THEN DECLARE B, DOUBLE, 3, nx, ny, 4  
DECLARE P, DOUBLE, 1, 20  
  
ZBFREAD "test1.zbf", B, P  
  
FOR j, 1, ny, 1  
  FOR i, 1, nx, 1  
    FORMAT 4.0  
    PRINT i, j,  
    FORMAT 12.6  
    IF (ip == 1)
```

```

        PRINT B(i, j, 1),
        PRINT B(i, j, 2),
        PRINT B(i, j, 3),
        PRINT B(i, j, 4)
    ELSE
        PRINT B(i, j, 1),
        PRINT B(i, j, 2)
    ENDIF
NEXT
NEXT

```

```

! save the beam
ZBFWRITE "TEST2.ZBF", B, P

! release the allocated memory
RELEASE B
RELEASE P

```

ZBFRESAMPLE

Purpose:

Resamples a ZBF file to a new width and point spacing.

Syntax:

```
ZBFRESAMPLE filename, nx, ny, wx, wy, decenterx, decentery
```

Discussion:

This keyword requires the name of the ZBF file and six numbers. The beam will be resampled and interpolated as required to create a new beam file with nx and ny points, of total width wx and wy, in the x and y directions respectively. The nx and ny values must be powers of 2; such as 32, 64, 128, etc. The decenterx and decentery values may be provided to optionally decenter the new beam relative to the old beam. If either nx or ny is zero, no change is made to the existing beam sampling. If either wx or wy is zero, no change is made to the existing beam width. The length units in the ZBF file are converted automatically to the current lens units. The resulting data is written back to the same file name.

See “ZBFCLR” on page 740 for comments that apply to all ZBF related keywords.

Example:

```
ZBFRESAMPLE "TEST.ZBF", 128, 128, 25.4, 25.4, 0, 0.4
```

ZBFSHOW

Purpose:

Displays a ZBF file in a viewer window.

Syntax:

```
ZBFSHOW filename
```

Discussion:

This keyword requires only the name of the ZBF file. This command will open a new window to display the beam file.

See “ZBFCLR” on page 740 for comments that apply to all ZBF related keywords.

Example:

```
ZBFSHOW "new beam data.ZBF"
ZBFSHOW N$
```

ZBFSUM

Purpose:

Sums either coherently or incoherently the data in two ZBF files and places the resulting data in a third ZBF file.

Syntax:

```
ZBFSUM coherent, filename1, filename2, outfile
```

Discussion:

This keyword requires an integer to indicate if the sum is coherent (any value other than zero) or incoherent (zero), and the names of three ZBF files. If an incoherent sum is performed, then the output data will be real valued only. If the two source files do not have the same number of data points, point spacing, and reference radii in both x and y directions, then the second source file listed is first scaled and interpolated, and the reference radii is adjusted, to match the first file before the sum is performed. The length units in the ZBF files are converted automatically to the current lens units. The outfilename may be the same as one of the source file names; in which case the original file is overwritten.

See “ZBFCLR” on page 740 for comments that apply to all ZBF related keywords.

Example:

```
ZBFSUM 1, "a.ZBF", "b.zbf", "coherent a plus b.zbf"  
ZBFSUM 0, "a.ZBF", "b.zbf", "incoherent a plus b.zbf"  
ZBFSUM 0, A$, B$, C$
```

ZBFTILT

Purpose:

Multiplies the data in a ZBF file by a complex phase factor to introduce phase tilt to the beam.

Syntax:

```
ZBFTILT filename, cx, cy, tx, ty
```

Discussion:

This keyword requires the name of the ZBF file and four numbers. The phase of the beam is modified by a phase angle given by $\theta = (x - cx)tx + (y - cy)ty$. The cx and cy values are the center of the phase tilt, and tx and ty are the slopes of the tilt in units of radians per lens unit length. The coordinates x and y refer to positions within the beam file, with the center coordinate (x = 0, y = 0) being at the point (nx/2 + 1, ny/2 + 1) where nx and ny are the number of points in the x and y directions. The length units in the ZBF file are converted automatically to the current lens units. The resulting data is written back to the same file name.

See “ZBFCLR” on page 740 for comments that apply to all ZBF related keywords.

Example:

```
ZBFTILT "TEST.ZBF", 0.0, 0.0, 0.0, 0.01237
```

ZBFWRITE

Purpose:

Writes electric field and beam property data arrays to a ZBF file.

Syntax:

```
ZBFWRITE filename, beamname, propertyname
```

Discussion:

This keyword requires the name of the ZBF file and the name of two arrays defined by a previous call to DECLARE. The beamname must be a 3 dimensional array, of minimum size (nx, ny, 2) for an unpolarized beam and minimum size (nx, ny, 4) for a polarized beam. The propertyname array must be a one dimensional array of minimum size 14. The following beam data must be placed in the specified propertyname array: nx, ny, dx, dy, waist_x, waist_y, position_x, position_y, rayleigh_x, rayleigh_y, wavelength (in lens units), total power, peak irradiance (power per area), the is_polarized flag (0 for no, 1 for yes), and the media index; the values are placed in array positions 1 through 15. The electric field data must be placed in the beamname array. The third dimension of the beamname array is 1 for Ex Real, 2 for Ex Imaginary, and if the beam is polarized, 3 for Ey Real, and 4 for Ey Imaginary.

See “ZBFCLR” on page 740 for comments that apply to all ZBF related keywords. See also “ZBFREAD” on page 741.

ZRDAPPEND

Purpose:

Appends data from one ZRD file onto the end of a second file. The first file is left unaltered. This command only makes sense to use if the two ZRD files were traced in the same system, as no attempt is made to modify object numbers or coordinates.

Syntax:

```
ZRDAPPEND infilename, outfilename
```

Discussion:

This keyword requires the name of two input ZRD files. For information on the ZRD format, see “Ray database (ZRD) files” on page 583. The file names may each be defined by either a literal string or a single string variable. The file names must include the full path name. Although any extensions may be specified, the extension ZRD is recommended for both file names. The input and output file names must be distinct. The input and output files must use the same ZRD format. See also ZRDSUM.

Example:

```
ZRDAPPEND "C:\TEMP\CART.ZRD", "C:\TEMP\HORSE.ZRD"
```

ZRDFILTER

Purpose:

Opens a ZRD ray database file, applies a filter, and saves the filtered subset of rays to a new ZRD file.

Syntax:

```
ZRDFILTER infilename, outfilename, filterstring
```

Discussion:

This keyword requires the name of an input ZRD file, the name of the output (filtered) ZRD file, and the filter string to apply. For information on the ZRD format, see “Ray database (ZRD) files” on page 583. For information on the filter string format see “The filter string” on page 589. The file names and filter string may each be defined by either a literal string or a single string variable. The file names must include the full path name. Although any extensions may be specified, the extension ZRD is recommended for both the input and output files. The input and output file names must be distinct.

Example:

```
ZRDFILTER "C:\TEMP\TEST1.ZRD", "C:\TEMP\TEST1_H4.ZRD", "H4"
```

ZRDPLAYBACK

Purpose:

This keyword reads a ZRD file, and adds ray amplitude data to one or all detectors. This command only makes sense to use if the ZRD file was generated with the currently loaded system, as no attempt is made to modify object numbers or coordinates.

Syntax:

```
ZRDPLAYBACK zrdfilename, surface, detector, clear, filterstring
```

Discussion:

This keyword requires the name of a ZRD file. For information on the ZRD format, see “Ray database (ZRD) files” on page 583. The file name may each be defined by either a literal string or a single string variable. The file name must not include a path and the file must be in the same folder as the current lens file. The surface argument is the surface number of the Non-Sequential surface, use 1 for Non-Sequential mode. The integer detector is either 0 for all detectors, or the object number of a specific single detector. If clear is set to anything other than zero, then the specified detector(s) are cleared prior to playback. The optional filterstring is applied to the rays before being played back.

Example:

```
ZRDPLAYBACK "C:\TEMP\MYRAYS.ZRD", 1, 0, 1, "H3 & H5"
```

ZRDSAVERAYS

Purpose:

Opens a ZRD ray database file, applies a filter, and saves the rays in the filtered subset that intersect a specified object number.

Syntax:

```
ZRDSAVERAYS infilename, outfilename, filterstring, object
```

Discussion:

This keyword requires the name of an input ZRD file, the name of the output source ray DAT file, the filter string to apply, and the object number. For information on the ZRD format, see "Ray database (ZRD) files" on page 583. Source ray files are described in "Binary Source File format" on page 547. For information on the filter string format see "The filter string" on page 589. The file names and filter string may each be defined by either a literal string or a single string variable. The file names must include the full path name. Although any extensions may be specified, the extension ZRD is recommended for the input file, and the extension DAT is recommended for the output file. The input and output file names must be distinct. If no filtering is desired, the filterstring argument should be replaced with two double quotes like this: "".

Example:

```
ZRDSAVERAYS "C:\TEMP\TEST1.ZRD", "C:\TEMP\TEST1_H6.DAT", "", 6
```

ZRDSUM

Purpose:

Concatenates two ZRD files into a third file. This command only makes sense to use if the two ZRD files were traced in the same system, as no attempt is made to modify object numbers or coordinates.

Syntax:

```
ZRDSUM infilename1, infilename2, outfilename
```

Discussion:

This keyword requires the name of two input ZRD files and the name of the output ZRD file. For information on the ZRD format, see "Ray database (ZRD) files" on page 583. The file names may each be defined by either a literal string or a single string variable. The file names must include the full path name. Although any extensions may be specified, the extension ZRD is recommended for all three file names. The input and output file names must be distinct. Both input files must use the same ZRD format, and the output file will use this same ZRD format. See also ZRDAPPEND.

Example:

```
ZRDSUM "C:\TEMP\TEST1.ZRD", "C:\TEMP\TEST2.ZRD", "C:\TEMP\TEST_SUM.ZRD"
```

Example Macro 1

The following sample macro will print out the image surface intercept coordinates of the chief ray at every defined field angle:

```
nfield = NFLD()
maxfield = MAXF()
n = NSUR()
FOR i, 1, nfield, 1
    hx = FLDX(i)/maxfield
    hy = FLDY(i)/maxfield
    PRINT "Field number ", i
    RAYTRACE hx,hy,0,0,PWAV()
    PRINT "X-field angle : ",FLDX(i)," Y-field angle : ", FLDY(i)
    PRINT "X-chief ray : ",RAYX(n), " Y-chief ray : ", RAYY(n)
    PRINT
NEXT
PRINT "All Done!"
```

The first line of the macro calls the NFLD() function which returns the number of defined field angles, and assigns it to the variable "numfield". The second line calls MAXF(), which returns the maximum radial field and stores the value in the variable "maxfield". The number of surfaces is then stored in the variable "n" by calling the function NSUR().

The FOR loop is then defined with i being the counter for the field position, starting at 1, with a maximum value of nfield, and an increment of 1 on each loop. The two macro lines which define "hx" and "hy" use the functions FLDX() and FLDY(), which return the x- and y- field values for the current field position number "i". The chief ray is then traced by the keyword RAYTRACE. Note the chief ray intersects the center of the pupil, which is why the two pupil coordinates are both zero. The PWAV() function returns the primary wavelength number, which is usually used for the chief ray. The various PRINT commands will output the chief ray coordinates on the image surface to the screen.

Example Macro 2

The following sample macro will estimate the RMS spot radius (on axis) of the current optical system. The macro traces many random rays through the system, and records the radial offset from the primary wavelength chief ray. The current wavelength weighting is applied to estimate the RMS spot radius.

```
PRINT "Primary wavelength is number ",
FORMAT .0
PRINT PWAV(),
FORMAT .4
PRINT " which is ", WAVL(PWAV()), " micrometers."
PRINT "Estimating RMS spot radius for each wavelength."

! How many random rays to trace to make estimate?
n = 100

! Initialize the timer
TIMER

! Store the number of surfaces for later use
ns = NSUR()

! Start at wavelength 1
weightsum = 0
wvrms = 0
FOR w, 1, NWAV(), 1
  rms = 0
  FOR i, 1, n, 1
    hx = 0
    hy = 0
    angle = 6.283185 * RAND(1)
    ! SQRT yields uniform distribution in pupil
    radius = SQRT(RAND(1))
    px = radius * COSI(angle)
    py = radius * SINE(angle)
    RAYTRACE hx, hy, px, py, w
    x = RAYX(ns)
    y = RAYY(ns)
    rms = rms + (x*x) + (y*y)
  NEXT
  rms = SQRT(rms/n)
  wvrms = wvrms + ( WWGT (w) * rms )
  weightsum = weightsum + WWGT(w)
  FORMAT .4
  PRINT "RMS spot radius for ", WAVL(w),
  FORMAT .6
  PRINT " is ", rms
NEXT
wvrms = wvrms / weightsum
PRINT "Wavelength weighted rms is ", wvrms
FORMAT .2
t = ETIM()
PRINT "Elapsed time was ",t," seconds."
```

Note the use of the trailing commas in the first two PRINT commands. These allow the data printed from the first three PRINT commands to appear on the same line. The intervening FORMAT commands changes the way

the numerical output is printed, even on the same line. The "!" character is used to indicate a comment, which is ignored when the macro is run.

Calling a Macro from within a Macro

To call another ZPL macro from within a ZPL macro, use the CALLMACRO keyword. Data may also be passed between the macros using the keywords CALLSETDBL and CALLSETSTR and the functions CALD and \$CALLSTR.

The first macro that is executed, usually from the ZPL Macro Dialog box, is the "parent" macro. The parent macro can call other macros, and these macros are called "child" macros. The parent macro creates a buffer of 51 numeric and 51 string values. The buffer can be used to set or retrieve numeric or string values, which allows macros to share data, or pass arguments to one another. The buffer is unique to the parent macro, and any child macros called by the parent. If more than one parent macro is executing in parallel (for example, in two windows each updating a different macro) each parent has their own buffer, and no data is shared between parents.

The easiest way to see how the parent and child macros work is by example. This example will use just two very simple macros, but there is no hard limit as to how many macros may be called or nested.

Consider two macros, called PARENT.ZPL and CHILD.ZPL The PARENT.ZPL macro is:

```
CALLSETDBL 1, 3.5
CALLSETSTR 1, "Hello World"
CALLMACRO CHILD.ZPL
PRINT CALD(1)
```

The CHILD.ZPL macro is:

```
PRINT "Executing child macro"
PRINT CALD(1)
A$ = $CALLSTR(1)
PRINT A$
CALLSETDBL 1, 7.11
```

When the PARENT.ZPL macro is executed, the macro uses CALLSETDBL to place the numeric value 3.5 in the parent's call numeric buffer in index position 1. The CALLSETSTR is then used to place the string "Hello World" in the parent's call string buffer in index position 1. Any index between 0 and 50, inclusive may be used for either the numeric or string buffers.

The PARENT.ZPL macro then executes the macro CHILD.ZPL. The child macro prints the message "Executing child macro" to the child's output window, which is not visible during execution. The child then extracts the numeric value from numeric buffer at index 1 using CALD(1), and the string value from the string buffer index 1 using \$CALLSTR(1). These values are also printed. Finally, the child macro overwrites the numeric value in index 1 with the new value 7.11, and the child macro completes.

Control now returns to the parent macro, which copies the child macro output window contents to the parent's output window, and the child's output window contents are discarded. Lastly, the modified value from the numeric call buffer at index 1 is printed. The output of this macro looks like this:

```
Executing child macro
3.5000
Hello World
7.1100
```

The sample macros PARENT.ZPL and CHILD.ZPL are included with ZEMAX in the Macros folder.

Introduction



This feature is only available in the EE edition of ZEMAX.

ZEMAX has a very powerful feature which allows another Windows program to establish a communication link with ZEMAX, and for this other program to get data from ZEMAX on the lens. The idea is that a program can use ZEMAX to trace rays through the lens, and then the data is sent to the other program for further analysis or computation. Extensions are more complicated than macros; for simple programming jobs see "ZEMAX PROGRAMMING LANGUAGE" on page 651.

There are three closely related ways this feature can be used to extend the capabilities of ZEMAX:

First, any stand-alone Windows program may be designed to establish a link with ZEMAX which is used to extract data about a lens that ZEMAX can provide- typically ray trace or other optical data. The program may use this data any way it sees fit.

Second, a stand alone program can do its analysis "hidden" from the user, and then the data generated may be displayed in the usual ZEMAX graphic or text window. When using this mode, the program is called an "Extension" to ZEMAX. Extensions appear as menu options under the "Extensions" menu in the ZEMAX program. The Extensions must be placed in the <program>\Extend folder (see "Folders" on page 68) for ZEMAX to run them.

Third, the stand alone program may be used to compute data that ZEMAX can optimize. In this mode, the program is called a "User Defined Operand" or UDO. These UDO's must be placed in the <program>\UDO folder (see "Folders" on page 68) for ZEMAX to run them.

Applications that require complex user interfaces, with many inputs, menus, and buttons are better suited to be stand alone applications that do not use ZEMAX to display the computed data. Applications that may be complex computationally, but require only a single dialog box type user interface, are better suited to be integrated with ZEMAX as Extensions. Extensions have the advantage of looking and acting just like regular ZEMAX analysis windows. Features implemented as Extensions may be viewed within ZEMAX, updated, have settings changed, and be printed, copied to the clipboard, or be panned and zoomed just like any ZEMAX analysis window.

UDO's are for applications that compute data that must be optimized by the ZEMAX optimizer. UDO's are described in detail in the chapter "Optimization" under "User defined operands".

Well written applications may even be designed to operate either as a stand alone program or as an Extension, as shown in the examples section.

The communication between application and ZEMAX is accomplished using Dynamic Data Exchange (DDE). DDE is a protocol defined within the Windows operating system for sharing data between programs. Two programs can establish a DDE link, with one program acting as the "server" and the other the "client". The client generally requests specific data from the server, and the server sends the data back to the client.

ZEMAX is intended to be the server, and any Windows program with the right modifications can be made a client. This chapter defines the DDE interface to ZEMAX, so that the user may write software programs which take advantage of the ZEMAX ray trace engine.

Requirements for writing Extensions

The purpose of this chapter and the DDE features is to assist experienced programmers in writing code that can interface to ZEMAX. This feature allows programmers to write third party software programs that work in conjunction with ZEMAX. The third party program may use ZEMAX as a ray trace engine, but provide an interface, graphics, input and output, and extended calculations that are customized for a particular type of analysis.

Extensions can be easily implemented using any program or programming language that supports DDE. Examples of such programs include Visual Basic, Excel, and MatLab, but there are many others and DDE is a fairly common interface between scientific programs. If you are writing an application in C, then using this feature requires a knowledge of Windows and C programming, with at least some mastery over message passing,

message loops, memory management, and global atoms, handles, and pointers. An excellent book that contains all the information needed to learn DDE programming is "Programming Windows, Fourth Edition" by Charles Petzold, Microsoft Press. The examples and code sections in the text of this chapter are not complete programs, but are code fragments intended only for illustration.

Although programming a Windows application is somewhat complex, the good news is that all of the difficult code has already been written, and is available for study and re-use. The details on establishing the communication link with ZEMAX have already been worked out, and the sample code may be "cut and paste" with only the code relating to the specific computation required needing change.

There is also a C source code file, called ZCLIENT, that vastly simplifies the DDE communication with ZEMAX by encapsulating all the communication in wrapper functions. ZCLIENT provides a number of functions that can be called to extract data from ZEMAX without requiring knowledge of DDE programming. ZCLIENT is described later in this chapter.

There are complete examples of client programs and Extensions, including source code, provided with ZEMAX and discussed at the end of this chapter.

The power of the Extension comes at a price, although it is a reasonable one. Use of the DDE features does require that the user have a suitable compiler or development tool that can generate 32 bit Windows executables. It is also assumed that the user can write the required code, and most importantly, ensure that the code is reliable and bug-free. To maximize speed, ZEMAX performs very little error checking on data returned by the client, and so buggy Extensions are capable of bringing ZEMAX to a crash, or placing ZEMAX in an infinite loop.

For this reason, technical support on the implementation of DDE or Extensions is strictly limited to illustrating that the provided sample files work correctly. If you need a ZEMAX Extension, and do not possess the desire or ability to write them yourself, please feel free to contact ZEMAX Technical Support for a quote on developing a custom program to meet your requirements. We have considerable experience in developing these types of programs, and can generally write Extensions at very competitive rates on short notice.

Establishing the link

To establish a DDE link with ZEMAX, the client program must broadcast a message to all top level windows that includes a reference to the application name, and the topic name. For ZEMAX, the application name is "ZEMAX", and the topic name can be any non-NULL string. If two copies of ZEMAX are running, the application name may also be either "ZEMAX1" or "ZEMAX2" to distinguish between the two. ZEMAX does not use the topic name, just the application name and the "item". The item indicates to ZEMAX what data is being requested. The various items ZEMAX supports are described later.

See the many sample DDE client codes for specific code examples.

Once the DDE link is established, any analysis that can be done based upon ray tracing may be performed by the server, and the data returned to the client.

Comments on the DDE timeout interval

DDE is based upon message-passing. To prevent either the client or the server from waiting forever for a response to a message that may never come, a maximum wait time must be specified. This wait time is specified in the client application. In the sample interface code ZCLIENT, the default wait time is defined by the source code line

```
#define DDE_TIMEOUT 50000
```

This `DDE_TIMEOUT` is in milliseconds, so 50000 ms is 50 seconds. It is critical that no DDE item command requires longer than this length of time to complete. If the computation requested exceeds this time, the message passing will fail and invalid data may result. The default time limit may be edited to a larger value, however, more care must be taken in programming reliable code as the client and server applications may appear to "hang" if the message passing is not done correctly.

Terminating the link

When all computations are complete, and the DDE link is to be terminated by the client, the client sends a termination message to the ZEMAX server. This message will terminate only the server portion of ZEMAX, not ZEMAX itself.

Extracting data from ZEMAX

ZEMAX supports a number of capabilities under DDE. Each function is given a name, called the "item", and this item is passed to the ZEMAX server using the WM_DDE_REQUEST message. Any data ZEMAX requires (for example to trace a ray) is passed to ZEMAX by encoding within the item name. Most items in ZEMAX have names that require no further encoding, but some items have coded values appended to the name. The values are appended to the item name and are delimited by commas.

ZEMAX responds to each item request with a WM_DDE_DATA message, and the message contains a pointer to a block of memory holding a string. Most data is passed from ZEMAX back to the client application in a string, this corresponds to CF_TEXT format. The client application must extract the string and free the memory for the string.

See the many sample DDE client codes for specific code examples.

The data items

Here are the supported data items:

CloseUDOData

This item is used to close the user defined operand buffer, which allows the ZEMAX optimizer to proceed. See "SetUDOIItem" on page 775 for details. The syntax for CloseUDOData is:

`CloseUDOData, buffercode`

See also GetUDOSystem and SetUDOIItem.

DeleteConfig

This item deletes an existing configuration in the multi-configuration editor. The syntax is:

`DeleteConfig, config`

See also InsertConfig.

DeleteMCO

This item deletes a multi-configuration operand in the multi-configuration editor. The syntax is:

`DeleteMCO, operand`

The operand argument is an integer between 1 and the current number of operands, inclusive. The return value is the new number of operands. See also InsertMCO.

DeleteMFO

This item deletes an optimization operand in the merit function editor. The syntax is:

`DeleteMFO, operand`

The operand argument is an integer between 1 and the current number of operands, inclusive. The return value is the new number of operands. See also InsertMFO.

DeleteObject

This item deletes an existing object. The syntax is:

`DeleteObject, surf, object`

See also InsertObject.

DeleteSurface

This item deletes an existing surface. The syntax is:

`DeleteSurface, surf`

See also InsertSurface.

ExportCAD

This item export lens data in IGES/STEP/SAT format. The syntax is:

ExportCAD, filename, filetype, num spline, first, last, rayslayer, lenslayer, exportdummy, usesolids, raypattern, numrays, wave, field, deletevignetted, dummythick, split, scatter, usepol, config

For details on this feature and the meaning of all parameters, see "EXPORTCAD" on page 676, and "Export IGES/SAT/STEP Solid" on page 257.

There is a complexity in using this feature via DDE. The export of lens data may take a long time relative to the timeout interval of the DDE communication. Therefore, calling this data item will cause ZEMAX to launch an independent thread to process the request. Once the thread is launched, the return value is the string "Exporting filename". However, the actual file may take many seconds or minutes to be ready to use. To verify that the export is complete and the file is ready, use the ExportCheck data item. ExportCheck will return 1 if the export is still running, or 0 if it has completed. Generally, the ExportCheck data item call will need to be placed in a loop which executes until ExportCheck returns 0. A typical loop test in C code might look like this:

```
/* check if the export is done */
still_working = 1;
while(still_working)
{
    // Delay for 200 milliseconds
    Sleep(200);
    PostRequestMessage("ExportCheck", szBuffer);
    if (szBuffer[0] == '1')
    {
        // Still working!
    }
    else
    {
        // All done!
        still_working = 0;
    }
}
```

If ExportCad returns "BUSY!" then the export thread is already running from a previous request, and the export will not be performed.

ExportCheck

This item indicates the status of the last executed ExportCAD data item. The return value is 1 if still running, or 0 if it has completed. See ExportCad above for a discussion of the use of this data item.

FindLabel

This item finds the surface that has the integer label associated with the specified surface. The syntax is

FindLabel, label

The returned value is the surface number of the first surface with the identical integer label, or -1 if no surface has the specified label. See also SetLabel and GetLabel.

GetAddress

This item extracts the address line 1, 2, 3 or 4 from the environment settings. The syntax is

GetAddress, 1

with similar commands for lines 2, 3, and 4. If the user has selected to hide the address box, then all strings return as "\r\n".

GetAperture

This item extracts surface aperture data. The syntax is

GetAperture, surf

The returned string is formatted as follows:

"type, min, max, xdecenter, ydecenter, aperturefilename"

This item returns the type as an integer code; 0 for no aperture, 1 for circular aperture, 2 for circular obscuration, 3 for spider, 4 for rectangular aperture, 5 for rectangular obscuration, 6 for elliptical aperture, 7 for elliptical obscuration, 8 for user defined aperture, 9 for user defined obscuration, and 10 for floating aperture. The min and max values have different meanings for the elliptical, rectangular, and spider apertures than for circular apertures; see "Aperture type and other aperture controls" on page 79 for details. See also SetAperture.

GetApodization

This item computes the intensity apodization of a ray from the apodization type and value. The syntax is:

GetApodization, px, py

Where px and py are the normalized pupil coordinates. The returned string is the intensity apodization.

GetAspect

This item extracts the graphic display aspect ratio and the width or height of the printed page in current lens units. For example, if the current aspect ratio is 3 x 4, the aspect ratio returned will be 0.75. Knowing the correct aspect ratio is required when drawing isometric plots. The format of the returned data is aspect, width. If the aspect ratio is greater than 1, then the plot is taller than it is wide, and the format of returned data is aspect, height.

The syntax is:

GetAspect, filename

where filename is the name of the temporary file associated with the window being created or updated. If the temporary file name is left off; then the default aspect ratio and width (or height) is returned.

GetBuffer

The GetBuffer item is used to retrieve client specific data from a window being updated. The syntax is:

GetBuffer, n, tempfile

where n is the buffer number, which must be between 0 and 15 inclusive; and tempfile is the name of the temporary file for the window being updated. The tempfile name is passed to the client when ZEMAX calls the client; see the discussion "How ZEMAX calls the client" for details. Note each window may have it's own buffer data, and ZEMAX uses the filename to identify the window for which the buffer contents are required. See also SetBuffer.

GetComment

This item extracts the surface comment, if any. The surface for which the comment data is returned is appended to the item name, for example, to get the comment for surface 5, the item name should be "GetComment, 5".

GetConfig

This item extracts the current configuration number, the number of configurations, and the number of multiple configuration operands.

The returned string is formatted as follows:

"currentconfig, numberconfig, numbermcooper"

See also SetConfig.

GetDate

This item extracts the current date and time, in the format selected by the user in the ZEMAX environment dialog box.

GetExtra

This item extracts the extra surface data. The syntax is

GetExtra, surf, column

The returned string is formatted as follows:

"value"

See also SetExtra.

GetField

This item extracts field data. The syntax is

GetField, n

where n is zero or the field number. If n is zero, the returned string is formatted as follows:

"type, number, max_x_field, max_y_field, normalization_method"

The parameter type is an integer; either 0, 1, 2, or 3 for angles in degrees, object height, paraxial image height, or real image height, respectively. The parameter number is the number of fields currently defined. The max_x_field and max_y_field values are used to normalize field coordinates, see "Normalized field coordinates" on page 55. The field normalization method is 0 for radial and 1 for rectangular.

If n is not zero, but corresponds to a valid field number, the returned string is formatted as follows:

"xfield, yfield, weight, vdx, vdy, vcx, vcy, van"

These 8 values are exponential format floating point numbers for the various field data. See also SetField.

GetFile

This item extracts the full name of the lens, including the drive and path. Extreme caution should be used if the file is to be tampered with; since at any time ZEMAX may read or write from/to this file.

GetFirst

This item extracts first order data about the lens. The returned string is formatted as follows:

"focal, pwfn, rwfn, pima, pmag"

Values in the string are the EFL, paraxial working F/#, real working F/#, paraxial image height, and paraxial magnification.

GetGlass

This item extracts some data about the glass on any surface. The syntax is

GetGlass, surf

If the specified surface is not valid, is not made of glass, or is gradient index, the returned string is empty. This data may be meaningless for glasses defined only outside of the FdC band. Otherwise, the returned string is formatted as follows:

"name, nd, vd, dpgf"

GetGlobalMatrix

This item returns the matrix required to convert any local coordinates (such as from a ray trace) into global coordinates. For details on the global coordinate matrix, see "Global Coordinate Reference Surface" on page 111.

The syntax for this item is:

GetGlobalMatrix, surf

The returned data string is in the format

"R11,R12,R13,R21,R22,R23,R31,R32,R33,Xo,Yo,Zo"

GetIndex

This item extracts index of refraction data for any surface. The syntax is

GetIndex, surf

If the specified surface is not valid, or is gradient index, the returned string is empty. Otherwise, the returned string is formatted as follows:

"n1, n2, n3..."

where the index values correspond to the index at each defined wavelength, in order.

GetLabel

This item retrieves the integer label associated with the specified surface. The syntax is

```
GetLabel, surf
```

The returned value is the label. See also SetLabel and FindLabel.

GetMetaFile

This item creates a Windows Metafile of any ZEMAX graphical analysis plot. The syntax is

```
GetMetaFile, "metafilename", type, "settingsfilename", flag
```

The metafilename must be in quotes and include the full path, name, and extension for the metafile.

The type argument is a 3 character case-sensitive label that indicates the type of analysis to be performed. The 3 letter labels are identical to those used for the button bar in ZEMAX. A list of codes may be found on the "Buttons" tab of the File, Preferences dialog box. The labels are case sensitive. If no label is provided or recognized, a 3D Layout plot will be generated.

If a valid file name is used for the "settingsfilename" argument, ZEMAX will use or save the settings used to compute the metafile graphic, depending upon the value of the flag parameter.

If the flag value is 0, then the default settings will be used. If the lens file has its own default settings, then those will be used; these are the settings stored in the "lensfilename.cfg" file. If no lens specific default settings exist, then the default settings for all ZEMAX files, stored in the file "ZEMAX.CFG" will be used, if any. If no previous settings have been saved for this or any other lens, then the default settings used are the "factory" defaults used by ZEMAX.

If the flag value is 1, then the settings provided in the settings file, if valid, will be used to generate the file. If the data in the settings file is in anyway invalid, then the default settings will be used to generate the file. The only valid settings files are those generated by ZEMAX, then renamed and saved to a new user defined file name. For example, on the Spot Diagram settings dialog, pressing "Save" will generate a "SPT.CFG" file with the saved settings. If this file is renamed to "MySPT.CFG", then the file, with a full path, may be used as the "settingsfilename" argument to GetTextFile.

If the flag value is 2, then the settings provided in the settings file, if valid, will be used and the settings box for the requested feature will be displayed. After the user makes any changes to the settings the graphic will then be generated using the new settings. The dialog boxes used to change the analysis settings use the data from the lens currently in the Lens Data Editor, so the item PushLens must be used prior to using flag = 2.

No matter what the flag value is, if a valid file name is provided for the settingsfilename, the settings used will be written to the settings file, overwriting any data in the file.

Only graphic, and not text files, are supported by GetMetaFile.

See also GetTextFile and OpenWindow.

GetMulticon

This item extracts data from the multi-configuration editor. The syntax is

```
GetMulticon, config, row
```

If the config number is greater than zero and less than or equal to the number of configurations, the returned string is formatted as follows:

```
"value, num_config, num_row, status, pickuprow, pickupconfig, scale, offset"
```

If the config number is zero, the multicon operand type data is returned as follows:

```
"operand_type, number1, number2, number3"
```

The numbers are used as described in "Summary of multi-configuration operands" on page 579. The status integer is 0 for fixed, 1 for variable, 2 for pickup, and 3 for thermal pickup. If status is 2 or 3, the pickuprow and pickupconfig values indicate the source data for the pickup solve.

See also SetMulticon.

GetName

This item extracts the name of the lens. The returned string is the current lens name as entered on the General data dialog box.

GetNSCData

This item gets the various data for NSC groups. The syntax is:

`GetNSCData, surf, code`

Surf refers to the surface number of the NSC group, use 1 if the program mode is Non-Sequential. Currently only code = 0 is supported, in which case the returned data is the number of objects in the NSC group.

GetNSCMatrix

This item returns the rotation and position matrices for NSC objects relative to the NSC surface origin.

The syntax for this item is:

`GetNSCMatrix, surf, object`

The returned data string is in the format

“R11,R12,R13,R21,R22,R23,R31,R32,R33,Xo,Yo,Zo”

GetNSCObjectData

This item gets the various data for NSC objects. The syntax is:

`GetNSCObjectData, surf, object, code`

Surf and object refer to the surface number and object number. The code is one of the integer values in the table below; the returned value is the data item indicated. See also `SetNSCObjectData`.

Code	Data returned by GetNSCObjectData
0	Object type name. (string)
1	Comment, which also defines the file name if the object is defined by a file. (string)
2	Color. (integer)
5	Reference object number. (integer)
6	Inside of object number. (integer)
The following codes set values on the Type tab of the Object Properties dialog.	
3	1 if object uses a user defined aperture file, 0 otherwise. (integer)
4	User defined aperture file name, if any. (string)
29	Sets the "Use Pixel Interpolation" checkbox. Use 1 for checked, 0 for unchecked.
The following codes set values on the Sources tab of the Object Properties dialog.	
101	Sets the source object random polarization. Use 1 for checked, 0 for unchecked.
102	Sets the source object reverse rays option. Use 1 for checked, 0 for unchecked.
103	Sets the source object Jones X value.
104	Sets the source object Jones Y value.
105	Sets the source object Phase X value.
106	Sets the source object Phase Y value.
107	Sets the source object initial phase in degrees value.

Code	Data returned by GetNSCObjectData
108	Sets the source object coherence length value.
109	Sets the source object pre-propagation value.
110	Sets the source object sampling method; 0 for random, 1 for Sobol sampling.
111	Sets the source object bulk scatter method; 0 for many, 1 for once, 2 for never.
The following codes set values on the Bulk Scatter tab of the Object Properties dialog.	
202	Sets the Mean Path value.
203	Sets the Angle value.
211-226	Sets the DLL parameter 1-16, respectively.

GetNSCObjectFaceData

This item gets the various data for NSC object faces. The syntax is:

`GetNSCObjectFaceData, surf, object, face, code`

Surf and object refer to the surface number and object number. Face is the face number. The code is one of the integer values in the table below; the returned value is the data item indicated. See also `SetNSCObjectFaceData`.

Code	Data returned by GetNSCObjectFaceData
10	Coating name. (string)
20	Scatter code. Use 0 for None, 1 for Lambertian, 2 for Gaussian, 3 for ABg, and 4 for user defined. (integer)
21	Scatter fraction. (double)
22	Number of rays to scatter. (integer)
23	Gaussian scatter sigma. (double)
24	Face is setting; use 0 for object default, 1 for reflective, 2 for absorbing. (integer)
30	ABg scatter profile name for reflection. (string)
31	ABg scatter profile name for transmission. (string)
40	User Defined Scatter DLL name. (string)
41-46	User Defined Scatter Parameter 1 - 6. (double)
60	User Defined Scatter data file name. (string)

GetNSCParameter

This item gets the parameter data for NSC objects. The syntax is:

`GetNSCParameter, surf, object, parameter`

Surf and object refer to the surface number and object number, and parameter is the integer parameter number. The returned string is the parameter value. See also `SetNSCParameter`.

GetNSCPosition

This item gets the position data for NSC objects. The syntax is:

`GetNSCPosition, surf, object`

Surf and object refer to the surface number and object number. The returned string is

"x, y, z, tilt-x, tilt-y, tilt-z, material"

See also SetNSCPosition.

GetNSCProperty

This item mimics the ZPL function NPRO (see "NPRO(surf, object, code, face)" on page 661). The syntax is:

GetNSCProperty surface, object, code, face

Surface and object refer to the surface number and object number, and code refers to the integer code used to identify specific types of NSC object data as described in "SETNSCPROPERTY" on page 721. The face argument is used by some codes to identify a particular object face. GetNSCProperty can return both text and numeric data. The return string format is

"data"

where data is either the string or numeric data retrieved.

See also SetNSCProperty

GetNSCSettings

This item extracts the maximum number of intersections, segments, nesting level, minimum absolute intensity, minimum relative intensity, glue distance, miss ray distance, and ignore errors flag used for NSC ray tracing. The data is returned in a string as follows:

"Max_Intersections, Max_Segments, Max_Nesting, Min_Absolute_Intensity, Min_Relative_Intensity, Glue_Distance, Miss_Ray_Length, Ignore_Errors"

Ignore_Errors is 1 if true, 0 if false. See also SetNSCSettings.

GetOperand

This item gets the operand data from the Merit Function Editor. The syntax is:

GetOperand, row, column

Row is the operand row number in the Merit Function Editor. Column is 1 for the operand type as a string, such as "EFFL", 2 for int1, 3 for int2, 4-7 for data1-data4, 8 for target, 9 for weight, 10 for value, 11 for percent contribution, and 12-13 for data5-data6. The returned string is the numeric or string value of the requested parameter. To update the merit function prior to calling GetOperand, use the Optimize item with the number of cycles set to -1. See also SetOperand and Optimize.

GetPath

This item extracts the full path name to the <data> folder, and the path name to the default folder for lenses. The returned strings are separated by a comma.

GetPolState

This item extracts the default polarization state set by the user. The data is formatted as follows:

"nIsPolarized, Ex, Ey, Phax, Phay"

If nIsPolarized is anything other than zero, then the default polarization state is unpolarized. Otherwise, the Ex, Ey, Phax, and Phay values are used to define the polarization state. Ex and Ey should each be normalized to a magnitude of unity, although this is not required. Phax and Phay are in degrees. See also SetPolState.

GetPolTrace

This item is very similar to GetTrace, with the added capability to trace polarized rays through the system. See the GetTrace documentation for details. The syntax for the GetPolTrace item is as follows:

GetPolTrace, wave, mode, surf, hx, hy, px, py, Ex, Ey, Phax, Phay

The arguments are identical to GetTrace, except for the additional Ex, Ey, Phax, and Phay arguments. Ex and Ey are the normalized electric field magnitudes in the x and y directions. The quantity $Ex^2 + Ey^2$ should have a value of 1.0 (with an important exception described below) although any values are accepted. Phax and Phay are the relative phase, in degrees.

If Ex, Ey, Phax, and Phay are all zero, and only in this case, then ZEMAX assumes an "unpolarized" ray trace is required. An unpolarized ray trace actually requires ZEMAX to trace two orthogonal rays and the resulting transmitted intensity be averaged. If any of the four values are not zero, then a single polarized ray will be traced.

For example, to trace the real unpolarized marginal ray to the image surface at wavelength 2, the item string would be

```
GetPolTrace, 2, 0, -1, 0.0, 0.0, 0.0, 1.0, 0, 0, 0, 0
```

For polarized rays, the data message comes back in the following format:

```
"error, intensity, Exr, Eyr, Ezr, Exi, Eyi, Ezi"
```

The integer error will be zero if the ray traced successfully, otherwise it will be a positive or negative number. If positive, then the ray missed the surface number indicated by error. If negative, then the ray total internal reflected (TIR) at the surface given by the absolute value of the error number. Always check to verify the ray data is valid before using the rest of the string!

The intensity will be the transmitted intensity. It is always normalized to an input electric field intensity of unity. The transmitted intensity accounts for surface, thin film, and bulk absorption effects, but does not consider whether or not the ray was vignetted. The Ex, Ey, and Ez values are the electric field components, with the r and i characters denoting the real and imaginary portions.

For unpolarized rays, the data message is:

```
"error, intensity"
```

Although GetPolTrace is easy to program and use, there is a significant disadvantage to using GetPolTrace: only one ray per DDE call is traced. The overhead of passing data via DDE can be large compared to tracing the ray, and so if a large number of rays are required to be traced, the execution may be relatively slow. For information on tracing large numbers of rays, see the discussion "Tracing large numbers of rays" elsewhere in this chapter. See also GetPolTraceDirect.

GetPolTraceDirect

GetPolTraceDirect provides the same direct access to the ZEMAX ray tracing engine that GetTraceDirect does, and allows for polarized ray tracing as well. See GetTraceDirect for important details on the direct access version of GetTrace. The syntax for the GetPolTraceDirect item is:

```
GetPolTraceDirect, wave, mode, startsurf, stopsurf, x, y, z, l, m, n, Ex, Ey, Phax, Phay
```

The Ex, Ey, Phax, and Phay parameters are exactly the same as defined for GetPolTrace.

The data message comes back in the same format as described in GetPolTrace.

GetPupil

This item extracts pupil data. The returned string is formatted as follows:

```
"type, value, ENPD, ENPP, EXPD, EXPP, apodization_type, apodization_factor"
```

The parameter type is an integer indicating the system aperture type, a number between 0 and 5, for entrance pupil diameter, image space F/#, object space NA, float by stop size, paraxial working F/#, or object cone angle, respectively. The value is the system aperture value, unless float by stop size is being used, in which case the value is the stop surface semi-diameter. The next 4 values are the entrance pupil diameter, entrance pupil position, exit pupil diameter, and exit pupil position, all in lens units. The apodization_type is an integer which is set to 0 for none, 1 for Gaussian, 2 for Tangential. The apodization_factor is the number shown on the general data dialog box.

GetRefresh

This data item causes ZEMAX to copy the lens data from the LDE into the stored copy of the server. The lens is then updated, which means ZEMAX recomputes all pupil positions, solves, and index data. If the lens can be updated, ZEMAX returns the string "0", otherwise, it returns "-1". If the GetUpdate returns "-1", no ray tracing can be performed.

All subsequent commands will now affect or be executed upon the newly copied lens data. The old lens data, if any, cannot be recovered. See also GetUpdate and PushLens.

GetSag

This item computes the sag of any surface. The item name is formatted as "GetSag, surf, x, y", where surf is the surface number, and x and y are the coordinates on the surface for which the sag is computed. The returned string is formatted as "sag, alternatesag". The values for x, y, and the sag are all in lens units.

GetSequence

This item returns the sequence number of the lens in the Server's memory, then the sequence number of the lens in the LDE, separated by a comma.

GetSerial

This item returns the serial number of the ZEMAX key.

GetSettingsData

This item retrieves the settings data used by a window. The data must have been previously stored by a call to SetSettingsData. The data may have been stored by a previous execution of the client program. The syntax is:

GetSettingsData, tempfile, number

The tempfile data is the name of the output file passed by ZEMAX to the client. ZEMAX uses this name to identify for which window the GetSettingsData request is for.

The number is the data number used by the previous SetSettingsData call. Currently, only number = 0 is supported. This number may be used to expand the feature in the future.

The return value is the string that was saved by a previous SetSettingsData command for the window and number.

See also "How ZEMAX calls the client" on page 781 and SetSettingsData.

GetSolve

This item returns data about any solve on any surface. The syntax is "GetSolve, surface, code" where code is an integer code indicating which surface parameter the solve data is for. The solve data is returned in the following formats, depending upon the code value.

GetSolve Code	Returned data format
0, curvature	solvetype, parameter1, parameter2
1, thickness	solvetype, parameter1, parameter2, parameter3
2, glass	solvetype (for solvetype = 0) solvetype, Index, Abbe, Dpgf (for solvetype = 1, model glass) solvetype, pickupsurf (for solvetype = 2, pickup) solvetype (for solvetype = all other values)
3, semi-diameter	solvetype, pickupsurf
4, conic	solvetype, pickupsurf
5-16, parameters 1-12	solvetype, pickupsurf, offset, scalefactor, pickupcolumn
17, parameter 0	solvetype, pickupsurf, offset, scalefactor, pickupcolumn
1001+, extra data values 1+	solvetype, pickupsurf, scalefactor

The solvetype is an integer code, and the parameters have meanings that depend upon the solve type; see the chapter "SOLVES" on page 473 for details. See also SetSolve.

GetSurfaceData

This item gets surface data. The syntax is:

GetSurfaceData, surf, code

or

GetSurfaceData, surf, code, arg2

Surf refers to the surface number. The code is one of the integer values in the table below; the returned value is the data item indicated. See also SetSurfaceData. Not all codes use the arg2 value.

Code	Date returned by GetSurfaceData
0	Surface type name. (string)
1	Comment. (string)
2	Curvature (not radius).
3	Thickness.
4	Glass. (string)
5	Semi-Diameter.
6	Conic.
7	Coating. (string)
8	Thermal Coefficient of Expansion (TCE).
9	User Defined Surface DLL name. (string)
20	Ignore surface flag, 0 for not ignored, 1 for ignored.
51	Before tilt and decenter order; 0 for Decenter then Tilt, 1 for Tilt then Decenter.
52-56	Before decenter x, decenter y, tilt x, tilt y, and tilt z values, respectively.
60	After status; 0 for explicit, 1 for pickup this surface, 2 for reverse this surface, 3 for pickup this surface - 1, 4 for reverse this surface - 1, etc.
61	After tilt and decenter order; 0 for Decenter then Tilt, 1 for Tilt then Decenter.
62-66	After decenter x, decenter y, tilt x, tilt y, and tilt z values, respectively.
70	Use Layer Multipliers and Index Offsets. Use 1 for true, 0 for false.
71	Layer Multiplier value. The coating layer number is defined by arg2.
72	Layer Multiplier status. Use 0 for fixed, 1 for variable, or n+1 for pickup from layer n. The coating layer number is defined by arg2.
73	Layer Index Offset value. The coating layer number is defined by arg2.
74	Layer Index Offset status. Use 0 for fixed, 1 for variable, or n+1 for pickup from layer n. The coating layer number is defined by arg2.
75	Layer Extinction Offset value. The coating layer number is defined by arg2.
76	Layer Extinction Offset status. Use 0 for fixed, 1 for variable, or n+1 for pickup from layer n. The coating layer number is defined by arg2.

For example, to get the conic data for surface 8, the item name should be "GetSurfaceData, 8, 6". To get the coating multiplier for layer 4 on surface 7, the item name should be "GetSurfaceData, 7, 71, 4".

See also SetSurfaceData and GetSurfaceParameter.

GetSurfaceDLL

This item extracts the name of the DLL if the surface is a user defined type. The surface for which the surface DLL name is returned is appended to the item name, for example, to get the DLL name for a user defined surface 5, the item name should be "GetSurfaceDLL, 5".

The returned string is formatted as follows:

"dllname, surfacename"

The dllname is the name of the defining DLL, the surfacename is the string displayed by the DLL in the surface type column of the Lens Data Editor.

GetSurfaceParameter

This item extracts surface parameter data. The surface and parameter number to be returned is appended to the item name, for example, to get the parameter 5 data for surface 8, the item name should be "GetSurfaceParameter, 8, 5".

See also GetSurfaceData and SetSurfaceParameter.

GetSystem

This item extracts system data.

The returned string is formatted as follows:

"numsurfs, unitcode, stopsurf, nonaxialflag, rayaimingtype, adjust_index, temp, pressure, globalrefsurf, need_save"

This item returns the number of surfaces, the unit code (0, 1, 2, or 3 for mm, cm, in, or M), the stop surf number, a flag to indicate if the system is non-axial symmetric (0 for axial, 1 if not axial), the ray aiming (0 for off, 1 for paraxial, 2 for real), the adjust index data to environment flag (0 if false, 1 if true), the current temperature and pressure, and the global coordinate reference surface number. The need_save value indicates whether the file has been modified

See also SetSystem and GetSystemAper.

GetSystemAper

This item extracts system aperture data.

The returned string is formatted as follows:

"type, stopsurf, aperture_value"

This item returns the system aperture type (0, 1, 2, 3, 4, or 5 for Entrance Pupil Diameter, Image Space F/#, Object Space NA, Float by Stop Size, Paraxial Working F/#, or Object Cone Angle, respectively), the stop surface number, and the system aperture value. If the aperture type is Float by Stop Size, the aperture value is the stop surface semi-diameter.

See also GetSystem and SetSystemAper.

GetSystemProperty

This item mimics the ZPL function SYPR (see "SYPR(code)" on page 667). The syntax is:

GetSystemProperty code

Code is an integer value that defines the specific system property requested. For a listing of these codes, see "SETSYSTEMPROPERTY, SYSP" on page 728. The return string format is

"data"

where data is either the string or numeric data retrieved.

See also SetSystemProperty

GetTextFile

This item creates a text file of any ZEMAX analysis that supports text. The syntax is

GetTextFile, "textfilename", type, "settingsfilename", flag

The textfilename must be in quotes and include the full path, name, and extension for the file to be created.

The type argument is a 3 character case-sensitive label that indicates the type of analysis to be performed. The 3 letter labels are identical to those used for the button bar in ZEMAX. A list of codes may be found on the "Buttons" tab of the File, Preferences dialog box. The labels are case sensitive. If no label is provided or recognized, a standard ray trace will be generated.

If a valid file name is used for the "settingsfilename" argument, ZEMAX will use or save the settings used to compute the text file, depending upon the value of the flag parameter.

If the flag value is 0, then the default settings will be used. If the lens file has it's own default settings, then those will be used; these are the settings stored in the "lensfilename.cfg" file. If no lens specific default settings exist, then the default settings for all ZEMAX files, stored in the file "ZEMAX.CFG" will be used, if any. If no previous settings have been saved for this or any other lens, then the default settings used are the "factory" defaults used by ZEMAX.

If the flag value is 1, then the settings provided in the settings file, if valid, will be used to generate the file. If the data in the settings file is in anyway invalid, then the default settings will be used to generate the file. The only valid settings files are those generated by ZEMAX, then renamed and saved to a new user defined file name. For example, on the Spot Diagram settings dialog, pressing "Save" will generate a "SPT.CFG" file with the saved settings. If this file is renamed to "MySPT.CFG", then the file, with a full path, may be used as the "settingsfilename" argument to GetTextFile.

If the flag value is 2, then the settings provided in the settings file, if valid, will be used and the settings box for the requested feature will be displayed. After the user makes any changes to the settings the file will then be generated using the new settings. The dialog boxes used to change the analysis settings use the data from the lens currently in the Lens Data Editor, so the item PushLens must be used prior to using flag = 2.

No matter what the flag value is, if a valid file name is provided for the settingsfilename, the settings used will be written to the settings file, overwriting any data in the file.

Only text, and not graphic files, are supported by GetTextFile.

See also GetMetaFile and OpenWindow.

GetTol

This item extracts tolerance data. The syntax is

```
GetTol, n
```

where n is zero or the tolerance operand number. If n is zero, the returned string is formatted as follows:

```
"number"
```

where number is the number of tolerance operands defined. If n is not zero, but corresponds to a valid tolerance operand number, the returned string is formatted as follows:

```
"toltype, int1, int2, min, max, int3"
```

GetTrace

This item requires that the client provide additional data. In order to trace a ray, ZEMAX needs to know the relative field and pupil coordinates, the wavelength, the mode (either real, mode = 0 or paraxial, mode = 1) as well as the surface to trace the ray to. All of this data is encoded and appended to the item name. This is accomplished by delimiting the various parameters with commas as follows:

```
GetTrace, wave, mode, surf, hx, hy, px, py
```

For example, to trace the real chief ray to surface 5 at wavelength 3, the item string would be

```
GetTrace, 3, 0, 5, 0.0, 1.0, 0.0, 0.0
```

Although this looks cumbersome, it is easy to program. ZEMAX receives the item, recognizes the "GetTrace" at the beginning of the name, strips off the rest of the string, then parses it, and traces the ray. Usually, the ray data is only needed at the image surface; setting the surface number to -1 will yield data at the image surface. Note 0 is reserved for the object surface.

The data message comes back in the following format:

```
"error, vigcode, x, y, z, l, m, n, l2, m2, n2, intensity"
```

The integer error will be zero if the ray traced successfully, otherwise it will be a positive or negative number. If positive, then the ray missed the surface number indicated by error. If negative, then the ray total internal reflected (TIR) at the surface given by the absolute value of the error number. Always check to verify the ray data is valid before using the rest of the string!

The parameter vigcode is the first surface where the ray was vignetted. Unless an error occurs at that surface or subsequent to that surface, the ray will continue to trace to the requested surface.

The x, y, and z values are the coordinates on the requested surface.

The l, m, and n values are the direction cosines after refraction into the media following the requested surface.

The l2, m2, and n2 values are the surface intercept direction normals at the requested surface.

The intensity is the relative transmitted intensity of the ray, including any pupil or surface apodization defined.

Although GetTrace is easy to program and use, there is a significant disadvantage to using GetTrace: only one ray per DDE call is traced. The overhead of passing data via DDE can be quite large compared to tracing the ray, and so if a large number of rays are required to be traced, the execution may be relatively slow. For information on tracing large numbers of rays, see the discussion "Tracing large numbers of rays" elsewhere in this chapter. See also GetTraceDirect.

GetTraceDirect

GetTraceDirect provides a more direct access to the ZEMAX ray tracing engine. Normally, rays are defined by the normalized field and pupil coordinates hx, hy, px, and py. ZEMAX takes these normalized coordinates and computes the object coordinates (x, y, and z) and the direction cosines to the entrance pupil aim point (l, m, and n; for the x-, y-, and z-direction cosines, respectively).

However, there are times when it is more appropriate to trace rays by direct specification of x, y, z, l, m, and n. The direct specification has the added flexibility of defining the starting surface for the ray anywhere in the optical system. This flexibility comes at the price of requiring the client program to carefully ensure that the starting ray coordinates are meaningful.

Like GetTrace, this item requires that the client provide additional data. In order to trace a ray, ZEMAX needs to know x, y, z, l, m, n, the wavelength, the mode (either real, mode = 0 or paraxial, mode = 1) as well as the starting and stopping surfaces to trace the ray from and to. All of this data is encoded and appended to the item name. This is accomplished by delimiting the various parameters with commas as follows:

```
GetTraceDirect, wave, mode, startsurf, stopsurf, x, y, z, l, m, n
```

The data message comes back in the following format:

```
"error, vigcode, x, y, z, l, m, n, l2, m2, n2, intensity"
```

where the parameters are exactly the same as for GetTrace, except for intensity. The intensity is the relative transmitted intensity of the ray, excluding any pupil apodization defined. Note GetTrace includes pupil apodization, GetTraceDirect does not. Both include surface apodizations.

Although GetTraceDirect is easy to program and use, there is a significant disadvantage to using GetTraceDirect: only one ray per DDE call is traced. The overhead of passing data via DDE can be quite large compared to tracing the ray, and so if a large number of rays are required to be traced, the execution may be relatively slow. For information on tracing large numbers of rays, see the discussion "Tracing large numbers of rays" elsewhere in this chapter. See also GetTrace.

GetUDOSystem

This item is used to load a particular lens from the optimization function memory into the ZEMAX server's memory. The only time this item name should be used is when implementing a User Defined Operand, or UDO. UDO's are described in detail in the chapter "Optimization". See that chapter for details. The syntax for GetUDOSystem is

```
GetUDOSystem, buffercode
```

where buffercode is an integer passed to the UDO on the command line. See also SetUDOData.

GetUpdate

This data item causes ZEMAX to update the lens, which means ZEMAX recomputes all pupil positions, solves, and index data. If the lens can be updated, ZEMAX returns the string "0", otherwise, it returns "-1". If the GetUpdate returns "-1", no ray tracing can be performed. To update the merit function, use the Optimize item with the number of cycles set to -1. See also GetRefresh and Optimize.

GetVersion

This item returns the version number of ZEMAX. This number is generally a 5 digit integer with no decimal place, such as 10000 for version 10.0. Later versions will have larger numbers.

GetWave

This item extracts wavelength data. The syntax is

GetWave, n

where n is zero or the wavelength number. If n is zero, the returned string is formatted as follows:

"primary, number"

The parameter primary is an integer which indicates which of the wavelengths is the primary one. The parameter number is the number of wavelengths currently defined. If n is not zero, but corresponds to a valid wavelength number, the returned string is formatted as follows:

"wavelength, weight"

These 2 values are exponential format floating point numbers which correspond to the value and weight for the specified wavelength. See also SetWave.

Hammer

Hammer calls the ZEMAX Hammer optimizer. For information on Hammer, see "The Hammer algorithm" on page 541. The syntax is

Hammer, n, algorithm

where n is the number of cycles to run, and algorithm is 0 for damped least squares and 1 for orthogonal descent. The return value is the final merit function. If the merit function value returned is 9.0E+009, the optimization failed, usually because the lens or merit function could not be evaluated. If n is less than 1, Hammer updates all operands in the merit function and returns the current merit function, and no optimization is performed. The number of cycles should be kept small enough to allow the algorithm to complete and return before the DDE communication times out, or an error will occur.

ImportExtraData

This item imports extra data and grid surface data values into an existing surface. The syntax is:

ImportExtraData, surf, filename

For details on how ZEMAX imports extra data, see "Extra Data" on page 93. For information on importing grid surface data, see "Importing grid data" on page 314.

InsertConfig

This item inserts a new configuration in the multi-configuration editor. The syntax is:

InsertConfig, config

The new configuration will be placed at the location indicated by the parameter config. See also DeleteConfig.

InsertMCO

This item inserts a new multi-configuration operand in the multi-configuration editor. The syntax is:

InsertMCO, operand

The operand argument is an integer between 1 and the current number of operands plus 1, inclusive. The return value is the new number of operands. See also DeleteMCO.

InsertMFO

This item inserts a new optimization operand in the merit function editor. The syntax is:

`InsertMFO, operand`

The operand argument is an integer between 1 and the current number of operands plus 1, inclusive. The return value is the new number of operands. See also `DeleteMFO`.

InsertObject

This item inserts a new NSC object. The syntax is:

`InsertObject, surf, object`

The new null object will be placed at the location indicated by the parameters `surf` and `object`. See also `SetNSCObjectData` to define data for the new surface and the `DeleteObject` item.

InsertSurface

This item inserts a new surface. The syntax is:

`InsertSurface, surf`

The new surface will be placed at the location indicated by the parameter `surf`. See also `SetSurface` to define data for the new surface and the `DeleteSurface` item.

LoadDetector

Loads the data saved in a file to an NSC Detector Rectangle, Detector Color, Detector Polar, or Detector Volume object. The syntax is:

`LoadDetector, surf, object, filename`

`Surf` is the surface number of the non-sequential group; use 1 when using non-sequential mode. `Object` is the object number of the detector object. The filename may include the full path, if no path is provided the path of the current lens file is used. The extension should be DDR, DDC, DDP, or DDV for Detector Rectangle, Color, Polar, and Volume objects, respectively. The return value is zero if the load was successful, otherwise an error code is returned.

LoadFile

Loads a ZEMAX file into the server. The syntax is:

`LoadFile, filename, append`

Note that loading a file does not change the data displayed in the LDE; the server process has a separate copy of the lens data. The file name to be loaded is appended to the `LoadFile` item name, and must include the full path. For example: "`LoadFile, C:\ZEMAX\SAMPLES\COOKE.ZMX`".

The value `append` may be omitted. If present and greater than zero, then the new file is appended to the current file starting at the surface number defined by the value `append`.

The returned string is the same as for the `GetUpdate` item; after updating the newly loaded lens file. If a value other than 0 is returned, the Update failed, if -999 is returned, the file could not be loaded. See also `GetPath`, `SaveFile`, and `PushLens`.

LoadMerit

Loads a ZEMAX .MF or .ZMX file and extracts the merit function and places it in the lens loaded in the server. The syntax is

`LoadMerit, filename`

Note that loading a merit function file does not change the data displayed in the LDE; the server process has a separate copy of the lens data. The file name to be loaded is appended to the `LoadMerit` item name, and must include the full path. For example: "`LoadMerit, C:\ZEMAX\SAMPLES\MyMerit.MF`". The returned string is formatted as follows:

"number, merit"

where `number` is the number of operands in the merit function and `merit` is the value of the merit function. If the merit function value is 9.00e+009, the merit function cannot be evaluated. See also `Optimize`, `SaveMerit`.

LoadTolerance

Loads a tolerance file previously saved with SaveTolerance places the tolerances in the lens loaded in the server. The syntax is:

```
LoadTolerance, filename
```

If no path is provided in the filename, the <data>\Tolerance folder is assumed. The returned string is formatted as follows:

"number"

where number is the number of tolerance operands loaded.

MakeGraphicWindow

This item notifies ZEMAX that graphic data has been written to a file and may now be displayed as a ZEMAX child window. The primary purpose of this item is to implement user defined features in a client application, that look and act like native ZEMAX features. The item string must be formatted as follows:

```
MakeGraphicWindow, filename, modulename, wintitle, textflag, settingsdata
```

or

```
SetSettingsData, 0, settingsdata
```

```
MakeGraphicWindow, filename, modulename, wintitle, textflag
```

The filename is the full path and file name to the temporary file that holds the graphic data. This must be the same name as passed to the client executable in the command line arguments, if any. The modulename is the full path and executable name of the client program that created the graphic data. The wintitle is the string which defines the title ZEMAX should place in the top bar of the window.

The textflag should be 1 if the client can also generate a text version of the data. Since the current data is a graphic display (it must be if the item is MakeGraphicWindow) ZEMAX wants to know if the "Text" menu option should be available to the user, or if it should be grayed out. If the text flag is 0, ZEMAX will gray out the "Text" menu option and will not attempt to ask the client to generate a text version of the data.

The settings data is a string of values delimited by spaces (not commas) which are used by the client to define how the data was generated. These values are only used by the client, not by ZEMAX. The settings data string holds the options and data that would normally appear in a ZEMAX "settings" style dialog box. The settings data should be used to recreate the data if required. Because the total length of a data item cannot exceed 255 characters, the data item SetSettingsData may be used prior to the call to MakeGraphicWindow to specify the settings data string rather than including the data as part of MakeGraphicWindow. See "How ZEMAX calls the client" on page 781 for more details on the settings data.

A sample item string might look like this:

```
MakeGraphicWindow,C:\TEMP\ZGF001.TMP,C:\ZEMAX\FEATURES\CLIENT.EXE,ClientWindow,1,  
0 1 2 12.55
```

This item indicates that ZEMAX should open a graphic window, display the data stored in the file C:\TEMP\ZGF001.TMP, and that any updates or setting changes can be made by calling the client module C:\ZEMAX\FEATURES\CLIENT.EXE. This client can generate a text version of the graphic, and the settings data string (used only by the client) is "0 1 2 12.55".

MakeTextWindow

This item notifies ZEMAX that text data has been written to a file and may now be displayed as a ZEMAX child window. The primary purpose of this item is to implement user defined features in a client application, that look and act like native ZEMAX features. The item string must be formatted as follows:

```
MakeTextWindow, filename, modulename, wintitle, settingsdata
```

or

```
SetSettingsData, 0, settingsdata
```

```
MakeTextWindow, filename, modulename, wintitle
```

The filename is the full path and file name to the temporary file that holds the text data. This must be the same name as passed to the client executable in the command line arguments, if any. The modulename is the full path and executable name of the client program that created the text data. The wintitle is the string which defines the title ZEMAX should place in the top bar of the window.

The settingsdata is a string of values delimited by spaces (not commas) which are used by the client to define how the data was generated. These values are only used by the client, not by ZEMAX. The settingsdata string holds the options and data that would normally appear in a ZEMAX "settings" style dialog box. The settingsdata should be used to recreate the data if required. Because the total length of a data item cannot exceed 255 characters, the data item SetSettingsData may be used prior to the call to MakeTextWindow to specify the settings data string rather than including the data as part of MakeTextWindow. See "How ZEMAX calls the client" on page 781 for more details on the settings data.

A sample item string might look like this:

```
MakeTextWindow,C:\TEMP\ZGF002.TMP,C:\ZEMAX\FEATURES\CLIENT.EXE,ClientWindow,6 5 4  
12.55
```

This item indicates that ZEMAX should open a text window, display the data stored in the file C:\TEMP\ZGF002.TMP, and that any updates or setting changes can be made by calling the client module C:\ZEMAX\FEATURES\CLIENT.EXE. The settingsdata string (used only by the client) is "6 5 4 12.55".

ModifySettings

ModifySettings is used to change specific options in ZEMAX settings files. The settings files are used by MakeTextWindow and MakeGraphicWindow. The item string syntax is:

```
ModifySettings, filename, type, value
```

The filename is the full name of the settings file, including the path. The type is a mnemonic that defines what option value is being modified. The valid values for type are as defined in the ZPL macro command MODIFY-SETTINGS that serves the same function as ModifySettings does for extensions. See "MODIFYSETTINGS" on page 697 for a complete list of the type codes. The return value is an integer code: 0 for no error, -1 for invalid file, -2 for incorrect version number, and -3 for a file access conflict.

NewLens

This item erases the current lens. The "minimum" lens that remains is identical to the lens in the Lens Data Editor when "File, New" is selected. No prompt to save the existing lens is given.

NSCCoherentData

This item returns data from an NSC detector. The syntax is

```
NSCCoherentData, surface, object, pixel, data
```

where the arguments are exactly the same as defined for the NSDC optimization operand described in "NSDC" on page 509. The return value is the detector data formatted as a string.

NSCDetectorData

This item returns data from an NSC detector. The syntax is

```
NSCDetectorData, surface, object, pixel, data
```

where the arguments are exactly the same as defined for the NSDD optimization operand described in "NSDD" on page 509. The return value is the detector data formatted as a string.

NSCTrace

This item traces rays from one or all NSC sources with various optional arguments. The syntax is

```
NSCTrace, surf, source, split, scatter, usepolar, ignore_errors, no_random_seed, save,  
savefilename, filter
```

where the arguments are exactly the same as defined for the NSTR macro command described in "NSTR" on page 702.

For information about timeout limits when executing potentially long calculations, see "Comments on the DDE timeout interval" on page 750.

OpenWindow

OpenWindow will open a new analysis window on the main ZEMAX screen. The syntax is

OpenWindow, Type

The type argument is a 3 character case-sensitive label that indicates the type of analysis to be performed. The 3 letter labels are identical to those used for the button bar in ZEMAX. A list of codes may be found on the "Buttons" tab of the File, Preferences dialog box. The labels are case sensitive. See also GetMetaFile.

OperandValue

OperandValue returns the value of any optimization operand, even if the operand is not currently in the merit function. The syntax is

OperandValue, type, int1, int2, data1, data2, data3, data4, data5, data6

The type argument is the name of the operand, such as EFFL. The remaining values are the usual values used by the Merit Function Editor; see "Optimization operands" on page 489. Trailing zero values are not required. For example, to retrieve the EFFL at wavelength 4 the syntax is:

OperandValue, EFFL, 0, 4

Optimize

Optimize calls the ZEMAX Damped Least Squares optimizer. The syntax is

Optimize, n, algorithm

where n is the number of cycles to run and algorithm is 0 for Damped Least Squares (the default) and 1 for Orthogonal Descent. For more information on these algorithms see "Performing an optimization" on page 526. The return value is the final merit function. If the merit function value returned is 9.0E+009, the optimization failed, usually because the lens or merit function could not be evaluated. If n is zero, the optimization runs in automatic mode. If n is less than zero (for example, n = -1), Optimize updates all operands in the merit function and returns the current merit function, and no optimization is performed.

For information about timeout limits when executing potentially long calculations, see "Comments on the DDE timeout interval" on page 750.

PushLens

The syntax is:

PushLens, flag

PushLens will take the lens currently loaded in the server's memory and push it into the Lens Data Editor. This operation requires the permission of the user running the ZEMAX program. To check or set this permission status, see the PushLensPermission data item described below.

The returned string is the same as for the GetUpdate item; after updating the newly pushed lens file. If a value other than 0 is returned, the Update failed, if -999 is returned, the lens could not be pushed into the LDE.

If the flag value is zero or is omitted, the open windows are not updated. If the flag value is 1, then all open analysis windows are updated.

See also PushLensPermission, GetPath, GetRefresh, LoadFile, and SaveFile.

PushLensPermission

The syntax is:

PushLensPermission

The return value is 1 if ZEMAX is set to accept PushLens commands, or 0 if extensions are not allowed to use PushLens. The PushLens data item can only succeed if the user has set the appropriate permissions in ZEMAX. To allow a DDE extension to use PushLens, the option "Allow Extensions To Push Lenses" must be checked. This option is found under File, Preferences on the Editors tab, and is described in "Editors" on page 71. The proper use of PushLens is to first call PushLensPermission. If the return value is 0, then the client should display a dialog box instructing the user to turn on the "Allow Extensions To Push Lenses" option before proceeding.

QuickFocus

This feature performs a quick best focus adjustment of the optical system. The syntax is:

`QuickFocus, mode, centroid`

The value mode should be 0, 1, 2, or 3 for RMS spot radius, spot x, spot y, or wavefront OPD. The value for centroid should be 0 or 1 to indicate the RMS should be referenced to the chief ray or image centroid, respectively. The "best" focus is chosen as a wavelength weighted average over all fields. The return value is always "OK".

ReleaseWindow

When ZEMAX calls the client to update or change the settings used by the client function, the menu bar is grayed out on the window to prevent multiple updates or setting changes from being requested simultaneously. Normally, when the client code sends the `MakeTextWindow` or `MakeGraphicWindow`, the menu bar is once again activated. However, if during an update or setting change, the new data cannot be computed, then the window must be released. The `ReleaseWindow` serves just this one purpose. If the user selects "Cancel" when changing the settings, the client code should send a `ReleaseWindow` item to release the lock out of the menu bar. If this command is not sent, the window cannot be closed, which will prevent ZEMAX from terminating normally. The `ReleaseWindow` item takes just one argument: the name of the temporary file. The syntax is:

`ReleaseWindow, filename`

The return value is zero if no window is using the filename, or a positive integer number if the file is being used.

SaveDetector

Saves the data currently on an NSC Detector Rectangle, Detector Color, Detector Polar, or Detector Volume object to a file. The syntax is:

`SaveDetector, surf, object, filename`

Surf is the surface number of the non-sequential group; use 1 when using non-sequential mode. Object is the object number of the detector object. The filename may include the full path, if no path is provided the path of the current lens file is used. The extension should be DDR, DDC, DDP, or DDV for Detector Rectangle, Color, Polar, and Volume objects, respectively. The return value is zero if the save was successful, otherwise an error code is returned.

SaveFile

Saves the lens currently loaded in the server to a ZEMAX file. The file name to be used for the save is appended to the `SaveFile` item name, and must include the full path. For example: "`SaveFile,C:\ZEMAX\SAMPLES\COOKE.ZMX`". The returned string is the same as for the `GetUpdate` item; after updating the newly saved lens file. If a value other than 0 is returned, the Update failed, if -999 is returned, the file could not be saved. See also `GetPath`, `GetRefresh`, `LoadFile`, and `PushLens`.

SaveMerit

Saves the current merit function to a ZEMAX .MF file. The syntax is

`SaveMerit, filename`

The file name must include the full path. For example: "`SaveMerit,C:\TEMP\MyMerit.MF`". The returned string is formatted as follows:

"number"

where number is the number of operands in the merit function. See also `Optimize`, `LoadMerit`.

SaveTolerance

Saves the tolerances of the current lens to a file. The syntax is:

`SaveTolerance, filename`

If no path is provided in the filename, the <data>Tolerance folder is assumed. The returned string is formatted as follows:

"number"

where number is the number of tolerance operands saved.

SetAperture

This item sets surface aperture data. The syntax is

```
SetAperture, surf, type, min, max, xdecenter, ydecenter, aperturefile
```

The returned string is formatted as follows:

```
"type, min, max, xdecenter, ydecenter, aperturefile"
```

This item uses an integer code for the surface aperture type; 0 for no aperture, 1 for circular aperture, 2 for circular obscuration, 3 for spider, 4 for rectangular aperture, 5 for rectangular obscuration, 6 for elliptical aperture, 7 for elliptical obscuration, 8 for user defined aperture, 9 for user defined obscuration, and 10 for floating aperture. The min and max values have different meanings for the elliptical, rectangular, and spider apertures than for circular apertures; see "Aperture type and other aperture controls" on page 79 for details.

If SetAperture is used to set user defined apertures or obscurations, the aperturefile must be the name of a file which lists the x, y, coordinates of the user defined aperture file in a two column format. For more information on user defined apertures, see "User defined apertures and obscurations" on page 80.

See also GetAperture.

SetBuffer

The SetBuffer item is used to store client specific data with the window being created or updated. The buffer data can be used to store user selected options, instead of using the settings data on the command line of the MakeTextWindow or MakeGraphicWindow items. The data must be in a string format. The syntax is:

```
SetBuffer,1,any text you want.....
```

There are 16 buffers provided, numbered 0 through 15, and each can be set using SetBuffer,0,.....; SetBuffer,1,.....etc. The text following the "SetBuffer,n," is the only text that is stored; and it may be a maximum of 240 characters.

Note the buffer data is not associated with any particular window until either the MakeTextWindow or MakeGraphicWindow items are issued. Once ZEMAX receives the MakeTextWindow or MakeGraphicWindow item, the buffer data is then copied to the appropriate window memory, and then may later be retrieved from that window's buffer using GetBuffer. See also GetBuffer.

SetConfig

This item switches the current configuration number and updates the system. The desired configuration is appended to the item name. For example, to switch to configuration 3, the item name is "SetConfig,3".

The returned string is formatted as follows:

```
"currentconfig, numberconfig, error"
```

The currentconfig is the new configuration number, which will be between 1 and the value of numberconfig. Normally, this will be the desired configuration requested in the item name, as long as it was a valid configuration number. The error code is the same as returned by the GetUpdate item, and will be zero if the new current configuration is traceable. See also GetConfig.

SetExtra

This item sets the extra surface data. The syntax is

```
SetExtra, surf, column, value
```

The returned string is formatted as follows:

```
"value"
```

See also GetExtra.

SetField

This item sets field data. The syntax is

```
SetField, 0, type, number, normalization
```

or

SetField, n, xf, yf, wgt, vdx, vdy, vcx, vcy, van

If the value for n is zero, then the field type, total number of fields, and field normalization type is set to the new values. The field type is 0 for angle, 1 for object height, 2 for paraxial image height, and 3 for real image height. The field normalization type is 0 for radial and 1 for rectangular.

If n is a valid field number (between 1 and the number of fields, inclusive) then the field x and y values, field weight, and vignetting factors are all set. See also *GetField*. The returned data is the same as for *GetField, n*.

SetFloat

This item sets all surfaces without surface apertures to have floating apertures. Floating apertures will vignette rays which trace beyond the semi-diameter. The syntax is

SetFloat

The returned value is "OK".

SetLabel

This item associates an integer label with the specified surface. The label will be retained by ZEMAX as surfaces are inserted or deleted around the target surface. The syntax is

SetLabel, surf, label

The returned value is the label. See also *GetLabel* and *FindLabel*.

SetMulticon

This item sets data in the multi-configuration editor. To set multi-configuration data, the syntax is

SetMulticon, config, row, value, status, pickuprow, pickupconfig, scale, offset

If the config number is zero, *SetMulticon* may be used to set the operand type and number data using this syntax:

SetMulticon, 0, row, operand_type, number1, number2, number3

The returned string is formatted as defined for "*GetMulticon*". The numbers are used as described in "Summary of multi-configuration operands" on page 579.

See also *GetMulticon*.

SetNSCObjectData

This item sets the various data for NSC objects. The syntax is:

SetNSCObjectData, surf, object, code, data

Surf and object refer to the surface number and object number. The code is one of the integer values defined in the table in the discussion for "*GetNSCObjectData*" on page 756. The returned value is identical to that returned by *GetNSCObjectData* after the new data is accepted. See also *SetNSCObjectFaceData*.

SetNSCObjectFaceData

This item sets the various data for NSC object faces. The syntax is:

SetNSCObjectFaceData, surf, object, face, code, data

Surf and object refer to the surface number and object number. Face is in the integer face number. The code is one of the integer values defined in the table in the discussion for "*GetNSCObjectFaceData*" on page 757. The returned value is identical to that returned by *GetNSCObjectFaceData* after the new data is accepted.

SetNSCPosition

This item sets the position data for NSC objects. The syntax is:

SetNSCPosition, surf, object, code, data

Surf and object refer to the surface number and object number. Code is 1-7 for x, y, z, tilt-x, tilt-y, tilt-z, and material, respectively. Data is either a number (for codes 1-6) or a string with the material name (for code 7). The returned value is that of *GetNSCPosition* after the new data has been applied. See also *GetNSCPosition*.

SetNSCParameter

This item sets the parameter data for NSC objects. The syntax is:

```
SetNSCParameter, surf, object, parameter, data
```

Surf and object refer to the surface number and object number. Parameter is the integer parameter number, and data is the new numeric value for that parameter. The returned value is that of GetNSCParameter after the new data has been applied. See also GetNSCParameter.

SetNSCProperty

This item mimics the ZPL keyword SETNSCPROPERTY (see "SETNSCPROPERTY" on page 721). The syntax is:

```
SetNSCProperty, surface, object, code, face, value
```

Surface and object refer to the surface number and object number, and code refers to the integer code used to identify specific types of NSC object data as described in "SETNSCPROPERTY" on page 721. The face argument is used by some codes to identify a particular object face. The value is either a string or a numeric value, depending upon the property being set. The return value is the same as for GetNSCProperty and corresponds to the newly set data.

See also GetNSCProperty

SetNSCSettings

This item sets the maximum number of intersections, segments, nesting level, minimum absolute intensity, minimum relative intensity, glue distance, miss ray distance, and ignore errors flag used for NSC ray tracing. The syntax is:

```
SetNSCSettings, MaxInt, MaxSeg, MaxNest, MinAbsI, MinRelI, GlueDist, MissRayDist, IgnoreErrors
```

The returned value is identical to that of GetNSCSettings. Since the MaxSeg value may require large amounts of RAM, verify that the new value was accepted by checking the returned string value.

See also GetNSCSettings.

SetOperand

This item sets the operand data in the Merit Function Editor. The syntax is:

```
SetOperand, row, column, value
```

Row is the operand row number in the Merit Function Editor. Column is 1 for type, 2 for int1, 3 for int2, 4-7 for data1-data4, 8 for target, 9 for weight, and 10-11 for data 5-6. The returned string is the same as for GetOperand. To set the operand type, use column = 1, and for the value use the name of the operand, such as "EFFL".

To update the merit function after calling SetOperand, use the Optimize item with the number of cycles set to -1. See also GetOperand and Optimize.

SetPolState

This item sets the default polarization state set by the user. The data is formatted as follows:

```
SetPolState, nIsPolarized, Ex, Ey, Phax, Phay
```

See GetPolState for a complete description.

SetSettingsData

This item sets the settings data used by a window in temporary storage before calling MakeGraphicWindow or MakeTextWindow. The data may be retrieved using GetSettingsData. The syntax is:

```
SetSettingsData, number, data
```

Currently, only number = 0 is supported. This number may be used to expand the feature in the future.

See also "How ZEMAX calls the client" on page 781 and GetSettingsData.

SetSolve

This item sets solve data and returns data about the new solve on any surface. The syntax is:

`SetSolve, surface, code, data...`

where the code is as defined in the discussion of “GetSolve” and the “data...” are the exact parameters as defined for GetSolve. The parameter list depends upon the code listed. The return string is the same as GetSolve returns for the newly defined solve data.

SetSurfaceData

This item sets surface data. The item name is formatted as follows:

`SetSurfaceData, surf, code, data`

or

`SetSurfaceData, surf, code, data, arg2`

Where the values for code are the same as documented for the GetSurfaceData item. To set the surface type to a user defined surface, send the new DLL name using code 9 rather by setting the surface type. The return string is the newly set surface data formatted as documented for the GetSurfaceData item. See also GetSurfaceData and SetSurfaceParameter.

SetSurfaceParameter

This item sets surface parameter data. The item name is formatted as follows:

`SetSurfaceParameter, surf, parm, newvalue`

where surf is the surface number and parm is the parameter number. The return value is identical to GetSurfaceParameter. See also SetSurfaceData and GetSurfaceParameter.

SetSystem

This item sets system data. The syntax is

`SetSystem, unitcode, stopsurf, rayaimingtype, useenvdata, temp, pressure, globalrefsurf`

The returned string is formatted as defined for GetSystem. See also GetSystem for details on the arguments. The value for useenvdata is ignored.

SetSystemAper

This item sets system aperture data. The syntax is

`SetSystemAper, type, stopsurf, aperture_value`

The returned string is formatted as defined for GetSystemAper. See GetSystemAper for details on the arguments.

SetSystemProperty

This item mimics the ZPL keyword SETSYSTEMPROPERTY. The syntax is:

`SetSystemProperty code, value1, value2`

Code is an integer value that defines the specific system property requested. For a listing of these codes, see “SETSYSTEMPROPERTY, SYSP” on page 728. The return string format is the same as for GetSystemProperty.

See also GetSystemProperty.

SetTol

This item sets tolerance operand data. The syntax is

`SetTol, n, col, value`

where n is the tolerance operand number, col is 1 for type, 2-4 for int1-int3, 5 for min, and 6 for max. The returned string is formatted as for GetTol.

See also GetTol.

SetUDOData

Obsolete, see SetUDOItem below.

SetUDOItem

This item is used to pass just one datum computed by the client program to the ZEMAX optimizer. The only time this item name should be used is when implementing a User Defined Operand, or UDO. UDO's are described in "Optimizing with externally compiled programs" on page 534. The syntax for SetUDOItem is

```
SetUDOItem, buffercode, datanumber, data
```

where buffercode is an integer passed to the UDO on the command line and data is the data item number being passed. After the last data item has been sent, the buffer must be closed using the CloseUDOData item before the optimization may proceed. A typical implementation may consist of the following series of item messages:

```
SetUDOItem, buffercode, 0, value0  
SetUDOItem, buffercode, 1, value1  
SetUDOItem, buffercode, 2, value2  
CloseUDOData, buffercode
```

See also GetUDOSystem.

SetVig

This item sets the vignetting factors. For more information, see "Vignetting factors" on page 115. No arguments are required. SetVig always returns "OK" as the return string.

SetWave

This item sets wavelength data. The syntax is

```
SetWave, 0, primary, number  
or
```

```
SetWave, n, wavelength, weight
```

If the value for n is zero, then the primary wavelength number and the total number of wavelengths is set to the new integer values. If n is a valid wavelength number (between 1 and the number of wavelengths, inclusive) then the wavelength in micrometers and the wavelength weight are both set. See also GetWave. The returned data is the same as for GetWave, n.

WindowMaximize

This item will maximize the main ZEMAX window or any analysis window ZEMAX currently displays. The syntax is

```
WindowMaximize, winnum
```

Use 0 for winnum to maximize the main ZEMAX window, otherwise, the integer window number.

WindowMinimize

This item will minimize the main ZEMAX window to an icon. See WindowMaximize for details.

WindowRestore

This item will restore the main ZEMAX window to it's previous size and position. See WindowMaximize for details.

Tracing large numbers of rays

If only a few rays need to be traced, then it is easy and fast enough to use DDE items such as GetTrace or GetPolTrace. However, if more than perhaps 100-500 rays need to be traced, it is far faster to use the array ray tracing technique.

Instead of tracing the rays one at a time, an array containing a list of all the rays to be traced is created, and then the entire array passed to ZEMAX at one time. ZEMAX then traces all the rays, and passes the entire array back to the client.

This method mimics the behavior of GetTrace, GetTraceDirect, GetPolTrace, and GetPolTraceDirect, with the added power that any number of rays to be traced at one time, although it is slightly more complex to program.

There are 2 steps required:

- 1) Fill an array of structures with the required data
- 2) Pass the array to ZEMAX

Step 1: Placing the ray data in the array

The data defining the list of rays to be traced is placed in an array of structures of the following form:

```
typedef struct
{
    double x, y, z, l, m, n, opd, intensity;
    double Exr, Exi, Eyr, Eyi, Ezr, Ezi;
    int wave, error, vigcode, want_opd;
}DDERAYDATA;
```

Inside the client program, an array of type DDERAYDATA is created with a declaration such as:

```
DDERAYDATA MyData[101];
```

There are four trace modes supported, numbered from 0 to 3, and they correspond to the actions of GetTrace (mode 0), GetTraceDirect (mode 1), GetPolTrace (mode 2), and GetPolTraceDirect (mode 3). Each mode requires a slightly different set of data to be placed in the DDERAYDATA structure. The mode is set by the opd parameter in array position zero.

Each ray requires 1 array element. If there are "n" rays to be traced, then the array structures 1 through n will be used to hold the data for these n rays. Array position zero is reserved for some header data that needs to be passed with the array.

Mode 0: Similar to GetTrace

Like GetTrace, the rays are defined by the relative field and pupil coordinates, the wavelength, the mode (either real, mode = 0 or paraxial, mode = 1) as well as the surface to trace the ray to. Before the rays can be traced, the array should be modified in array positions 1 through n to contain the definitions for each ray. The field, pupil, and wavelength data are all stored "in place" in the array as follows:

```
MyData[i].x = hx;
MyData[i].y = hy;
MyData[i].z = px;
MyData[i].l = py;
MyData[i].m = 0.0; /* ignored on input */
MyData[i].n = 0.0; /* ignored on input */
MyData[i].opd = 0.0; /* ignored on input */
MyData[i].intensity = 1.0; /* initial intensity */
MyData[i].Exr = 0.0; /* on output, x cosine of the surface normal */
MyData[i].Exi = 0.0; /* ignored */
MyData[i].Eyr = 0.0; /* on output, y cosine of the surface normal */
MyData[i].Eyi = 0.0; /* ignored */
MyData[i].Ezr = 0.0; /* on output, z cosine of the surface normal */
MyData[i].Ezi = 0.0; /* ignored */
MyData[i].wave = wavenumber;
MyData[i].error = 0; /* must initially be zero */
MyData[i].vigcode = 0; /* must initially be zero */
MyData[i].want_opd = 0; /* 0 if OPD data is not needed, 1 if it is */
```

Note the argument i should be between 1 and n, the number of rays. The array may be of any size, however, it is bad programming practice to trace too many rays at once, because it makes it difficult to accept a user interrupt if the ray tracing is to be terminated. The initial intensity will be scaled by any pupil apodization or surface apodization that occurs. Note mode 0 considers pupil apodization, while mode 1 does not.

The array position 0 is reserved for other data that applies to all the rays. The data in array position 0 is used to tell ZEMAX the mode, the number of rays, and the surface to trace the rays to. The data is placed in the array as follows:

```
MyData[0].x = 0.0; /* ignored */
MyData[0].y = 0.0; /* ignored */
MyData[0].z = 0.0; /* ignored */
MyData[0].l = 0.0; /* ignored */
MyData[0].m = 0.0; /* ignored */
MyData[0].n = 0.0; /* ignored */
```



```

MyData[0].opd = 0; /* sets mode 0 */
MyData[0].intensity = 0.0; /* ignored */
MyData[0].Exr = 0.0; /* ignored */
MyData[0].Exi = 0.0; /* ignored */
MyData[0].Eyr = 0.0; /* ignored */
MyData[0].Eyi = 0.0; /* ignored */
MyData[0].Ezr = 0.0; /* ignored */
MyData[0].Ezi = 0.0; /* ignored */
MyData[0].wave = mode; /* 0 for real rays, 1 for paraxial rays */
MyData[0].error = numrays; /* the number of rays in the array */
MyData[0].vigcode = 0; /* must initially be zero */
MyData[0].want_opd = lastsurf; /* -1 for image, or any valid surface number */

```

The number of rays must be placed in the error variable. Note that any declaration for an array of DDERAY-DATA types must be 1 bigger than the number of rays, since array position 0 is reserved for the data described above. All rays in the array must be traced to the same surface using the same mode, but they all may have different wavelengths and field and pupil coordinates.

If the variable want_opd is 0, then the usual x, y, z, l, m, and n data will be computed for the ray. If want_opd is anything other than zero, then the OPD for the ray will be returned along with the usual data. Computing the OPD takes additional time beyond that for regular ray tracing, and ZEMAX only performs this additional computation if requested to do so.

Computing the OPD is also more complicated for the client program, because OPD means optical path difference; which means two rays must be traced rather than just one. To compute OPD for an arbitrary ray, ZEMAX must trace the chief ray, then the arbitrary ray, then subtract the phase of the two to get the OPD. Rather than trace the same chief ray over and over, which is slow, usually the chief ray is traced once, and then the phase of the chief ray is subtracted from each subsequent ray. ZEMAX uses the sign of the want_opd parameter to specify which calculation should be done. If want_opd is less than zero (such as -1) then the both the chief ray and specified ray are requested, and the OPD is the phase difference between the two in waves of the current wavelength. If want_opd is greater than zero, then the most recently traced chief ray data is used. Therefore, the want_opd flag should be -1 whenever the chief ray changes; and +1 for all subsequent rays which do not require the chief ray be traced again. Generally the chief ray changes only when the field coordinates or wavelength changes. This method is much faster if there are many rays being traced from the same field point, as is the case for many optical analysis calculations.

The opd can only be computed if the last surface is the image surface, otherwise, the opd value will be zero.

Mode 1: Similar to GetTraceDirect

Like GetTraceDirect, the rays are defined by x, y, z, l, m, and n coordinates on any starting surface, as well as the wavelength, the mode (either real, mode = 0 or paraxial, mode = 1) and the surface to trace the ray to. Before the rays can be traced, the array should be modified in array positions 1 through n to contain the definitions for each ray. The field, pupil, and wavelength data are all stored "in place" in the array as follows:

```

MyData[i].x = x;
MyData[i].y = y;
MyData[i].z = z;
MyData[i].l = l;
MyData[i].m = m;
MyData[i].n = n;
MyData[i].opd = 0.0; /* ignored on input */
MyData[i].intensity = 1.0; /* initial intensity */
MyData[i].Exr = 0.0; /* on output, x cosine of the surface normal */
MyData[i].Exi = 0.0; /* ignored */
MyData[i].Eyr = 0.0; /* on output, y cosine of the surface normal */
MyData[i].Eyi = 0.0; /* ignored */
MyData[i].Ezr = 0.0; /* on output, z cosine of the surface normal */
MyData[i].Ezi = 0.0; /* ignored */
MyData[i].wave = wavenumber;
MyData[i].error = 0.0; /* must initially be zero */
MyData[i].vigcode = 0.0; /* must initially be zero */
MyData[i].want_opd = 0; /* ignored */

```

Note the argument *i* should be between 1 and *n*, the number of rays. The array may be of any size, however, it is bad programming practice to trace too many rays at once, because it makes it difficult to accept a user interrupt if the ray tracing is to be terminated. The initial intensity will be scaled by any surface apodization that occurs. Note mode 0 considers pupil apodization, while mode 1 does not.

The array position 0 is reserved for other data that applies to all the rays. The data in array position 0 is used to tell ZEMAX the mode, the number of rays, and the surface to trace the rays to. The data is placed in the array as follows:

```
MyData[0].x = 0.0; /* ignored */
MyData[0].y = 0.0; /* ignored */
MyData[0].z = 0.0; /* ignored */
MyData[0].l = 0.0; /* ignored */
MyData[0].m = 0.0; /* ignored */
MyData[0].n = 0.0; /* ignored */
MyData[0].opd = 1; /* sets mode 1 */
MyData[0].intensity = 0.0; /* ignored */
MyData[0].Exr = 0.0; /* ignored */
MyData[0].Exi = 0.0; /* ignored */
MyData[0].Eyr = 0.0; /* ignored */
MyData[0].Eyi = 0.0; /* ignored */
MyData[0].Ezr = 0.0; /* ignored */
MyData[0].Ezi = 0.0; /* ignored */
MyData[0].wave = mode; /* 0 for real rays, 1 for paraxial rays */
MyData[0].error = numrays; /* the number of rays in the array */
MyData[0].vigcode = startsurf; /* the surface on which the coordinates start */
MyData[0].want_opd = lastsurf; /* -1 for image, or any valid surface number */
```

The behavior of this mode is otherwise the same as in mode 0. Computation of the OPD is not permitted in this mode.

Mode 2: Similar to GetPolTrace

This mode is very similar to GetPolTrace. Before the rays can be traced, the array should be modified in array positions 1 through *n* to contain the definitions for each ray. The field, pupil, and wavelength data are all stored "in place" in the array as follows:

```
MyData[i].x = hx;
MyData[i].y = hy;
MyData[i].z = px;
MyData[i].l = py;
MyData[i].m = 0.0; /* ignored */
MyData[i].n = 0.0; /* ignored */
MyData[i].opd = 0.0; /* ignored */
MyData[i].intensity = 1.0; /* initial intensity */
MyData[i].Exr = Exr /* Electric field X real */
MyData[i].Exi = Exi /* Electric field X imaginary */
MyData[i].Eyr = Eyr /* Electric field Y real */
MyData[i].Eyi = Eyi /* Electric field Y imaginary */
MyData[i].Ezr = Ezr /* Electric field Z real */
MyData[i].Ezi = Ezi /* Electric field Z imaginary */
MyData[i].wave = wavenumber;
MyData[i].error = 0.0; /* must initially be zero */
MyData[i].vigcode = 0.0; /* must initially be zero */
MyData[i].want_opd = 0; /* ignored */
```

Note the argument *i* should be between 1 and *n*, the number of rays. The array may be of any size, however, it is bad programming practice to trace too many rays at once, because it makes it difficult to accept a user interrupt if the ray tracing is to be terminated.

The intensity parameter is used to return the relative transmitted intensity of the ray. If all six of the electric field values are zero; then ZEMAX will use the *Ex* and *Ey* values provided in array position 0 to determine the electric field. Otherwise, the electric field is defined by these six values. The defined electric field vector must be orthogonal to the ray vector or incorrect ray tracing will result. Even if these six values are defined for each ray, values for *Ex* and *Ey* in the array position 0 must still be defined, otherwise an unpolarized ray trace will result.

The array position 0 is reserved for other data that applies to all the rays. The data in array position 0 is used to tell ZEMAX the mode, the number of rays, and the surface to trace the rays to. The data is placed in the array as follows:

```
MyData[0].x = Ex; /* Electric field amplitude in x */
MyData[0].y = Ey; /* Electric field amplitude in y */
MyData[0].z = Phax; /* Phase in degrees for Ex */
MyData[0].l = Phay; /* Phase in degrees for Ey */
MyData[0].m = 0.0; /* ignored */
MyData[0].n = 0.0; /* ignored */
MyData[0].opd = 2; /* sets mode 2 */
MyData[0].intensity = 0.0; /* ignored */
MyData[0].Exr = 0.0; /* ignored */
MyData[0].Exi = 0.0; /* ignored */
MyData[0].Eyr = 0.0; /* ignored */
MyData[0].Eyi = 0.0; /* ignored */
MyData[0].Ezr = 0.0; /* ignored */
MyData[0].Ezi = 0.0; /* ignored */
MyData[0].wave = mode; /* 0 for real rays, 1 for paraxial rays */
MyData[0].error = numrays; /* the number of rays in the array */
MyData[0].vigcode = 0; /* must initially be zero */
MyData[0].want_opd = lastsurf; /* -1 for image, or any valid surface number */
```

If all size of the electric field values are zero; then ZEMAX will use the

The behavior of this mode is otherwise the same as in mode 0.

Mode 3: Similar to GetPolTraceDirect

Like GetPolTraceDirect, the rays are defined by x, y, z, l, m, and n coordinates on any starting surface, as well as the wavelength, the mode (either real, mode = 0 or paraxial, mode = 1) and the surface to trace the ray to. Before the rays can be traced, the array should be modified in array positions 1 through n to contain the definitions for each ray. The field, pupil, and wavelength data are all stored "in place" in the array as follows:

```
MyData[i].x = x;
MyData[i].y = y;
MyData[i].z = z;
MyData[i].l = l;
MyData[i].m = m;
MyData[i].n = n;
MyData[i].opd = 0.0; /* ignored */
MyData[i].intensity = 1.0; /* initial intensity */
MyData[i].Exr = Exr /* Electric field X real */
MyData[i].Exi = Exi /* Electric field X imaginary */
MyData[i].Eyr = Eyr /* Electric field Y real */
MyData[i].Eyi = Eyi /* Electric field Y imaginary */
MyData[i].Ezr = Ezr /* Electric field Z real */
MyData[i].Ezi = Ezi /* Electric field Z imaginary */
MyData[i].wave = wavenumber;
MyData[i].error = 0.0; /* must initially be zero */
MyData[i].vigcode = 0.0; /* must initially be zero */
MyData[i].want_opd = 0; /* ignored */
```

Note the argument i should be between 1 and n, the number of rays. The array may be of any size, however, it is bad programming practice to trace too many rays at once, because it makes it difficult to accept a user interrupt if the ray tracing is to be terminated.

The intensity parameter is used to return the relative transmitted intensity of the ray. If all six of the electric field values are zero; then ZEMAX will use the Ex and Ey values provided in array position 0 to determine the electric field. Otherwise, the electric field is defined by these six values. The defined electric field vector must be orthogonal to the ray vector or incorrect ray tracing will result. Even if these six values are defined for each ray, values for Ex and Ey in the array position 0 must still be defined, otherwise an unpolarized ray trace will result.

The array position 0 is reserved for other data that applies to all the rays. The data in array position 0 is used to tell ZEMAX the mode, the number of rays, and the surface to trace the rays to. The data is placed in the array as follows:

```

MyData[0].x = Ex; /* Electric field amplitude in x */
MyData[0].y = Ey; /* Electric field amplitude in y */
MyData[0].z = Phax; /* Phase in degrees for Ex */
MyData[0].l = Phay; /* Phase in degrees for Ey */
MyData[0].m = 0.0; /* ignored */
MyData[0].n = 0.0; /* ignored */
MyData[0].opd = 3; /* sets mode 3 */
MyData[0].intensity = 0.0; /* ignored */
MyData[i].Exr = 0.0; /* ignored */
MyData[i].Exi = 0.0; /* ignored */
MyData[i].Eyr = 0.0; /* ignored */
MyData[i].Eyi = 0.0; /* ignored */
MyData[i].Ezr = 0.0; /* ignored */
MyData[i].Ezi = 0.0; /* ignored */
MyData[0].wave = mode; /* 0 for real rays, 1 for paraxial rays */
MyData[0].error = numrays; /* the number of rays in the array */
MyData[0].vigcode = startsurf; /* the surface on which the coordinates start */
MyData[0].want_opd = lastsurf; /* -1 for image, or any valid surface number */

```

The behavior of this mode is otherwise the same as in mode 1.

Mode 5: For tracing non-sequential rays

Mode 5 is used to define and trace a ray inside of a non-sequential group. Unlike the other modes; there is no text based command for tracing NSC rays, and only 1 ray at a time may be traced. However, rays may split or scatter into multiple paths, and mode 5 will trace them all and return the entire tree of ray data.

To define the starting ray, use array position zero. The data is placed in the array as follows:

```

MyData[0].x = 0.0; /* starting x coordinate */
MyData[0].y = 0.0; /* starting y coordinate */
MyData[0].z = 0.0; /* starting z coordinate */
MyData[0].l = 0.0; /* starting x direction cosine */
MyData[0].m = 0.0; /* starting y direction cosine */
MyData[0].n = 1.0; /* starting z direction cosine */
MyData[0].opd = 5+nMaxSegments; /* sets mode 5, see comments below */
MyData[0].intensity = 1.0; /* initial intensity */
MyData[0].Exr = 0.0; /* initial E field if doing pol ray tracing */
MyData[0].Exi = 0.0;
MyData[0].Eyr = 0.0;
MyData[0].Eyi = 0.0;
MyData[0].Ezr = 0.0;
MyData[0].Ezi = 0.0;
MyData[0].wave = 0; /* wavelength number, use 0 for randomly selected by weight */
MyData[0].error = 1; /* NSC group surface, 1 if the program mode is NSC */
MyData[0].vigcode = 0; /* controls polarization, split, and scatter */
MyData[0].want_opd = 0; /* the inside of flag, use 0 if ray is not inside anything */

```

Only array position zero is used to define the starting ray. ZEMAX will trace the ray through the NSC group and fill the remaining array positions with ray data. Since rays may split or scatter into multiple child rays, a very large number of segments may be returned. The nMaxSegments number is the maximum allowed size of the array for ZEMAX to return. The array should also be large enough to hold at least nMaxSegments elements. The nMaxSegments value ZEMAX is using for the current optical system may be determined by the GetNSCSettings command.

If polarization ray tracing is used, the initial electric field values must be provided. The user application must ensure the defined vector is orthogonal to the ray propagation vector, and that the resulting intensity matches the starting intensity value, otherwise incorrect ray tracing results will be produced.

The integer vigcode value determines if polarization ray tracing, splitting, and scattering are to be used. To determine the vigcode value, use 0 for no polarization, 1 for polarization, plus 2 if splitting is used, plus 4 if scattering is used. The resulting integer will be between 0 and 7 inclusive. Note if ray splitting is to be used, polarization must be used as well.

The want_opd flag indicates where the ray starts. If zero, then the ray is assumed to not start in any object. Otherwise, the ray starts inside the media defined by the specified object number. This is identical to the "inside of" flag set on the NSC Editor in ZEMAX.

The array is passed to ZEMAX just like the other modes, see “Step 2: Pass the array to ZEMAX” on page 781.

When the ray tracing is complete, the array will contain the entire ray tree. Array element zero also contains the total number of segments stored in the array element `want_opd`:

```
nNumRaySegments = MyData[0].want_opd;
```

The other values are interpreted as follows:

```
i = the segment number
MyData[i].wave = segment level
MyData[i].want_opd = segment parent
MyData[i].vigcode = inside of object number
MyData[i].error = hit object number
MyData[i].x, .y, .z, .l, .m, .n = the ray coordinates and cosines
MyData[i].intensity = the ray intensity
MyData[i].opd = the path length to the hit object
```

See also `GetNSCSettings` and `SetNSCSettings`.

Step 2: Pass the array to ZEMAX

Once all the rays are defined, the array must be sent to the ZEMAX server. The easy way to do this is using the `ZCLIENT PostArrayTraceMessage` function:

```
PostArrayTraceMessage(szBuffer, RD);
```

The `RD` variable is the pointer to the `DDERAYDATA` array holding the list of rays. The ZEMAX server will trace all the rays, then pass the data back in the same array. For more details on the code sample, see the section “a sample client program” at the end of the chapter.

Note that an extension should never request that more rays be traced at one time than can be traced within the timeout period specified in the `zclient.c` code. If more rays are traced in one call to `PostArrayTraceMessage` than can be traced in the time allowed by `DDE_TIMEOUT`, an error will occur, and `PostArrayTraceMessage` will return -1. If the ray trace is successful, `PostArrayTraceMessage` will return 0. It is good programming practice to trace rays in small groups of 1000-5000 rays at a time. Once this many rays are traced, the overhead of multiple calls using `PostArrayTraceMessage` becomes negligible.

For information about timeout limits when executing potentially long calculations, see “Comments on the DDE timeout interval” on page 750.

How ZEMAX calls the client

ZEMAX looks in the `<program>\Extend` folder (see “Folders” on page 68) for Extension executables. Each program in this folder ending in the usual `.EXE` extension is assumed to be a valid ZEMAX Extension program. The name of the executable is placed in the pull down menu under “Extensions” on the main menu bar.

When any of the listed Extension names is selected from the menu, ZEMAX executes the Extension program with the following syntax:

```
programname textflag optionsflag tempfile {settings data}
```

The `programname` is the full path to the executable, for example, “`C:\ZEMAX\EXTENSIONS\DDE_DEMO.EXE`”.

The `textflag` is initially always 0. The convention ZEMAX uses for Extensions is the same as for internal ZEMAX features: if the feature generates both a graphic and a text format for output, show the graphic version first. If the feature does not have a graphic display, then show a text window. The Extension must generate a graphic file if it can and is requested to do so. The user can click on “text” to look at the text version in a new window; if this happens ZEMAX will call the Extension again with the text flag set to 1. If the Extension only generates text, then it can generate a text file and open it by sending the `MakeTextWindow` item. ZEMAX will never send the `textflag` set to 1 unless the Extension specifically states it can support a text mode; this is why the `MakeGraphicWindow` returns a flag to indicate if the feature supports text, and why ZEMAX always ask for a graphic mode first.

The `optionsflag` is set to 0 if the user does not wish to see the “settings” box; or 1 if the user does. There are two ways this flag may be set; either the user clicked on “Settings” on a window already open, or, the user has selected “Show Options First” from the Environment options within ZEMAX. In this latter case the settings box may need to be shown before the calculation is done the first time. If this flag is set the Extension must prompt

the user for setting options using a dialog box or other interface mechanism. If there are no user defined options, then a message box stating "This window has no options" should be displayed.

The tempfile is the fully qualified path and file name of the temporary file to hold the data generated by the Extension. Initially, this file will not exist, and the Extension needs to create it, write either text or graphic data to the file, and then call MakeTextWindow or MakeGraphicWindow as appropriate. If the user is updating an existing window, then the file may already exist, and it should be overwritten. The file should never be opened and read as the data may not be what the Extension originally put there. ZEMAX does some translation of the Extension created graphic file in particular. No other file should be used for output, because ZEMAX expects windows which are being updated to retain their original temporary file names. Only one window may be created by a single call to the Extension.

When ZEMAX receives a MakeTextWindow or MakeGraphicsWindow request from the Extension, one of the arguments encoded in the atom name is a space delimited string of Extension defined values. This string should be used to hold the "settings data" for the settings the Extension requires, such as the surface number, field, wavelength, or other data for which the data in the window was computed. When ZEMAX updates a window, it will call the Extension and provide the settings data string back to the Extension.

It is important that the Extension does not attempt to store the settings data itself, either statically or in a file. The reason is that ZEMAX may have multiple text and graphic windows open at once, all using the same Extension code to generate the data displayed. For example, a Extension that computes and displays the power of a single surface may be used by several windows simultaneously, each for a different surface. Only ZEMAX "knows" which window is using which setting data; and so ZEMAX passes this data on the command line. On the initial call to the Extension, the settings data string is not supplied, and the Extension must determine it's own default settings.

Because DDE data items cannot be longer than 255 characters, it is a good idea to use the SetSettingsData rather than return the settings data appended to the end of the MakeGraphicsWindow or MakeTextWindow commands. If the extension uses SetSettingsData to store the settings data with the window, then GetSettingsData should be used to retrieve the data whenever needed.

Generating a text window

Generating a text window for display by ZEMAX is accomplished with the MakeTextWindow item. First, open the filename passed to the Extension using the typical C language fopen function, with "wt" as the mode for "write text". Then, output data to the file, close the file, then send the MakeTextWindow item to ZEMAX. ZEMAX will read the data and display it without modification or interpretation. A test code might look like:

```
FILE *out;
out = fopen(szFile, "wt");
fputs("Write this line\n", out);
fclose(out);
```

Do not forget to close the file!

Generating a graphic window

The Extension generates a graphic window by writing out graphic format data to the temporary file name provided on the command line. The graphic data is created using a script language similar to that used by the ZPL macro language or the annotation feature.

All ZEMAX graphics are written on a virtual screen whose coordinates are 1.00 wide by 1.00 high, with the 0, 0 origin at the lower left corner. The screen aspect ratio is normally 4 wide by 3 high, or 5 wide by 3 high, depending upon user preference (See the File Menu chapter under Environments for details). Therefore, the virtual pixels being drawn are wider than they are tall.

Each line in the script file should be followed by a newline "\n" character. After writing all of the graphic commands to the file, the file should be closed for writing.

The following graphic script commands are supported.

ADDRESS

Syntax: ADDRESS

The address command will cause the "address box" in the lower right edge of the plot to appear, unless the user has selected this box not to appear.

BOX

Syntax: BOX x1 y1 x2 y2

The BOX command will draw four lines, connecting the four points of a rectangle. The coordinates are the lower left and upper right coordinates of the rectangle.

DATA

Syntax: DATA n "string"

The DATA command accepts an integer argument between 0 and 5, and a text string. The integer argument indicates on which row of the data box the string should be written. Normally, the lens name goes on row 0; and the date on row 1. The LENSNAME and DATE commands do this automatically. The DATA command allows any string to be placed on any of the lines.

DATE

Syntax: DATE

The current data and time is normally written on the second line of the data box below the plot title. This command has no arguments, and will print the date and time (depending upon the User's Environment preferences) on this second line.

FRAME

Syntax: FRAME

Either a FRAME or a NOFRAME command must be the first line in the script file. FRAME will cause a standard ZEMAX graphics frame to be drawn, with a rectangle around the entire perimeter, and two ruling lines at about 0.2 in the y direction. The plot title can be written centered inside these ruling lines using the TITLE command. See also NOFRAME.

GRID

Syntax: GRID x1 y1 x2 y2 nx ny

The GRID command creates a grid of nx by ny grey boxes. This command is typically used to make the background grid lines that mimics the style of most ZEMAX graphics. The coordinates are the lower left and upper right coordinates of the rectangle. See also BOX.

LENSNAME

Syntax: LENSNAME

The current lens name is normally written on the first line of the data box below the plot title. This command has no arguments, and will print the lens name on this line.

LINE

Syntax: LINE x1 y1 x2 y2

This is the most basic graphic command; it draws a line connecting the two points. Note that the coordinates should be between 0 and 1.0.

NOFRAME

Syntax: NOFRAME

Either a FRAME or a NOFRAME command must be the first line in the script file. NOFRAME will cause a standard ZEMAX graphics frame to be drawn, but without the two ruling lines. See also FRAME.

PENSTYLE

Syntax: PENSTYLE ncolor nstyle nwidth

The PENSTYLE command changes the current pen color, style, and width. The integer ncolor must be 0 for black, or between 1 and 12, for the corresponding pen color. The colors may be defined by the user in the File, Environments dialog. The nstyle integer can be 0 for a solid line, or an integer between 1 and 4 for various styles of dashed lines. The integer nwidth is the relative width of the pen. The default value is 1. A value of 2 will make lines twice as wide as the default width. If the pen style is not a solid line (nstyle = 0) then the pen width must be 1 pixel; Windows does not support thick lines which are not solid. All subsequent text and lines drawn will be in the new color, style, and width.

TEXT

Syntax: TEXT "string" x y angle width height

The text command is used to place text data on the graphic. The x and y coordinates are the starting coordinates for the string. The angle is in degrees. The width and height values affect the font size. The default value of 10 for each yields the same font size ZEMAX normally uses; which is 1/70 of the screen width and 1/40 of the screen height. Numbers larger or smaller than 10 will yield proportionally scaled font sizes.

TITLE

Syntax: TITLE "string"

If FRAME was used to define the graphic rather than NOFRAME, then the plot title would normally be centered within the ruling lines along the bottom 20% of the plot. The "string" provided will be displayed in this position by the TITLE command.

A sample Extension program

Within the <program>\Extend folder (see "Folders" on page 68) is a sample application called DDE_DEMO which illustrates the use of DDE with ZEMAX. The source code is included. DDE_DEMO can operate in either of two modes: stand alone, or as an Extension.

If run directly from a command prompt, Windows Explorer, or by double-clicking on the DDE_DEMO icon, it will run as a stand alone application, conducting the DDE in the background. The data it generates will be displayed in the DDE_DEMO window. To run DDE_DEMO in stand alone mode, first run ZEMAX, load a lens file, then launch DDE_DEMO by double clicking on the DDE_DEMO icon. Try loading different lens files, or editing the one in the LDE, and watch the DDE_DEMO window to monitor the changes. DDE_DEMO will list out some of the system, surface, field, and wavelength data, and will trace a few rays, and list the x, y, and z coordinates of the rays on the screen. The program automatically updates itself every few seconds.

To run DDE_DEMO as an Extension, run ZEMAX, load a lens file, then select Extensions, then DDE_DEMO. Note that any .EXE application placed in the <program>\Extend folder (see "Folders" on page 68) will be listed in the ZEMAX Extensions menu; whether or not it is a valid stand alone or Extension capable of communicating with ZEMAX. DDE_DEMO will execute in a "hidden" window, establish the DDE link with ZEMAX, extract the system and ray data, write the data to a file, notify ZEMAX that the file is ready (using the MakeTextWindow item) and then close. The client's job is finished once it has generated the data; and it should terminate itself when done.

When ZEMAX receives the MakeTextWindow item, it reads the contents of the file and displays the data in a standard ZEMAX text window. If the window is now "updated", or the "settings" option is selected, ZEMAX will call the client program again and ask for the data to be computed again.

Note that the exact same program can act as both a stand alone client and as an Extension. The secret to the bimodal operation is that if ZEMAX calls the client, it passes command line arguments to the client telling it what type of data to create and where to write the data to. If no command line arguments are provided, the client can assume it was launched from the operating system, and therefore it should run in stand alone mode. The DDE exchange and subsequent processing is the same either way; only a few key lines of code are required to determine whether the window should be visible or hidden, and whether the data is written to a file or the client's own screen.

The DDE_DEMO code

The source code for DDE_DEMO is provided. It contains the basic algorithms, variables, declarations, and Windows message handling required of all client programs. There are portions of the code which illustrate the use of WM_DDE_POKE for tracing arrays of rays as well.

DDE_DEMO is a rather difficult program to use as a starting point for an Extension, although it does show some high level techniques. A simpler way to implement ZEMAX Extensions is described below.

A simplified technique for implementing Extensions

Obviously, DDE programming for Windows is not for the faint of heart; it involves numerous programming skills and significant familiarity with message loops, DDE, pointers, atoms, and global handles.

However, most of the code used to communicate with ZEMAX is standard "boiler plate" code which is common to any DDE client program. This code has all been written already, and has been placed into a single source code module program called ZCLIENT. ZCLIENT is provided in source code form with ZEMAX, and may be freely copied (as long as the copyright is retained), and used in new Extensions.

The ZCLIENT program

ZCLIENT handles all of the communication with ZEMAX transparently. Embedded within ZCLIENT is a call to a single function given the name "UserFunction". This function is placed in a separate C program provided by the user, and is compiled along with ZCLIENT to create the Extension executable.

When ZEMAX calls the Extension, execution begins within ZCLIENT. ZCLIENT establishes the DDE communication, then calls UserFunction. Within UserFunction, two functions provided by ZCLIENT, PostRequestMessage and GetString, are all that are usually required to get the data needed from ZEMAX. The data is then formatted as either a text or a graphic, and sent back to ZEMAX for display.

The syntax for PostRequestMessage is as follows:

```
PostRequestMessage(szItemName, szBuffer);
```

szItemName contains the name of the item desired; the available item names are described earlier in this chapter. The desired data is then sent back in the string szBuffer. This function returns 0 if the data transfer was successful, otherwise, it returns -1. ZEMAX normally provides comma delimited data, and the function GetString extracts the individual items. The syntax for GetString is:

```
GetString(szBuffer, nItem, szSubString);
```

The szBuffer string is the buffer returned by PostRequestMessage. The integer nItem is the number of the item desired; 0 for the first item, 1 for the second, etc. The character szSubString holds the string from the nItem position.

For example, to get the name of the lens, the code would be:

```
PostRequestMessage("GetName", szBuffer);  
GetString(szBuffer, 0, szLensName);
```

The string szLensName now contains the name of the lens. Because "GetName" only returns 1 item, this code could be shortened to just one line:

```
PostRequestMessage("GetName", szLensName);
```

There is also a modified version of PostRequestMessage for doing fast tracing of large numbers of rays. The function is PostArrayTraceMessage. This function accepts two arguments: the buffer string, and the address of the array of ray data. A sample array ray trace code using PostArrayTraceMessage is provided, it is called ARR_DEMO.C.

To trace rays using PostArrayTraceMessage, the ray list must first be defined, then a single call to PostArrayTraceMessage must be made. Here is a sample of the technique:

```
/* Fill RD array position 0 with the header data. */  
RD[0].x = 0.0;  
RD[0].y = 0.0;  
RD[0].z = 0.0;  
RD[0].l = 0.0;  
RD[0].m = 0.0;  
RD[0].n = 0.0;  
RD[0].opd = 0.0; /* this is where we set the mode, mode 0 is like GetTrace */  
RD[0].intensity = 0.0;  
RD[0].wave = 0;  
RD[0].error = 25; /* trace 25 rays */  
RD[0].vigcode = 0;  
RD[0].want_opd = -1;
```

```

/* Define the 25 rays. Obviously, you can define any rays you want... */
k = 0;
for (i = -2; i <= 2; i++)
    {
    for (j = -2; j <= 2; j++)
        {
        k++;
        RD[k].x = 0.0;
        RD[k].y = 0.0;
        RD[k].z = (double) i / 4.0;
        RD[k].l = (double) j / 4.0;
        RD[k].m = 0.0;
        RD[k].n = 0.0;
        RD[k].opd = 0.0;
        RD[k].intensity = 0.0;
        RD[k].wave = 1;
        RD[k].error = 0;
        RD[k].vigcode = 0;
        RD[k].want_opd = 0;
        }
    }
/* Now go get the data */
PostArrayTraceMessage(szBuffer, RD);

/* The output data is now stored in RD, ready to use! */
/* Note just one line of code to trace all rays. */

```

Using ZCLIENT is vastly easier than writing all the DDE communication code yourself. This sample array trace is taken from ARR_DEMO, one of the sample DDE codes that use ZCLIENT.

For information about timeout limits when executing potentially long calculations, see "Comments on the DDE timeout interval" on page 750.

A sample Extension program using ZCLIENT

The source code for a complete Extension which uses ZCLIENT is provided. The Extension is called PHASPLOT, and it uses ZCLIENT.C as well as PHASPLOT.C. PHASPLOT generates either a graph or a text listing of the phase of a binary 2 surface. The code illustrates how to make a graphic and a text window, how to get data out of ZEMAX, and how to format output for graphics and text display.

PHASPLOT also has a ZEMAX style "settings" box, the source code for which is included in PHASPLOT.C in the "SurDlgProc" function; which uses the PHASPLOT.RC file to define the dialog box appearance.

Writing Extensions from the PhasePlot template

The best way to learn how to write Extensions is to start with PhasePlot, and start editing! Make a copy of PhasePlot.C and PhasePlot.RC, and rename the copies to something else (MY_CODE.C, MY_CODE.RC). Then create a new project using any of the good commercially available C compilers (Borland and MicroSoft both make excellent tools) and load these 3 files into the project. Recompile and build the project into a test EXE to make sure everything is in place.

Then, edit the code within UserFunction and the settings dialog function to suit your requirements.

Getting analysis data using GetTextFile and GetMetaFile

One often overlooked capability of Extensions is the ability to create text and graphic data files from analysis already supported by ZEMAX by using a single call to zclient. For example, to generate a text listing of the spot diagram, use the command

```
PostRequestMessage("GetTextFile, \"C:\\OUTPUT.TXT\", Spt, , 0", szBuffer);
```

The spot diagram data will be placed in the text file "OUTPUT.TXT" in the specified folder. For more information, see "GetTextFile" on page 762 and "GetMetaFile" on page 755.

Symbols

µm 104

A

Abbe number

Tolerance on 553

Abbe V number

Definition for 601

For MILNUM glasses 600

Operand 505, 507

ABCD surfaces 280

Aberration

Longitudinal 189

Pupil 129

Ray fans 127

Wavefront 148

Aberration coefficients 194, 196, 520

Annular Zernike 202

Extended Zernike 200

ABg model 367

Catalog of data 250

Wavelength scaling of data 251

Aborting

Editing of cell values 42

Long computations 46

Absorbers 353

Absorption 594, 612

Active configuration

Defined 51

Active Cursor 43

Address box

Preferences 68

Adjust Index Data To Environment 108

Adjustable columns 42

Afocal Image Space 103, 579

Afocal mode units 105

Air pressure 603

Airy disk

Defined 130

Alternate Even surface 281

Alternate Odd surface 281

Alternate surface intersection 281

Analysis units 105

Anamorphic asphere surface 334

Angle of incidence 517

Angular distributions 463, 464, 466

Angular magnification

Defined 51

Operand 491

Angular Spectrum 644

Angular spectrum 629

Animation 266

Anisotropic media 357

Annotate

Graphics 43

Annotations

On graphic windows 44

Annular aspheric surface 334

Annulus object 398

ANSI 71

Anti-nodal planes 51

Anti-principal planes 51

Aperture

System aperture value 102

Aperture checking

With footprint diagram 188

Aperture type 79

Apertures 79

Floating 80

Removing 254

Soft edged 336

Surface aperture 60

System aperture 60, 101

User defined 80

Aplanatic condition

Solving for 475

Apodization 579

Cosine cubed 102

Defined 51

Gaussian 103, 488

Apodization factor 579

Defined 51

Specifying 103

Apodizations

User defined 51, 333, 334

Appending

Lens files 67

AR

Codes 589

Archive file 65, 66

Arms

For pupil integration 486

Array

Of diodes 450

Array object 399

Array of lenses 394, 395

Array Ring 401

Array variables

In ZPL macros 653

ZPL macros 733, 739

Arrays

Of lenses 335

Of sources 358

- Arrays of objects 96
- Arrow
 - Drawing on plots 44
- Arrow source
 - Size of 114
- ASCII 71
- Aspect ratio 43, 45
 - Printing 69
- Aspheric object 415, 416, 417, 428
- Aspheric surface object 402, 405, 446
 - Hologram 422
- Aspheric surfaces 298
 - High order terms 298, 301, 323
 - Superconic 326
- Asterisk
 - Hyperhemispheric surfaces 59
 - In data cells 71
 - Indicates surface apertures 75
- Astigmatism
 - Calculating 194
 - Operand 492
- Athermal Glass Map 210
- Athermal lenses
 - Analyzing and optimizing 608
- Atmospheric refraction 281
- Atmospheric turbulence 313
- Axial color
 - Operand 492, 493
- Axial symmetry 487
- Axicon 326
- Axicon surface object 403

B

- Back focal length
 - Defined 51
- Beam footprint 188
- Beam splitter
 - Defining coatings for 364
- Beamsplitter
 - Polarization modeling 625
- Beep command 671
- Beer's Law 625
- Bending
 - Using fold mirrors 255
- Best fit glass 271
- Best fit sphere 208
- Biconic 278
 - Surface type 282
- Biconic lens object 403
- Biconic object 403, 404
- Bicubic spline 314
- Binary 1 object 407
- Binary 2 object 408

- Binary optics 283, 284, 340, 342
 - Sign conventions for 284
- Biocular analysis 180, 182
- Birefringent
 - NSC rays 357
- Birefringent materials 287
- Birefringent waveplates 357
- Bitmap images 170
- Bitmaps
 - As sources 435
 - Exporting 48, 675, 676, 677
- Black Body 361
- Black Box 260, 278
- BMP 170
 - Creating 43
- Boolean
 - Operations on objects 97
- Boolean object 409
- Boresight error 493
 - Tolerancing 545, 561
- Boundary conditions
 - On (center) thickness 495
 - On curvature 495
- Boundary operands
 - Understanding 524
- Brightness Enhancement Filter 414
- BSDF 370
- BSDF data 369
- Bubbles 436
- Bugs
 - Reporting 39
- Bulk absorption 364, 625
- Bulk scatter 358
- Bulk scattering 356
- Bulk transmission 625
- Button bar
 - Defining 74

C

- CAD files
 - Importing 423
- Candela 105
- Capping
 - GRADIUM profiles 311
- Capturing
 - Screen images 48
- Cardinal planes 51
- Cardinal points 51, 209, 272
- Cell
 - Editing contents 42
- Cell width
 - Changing 70, 71

- Celsius
 - Defining temperature 603
- Center thickness
 - Operand 495, 505, 506, 507, 508
- Centroid 130
 - Determining coordinates of 131
 - Encircled energy 158
 - Optimization 512, 521
- Check for updates 71
- Checking for 71
- Chief ray
 - Definition of 51
 - Solve 475, 476
- Chief ray solve
 - On coordinate breaks 479
- Chinese 71
- Chord Tolerance 410, 423
- Chromatic aberration 191
- Chromaticity 361
- CIE 1931 361
- Circular 79
- Clipboard
 - Copying graphics to 43
 - Copying text to 45
 - Cut, Copy, and Paste 87, 88, 90, 92, 94
 - Selecting metafile format 69
 - Using 48
- Cloning windows 43, 45
- Coating file 109, 249, 613
- Coating multiplier 86, 493, 494, 617, 621
- Coatings
 - Adding to surfaces 249
 - Defined by tables 619
 - Editing 249
 - Loading coating file 249
 - Multi-configuration control of 580
 - Optical 219
 - Properties of 219
 - Tapered 615, 617
- Coherence length 358, 373
- Coherent image 172
- Coherent imaging 174
- COLOR
 - ZPL macros 672
- Color
 - True color PSF 147
- Color of rows 77, 354
- Color slides 435
- Colors
 - Changing default 69
 - Defining 73
 - Of objects 363
 - Of surfaces 77
- Column resizing 42
- Coma
 - Calculating 194
 - Operand 495
- Comments 76
 - In the merit function using BLNK 493
 - In the tolerance data using COMM 555
- Company name
 - On plots 68
- Compensators
 - For tolerancing 555
- Compound parabolic concentrator (CPC) 394
- Cone angle 101
- Cone object 411, 412
- Configuration
 - Changing 116, 584
 - Selecting on windows 43
 - Zoom position 116
- Conic
 - Description of 326
 - Tolerance on 546
- Conjugate
 - Surface type 293
- Conjugate surfaces 236
- Conrady
 - Dispersion formula 591
- Consider objects 354
- Consider objects list 363, 364
- Constant
 - Operand 495
- Construction optics
 - Constraining 494
 - For holograms 320
- Contours
 - Setting 151
- Convergence 180, 182
- Convert semi-diameters
 - to apertures 249, 253, 254
- Converting to NSC groups 267
- Coordinate axes
 - Definition of 52
- Coordinate breaks
 - And paraxial rays 110
 - Don't print data for 113
 - Reversing with "order" 294
 - Solves for 479
 - Surface tilts and decenters 84
 - Surface type 293
- Coordinate reference
 - Drawing 362
- Coordinate Return 85
- Coordinate system
 - Conventions for 52

- Copy
 - Operand data 89, 90, 92
 - Surface data 87, 95
- Corner object 432, 444
- Cosine lamp 453
- Cosine source 453
- Cosine space 131, 463, 464, 466
- Cosine surface 322
- Coupling efficiency 165, 169, 206, 499, 500, 657, 668
- CPC 394, 411
- CPU's
 - Number of 378, 526, 541
- CR
 - Codes 589
- Creating POB files 96
- Cross sections
 - Point Spread Functions 144, 145
- Cubic spline
 - Extended surface type 299
- Cubic spline surfaces 294
- Curvature 519
 - Boundary conditions on 495
 - Operand 495
- Cut
 - Surface data 87
- Cylinder arrays 335
- Cylinder Fresnel surface 295
- Cylinder lens object 426, 427, 441, 442
- Cylinder pipe object 413
- Cylinder volume 496
- Cylinder volume object 413, 414, 415
- Cylindrical lens
 - Paraxial model for 322
- Cylindrical surfaces 282, 328, 329
 - Gratings placed on 302, 329

D

- D65 361
- Damped Least Squares 526
- Data folder 68
- Date
 - Printing on plots and text 70, 71
- DDE 749
- Debugging
 - Tolerancing 577
- Decenters
 - Tolerancing 549
- Decenters and tilts
 - On surfaces 84
 - With coordinate break surfaces 293
- Decimals
 - Changing default 70, 71
- Default merit function 483
 - Start at row 497
- Defaults, changing 68
 - defined 345
- Delano diagrams 190
- Deleting
 - Operands 88, 90, 92, 94
 - Surfaces 75, 87
 - Surfaces in ZPL macros 674, 678
- Density
 - Discussion of 272
 - Of lenses 272
- Detector object 354, 463, 464, 466, 468, 470
- Detector viewer 384
- Detector volume 470
- Dialogs 41
- Diameter
 - Operand 497
- Diameter to thickness ratio
 - Constraining 506, 508
- Diattenuation
 - Of coatings 220
- Dichroism 291
- Diffraction 627
- Diffraction grating 414
 - Multiple orders 357
- Diffraction grating object 414
- Diffraction gratings 296, 323
 - On cylinders 329
 - On toroids 329
- Diffraction Imaginary Transfer Function 134, 135, 136, 138, 139, 140, 141, 142, 145, 146, 148
- Diffraction limited
 - Meaning of 52
- Diffraction Modulation Transfer Function 134, 135, 136, 138, 139, 140, 141, 142, 145, 146, 147, 148
- Diffraction Phase Transfer Function 134, 135, 136, 138, 139, 140, 141, 142, 145, 146, 148
- Diffraction Real Transfer Function 134, 135, 136, 138, 139, 140, 141, 142, 145, 146, 148
- Diffractional optics 313
- Digital Mirror Device (DMD) 427
- Diode 450
- Dipvergence 180, 182
- Directivity 384
- Dispersion 209, 210, 211, 212, 243
 - Table data for 601
- Dispersion coefficients 590

- Distortion 182, 185
 - Calculating 194
 - Definition of 183
 - F-Theta 183
 - Grid 184
 - Operand 497
 - Representing 170
- Divergence 180, 182
- DLL 331
- dn/dt 607
- Double pass 264
- dPgF 589
- Drawing
 - Element 122, 123
- Dual zone 287
- DXF 234, 235
- Dynamic Data Exchange (DDE) 749
- Dynamic Link Library 331

E

- Edge margin 111
- Edge Response Function
 - Optimization operand for 498
- Edge Spread Function 145, 160
- Edge thickness 499
 - Controlling 476
 - Defined 52
 - Determining 271
 - Operand 499, 506, 508, 521
 - X- direction 522, 523
- Editing
 - Spreadsheet values 42
- Editor windows 41
 - Menu 75
 - Preferences 71
- Editors
 - Preferences 70
- Effective F/# 187, 188, 498
- Effective focal length
 - Defined 53
 - Displaying 271
 - Operand 498
 - Single element 271
- Element drawings 122, 123
 - ISO 10110 123
- Element power 271
 - Solve 475
- Ellipse
 - Defining 326
- Elliptical 79
- Elliptical object 415
- Elliptical surface
 - Defining 293

- Encircled energy 158
- END
 - ZPL macros 687
- English 71
- Ensquared energy 158, 159, 161
- Entrance pupil diameter
 - Defined 53
 - Operand 498
 - Setting 101
- Entrance pupil position 498
 - Defined 53
 - Operand 498
- Entry ports 347
- Environment
 - Defining 603
 - Temperature and pressure 603
- EQUA
 - Operand 498
- Equal operand 498
- Escape
 - Aborting long computations 46
- EULUMDAT 259, 463
- EULUMDAT format 454
- Even Asphere 278
- Even Asphere lens object 415
- Even Aspheric 415
- Exclude Substitution 588
- Exit ports 347
- Exit pupil
 - Significance of 110
- Exit pupil diameter
 - Defined 53
- Exit pupil position
 - Defined 53
 - Operand 499
- Exponential notation 71
- Exporting
 - CAD format files 258, 676, 751
 - Graphics and text 48
- Extended 2
 - Constants of dispersion formula 592
- Extended 3
 - Constants of dispersion formula 592
- Extended cubic spline
 - Surface type 299
- Extended polynomial
 - Surface type 301
- Extended Polynomial Lens object 416
- Extended Polynomial Surface object 417
- Extended sources 165, 170
 - Encircled energy computations 161
- Extensions 275

Extra data
 Defined 53, 277
 Operand 522
 Optimization of 530
 ZPL Macros 679
Extraordinary rays 287
Extruded object 417

F

F/#
 Capped at 10,000 112
 Curvature solve 473
 Displaying 271
 Effective 188
 For fields, wavelengths 272
 Method of computing 112
 Paraxial working 58
 Working 62
Fabrication support 208
Faceted object 439, 440
Faceted surfaces 304
Far field 629
Far field pattern 453
Far-field cosine space 463, 464, 466
Fast Trace option 354
Fiber coupling efficiency 165, 169, 206, 499, 500, 657, 668
Fiber optic object 443
Field angles
 Conventions for 53
Field curvature 180, 181, 182, 231, 233, 499
 Calculating 194
 Operand 499
Field curvature plot
 Meaning of 181, 182, 183
 Offsets on axis 183, 190
Field normalization 55, 115
Field positions
 Defining 115
Field type
 Multi-configuration control of 580
File Encoding 71
File menu 41, 65
File names
 On plots 68
File Prefix
 For Monte Carlo files 563
Files
 Finding names in ZPL macros 677
Filter string 378, 386
Filters
 User defined neutral density 333, 334
Fitting 361
 Test plates 242
Fletch
 Size of 70
Fletch Rays 119, 121
Flip
 Lens elements 254, 262
Float by stop size 101
 Defined 53
Floating apertures 80
Fluorescence 356, 358
Focal length
 Scaling by 263
Focal mode 103
Fold mirrors
 Inserting 96, 255
Folder names
 Preferences 68
Folders 68
Font size
 Changing 70, 71
 In text windows 70
 Setting in ZPL macros 731
 When printing 72
Footprint
 Beam 170
Footprint diagram 188
FOR loop
 ZPL macros 678
Foucault analysis 150
 As optimization operand 500
FR
 Codes 589
Fraunhofer diffraction 632
Freeform Z object 500
Fresnel
 General polynomial form 304
 Zone plates 343
Fresnel diffraction 631
Fresnel lens 418, 419, 650
Fresnel lens object 418
Fresnel lenses 440
Fresnel number 628
Fresnel object 441
Fresnel surfaces 295, 299, 303, 304
Fringes
 Test wavelength 557
F-theta
 DISC operand 496
F-Theta distortion 181, 183
Full width half max 161

G

Gases

Thermal modeling 595

Gaussian Beam

Analysis 222

Operands 501

Paraxial analysis 222

Physical optics analysis 627

Skew analysis 225

Gaussian factor

Optimization 488

Gaussian Quadrature 485

Gaussian scattering 366

Gaussian source 453

General data 101

Geometric MTF 139

Vs. Field 140

Geometry errors 378

Create source ray for 98

Ghost reflections 54, 263, 376, 386

Glass

Best fit 271

Codes 589

Constraining 529

Cost factors 589

Dispersion formulas for 590

Entering type 54

Excluding from optimization 588

Gradient index 305, 306, 308, 309, 312, 356

Model 271

Optimizing 528, 541

Optimizing with RGLA 518

Pick up solves 477

Removing in multi-configurations 584

Substitution 477, 529, 542, 588

Table data for 601

TCE 589

Templates 240

Transmittance of 211, 212

ZPL macros 669

Glass catalog 209, 210, 211, 212, 243, 244, 250, 587

Adding new glasses to 590, 592

Copying or moving 589

Creating new 589

Managing multiple 587

Melt data 593

Obsolete glasses 599

Selecting 587

Transmission data 594

Glass Substitution 686

Global coordinates 111, 205, 754, 756

From ZPL macros 664

Multi-configuration control of 581

Operand 501, 517

Reference surface 77, 111

Global optimization 239, 537, 539

Glue distance 114

GOSUB 687

Gouy shift 224, 634

gradient fill 73

Gradient index 305, 306, 308, 309, 311, 312

Boundary operands for 532

Operand 503

Vignetting 113

Gradient Lens Corporation 309

Gradient Top 73

Graphic frames 70

Graphics

Colors 73

Copying to clipboard 43

Exporting using Clipboard 48

Preferences 69

Grating mode 324

Gratings 296, 323

Elliptical 296, 297

Multiple orders 357

Toroidal 302, 329

Grid phase surface 313

Grid Sag object 421

Grid sag surface 313

Guard band 638

H

Hammer optimization 239

Hardware lock

Policy regarding 38

Harvey-Shack scattering 367

Heads-Up Displays (HUD) 180, 182

Help

Button 46

Menu option 41

Henry-Greenstein

Scattering 371

Hermite-Gaussian 638

Herzberger

Dispersion formula 591

Hexapolar rings

Defined 54

Hidden line drawings 118, 120, 125, 126, 230, 231

Hiding

Edges of surfaces 78

Rays on layout plots 78

Highlight color 70, 73

- Highlighting layouts 70
- Holograms 316, 317
 - Aspheric non-sequential 422
 - General polynomial form 283, 284
 - Multiple orders 357
 - Non-sequential solid 421
 - Optically fabricated 319
 - Optimizing construction optics 320
 - Toroidal 329
 - Used in reflection 316
 - Using Zernike coefficients 340, 342

- Huygens
 - Point Spread Function 146, 147

- Hx
 - Defined 55

- Hy
 - Defined 55

- Hyperbola
 - Defining 326

- Hyperhemispheres
 - Preventing with HHCN 504

- Hyperhemispheric 77
 - Described 59
 - Testing for 504

- Hyperhemispheric lens object 436

- Hyperhemispheric object 404

- Hyperhemispheric surface object 437

I

- Ideal coatings 619

- IES lamp data 457

- IESNA 259, 457, 463

- IGES files

 - Exporting 233, 234, 235, 257, 258

 - Importing 423

- Ignore objects 354

- Ignore surface 78

- Ignoring objects 354

- Illumination 170

- Image analysis 165, 170, 172, 177

 - Aborting long computations 46

 - As optimization operand 504

 - Editing IMA files 166, 173, 178, 180

 - IMA file format 167

- Image height

 - Field definition 53

- Image plane

 - Coordinates 131

- Image space F/#

 - Operand 504, 505

 - Setting 101

- Images

 - Appearance of 165, 170

- Importance Sampling 370

- Imported object 423

- Incoherent image 172

- Incoherent imaging 174

- Increment

 - tolerancing 561

- Index data

 - Operand 504

 - Optimizing 528

- Index of refraction 592, 600

 - Determining 271

 - Dispersion plot 209

 - Table of values 272

 - Temperature and pressure effects 603

- Inequality

 - Operand 512

- Initial Phase 358

- Input

 - ZPL macros 692

- Input focus 42

- Inserting

 - Operands 88, 89, 92, 93, 94

 - Surfaces 75, 87

 - Surfaces in ZPL macros 674, 675, 693, 694

- Inside of

 - Placing sources 449

- Integer

 - Rounding to using INTE 659

- Intensity 105

- Interface Language 71

- Interferogram

 - Data for surface deformation 340, 341, 343

 - Drawing 148

- Interferometers

 - Using data from 340, 342

- Internet 71

- Invariant 196

- Irradiance 105

- Irregular

 - Surface type 317

- Irregularity 550, 551

- ISO drawings 123

J

- Japanese 71

- Jones matrix 623

- Jones matrix object 395

- Jones matrix surface 318

- Jones vector 609

- JPG

 - Creating 43

- Jx, Jy 358

K

Keyboard

- Shortcuts 47

- Use of 46

Keyboard shortcuts 42, 46

Kinoforms 283, 284

Knife edge test 150

Korean 71

L

Lagrange invariant 196

- Operand 505

Lagrange multiplier 489

Lambertian 366

Lamp data in IES format 457

Landscape orientation 72

Language 71

Last zoom 43

Lateral color 190

- Operand 505

Layout

- Element drawing 122, 123

Layouts

- Scale bar width 73

- Three-dimensional 118, 120, 125, 126, 230, 231

- Two-dimensional 117

Lens catalog 246, 248

Lens Data Editor

- Printing contents of 271

Lens design

- References on 37

Lens element

- Reversing 262

- Tilting and decentering 254

Lens files

- Deleting 67

- Inserting 67

- Loading 65, 66, 67, 695, 696, 717

- Opening 65, 66

- Recently used list 74

- Saving 65, 717

Lens mounts

- Thermal expansion of 603

Lens object 436

Lens search 246, 248

Lens units

- Defined 54

- Specifying 104

Lenses

- Plotting to scale 118

Lenslet array 394, 395

Lenslet object 426, 427

Lightpipes 394, 395, 411, 412, 433

Line Spread Function 145, 160

Linear interpolation 314

Liquids

- Thermal modeling 595

Load and Save buttons 46

Loading

- Lens files 67

Lock window 43, 45, 697

Long computations

- Aborting from 46

Longitudinal aberration 187, 189, 190

- Operand 505

Longitudinal aberrations

- Calculating 194

Longitudinal color

- Calculating 194

Lost energy 378, 379

LPTD

- Operand 532

Ltk 589

Luneberg lens 337

M

Macro solves 479

Macros

- ZPL macros 651

Magnification

- Displaying 271

- Operand 491, 514

Main window 41

Marginal ray

- Definition of 54

- Solve 474, 475, 476

Margins 111

- Setting for printing 72

Mass

- Of lenses 272, 521

Maximum aspheric deviation 208

Mean free path 371

Melt data

- fitting 593

Melt Frequency 589

Merit function
 Default 483
 Defining 483
 Formula for 489
 Gaussian apodization of 488
 Ignoring operands in 498
 Loading and saving 89
 Loading from ZPL 696, 718
 Modifying 488
 Percent contribution of components 524
 Printing 240, 242
 ZPL macros 689, 703

Meta materials 588

Metafile
 Creating 43
 Windows clipboard format 69

Method to compute F/# 112

MicroElectroMechanical System (MEMS) 395

Micrometers 104

Microns 104

MIL number glasses 600

Mirrors
 Conventions for 60
 Deleting fold 96, 255
 Inserting fold 255
 Specifying 54, 76
 Substrate drawing 79

Model glass 271, 477, 528, 601

Modulation Transfer Function 134, 135, 136, 138, 139, 140, 141, 142, 145, 146, 148

Monte Carlo ray tracing 377

Mouse
 Shortcuts 44, 46

M-squared 223, 515, 646

MTF
 Aborting from long computation 46
 Calculating in macros 681, 682
 Geometric approximation 139
 Operand 506
 Optimization of 525
 Visualizing 165, 170
 Vs. Field 135, 136, 139, 140, 141
 Vs. Focus 134

MTF units 106

Multi-configuration 579
 Operand values 505
 Optimization and 584
 Tolerancing 577

Multimode 638, 642

Multi-mode fiber coupling 165, 169

Multiple CPU's
 Parallel processing 378, 526, 541

Multiple orders 357

Multi-threading 378, 526, 541

N

Narcissus
 YNI contribution 523

Near field 629

Nesting object limits 351

Nesting sources 449

Neural networks 539

Neutral density filters
 With user defined surfaces 333, 334

Nodal points 51, 272

Non-paraxial system
 Defined 55

Non-sequential 393, 449
 Converting surfaces to 267
 Defined 55
 Multi-configuration operands for 581
 Object definitions 393, 449
 Optimization of systems 531
 Optimization operands for 509, 512
 Ray tracing 345
 Ray tracing limits 113
 Restricting 113

Non-sequential program mode 67

Normal distribution 567

Normalize Coherent Power 467

Normalized coordinates 488, 709, 712
 Defined 57

Note files
 Editing 122

Notes 106

NSC 345
 With ports 346
 without ports 346

NSG America, Inc. 312

Null object 428

NURBS 97, 409

NURBS surfaces 324, 330

O

Object
 User defined 444

Object array 399

Object plane
 Tilting 327

Object space N.A.
 Setting 101

Object space numerical aperture
 Defined 57
 Operand 512

- Objects 393, 449
 - Arrays of 96
 - Ignoring 354
 - Inserting and deleting 94
 - Overlapping 351
 - Replicating 96
- Objects of revolution 439
- Obscuration type 79
- Obscurations 79, 84
 - Circular 79
 - Elliptical 79
 - Rectangular 79
 - User defined 80
 - Vignetting from 188
- Obsolete
 - glasses 599
- Odd Asphere 279
- Odd Asphere lens object 428
- Offense against the sine condition 513
- Ogive surface 339
- Opacity
 - Of objects 363
 - Of surfaces 77
- OPD
 - Optical path differences 512
- OpenGL 120
 - Shaded model graphics 120, 126
- Opening
 - Lens files 65, 66
- Operand data
 - Cut, Copy, and Paste 88, 90
- Operands
 - Defining complex 527
 - User defined 533
- Optical design
 - References on 37
- Optical invariant
 - Lagrange invariant 505
- Optical path differences
 - Operand 512
 - Referencing 110
 - Solves 476
- Optical transfer function 135, 138
- Optically fabricated hologram 319

- Optimization 239, 483
 - Global 239, 539
 - Global minimum 537
 - Global search 539
 - Hammer 239, 539
 - Merit functions for 483
 - Of glass 528
 - Operands 488, 527
 - Requirements 483
 - Selecting variables for 483
 - Suggestions for use 536
 - With external programs 534
 - With multi-configurations 584
 - ZPL macros 533

- Optimization operands
 - Plot of 212, 214

- Orders
 - Splitting rays by 357

- Orientation
 - Graphics 72

- Orientation indicator 119
 - Size of 70

- Orientation matrix, POP results 227

- Orthogonal Descent 526

- Output Pilot Radius 644

- Overlapping objects 351

- Overlaying graphics 43, 563

P

- Pan and zoom 45

- Parabasal
 - Discussion of 57

- Parabola
 - Defining 326

- Parallel processing
 - Multiple CPU's 378, 526, 541

- Parameter data
 - Defined 57, 277
 - Tolerancing 552

- Parameter solve
 - On coordinate breaks 479

- Paraxial 279
 - Definition of 57
 - Working F/# 58

- Paraxial focus 476

- Paraxial image height
 - Defined 58

- Paraxial lenses 321, 322

- Paraxial magnification
 - Defined 58

- Paraxial ray data
 - Operands 513

- Paraxial rays
 - And coordinate breaks 110
 - Tracing 205
- Paraxial surfaces 321, 322
- Paraxial XY 279
- Partially coherent image 172
- Paste
 - Operand data 89, 90, 92
 - Surface data 87, 89, 90, 92, 93, 95
- Pasting
 - Graphics and text 48
- Paths
 - For data files 68
- PAUSE
 - ZPL macros 705
- Peak to valley
 - Wavefront error 196
- Pen
 - Changing width of 72
- Pen color
 - ZPL macros 672
- Pen colors 73
- Pen width 69, 70, 72
- Periodic surface 322
- Petzval
 - Operand 514
- Phase
 - Point Spread Function 142
- Phase of a surface 152, 153
- Photometric units 105
- Physical optics 627
- Pick up solves 473
- Pilot Beam 644
- Pilot beam 633
- Pilot Radius 644
- Pinhole 644
- Pinhole aperture 650
- Pixel Interpolation 354
- Placing sources 449
- Plot width 72
- Plots
 - Changing pen width 72
- Plotting
 - All zoom positions 120
 - Optimization operands 212, 214
- POB files
 - Creating 96
- POB format 429
- Point Spread Function
 - Cross sections 144, 145, 147
 - FFT 141, 146
 - For NSC rays 466
 - Huygens 141, 146
- Point spread function 141, 146, 147
- Polarization 637
 - Defining global state 108
 - In non-sequential ray tracing 364
 - NSC sources 358
 - Partial 358
 - Phase of detectors 464, 466
 - Ray trace 216, 225, 229
 - Reflectivity, Transmission, Absorption 219
 - ZPL macros 708, 709
- Polygon object 429
- Polygon objects
 - Defining 429
- Polynomial
 - Extended surface type 301
- Polynomial surfaces 323
- POP orientation matrix 227
- Pop-up screens
 - Configuring 68
- Portrait orientation 72
- Ports in non-sequential tracing 346
- Power
 - Function 663
 - Maintaining with glass change 76
 - Operand 516
 - Solve 475
 - Watts 345
- PR
 - Codes 589
- Preferences 68
 - Address box 68, 69, 70
 - Editors 71
 - Folder names 68
 - Printing 72
- Prefix
 - Monte Carlo files 563
- Pre-Propagation 358
- Prescription Data 271
- Pressure
 - Of each surface 606
 - Setting 108
- Pressure effects 603
- Primary wavelength 116
 - Defined 58
 - Multi-configuration 582
- Principal planes 51, 272

Printing 42

- ASCII from ZPL macros 711
- Changing pen widths 72
- Coating definitions 249
- Contents of a window 271
- Font size 72
- Graphic windows 43
- Landscape mode 73
- Merit function 240
- Orientation of 72
- Portrait mode 73
- Preferences 72
- Prescription data 271
- Spreadsheet data 271
- Surface data 271
- System data 271
- Text margins 73
- Text windows 45, 271
- Tolerance data 242
- Using ZPL macros 710
- Windows in ZPL macros 711

Prisms 429

Probing rays 636

Profiles

- Scattering 109, 355

Program folder 68

Program mode 67

Progressive Addition Lenses 192

Propagation report

- For POP analysis 227, 648

Pupil aberration

- Controlling 106

Pupil position solve 477

Pupil positions

- Displaying 271

Pupil shift 107

Px

- Defined 57

Py

- Defined 57

Pyramid object 435

Q

Qbfs 323

Qcon 323

QSUM

- Operand 516

Q-Type 323

Quadratic sum

- Operand 516

Quick adjust 265

Quick focus 265

Quitting ZEMAX 74

R

Radii

- Entering 76

Radius

- Curvature 495

Radius of curvature

- Conventions for 58

Random number generator

- Seeding 557

Random Polarization 358

Ray aiming 77, 106

- While tolerancing 560

- ZPL macros 719

Ray database files 382

Ray fans 127, 129

Ray splitting 345, 376

Ray tracing 205

Rayleigh range 223, 632, 633, 634, 639

Raylist format 120

Raylists 120

Real propagation 58

Real ray

- Tracing 205

Recovering lost files 100

Rectangular 79

Rectangular object 433

Rectangular pipe object 433

Rectangular volume grating object 435

Rectangular volume object 435

Redo 100

Reference object 449

Reference objects

- Modifying 98

Reference sphere 110

References 37

Referencing OPD

- Described 110

Reflections

- Ghost 263

Reflectivity

- Of coatings 219

Reflectors

- Specifying 54, 76

Refractive index

- Index of refraction 271

Relative Cost 589

Relative cost of glasses 589

Relative illumination 187, 518

Replicate 399

Replicating objects 96

Report graphics 273

Resize columns 42

- Resolution
 - Estimating 165, 170
- Resolving
 - Two field points 131
- Retardance
 - Of coatings 221
- RETURN
 - ZPL macros 687
- Return
 - To prior surface coordinate 85
- Reverse Rays 358
- Reversing
 - Lens elements 254, 262
- RGLA
 - Operand 529
- Rings 486
- RMS
 - Spot size 131, 154, 155, 156, 157, 483
 - Wavefront error 154, 155, 156, 157, 196, 483
 - X and y spot size 131
- Robust
 - Ray aiming mode 107
- Roofs 434
- Rotating
 - 3D layout plots 119
- Rotation matrix 111, 272, 350
- Row color 77, 354
- Russian 71

S

- Sag
 - Constraining 519, 520
 - Operand 519, 520
 - ZPL macros 665
- Sag of a surface 151, 152, 153
- Sag table 208
- Sagittal
 - Definition of 59
- Sampling
 - Hexapolar rings 54
- Sampling Method 358
- SAT files
 - Exporting 258
 - Importing 423
- Save and Load buttons 46
- Saving
 - Lens files 65
- Scale bar
 - On layout plots 117
 - Setting accurately 73
- Scale factor 118
- Scale lens 262

- Scaling
 - On printers 73
- Scattering 364
 - ABg 367
 - Catalog of data 250
 - Gaussian 366
 - Harvey-Shack 367
 - Lambertian 366
 - On sequential surfaces 84
 - On spot diagrams 130, 141
 - User defined 368
 - Viewer 252
 - ZPL macros 718
- Scattering profile file 109
- Scattering profiles 355
- Schott
 - Constants of dispersion formula 590, 592
- Screen capture 48
- Scripted tolerancing 570
- Searching
 - For stock lenses 246, 248
- Seidel coefficients 194, 196, 200, 202
- SELFOC lenses 311, 312
- Sellmeier
 - Dispersion formula 590
- Semi diameter margin 111
- Semi-diameter
 - Accuracy of 112
- Semi-diameters
 - Accuracy of 59
 - Converting to apertures 253, 254
 - Defined 59
 - Entering 76
- Sequential program mode 67
- Sequential ray tracing 345
 - Defined 59
- Session files 65
- Set to default 483
- Set vignetting
 - On field dialog box 115
 - ZPL macros 733
- Variables 86
- Settings
 - Saving with the lens 46
- SETVECSIZE
 - ZPL macros 731, 732, 733
- Shaded model 120, 126
- Shear distance 313
- Shortcuts
 - Keyboard 41, 42
 - Summary of 47
- Sign conventions
 - For binary optics 284
 - For radius of curvature 58

- Simple Ray Splitting 114
- Simulated annealing 539
- Single mode fiber coupling 206
- Slide object 435
- Slider control 267
- SMIA-TV Distortion 186
- Sobol sampling 358, 723, 757
- Soft apertures 336
- Solid modeling 118, 120, 125, 126, 230, 231
- Solves 473
 - Aplanatic 475, 477
 - Center of curvature 477
 - Chief ray angle 475
 - Chief ray height 476
 - Chief ray normal 475
 - Compensator 477
 - Concentric 475, 477
 - Edge thickness 476
 - Marginal ray angle 474
 - Marginal ray height 476
 - Marginal ray normal 475
 - Optical path difference 476
 - Pick up 475, 476, 477, 478, 585
 - Position 477
 - Setting and removing 86
 - Suggestions for use 479
 - Thermal Pick Up 605
 - Turning off 473
 - ZPL Macro 95, 479, 734
- Sorting
 - Tolerance data 92
- Source
 - EULUMDAT 454
 - IES 457
- Source array 358
- Source arrow
 - Size of 114
- Source bitmap 435
- Source diode 450
- Source object
 - Ellipse 453, 459
 - Filament 454
 - File 454
 - Gaussian 457
 - Point 459
 - Rectangle 459
 - Tube 459
 - Volume cylinder 460
 - Volume ellipse 460
 - Volume rectangle 460
- Source units 105
- Sources 449
 - Extended 165, 170
 - User defined 454
- Special characters 60
- Spectrum 361
- Specular ray 365
- Spherical aberration
 - Calculating 194
 - Operand 520
- Spherical object 436
- Spherochromatism
 - Operand 520
- Spider
 - Spider aperture 79
- Spin Mode 43, 44, 70
- Spiral of Archimedes 401
- Spline
 - Bicubic using grid surface 314
 - Linear interpolation using grid surface 314
- Spline surfaces 294, 324, 330
- Spot diagrams 129, 132
 - Configuration matrix 132
 - Full field 131
 - Matrix 132
- Spot size 154, 155, 156, 157
- Spreadsheet
 - Printing Lens Data Editor contents 271
- Spreadsheet editor
 - Cursor controls for 75
 - Shortcuts 75
- Square root
 - Operand 505, 520
- Square wave MTF
 - Operand 506
- SR
 - Codes 589
- Standard deviation
 - Wavefront error 197
- Standard lens object 436
- Standard surface object 437
- Status bar
 - Modifying 74
- STEP files
 - Importing 423
- STL format 437
- STL Object 437
- STL objects
 - Defining 438
- Stock lens catalog 246, 248
- Stock lenses
 - Using search to find 246, 248
- Stop surface
 - Defining 77
 - Zooming in multi-configurations 583
- ZPL macros 737

- STP files
 - Exporting 258
 - Importing 423
- Straddling objects 351
- Strehl ratio 52, 60, 146, 147, 154, 155, 156, 157, 198, 520
 - ZPL macros 686
- String length
 - ZPL macros 663, 665
- String operators 656
- String variables
 - ZPL macros 655
- Strings
 - Comparing 662, 665, 666
- Subaperture data 197
- Subroutines
 - ZPL macros 687
- Substitute
 - Glass 477
- Substitution
 - Templates for 240
- Substrate
 - On mirrors 79
- Superconic
 - Surface type 326
- Supergaussian 450
- Surface apertures
 - Defined 60
 - Multi-configuration values 579
 - Specifying 79, 84
- Surface data
 - Cut, Copy, and Paste 87, 94
- Surface numbering 75
- Surface Phase 152, 153
- Surface power 271
- Surface Sag 151, 152, 153
- Surface tilts and decenters 84
- Surface types 277
 - ABCD 280
 - Alternate intersection 281
 - Conic 77, 277
 - Coordinate breaks 293
 - Cubic spline 294
 - Diffraction gratings 296, 323
 - Fresnel 303
 - Holograms 316
 - NURBS 324, 330
 - Plane 77, 277
 - Selecting 77, 277

- Surfaces
 - Deleting 75
 - Ignoring 78, 581
 - Inserting 75
 - Miscellaneous data 271
 - Reversing order of 254, 262
- Swept object 438
- System aperture
 - Defined 60
 - Multi-configuration value 579
 - Specifying 101
- System data
 - ZPL Macros 680, 683
- System menu 101

T

- Table coatings 619
- Table glasses 601
- Tabulated Fresnel object 441
- Tabulated object 439, 440
- Talbot imaging 650
- Tangential
 - Definition of 59
- Tapered coatings 615, 617
- TCE 589
 - Adding data 607
 - Data defined 605
 - In glass catalogs 589
- Technical support 39
- Telecentric mode 103
- TEM(0,0) 638
- Temperature
 - Of each surface 606
 - Reference 589, 604
 - Setting 108
 - Surface 605
 - System 108, 603
- Temperature effects 603
- Templates for glass substitution 240
- Test plate fit files 243
- Test plate fitting 242
- Test wavelength 557
 - Tolerancing 555
- Text
 - Copying to clipboard 45
- Text data
 - Exporting 683
- Text windows
 - Extracting data from 658
- Thermal effects 603
 - Auto setup 90
- Thermal modeling
 - Of gases and liquids 595

- Thickness
 - Boundary conditions on 495
 - Conventions for 60
 - Operand 522
 - Solves on 476
- Thicknesses
 - Entering 76
- Thin films 219
 - Coatings 219
- Thin lens
 - Cylindrical 322
 - Surface type 321
- Threading
 - Turning off 113
- Through focus
 - Geometric MTF 140
- Tilted surface 317, 327
- Tilts
 - Tolerancing 549
- Tilts and decenters
 - On surfaces 84
 - With coordinate break surfaces 293
- Time
 - Printing on plots and text 70, 71
- TIR 348, 349
 - Frustrated 625
 - Total Internal Reflection 60
- TIS 253, 366, 370
- Tolerance data
 - Cut, Copy, and Paste 92
 - Loading and saving 93
- Tolerance sensitivity
 - Optimizing for 530
- Tolerancing 242, 545
 - Debugging 577
 - Defining compensators for 559
 - Increment 561
 - Irregularity 550
 - Using scripts 570
- Toolbar
 - Defining 74
- Tools menu 239, 275, 345
- Top Hat 640
- Toroidal 280
- Toroidal faceted object 440
- Toroidal Grating 280
- Toroidal lens
 - Paraxial model for 322
- Toroidal lens object 404, 441, 442
- Toroidal surfaces 282, 328, 329
 - Gratings placed on 302, 329
 - Thermal effects 606
- Torus surface 434, 443
- Torus volume 434, 443
- Total Indicator Runout
 - Tolerancing 549
- Total Integrated Scatter 366
- Total integrated scatter 251
- Total Internal Reflection
 - ZPL Macros 712, 713
- Total track
 - Defined 60
 - Operand 521
- Transfer Functions 134, 135, 136, 138, 139, 140, 141, 142, 145, 146, 147, 148
- Transmission 218, 594
 - Internal 211, 212
 - Of coatings 220
- Transmission data
 - Glass catalogs 594
- Transparent objects
 - Opacity of 363
- Transparent surfaces
 - Opacity of 77
- Transverse aberrations
 - Calculating 194
- Transverse color
 - Calculating 194
- Triangle object 444
- Triangular corner object 444
- Tristimulus 360, 463, 464
- True color
 - PSF 147

U

- UDA Scale 81, 354
- Undo 100
 - Disabling 71
- Uniform distribution 567
- Uniform power spectrum 361
- Units 104
 - Determining in ZPL macros 667
 - Setting lens units 104
 - Source, irradiance, and flux 105
- Universal plot 212, 214
- Unlock window 43, 45, 738
- Unpolarized 109
- Unzoom 44
- Update
 - Described 101
- Update all
 - Described 101
- Updates 71
- Use polarization
 - Polarization state 108

- User defined
 - Apertures and obscurations 80, 364
 - Optimization operands 534
 - Surface 331
 - Surface apodization 333, 334
- User defined object 444
- User defined sources 454
- User defined surface
 - Diffractive optics 333
 - DLL 77
 - Gradient index 333
 - Lenslet arrays 333
 - Refractors and reflectors 332
 - Using linked code 331
 - Using look up table 313

V

- Vacuum
 - Index of refraction 603
 - Modeling 604
- Validate tolerances 93
- Variable line space gratings 339
- Variables
 - Defining within ZPL macros 732
 - Removing 240
 - Setting and removing 86
 - ZPL macros 716
- Variance
 - Wavefront error 196
- Vector variables
 - Setting the size of 733
- Vignetting 79, 84, 128, 188
 - ZPL macros 712, 713, 733
- Vignetting factors 495
 - Adjusting during optimization 520
 - Defined 61
 - Field curvature and distortion 184
 - Multi-configuration values 580
 - Removing 254
 - Setting 115
- Virtual propagation 62
- Volume
 - Of an element 271
 - Operand 522
- Volume detector 470
- Volume sources 449
- Volumes
 - Discussion of 272
- Volumetric scattering 356
- Voxels 410, 423, 470

W

- Watts 345

- Wavefront aberrations 194
- Wavefront error 154, 155, 156, 157
 - Peak to valley 196
 - RMS 196
- Wavefront propagation 627
- Wavelength Shift 356
- Wavelengths
 - Conventions for 62, 603
 - Weights for 116
- Waveplate Mode
 - Birefringent media 357
- Wedge object 435
- Weight
 - Configuration 580
 - Of lenses 272
- Wide angle apertures
 - In object space 101
- Window size
 - Changing default 69, 70
- Windows
 - Editor properties 71
 - Graphic 41
 - Resizing and moving 41
 - Shortcuts 47
 - Text 41
 - Types of 41
- Wire frame model 231
- Working F/# 101
 - Defined 62
 - Method of computing 112
 - Operand 522
 - Paraxial 58

Y

- YNU tracing 205
- Y-Ybar diagrams 190

Z

- ZBF beam file
 - Decentering 742
 - Properties of 741, 743
 - Resampling 742
- ZBF beam files
 - Macro commands related to 740
- ZBF file format 640, 641
- ZBF files
 - Macro commands for 692, 742
- ZEMAX Beam File 640, 642
- ZEMAX Encrypted Coating (ZEC) file 249
- ZEMAX Programming Language
 - ZPL macros 651
- Zernike
 - Deformed surfaces 340, 341, 343

- Zernike coefficients
 - Calculation of 200, 202
 - Fringe 198
 - Optimization of 523
 - Standard 200
- Zernike phase surface
 - Optimizing coefficients 340
- Zernike surface 405, 446
 - Optimizing coefficients 299, 340
- ZOF (ZEMAX Object Format) 424
- Zone plates 343
- Zoom and pan 45
- Zoom in, out 43
- Zoom lenses
 - Maintaining length of 477
 - Maximum semi-diameter solve 478
 - Multi-configuration feature 579
 - Operand 524
- Zoom Mode 43, 70
- Zoom positions
 - Changing 116
 - On plots 68
 - Plotting all 120
 - Showing number on plots 68
- ZPL Macro solves 95, 479, 734
- ZPL macros 275, 276, 651
 - Editing 651
 - Optimization of 533
 - Printing 651
 - Status message 651
 - String variables 655
- ZRD file format 382