

C Language Reference Manual

007-0701-130

St. Peter's Basilica image courtesy of ENEL SpA and InfoByte SpA. Disk Thrower image courtesy of Xavier Berenguer, Animatica.

Copyright © 1999 Silicon Graphics, Inc. All Rights Reserved. This document or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Silicon Graphics, Inc.

LIMITED AND RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in the Rights in Data clause at FAR 52.227-14 and/or in similar or successor clauses in the FAR, or in the DOD, DOE or NASA FAR Supplements. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Silicon Graphics, Inc., 1600 Amphitheatre Pkwy., Mountain View, CA 94043-1351.

Autotasking, CF77, CRAY, Cray Ada, CraySoft, CRAY Y-MP, CRAY-1, CRInform, CRI/*TurboKiva*, HSX, LibSci, MPP Apprentice, SSD, SUPERCLUSTER, UNICOS, X-MP EA, and UNICOS/mk are federally registered trademarks and Because no workstation is an island, CCI, CCMT, CF90, CFT, CFT2, CFT77, ConCurrent Maintenance Tools, COS, Cray Animation Theater, CRAY APP, CRAY C90, CRAY C90D, Cray C++ Compiling System, CrayDoc, CRAY EL, CRAY J90, CRAY J90se, CrayLink, Cray NQS, Cray/REELibrarian, CRAY S-MP, CRAY SSD-T90, CRAY SV1, CRAY T90, CRAY T3D, CRAY T3E, CrayTutor, CRAY X-MP, CRAY XMS, CRAY-2, CSIM, CVT, Delivering the power . . ., DGauss, Docview, EMDS, GigaRing, HEXAR, IOS, ND Series Network Disk Array, Network Queuing Environment, Network Queuing Tools, OLNET, RQS, SEGLDR, SMARTE, SUPERLINK, System Maintenance and Remote Testing Environment, Trusted UNICOS, and UNICOS MAX are trademarks of Cray Research, Inc., a wholly owned subsidiary of Silicon Graphics, Inc.

DynaText and DynaWeb are registered trademarks of Inso Corporation. Gaussian is a trademark of Gaussian, Inc. OpenMP is a trademark of the OpenMP Architecture Review Board. Portions of this publication may have been derived from the OpenMP Language Application Program Interface Specification. Silicon Graphics and the Silicon Graphics logo are registered trademarks of Silicon Graphics, Inc. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited. X/Open is a trademark of X/Open Company Ltd. The X device is a trademark of the Open Group.

The UNICOS operating system is derived from UNIX® System V. The UNICOS operating system is also based in part on the Fourth Berkeley Software Distribution (BSD) under license from The Regents of the University of California.

New Features

This revision of the *C Language Reference Manual* supports the 7.3 release of the MIPSpro compiler. See the `cc(1)` man page for changes or additions to command-line options.

This manual contains the following additions and deletions:

- Information about OpenMP C/C++ Application Program Interface (API) multiprocessing directives has been added.
- Information about the auto-parallelizing option (APO) has been added; this information was taken from the *MIPSpro Auto-Parallelizing Option Programmer's Guide*, which is no longer being revised. Detailed information can be found in the *MIPSpro C and C++ Pragmas* manual.
- Details about the multiprocessing C/C++ compiler `#pragma` directives have been removed. See the *MIPSpro C and C++ Pragmas* manual for information about `#pragma` directives.

Record of Revision

<i>Version</i>	<i>Description</i>
7.3	April 1999 This document has been updated to support the MIPSpro 7.3 release.

Contents

	<i>Page</i>
About This Manual	xv
Related Publications	xv
Obtaining Publications	xv
Conventions	xvi
Reader Comments	xvii
An Overview of ANSI C [1]	1
ANSI C	1
Strictly Conforming Programs	1
Name Spaces	1
Compiling ANSI Programs	2
Guidelines for Using ANSI C	2
Compiling Traditional C Programs	3
Helpful Programming Hints	3
Recommended Practices	3
Practices to Avoid	4
C Language Changes [2]	7
Preprocessor Changes	7
Replacement of Macro Arguments in Strings	8
Token Concatenation	9
Changes in Disambiguating Identifiers	10
Scoping Differences	10
Name Space Changes	12
Changes in the Linkage of Identifiers	12

	<i>Page</i>
Types and Type Compatibility	13
Type Promotion in Arithmetic Expressions	14
Type Promotion and Floating Point Constants	15
Compatible Types	16
Argument Type Promotions	17
Mixed Use of Functions	18
Function Prototypes	18
External Name Changes	19
Changes in Function Names	19
Changes in Linker-Defined Names	20
Data Area Name Changes	20
Standard Headers	21
Lexical Conventions [3]	23
Comments	23
Identifiers	23
Keywords	23
Constants	24
Integer Constants	24
Character Constants	25
Special Characters	25
Trigraph Sequences (ANSI C Only)	26
Floating Constants	27
Enumeration Constants	27
String Literals	28
Operators	28
Punctuators	29
Meaning of Identifiers [4]	31

	<i>Page</i>
Disambiguating Names	31
Scope	31
Block Scope	31
Function Scope	32
Function Prototype Scope	32
File Scope	32
Name Spaces	33
Name Space Discrepancies Between Traditional and ANSI C	33
Linkage of Identifiers	34
Linkage Discrepancies Between Traditional and ANSI C	36
Storage Duration	37
Object Types	38
Character Types	38
Integer and Floating Point Types	38
Derived Types	40
void Type	40
Objects and lvalues	40
Operator Conversions [5]	43
Conversions of Characters and Integers	43
Conversions of Float and Double	43
Conversion of Floating and Integral Types	44
Conversion of Pointers and Integers	44
Conversion of unsigned Integers	44
Arithmetic Conversions	45
Integral Promotions	45
Usual Arithmetic Conversions	45
Traditional C Conversion Rules	45

	<i>Page</i>
ANSI C Conversion Rules	46
Conversion of Other Operands	47
Conversion of lvalues and Function Designators	47
Conversion of void Objects	47
Conversion of Pointers	47
Expressions and Operators [6]	49
Precedence and Associativity Rules in C	49
Primary Expressions	51
Postfix Expressions	51
Subscripts	52
Function Calls	52
Structure and Union References	54
Indirect Structure and Union References	54
postfix ++ and postfix --	54
Unary Operators	55
Address-of and Indirection Operators	55
Unary + and - Operators	56
Unary ! and ~ Operators	56
Prefix ++ and -- Operators	56
sizeof Unary Operator	57
Cast Operators	57
Multiplicative Operators	58
Additive Operators	59
Shift Operators	60
Relational Operators	60
Equality Operators	61
Bitwise AND Operator	62
Bitwise Exclusive OR Operator	62

	<i>Page</i>
Bitwise Inclusive OR Operator	62
Logical AND Operator	62
Logical OR Operator	63
Conditional Operator	63
Assignment Operators	64
Assignment Using = (Simple Assignment)	65
Compound Assignment	65
Comma Operator	65
Constant Expressions	66
Integer and Floating Point Exceptions	67
Declarations [7]	69
Storage Class Specifiers	70
Type Specifiers	71
Structure and Union Declarations	72
Bitfields	75
Enumeration Declarations	76
Type Qualifiers	77
Declarators	78
Meaning of Declarators	78
Pointer Declarators	79
Qualifiers and Pointers	79
Pointer-related Command Options	80
Array Declarators	80
Function Declarators and Prototypes	82
Prototyped Functions Summarized	84
Restrictions on Declarators	85
Type Names	85
Implicit Declarations	86

	<i>Page</i>
typedef	87
Initialization	88
Initialization of Aggregates	89
Examples of Initialization	90
Statements [8]	93
Expression Statement	93
Compound Statement or Block	93
Selection Statements	94
if Statement	94
switch Statement	95
Iteration Statements	96
while Statement	96
do Statement	96
for Statement	96
Jump Statements	97
goto Statement	97
continue Statement	98
break Statement	98
return Statement	98
Labeled Statements	99
External Definitions [9]	101
External Function Definitions	101
External Object Definitions	102
Multiprocessing Directives [10]	103
OpenMP C/C++ API Multiprocessing Directives	104
Using Parallel Regions	104
Coding Rules of #pragma Directives	105
Parallel Regions	106

	<i>Page</i>
Parallel Reduction Operations in C and C++	109
Restrictions on the Reduction Clause	109
Performance Considerations	110
Reduction on User-Defined Types in C++	111
Reduction Example	111
Restrictions for the C++ Compiler	112
Restrictions on pfor	113
Restrictions on Exception Handling	113
Scoping Restrictions	115
Multiprocessing Advanced Features [11]	117
Run-time Library Routines	117
mp_block and mp_unblock	117
mp_setup, mp_create, and mp_destroy	117
mp_blocktime	118
mp_numthreads, mp_suggested_numthreads, mp_set_numthreads	118
mp_my_threadnum	119
mp_setlock, mp_unsetlock, mp_barrier	119
mp_set_slave_stacksize	119
Run-time Environment Variables	120
MP_SET_NUMTHREADS, MP_BLOCKTIME, MP_SETUP	120
MP_SUGNUMTHD, MP_SUGNUMTHD_MIN, MP_SUGNUMTHD_MAX, MP_SUGNUMTHD_VERBOSE	121
MP_SCHEDTYPE, CHUNK	121
MP_SLAVE_STACKSIZE	122
MPC_GANG	122
Communicating Between Threads Through Thread Local Data	122
Synchronization Intrinsics	125
Atomic fetch-and-op Operations	126
Atomic op-and-fetch Operations	127

	<i>Page</i>
Atomic compare-and-swap Operation	128
Atomic synchronize Operation	128
Atomic lock and unlock Operations	129
Atomic lock-test-and-set Operation	129
Atomic lock-release Operation	129
Example of Implementing a Pure Spin-Wait Lock	130
Parallel Programming on Origin Servers [12]	131
Performance Tuning of Parallel Programs on an Origin2000 Server	131
Improving Program Performance	132
Types of Data Distribution	135
Regular Data Distribution	135
Data Distribution With Reshaping	135
Implementation of Reshaped Arrays	135
Regular Versus Reshaped Data Distribution	138
Choosing Between Multiple Options	139
Explicit Placement of Data	140
Affinity Scheduling	140
Data and Thread Affinity	141
Data Affinity	141
Data Affinity for Redistributed Arrays	142
Data Affinity for a Formal Parameter	142
Data Affinity and the <code>#pragma pfor nest</code> Clause	142
Directives for Performance Tuning on Origin2000	143
Query Intrinsics for Distributed Arrays	143
Optional Environment Variables and Compile-Time Options	145
Multiprocessing Environment Variables	145
Compile-Time Options	146

	<i>Page</i>
Examples	147
Distributing a Matrix	147
Using Data Distribution and Data Affinity Scheduling	148
Parameter Passing	149
Redistributed Arrays	150
Irregular Distributions and Thread Affinity	151
The Auto-Parallelizing Option (APO) [13]	153
Appendix A Implementation-Defined Behavior	155
Translation (F.3.1)	155
Environment (F.3.2)	156
Identifiers (F.3.3)	157
Characters (F.3.4)	157
Integers (F.3.5)	159
Floating Point (F.3.6)	160
Arrays and Pointers (F.3.7)	161
Registers (F.3.8)	162
Structures, Unions, Enumerations, and Bitfields (F.3.9)	162
Qualifiers (F.3.10)	164
Declarators (F.3.11)	164
Statements (F.3.12)	165
Preprocessing Directives (F.3.13)	165
Library Functions (F.3.14)	166
Signals	167
Signal Notes	171
Diagnostics	174
Streams and Files	175
Temporary Files	177

	<i>Page</i>
errno and perror	177
Memory Allocation	184
abort Function	184
exit Function	184
getenv Function	185
system Function	185
strerror Function	185
Time Zones and the clock Function	185
Locale-Specific Behavior (F.4)	186
Common Extensions (F.5)	186
Environment Arguments (F.5.1)	187
Specialized Identifiers	187
Lengths and Cases of Identifiers	187
Scopes of Identifiers (F.5.4)	187
Writable String Literals (F.5.5)	188
Other Arithmetic Types (F.5.6)	188
Function Pointer Casts (F.5.7)	188
Non-int Bit-Field Types (F.5.8)	188
fortran Keyword (F.5.9)	189
asm Keyword (F.5.10)	189
Multiple External Definitions (F.5.11)	189
Empty Macro Arguments (F.5.12)	189
Predefined Macro Names (F.5.13)	190
Extra Arguments for Signal Handlers (F.5.14)	190
Additional Stream Types and File-Opening Modes (F.5.15)	190
Defined File Position Indicator (F.5.16)	190
Appendix B lint-style Comments	193

	<i>Page</i>
Appendix C Built-in Functions	195
Index	197
Figures	
Figure 1. Program Execution	108
Figure 2. Origin2000 Memory Hierarchy	132
Figure 3. Cache Behavior and Solutions	134
Figure 4. Implementation of block Distribution	136
Figure 5. Implementation of cyclic(1) Distribution	137
Figure 6. Implementation of block-cyclic Distribution	138
Tables	
Table 1. Effect of Compilation Options on Floating Point Conversions	15
Table 2. Using <code>__STDC__</code> to Affect Floating Point Conversions	16
Table 3. Effect of Compilation Mode on Names	21
Table 4. Reserved Keywords	23
Table 5. Escape Sequences for Nongraphic Characters	25
Table 6. Trigraph Sequences	27
Table 7. Storage Class Sizes	39
Table 8. Precedence and Associativity Examples	49
Table 9. Operator Precedence and Associativity	50
Table 10. Examples of Type Names	86
Table 11. Multiprocessing C/C++ Compiler Directives	103
Table 12. Loop Nest Optimization Directives Specific to the Origin2000 Server	143
Table 13. Integer Types and Ranges	159
Table 14. Ranges of floating point Types	161
Table 15. Alignment of Structure Members	163
Table 16. Signals	168
007-0701-130	xiii

	<i>Page</i>
Table 17. Valid Codes in a Signal-Catching Function	170
Table 18. lint-style Comments	193
Table 19. Built-in Functions	195

About This Manual

This manual contains a summary of the syntax and semantics of the C programming language as implemented on Silicon Graphics workstations. It documents previous releases of the Silicon Graphics C compiler as well as the American National Standards Institute (ANSI) C compiler.

The Silicon Graphics compiler system supports three modes of compilation: the old 32-bit mode (`-o32` or `-32` option), the new 32-bit mode (`-n32` option), and the 64-bit mode (`-64` option). For information on compilation modes and general compiler options for the old 32-bit mode, see the `o32(5)` man page and the *MIPS Compiling and Performance Tuning Guide*. For information on the new 32-bit mode and 64-bit mode, see the `cc(1)` man page and the *MIPSpro Compiling and Performance Tuning Guide*.

The term “traditional C” refers to the dialect of C described in the first edition of *The C Programming Language* by Kernighan and Ritchie.

Related Publications

The following documents contain additional information that may be helpful:

- *MIPSpro C and C++ Pragmas* describes the `#pragma` directives supported by the MIPSpro C and C++ compilers.
- *MIPS Compiling and Performance Tuning Guide* and *MIPSpro Compiling and Performance Tuning Guide* describe the compiler system, Dynamic Shared Objects (DSOs), and programming tools and interfaces. They also explain ways to improve program performance.
- *Topics in IRIX Programming* presents information about internationalizing an application; working with fonts; file and record locking; and interprocess communication.

Obtaining Publications

The *User Publications Catalog* describes the availability and content of all Cray hardware and software documents that are available to customers. Customers who subscribe to the Cray Inform (CRInform) program can access this information on the CRInform system.

To order a document, call +1 651 683 5907. Silicon Graphics employees may send electronic mail to `orderdisk@sgi.com` (UNIX system users).

Customers who subscribe to the CRInform program can order software release packages electronically by using the `Order Cray Software` option.

Customers outside of the United States and Canada should contact their local service organization for ordering and documentation information.

Conventions

The following conventions are used throughout this document:

<u>Convention</u>	<u>Meaning</u>																				
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.																				
<code>manpage(x)</code>	Man page section identifiers appear in parentheses after man page names. The following list describes the identifiers: <table><tbody><tr><td>1</td><td>User commands</td></tr><tr><td>1B</td><td>User commands ported from BSD</td></tr><tr><td>2</td><td>System calls</td></tr><tr><td>3</td><td>Library routines, macros, and opdefs</td></tr><tr><td>4</td><td>Devices (special files)</td></tr><tr><td>4P</td><td>Protocols</td></tr><tr><td>5</td><td>File formats</td></tr><tr><td>7</td><td>Miscellaneous topics</td></tr><tr><td>7D</td><td>DWB-related information</td></tr><tr><td>8</td><td>Administrator commands</td></tr></tbody></table>	1	User commands	1B	User commands ported from BSD	2	System calls	3	Library routines, macros, and opdefs	4	Devices (special files)	4P	Protocols	5	File formats	7	Miscellaneous topics	7D	DWB-related information	8	Administrator commands
1	User commands																				
1B	User commands ported from BSD																				
2	System calls																				
3	Library routines, macros, and opdefs																				
4	Devices (special files)																				
4P	Protocols																				
5	File formats																				
7	Miscellaneous topics																				
7D	DWB-related information																				
8	Administrator commands																				
<i>variable</i>	Some internal routines (for example, the <code>_assign_asgcmd_info()</code> routine) do not have man pages associated with them. Italic typeface denotes variable entries and words or concepts being defined.																				

user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.
[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.

The default shell in the UNICOS and UNICOS/mk operating systems, referred to as the *standard shell*, is a version of the Korn shell that conforms to the following standards:

- Institute of Electrical and Electronics Engineers (IEEE) Portable Operating System Interface (POSIX) Standard 1003.2-1992
- X/Open Portability Guide, Issue 4 (XPG4)

The UNICOS and UNICOS/mk operating systems also support the optional use of the C shell.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please tell us. Be sure to include the title and part number of the document with your comments.

You can contact us in any of the following ways:

- Send electronic mail to the following address:

`techpubs@sgi.com`

- Send a facsimile to the attention of "Technical Publications" at fax number +1 650 932 0801.
- Use the Suggestion Box form on the Technical Publications Library World Wide Web page:

`http://techpubs.sgi.com/library/`

- Call the Technical Publications Group, through the Technical Assistance Center, using one of the following numbers:

For Silicon Graphics IRIX based operating systems: 1 800 800 4SGI

For UNICOS or UNICOS/mk based operating systems or CRAY Origin2000 systems: 1 800 950 2729 (toll free from the United States and Canada) or +1 651 683 5600

- Send mail to the following address:

Technical Publications
Silicon Graphics, Inc.
1600 Amphitheatre Pkwy.
Mountain View, California 94043-1351

We value your comments and will respond to them promptly.

An Overview of ANSI C [1]

This chapter briefly discusses the scope of the new standard and lists some programming practices to avoid and some to use.

1.1 ANSI C

The ANSI standard on the C programming language is designed to promote the portability of C programs among a variety of data-processing systems. To accomplish this, the standard covers three major areas: the environment in which the program compiles and executes, the semantics and syntax of the language, and the content and semantics of a set of library routines and header files.

1.1.1 Strictly Conforming Programs

Strictly conforming programs adhere to the following guidelines:

- They use only those features of the language defined in the standard.
- They do not produce output dependent on any ill-defined behavior. Ill-defined behavior includes implementation-defined, undefined, and unspecified behavior which refers to areas that the standard does not specify.
- They do not exceed any minimum limit.

This ANSI C environment is designed to be a conforming hosted implementation, which will accept any strictly conforming program. Extensions are allowed only if the behavior of strictly conforming programs is not altered.

1.1.2 Name Spaces

Besides knowing which features of the language and library you can rely on when writing portable programs, you must be able to avoid naming conflicts with support routines used for the implementation of the library. To avoid such naming conflicts, ANSI divides the space of available names into a set reserved for the user and a set reserved for the implementation. Any name is in the user's name space if it meets these three requirements (this rule is given for simplicity; the space of names reserved for the user is actually somewhat larger than this):

- It does not begin with an underscore

- It is not a keyword in the language
- It is not reserved for the ANSI library

Strictly conforming programs may not define any names unless they are in the user's namespace. New keywords as well as those names reserved for the ANSI library are discussed in Section 2.6, page 21.

1.1.3 Compiling ANSI Programs

To provide the portable clean environment dictated by ANSI while retaining the many extensions available to Silicon Graphics users, two modes of compilation are provided for ANSI programs. Each of these switches to the `cc` command invokes the ANSI compiler:

- `-ansi` Enforces a pure ANSI environment, eliminating Silicon Graphics extensions. The ANSI symbol indicating a pure environment (`__STDC__`) is defined to be 1 for the preprocessor. Use this mode when compiling strictly conforming programs, because it guarantees purity of the ANSI namespace.
- `-xansi` Adds Silicon Graphics extensions to the environment. This mode is the default. The ANSI preprocessor symbol (`__STDC__`) is defined to be 1. The symbol to include extensions from standard headers (`__EXTENSIONS__`) is also defined, as is the symbol to inline certain library routines that are directly supported by the hardware (`__INLINE_INTRINSICS`.) Note that when these library routines are made to be intrinsic, they may no longer be strictly ANSI conforming (for example, `errno` may not be set correctly).

1.1.4 Guidelines for Using ANSI C

The following are some key facts to keep in mind when you use ANSI C:

- Use only `-lc` and/or `-lm` to specify the C and/or math libraries. These switches ensure the incorporation of the ANSI version of these libraries.
- Use the switch `-fullwarn` to receive additional diagnostic warnings that are suppressed by default. Silicon Graphics recommends using this option with the `-woff` option to remove selected warnings during software development.

- Use the switch `-wlint` (`-o32` mode only) to get `lint`-like warnings about the compiled source. This option provides `lint`-like warnings for ANSI and `-cckr` modes and can be used together with the other `cc` options and switches.
- Remember that the default compilation mode is shared and the libraries are shared.

1.1.5 Compiling Traditional C Programs

To compile code using traditional C (that is, non-ANSI), use the switch `-cckr`. The dialect of C invoked by `[-cckr]` is referred to interchangeably as `-cckr`, “the previous version of Silicon Graphics C,” and “traditional C” in the remainder of this document.

You can find complete information concerning ANSI and non-ANSI compilation modes in the `cc(1)` online reference page.

1.2 Helpful Programming Hints

Although the ANSI Standard has added only a few new features to the C language, it has tightened the semantics of many areas. In some cases, constructs were removed that were ambiguous, no longer used, or obvious hacks. The next two sections give two lists of programming practices. The first section recommends practices that you can use to ease your transition to this new environment. The second section lists common C coding practices that cause problems when you use ANSI C.

1.2.1 Recommended Practices

Follow these recommendations as you code:

- Always use the appropriate header file when declaring standard external functions. Avoid embedding the declaration in your code. This avoids inconsistent declarations for the same function.
- Always use function prototypes, and write your function prologues in function prototype form.
- Use the `offsetof()` macro to derive structure member offsets. The `offsetof()` macro is in `<stddef.h>`.
- Always use casts when converting.

- Be strict with your use of qualified objects, such as with `volatile` and `const`. Assign the addresses of these objects only to pointers that are so qualified.
- Return a value from all return points of all non-void functions.
- Use only structure designators of the appropriate type as the structure designator in `.` and `->` expressions (that is, ensure that the right side is a member of the structure on the left side).
- Always specify the types of integer bitfields as `signed` or `unsigned`.

1.2.2 Practices to Avoid

Avoid the following as you code:

- Never mix prototyped and nonprototyped declarations of the same function.
- Never call a function before it has been declared. This may lead to an incompatible implicit declaration for the function. In particular, this is unlikely to work for prototyped functions that take a variable number of arguments.
- Never rely on the order in which arguments are evaluated. For example, what is the result of the code fragment `foo(a++, a, ...)`?
- Avoid using expressions with side effects as arguments to a function.
- Avoid two side effects to the same data location between two successive sequence points (for example, `x=++x;`).
- Avoid declaring functions in a local context, especially if they have prototypes.
- Never access parameters that are not specified in the argument list unless using the `stdarg` facilities. Use the `stdarg` facilities only on a function with an unbounded argument list (that is, an argument list terminated with `...`).
- Never cast a pointer type to anything other than another pointer type or an integral type of the same size (`unsigned long`), and vice versa. Use a union type to access the bit-pattern of a pointer as a nonintegral and nonpointer type (that is, as an array of `chars`).
- Do not hack preprocessor tokens (for example, `FOO/**/BAR`).
- Never modify a string literal.

- Do not rely on search rules to locate `include` files that you specify with quotation marks.

C Language Changes [2]

This chapter describes changes to the C language, which include the following:

- Section 2.1, page 7, discusses two changes in the way the preprocessor handles string literals and tokens.
- Section 2.2, page 10, covers the four characteristics ANSI C uses to distinguish identifiers.
- Section 2.3, page 13, describes ANSI C changes to type promotions and type compatibility.
- Section 2.4, page 18, explains how ANSI C handles function prototyping.
- Section 2.5, page 19, discusses the changes in function, linker-defined, and data area names.
- Section 2.6, page 21, lists standard header files.

2.1 Preprocessor Changes

When compiling in an ANSI C mode (which is the default unless you specify `[-cckr]`), ANSI-standard C preprocessing is used. The preprocessor is built into the compiler and is functionally unchanged from the version appearing on IRIX™ Release 3.10.

The 3.10 version of the compiler had no built-in preprocessor and used two standalone preprocessors, for `-cckr` (`cpp(1)`) and ANSI C (`acpp(5)`) preprocessing, respectively. If you compile using the `-o32` option, you can activate `acpp` or `cpp` instead of the built-in preprocessor by using the `-oldcpp` option, and `acpp` in `-cckr` mode by using the `-acpp` option. Silicon Graphics recommends that you always use the built-in preprocessor, rather than `cpp` or `acpp`, because these standalone preprocessors may not be supported in future releases of the compilers.

`acpp` is a public domain preprocessor and its source is included in `/usr/src/gnu/acpp`.

Traditionally, the C preprocessor performed two functions that are now illegal under ANSI C. These functions are the substitution of macro arguments within string literals and the concatenation of tokens after removing a null comment sequence.

2.1.1 Replacement of Macro Arguments in Strings

Suppose you define two macros, *IN* and *PLANT*, as shown in this example:

```
#define IN(x)    `x`  
#define PLANT(y) "placing y in a string"
```

Later, you invoke them as follows:

```
IN(hi)  
PLANT(foo)
```

Compiling with `-cckr` makes these substitutions:

```
`hi`  
"placing foo in a string"
```

However, because ANSI C considers a string literal to be an atomic unit, the expected substitution does not occur. So, ANSI C adopted an explicit preprocessor sequence to accomplish the substitution.

In ANSI C, adjacent string literals are concatenated. Therefore, this is the result:

"abc" "def" becomes "abcdef".

This concatenation led to a mechanism for quoting a macro argument. When a macro definition contains one of its formal arguments preceded by a single #, the substituted argument value is quoted in the output. The simplest example of this is as follows:

Macro:	Invoked as:	Yields:
#define STRING_LITERAL(a) # a	STRING_LITERAL(foo)	``foo``

In conjunction with the rule of concatenation of adjacent string literals, the following macros can be defined:

Macro:	Invoked as:	Yields:
#define ARE(a,c) # a "are" # c	ARE(trucks,big)	"trucks" "are" "big" or ``trucks are big``

Blanks prepended and appended to the argument value are removed. If the value has more than one word, each pair of words in the result is separated by a single blank. Thus, the `ARE` macro could be invoked as follows:

Macro:	Invoked as:	Yields:
<code>#define ARE(a,c) # a "are" # c</code>	<code>ARE(fat cows, big)</code> or <code>ARE(fat cows, big)</code>	<code>"fat cows are big"</code>

Avoid enclosing your macro arguments in quotes, because these quotes are placed in the output string. For example:

```
ARE (`fat cows`, `big`) becomes ``\`fat cows\` are
\`big\`''
```

No obvious facility exists to enclose macro arguments with single quotes.

2.1.2 Token Concatenation

When compiling `[-cckr]`, the value of macro arguments can be concatenated by entering

```
#define glue(a,b) a/**/b
glue(FOO,BAR)
```

The result yields `FOOBAR`.

This concatenation does not occur under ANSI C, because null comments are replaced by a blank. However, similar behavior can be obtained by using the `##` operator in `-ansi` and `-xansi` mode. `##` instructs the precompiled to concatenate the value of a macro argument with the adjacent token, as illustrated by the following example:

This code:	Yields:
<code>#define glue_left(a) GLUED ## a</code> <code>#define glue_right(a) a ## GLUED</code> <code>#define glue(a,b) a ## b</code> <code>glue_left(LEFT)</code>	<code>GLUEDLEFT</code>

This code:	Yields:
<code>glue_right(RIGHT)</code>	<code>RIGHTGLUED</code>
<code>glue(LEFT,RIGHT)</code>	<code>LEFTRIGHT</code>

Furthermore, the resulting token is a candidate for further replacement. Note what happens in this example:

This code:	Yields:
<code>#define HELLO ``hello``</code>	
<code>#define glue(a,b) a ## b</code>	
<code>glue(HEL,LO)</code>	<code>``hello``</code>

2.2 Changes in Disambiguating Identifiers

Under ANSI C, an identifier has four disambiguating characteristics: its scope, linkage, name space, and storage duration. Each of these characteristics was used in traditional C, either implicitly or explicitly. Except in the case of storage duration, which is either static or automatic, the definitions of these characteristics chosen by the standard differ in certain ways from those you may be accustomed to, as detailed in Section 2.2.1, page 10, Section 2.2.2, page 12, and Section 2.2.3, page 12. For a discussion of the same material with a different focus, see Section 4.1, page 31.

2.2.1 Scoping Differences

ANSI C recognizes four scopes of identifiers: the familiar file and block scopes and the new function and function prototype scopes.

- Function scope includes only labels. As in traditional C, labels are valid until the end of the current function.
- Block scope rules differ from traditional C in one significant instance: the outermost block of a function and the block that contains the function arguments are the same under ANSI C.

For example, when compiling the following code, ANSI C complains of a redeclaration of *x*, whereas traditional C quietly hides the argument *x* with the local variable *x*, as if they were in distinct scopes:


```

int f(x);
int x;
{
    int x;
    x = 1;
}

```

- Function prototype scope is a new scope in ANSI C. If an identifier appears within the list of parameter declarations in a function prototype that is not part of a function definition, it has function prototype scope, which terminates at the end of the prototype. This allows any dummy parameter names appearing in a function prototype to disappear at the end of the prototype.

Consider the following example:

```

char * getenv (const char * name);
int name;

```

The `int` variable `name` does not conflict with the parameter `name` because the parameter went out of scope at the end of the prototype. However, the prototype is still in scope.

- File scope applies to identifiers appearing outside of any block, function, or function prototype.

One last discrepancy in scoping rules between ANSI and traditional C concerns the scope of the function `foo()` in the following example:

```

float f;
func0() {
    extern float foo() ;
    f = foo() ;
}
func1() {
    f = foo() ;
}

```

In traditional C, the function `foo()` would be of type `float` when it is invoked in the function `func1()`, because the declaration for `foo()` had file scope, even though it occurred within a function. ANSI C dictates that the declaration for `foo()` has block scope. Thus, there is no declaration for `foo()` in scope in `func1()`, and it is implicitly typed `int`. This difference in typing between the explicitly and implicitly declared versions of `foo()` results in a redeclaration error at compile time, because they both are linked to the same

external definition for `foo()` and the difference in typing could otherwise produce unexpected behavior.

2.2.2 Name Space Changes

ANSI C recognizes four distinct name spaces: one for tags, one for labels, one for members of a particular struct or union, and one for everything else. This division creates two discrepancies with traditional C:

- In ANSI C, each struct or union has its own name space for its members. This is a pointed departure from traditional C, in which these members were nothing more than offsets, allowing you to use a member with a structure to which it does not belong. This usage is illegal in ANSI C.
- Enumeration constants were special identifiers in versions of Silicon Graphics C prior to IRIX Release 3.3. In ANSI C, these constants are simply integer constants that can be used wherever they are appropriate. Similarly, in ANSI C, other integer variables can be assigned to a variable of an enumeration type with no error.

2.2.3 Changes in the Linkage of Identifiers

An identifier's linkage determines which of the references to that identifier refer to the same object. This terminology formalizes the familiar concept of variables declared `extern` and variables declared `static` and is a necessary augmentation to the concept of scope.

```
extern int mytime;  
static int yourtime;
```

In the previous example, both `mytime` and `yourtime` have file scope. However, `mytime` has external linkage, while `yourtime` has internal linkage. An object can also have no linkage, as is the case of automatic variables.

The preceding example illustrates another implicit difference between the declarations of `mytime` and `yourtime`. The declaration of `yourtime` allocates storage for the object, whereas the declaration of `mytime` merely references it.

If `mytime` is initialized as follows, storage is allocated:

```
int mytime = 0;
```

In ANSI C terminology, a declaration that allocates storage is referred to as a definition. This is different from traditional C.

In traditional C, neither of the following declarations was a definition:

```
extern int bert;  
int bert;
```

In effect, the second declaration included an implicit `extern` specification. This is not true in ANSI C.

Note: Objects with external linkage that are not specified as `extern` at the end of the compilation unit are considered definitions, and, in effect, initialized to zero. (If multiple declarations of the object are in the compilation unit, only one needs the `extern` specification.)

The effect of this change is to produce “multiple definition” messages from the linker when two modules contain definitions of the same identifier, even though neither is explicitly initialized. This is often referred to as the strict ref/def model. A more relaxed model can be achieved by using the `-common` compiler flag.

The ANSI C linker issues a warning when it finds redundant definitions, indicating the modules that produced the conflict. However, the linker cannot determine whether the definition of the object is explicit. If a definition is given with an explicit initialization, and that definition is not the linker’s choice, the result may be incorrectly initialized objects. This is illustrated in the following example:

```
module1.c:  
    int ernie;  
module2.c:  
    int ernie = 5;
```

ANSI C implicitly initializes `ernie` in `module1.c` to zero. To the linker, `ernie` is initialized in two different modules. The linker warns you of this situation, and chooses the first such module it encounters as the true definition of `ernie`. This module may or may not contain the explicitly initialized copy.

2.3 Types and Type Compatibility

Historically, C has allowed free mixing of arithmetic types in expressions and as arguments to functions. (Arithmetic types include integral and floating point types. Pointer types are not included.) C’s type promotion rules reduced the number of actual types used in arithmetic expressions and as arguments to the following three: `int`, `unsigned`, and `double`. This scheme allowed free

mixing of types, but in some cases forced unnecessary conversions and complexity in the generated code.

One ubiquitous example of unnecessary conversions is when float variables were used as arguments to a function. C's type promotion rules often caused two unwanted, expensive conversions across a function boundary.

ANSI C has altered these rules somewhat to avoid the unnecessary overhead in many C implementations. This alteration, however, may produce differences in arithmetic and pointer expressions and in argument passing. For a complete discussion of operator conversions and type promotions, see Chapter 5, page 43.

2.3.1 Type Promotion in Arithmetic Expressions

Two differences are noteworthy between ANSI and traditional C. First, ANSI C relaxes the restriction that all floating point calculations must be performed in double precision. In the following example, pre-ANSI C compilers are required to convert each operand to double, perform the operation in double precision, and truncate the result to float:

```
extern float f, f0, f1;
addf() {
    f = f0 + f1;
}
```

These steps are not required in ANSI C. In ANSI C, the operation can be done entirely in single-precision. (In traditional C, these operations were performed in single-precision if the [-float] compiler option was selected.)

The second difference in arithmetic expression evaluation involves integral promotions. ANSI C dictates that any integral promotions be "value-preserving." Traditional C used "unsignedness-preserving" promotions. Consider the following example:

```
unsigned short us = 1, them = 2;
int i;
test() {
    i = us - them;
}
```

ANSI C's value-preserving rules cause each of `us` and `them` to be promoted to `int`, which is the expression type. The unsignedness-preserving rules, in traditional C, cause `us` and `them` to be promoted to unsigned. The latter case yields a large unsigned number, whereas ANSI C yields -1. The discrepancy in

this case is inconsequential, because the same bit pattern is stored in the integer `i` in both cases, and it is later interpreted as `-1`.

However, if the case is altered slightly, as in the following example, the result assigned to `f` is quite different under the two schemes:

```
unsigned short us = 1, them = 2;
float f;
test() {
    f = us - them;
}
```

If you use the `-wlint` option, the compiler will warn about the implicit conversions from `int` or `unsigned` to `float`.

For more information on arithmetic conversions, see Section 5.6, page 45.

2.3.2 Type Promotion and Floating Point Constants

The differences in behavior of ANSI C floating point constants and traditional C floating point constants can cause numerical and performance differences in code ported from the traditional C to the ANSI C compiler.

For example, consider the result type of the following computation:

```
#define PI 3.1415926
float a, b;

b = a * PI;
```

The result type of `b` depends on which compilation options you use. Table 1, page 15, lists the effects of various options.

Table 1. Effect of Compilation Options on Floating Point Conversions

Compilation Option	PI Constant Type	Promotion Behavior
<code>-cckr</code>	double	<code>(float)((double)a * PI)</code>
<code>-cckr -float</code>	float	<code>a * PI</code>
<code>-xansi</code>	double	<code>(float)((double)a * PI)</code>
<code>-ansi</code>	double	<code>(float)((double)a * PI)</code>

Each conversion incurs computational overhead.

The `-float` flag has no effect if you also specify `-ansi` or `-xansi`. To prevent the promotion of floating constants to double (and promoting the computation to a double precision multiply) you must specify the constant as a single precision floating point constant. In the previous example, you would use the following statement:

```
#define PI 3.1415926f    /* single precision float */
```

Traditional C (compiled with the `-cckr` option) does not recognize the `float` qualifier, `f`, however. Instead, write the constant definition as follows:

```
#ifdef __STDC__
#define PI 3.1415926f
#else
#define PI 3.1415926
#endif
```

If you compile with the `-ansi`, `-ansiposix` or `-xansi` options, `__STDC__` is automatically defined, as though you used `-D__STDC__= 1` on your compilation line. Therefore, with the last form of constant definition noted above, the calculation in the example is promoted as described in Table 2, page 16.

Table 2. Using `__STDC__` to Affect Floating Point Conversions

Compilation Option	PI Constant Type	Promotion Behavior
<code>-cckr</code>	double	<code>(float)((double)a * PI)</code>
<code>-cckr -float</code>	float	<code>a * PI</code>
<code>-xansi</code>	float	<code>a * PI</code>
<code>-ansi</code>	float	<code>a * PI</code>

2.3.3 Compatible Types

To determine whether or not an implicit conversion is permissible, ANSI C introduced the concept of compatible types. After promotion, using the appropriate set of promotion rules, two non-pointer types are compatible if they have the same size, signedness, and integer or float characteristic, or, in the case of aggregates, are of the same structure or union type. Except as discussed in

the previous section, no surprises should result from these changes. You should not encounter unexpected problems unless you are using pointers.

Pointers are compatible if they point to compatible types. No default promotion rules apply to pointers. Under traditional C, the following code fragment compiled silently:

```
int *iptr;
unsigned int *uiptr;
foo() {
    iptr = uiptr;
}
```

Under ANSI C, the pointers `iptr` and `uiptr` do not point to compatible types (because they differ in unsignedness), which means that the assignment is illegal. Insert the appropriate cast to alleviate the problem. When the underlying pointer type is irrelevant or variable, use the wildcard type `void *`.

2.3.4 Argument Type Promotions

ANSI C rules for the promotion of arithmetic types when passing arguments to a function depend on whether or not a prototype is in scope for the function at the point of the call. If a prototype is not in scope, the arguments are converted using the default argument promotion rules: `short` and `char` types (whether signed or unsigned) are passed as `ints`, other integral quantities are not changed, and floating point quantities are passed as `doubles`. These rules are also used for arguments in the variable-argument portion of a function whose prototype ends in ellipses (...).

If a prototype is in scope, an attempt is made to convert each argument to the type indicated in the prototype prior to the call. The types of conversions that succeed are similar to those that succeed in expressions. Thus, an `int` is promoted to a `float` if the prototype so indicates, but a pointer to `unsigned` is not converted to a pointer to `int`. ANSI C also allows the implementation greater freedom when passing integral arguments if a prototype is in scope. If it makes sense for an implementation to pass short arguments as 16-bit quantities, it can do so.

Use of prototypes when calling functions allows greater ease in coding. However, due to the differences in argument promotion rules, serious discrepancies can occur if a function is called both with and without a prototype in scope. Make sure that you use prototypes consistently and that any prototype is declared to be in scope for all uses of the function identifier.

2.3.5 Mixed Use of Functions

To reduce the chances of problems occurring when calling a function with and without a prototype in scope, limit the types of arithmetic arguments in function declarations. In particular, avoid using `short` or `char` types for arguments; their use rarely improves performance and may raise portability issues if you move your code to a machine with a smaller word size. This is because function calls made with and without a prototype in scope may promote the arguments differently. In addition, be circumspect when typing a function argument `float`, because you can encounter difficulties if the function is called without a prototype in scope. With these issues in mind, you can quickly solve the few problems that may arise.

2.4 Function Prototypes

Function prototypes are not new to Silicon Graphics C. In traditional C, however, the implementation of prototypes was incomplete. In one case, a significant difference still exists between the ANSI C and the traditional C implementations of prototypes.

You can prototype functions in two ways. The most common method is simply to create a copy of the function declaration with the arguments typed, with or without identifiers for each, such as either of the following:

```
int func(int, float, unsigned [2]);
int func(int i, float f, unsigned u[2]);
```

You can also prototype a function by writing the function definition in prototype form:

```
int func(int i, float f, unsigned u[2])
{
    < code for func >
}
```

In each case, a prototype is created for `func()` that remains in scope for the rest of the compilation unit.

One area of confusion about function prototypes is that you must write functions that have prototypes in prototype form. Unless you do this, the default argument promotion rules apply.

ANSI C elicits an error diagnostic for two incompatible types for the same parameter in two declarations of the same function. Traditional C elicits an error diagnostic when the incompatibility may lead to a difference between the

bit-pattern of the value passed in by the caller and the bit-pattern seen in the parameter by the callee.

In the following example, the function `func()` is declared twice with incompatible parameter profiles:

```
int func (float);
int func (f)
float f;
{ ... }
```

The parameter `f` in `func()` is assumed to be type `double`, because the default argument promotions apply. Error diagnostics in traditional C and ANSI C are elicited about the two incompatible declarations for `func()`.

The following two situations produce diagnostics from the ANSI C compiler when you use function prototypes:

- A prototyped function is called with one or more arguments of incompatible type. (Incompatible types are discussed in Section 2.3, page 13.)
- Two incompatible (explicit or implicit) declarations for the same function are encountered. This version of the compiler scrutinizes duplicate declarations carefully and catches inconsistencies.

Note: When you use `-cckr` you do not get warnings about prototyped functions, unless you specify `-prototypes`.

2.5 External Name Changes

Many well-known UNIX external names that are not covered by the ANSI C standard are in the user's name space. These names fall into three categories:

1. Names of functions in the C library
2. Names defined by the linker
3. Names of data areas with external linkage

2.5.1 Changes in Function Names

Names of functions that are in the user's name space and are referenced by ANSI C functions in the C library are aliased to counterpart functions whose names are reserved. In all cases, the new name is formed simply by prefixing an underbar to the old name. Thus, although it was necessary to change the

name of the familiar UNIX C library function `write()` to `_write()`, the function `write()` remains in the library as an alias.

The behavior of a program may change if you have written your own versions of C library functions. If, for example, you have your own version of `write()`, the C library continues to use its version of `_write()`.

2.5.2 Changes in Linker-Defined Names

The linker is responsible for defining the standard UNIX symbols `end`, `etext`, and `edata`, if these symbols are unresolved in the final phases of linking. (See the `end(3C)` reference page for more information.) The ANSI C linker has been modified to satisfy references for `_etext`, `_edata`, and `_end` as well. The ANSI C library reference to `end` has been altered to `_end`.

This mechanism preserves the ANSI C name space, while providing for the definition of the non-ANSI C forms of these names if they are referenced from existing code.

2.5.3 Data Area Name Changes

The names of several well-known data objects used in the ANSI C portion of the C library were in the user's name space. These objects are listed in Table 3, page 21. These names were moved into the reserved name space by prefixing their old names with an underscore. Whether these names are defined in your environment depends on the compilation mode you are using (the default is `-xansi`).

Table 3, page 21, shows the effect of compilation mode on names and indicates whether or not these well-known external names are visible when you compile code in the various modes. The left column has three sets of names. Determine which versions of these names are visible by examining the corresponding column under your compilation mode.

Table 3. Effect of Compilation Mode on Names

Name	-cckr	-xansi	-ansi
environ	environ and _environ aliased	environ and _environ aliased	only _environ visible
timezone, tzname, altzone, daylight	unchanged	#define to ANSI C name if using <time.h>	_timezone, _tzname, _altzone, _daylight
sys_nerr, sys_errlist	unchanged	identical copies with names _sys_nerr, _sys_errlist	identical copies with names _sys_nerr, _sys_errlist

Definitions of some of the terms used in Table 3, page 21, are as follows:

- “aliased” means the two names access the same object.
- “unchanged” means the well-known version of the name is unaltered.
- “identical copies” means that two copies of the object exist—one with the well-known name and one with the ANSI C name (prefixed with an underbar). Applications should not alter these objects.
- “#define” means that a macro is provided in the indicated header to translate the well-known name to the ANSI C counterpart. Only the ANSI C name exists. You should include the indicated header if your code refers to the well-known name. For example, the name `tzname` is:
 - Unchanged when compiling `-cckr`
 - Converted to the reserved ANSI C name (`_tzname`) by a macro if you include `<time.h>` when compiling `-xansi`
 - Available only as the ANSI C version (`_tzname`) if compiling `-ansi` (the default is `-xansi`)

2.6 Standard Headers

Functions in the ANSI C library are declared in a set of standard headers. This set is self-consistent and is free of name space pollution, when compiling in the pure ANSI mode. Names that are normally elements of the user’s name space but are specifically reserved by ANSI are described in the corresponding standard header. Refer to these headers for information on both reserved names

and ANSI library function prototypes. The following list contains the set of standard headers:

```
<assert.h>
<ctype.h>
<errno.h>
<float.h>
<limits.h>
<locale.h>
<math.h>
<setjmp.h>
<signal.h>
<stdio.h>
<stddef.h>
<stdarg.h>
<string.h>
<stdlib.h>
<sys/errno.h>
<sys/signal.h>
<time.h>
```

Lexical Conventions [3]

This chapter covers the C lexical conventions including comments and tokens. A token is a series of contiguous characters that the compiler treats as a unit.

Blanks, tabs, newlines, and comments are collectively known as “white space.” White space is ignored except as it serves to separate tokens. Some white space is required to separate otherwise adjacent identifiers, keywords, and constants.

If the input stream has been parsed into tokens up to a given character, the next token is taken to include the longest string of characters that could possibly constitute a token.

3.1 Comments

The `/*` characters introduce a comment; the `*/` characters terminate a comment. They do not indicate a comment when occurring within a string literal. Comments do not nest. Once the `/*` introducing a comment is seen, all other characters are ignored until the ending `*/` is encountered.

3.2 Identifiers

An identifier, or name, is a sequence of letters, digits, and underscores (`_`). The first character cannot be a digit. Uppercase and lowercase letters are distinct. Name length is unlimited. The terms identifier and name are used interchangeably.

3.3 Keywords

The identifiers listed in Table 4, page 23, are reserved for use as keywords and cannot be used for any other purpose.

Table 4. Reserved Keywords

Keywords					
auto	default	float	register	struct	volatile
break	do	for	return	switch	while

Keywords				
case	double	goto	short	typedef
char	else	if	signed	union
const	enum	int	sizeof	unsigned
continue	extern	long	static	void

Traditional C reserves and ignores the `fortran` keyword.

3.4 Constants

The four types of constants are integer, character, floating, and enumeration. Each constant has a type, determined by its form and value.

In this section's discussions of the various types of constants, a unary operator preceding the constant is not considered part of it. Rather, such a construct is a constant-expression (see Section 6.19, page 66). Thus, the integer constant `0xff` becomes an integral constant expression by prefixing a minus sign, for instance, `-0xff`. The effect of the `-` operator is not considered in the discussion of integer constants.

As an example, the integer constant `0xffffffff` has type `int` in traditional C, with value `-1`. It has type `unsigned` in ANSI C, with value `232 - 1`. This discrepancy is inconsequential if the constant is assigned to a variable of integral type (for example, `int` or `unsigned`), as a conversion occurs. If it is assigned to a double, however, the value differs as indicated between traditional and ANSI C.

3.4.1 Integer Constants

An integer constant consisting of a sequence of digits is considered octal if it begins with `0` (zero). An octal constant consists of the digits `0` through `7` only. A sequence of digits preceded by `0x` or `0X` is considered a hexadecimal integer. The hexadecimal digits include `[aA]` through `[fF]`, which have values of `10` through `15`.

The suffixes `[lL]` traditionally indicate integer constants of type `long`. These suffixes are allowed, but are superfluous, because `int` and `long` are the same size in `-o32` and `-n32` modes. The `ll`, `LL`, `lL`, and `Ll` suffixes indicate a long long constant (a 64-bit integral type). Note that `long long` is not a

strict ANSI C type, and a warning is given for `long long` constants in `-ansi` and `-ansiposix` modes. The following are examples of `long long`:

```
12345LL
12345ll
```

In ANSI C, an integer constant can be suffixed with `uU`, in which case its type is `unsigned`. (One or both of `uU` and `lL` can appear.) An integer constant also has type `unsigned` if its value cannot be represented as an `int`. Otherwise, the type of an integer constant is `int`. The following are examples of `unsigned long long`:

```
123456ULL
123456ull
```

3.4.2 Character Constants

A character constant is a character enclosed in single quotation marks, such as `'x'`. The value of a character constant is the numerical value of the character in the machine's character set. An explicit new-line character is illegal in a character constant. The type of a character constant is `int`.

In ANSI C, a character constant can be prefixed by `L`, in which case it is a wide character constant. For example, a wide character constant for `'z'` is written `L'z'`. The type of a wide character constant is `wchar_t`, which is defined in the `stddef.h` file.

3.4.3 Special Characters

Some special and nongraphic characters are represented by the escape sequences shown in Table 5, page 25.

Table 5. Escape Sequences for Nongraphic Characters

Character Name	Escape Sequence
newline	<code>\n</code>
horizontal tab	<code>\t</code>
vertical tab	<code>\v</code>
backspace	<code>\b</code>

Character Name	Escape Sequence
carriage return	<code>\r</code>
form feed	<code>\f</code>
backslash	<code>\\</code>
single quote	<code>\'</code>
double quote	<code>\"</code>
question mark	<code>\?</code>
bell (ANSI C only)	<code>\a</code>

The `\ddd` escape sequence consists of the backslash followed by 1, 2, or 3 octal digits that specify the value of the desired character. A special case of this construction is `\0` (not followed by a digit), which indicates the ASCII character NUL.

In ANSI C, `\x` indicates the beginning of a hexadecimal escape sequence. The sequence is assumed to continue until a character is encountered that is not a member of the hexadecimal character set 0,1, ... 9, [aA], [bB], ... [fF]. The resulting unsigned number cannot be larger than a character can accommodate (decimal 255).

If the character following a backslash is not one of those specified in this section, the behavior is undefined.

3.4.3.1 Trigraph Sequences (ANSI C Only)

The character sets of some older machines lack certain members that have come into common usage. To allow the machines to specify these characters, ANSI C defined an alternate method for their specification, using sequences of characters that are commonly available. These sequences are termed trigraph sequences. Nine sequences are defined; each consists of three characters beginning with two question marks. Each instance of one of these sequences is translated to the corresponding single character. Other sequences of characters, perhaps including multiple question marks, are unchanged. Each trigraph sequence with the single character it represents is listed in Table 6, page 27.

Table 6. Trigraph Sequences

Trigraph Sequence	Single Character
??=	#
??([
??/	\
??)]
??'	^
??<	{
??!	
??>	}
??-	~

3.4.4 Floating Constants

A floating constant consists of an integer part, a decimal point, a fraction part, an [eE], and an optionally signed integer exponent. The integer and fraction parts both consist of a sequence of digits. Either the integer part or the fraction part (but not both) can be missing. Either the decimal point or the [eE] and the exponent (not both) can be missing.

In traditional C, every floating constant has type `double`.

In ANSI C, floating constants can be suffixed by either [fF] or [lL]. Floating constants suffixed with [fF] have type `float`. Those suffixed with [lL] have type `long double`, which has greater precision than `double` in `-n32` and `-64` modes and a precision equal to `double` in `-o32` mode.

3.4.5 Enumeration Constants

Names declared as enumerators have type `int`. For a discussion of enumerators, see Section 7.5, page 76. For information on the use of enumerators in expressions, see Section 4.2.2, page 38.

3.5 String Literals

A string literal is a sequence of characters surrounded by double quotation marks, as in `''...''`. A string literal has type `array of char` and is initialized with the given characters. The compiler places a null byte (`\0`) at the end of each string literal so that programs that scan the string literal can find its end. A double-quotation character (`''`) in a string literal must be preceded by a backslash (`\`). In addition, the same escapes as those described for character constants can be used. (See Section 3.4.2, page 25, for a list of escapes.) A backslash (`\`) and the immediately following newline are ignored. Adjacent string literals are concatenated.

In traditional C, all string literals, even when written identically, are distinct.

In ANSI C, identical string literals are not necessarily distinct. Prefixing a string literal with `L` specifies a wide string literal. Adjacent wide string literals are concatenated.

As an example, consider the sentence "He said, 'Hi there.'" This sentence could be written with three adjacent string literals:

```
"He said, " "\'Hi " "there.\' "
```

3.6 Operators

An operator specifies an operation to be performed. The operators `[]`, `()`, and `? :` must occur in pairs, possibly separated by expressions. The operators `#` and `##` can occur only in preprocessing directives.

operator can be one of the following:

```
[ ]( ). ->  
++ - - & * + - ~ ! sizeof  
/ % << >> < > <= >= == != ^ | && ||  
? :  
= *= /= %= += -= <<= >>= &= ^= |=  
, # ##
```

Individual operations are discussed in Chapter 6, page 49.

3.7 Punctuators

A punctuator is a symbol that has semantic significance but does not specify an operation to be performed. The punctuators [], (), and { } must occur in pairs, possibly separated by expressions, declarations, or statements. The punctuator # can occur only in preprocessing directives.

punctuator; one of the [] () { } * , : = ; ... #
following:

Some operators, determined by context, are also punctuators. For example, the array index indicator [] is a punctuator in a declaration (see Chapter 7, page 69), but an operator in an expression (see Chapter 6, page 49).

Meaning of Identifiers [4]

Traditional C formally based the interpretation of an identifier on two of its attributes: storage class and type. The storage class determined the location and lifetime of the storage associated with an identifier; the type determined the meaning of the values found in the identifier's storage. Informally, name space, scope, and linkage were also considered.

ANSI C formalizes the practices of traditional C. An ANSI C identifier is disambiguated by four characteristics: its scope, name space, linkage, and storage duration. The ANSI C definitions of these terms differ somewhat from their interpretations in traditional C.

Storage-class specifiers and their meanings are described in Chapter 7, page 69. Storage-class specifiers are discussed in this chapter only in terms of their effect on an object's storage duration and linkage.

You can find a discussion of focusing on changes to the language in Section 2.2, page 10, and Section 2.3, page 13.

4.1 Disambiguating Names

This section discusses the ways C disambiguates names: scope, name space, linkage, and storage class.

4.1.1 Scope

The region of a program in which a given instance of an identifier is visible is called its scope. The scope of an identifier usually begins when its declaration is seen, or, in the case of labels and functions, when it is implied by use. Although it is impossible to have two declarations of the same identifier active in the same scope, no conflict occurs if the instances are in different scopes. Of the four kinds of scope, two—file and block—are traditional C scopes. Two other kinds of scope—function and function prototype—are implied in traditional C and formalized in ANSI C.

4.1.1.1 Block Scope

Block scope is the scope of automatic variables (variables declared within a function). Each block has its own scope. No conflict occurs if the same identifier is declared in two blocks. If one block encloses the other, the

declaration in the enclosed block hides that in the enclosing block until the end of the enclosed block is reached. The definition of a block is the same in ANSI C and traditional C, with one exception, illustrated by the example below:

```
int f(x);
int x;
{
    int x;
    x = 1;
}
```

In ANSI C, the function arguments are in the function body block. Thus, ANSI C will issue an error of a “redeclaration of *x*.”

In traditional C, the function arguments are in a separate block that encloses the function body block. Thus, traditional C would quietly hide the argument *x* with the local variable *x*, because they are in distinct blocks.

ANSI C and traditional C differ in the assignment of block and file scope in a few instances. See Section 4.1.1.4, page 32, for more details.

4.1.1.2 Function Scope

Only labels have function scope. Function scope continues until the end of the current function.

4.1.1.3 Function Prototype Scope

If an identifier appears within the list of parameter declarations in a function prototype that is not part of a function definition (see Section 7.7.4, page 82), it has function prototype scope, which terminates at the end of the prototype. This termination allows any dummy parameter names appearing in a function prototype to disappear at the end of the prototype.

4.1.1.4 File Scope

Identifiers appearing outside of any block, function, or function prototype, have file scope. This scope continues to the end of the compilation unit. Unlike other scopes, multiple declarations of the same identifier with file scope can exist in a compilation unit, so long as the declarations are compatible.

Whereas ANSI C assigns block scope to all declarations occurring inside a function, traditional C assigns file scope to such declarations if they have the storage class `extern`. This storage class is implied in all function declarations,

whether the declaration is explicit (as in `int foo();`) or implicit (if there is no active declaration for `foo()` when an invocation is encountered, as in `f = foo();`). For a further discussion of this discrepancy, with examples, see Section 2.2.1, page 10.

4.1.2 Name Spaces

In certain cases, the purpose for which an identifier is used may disambiguate it from other uses of the same identifier appearing in the same scope. This is true, for example, for tags and is used in traditional C to avoid conflicts between identifiers used as tags and those used in object or function declarations. ANSI C formalizes this mechanism by defining certain name spaces. These name spaces are completely independent. A member of one name space cannot conflict with a member of another.

ANSI C recognizes the following four distinct name spaces:

Tags	<code>struct</code> , <code>union</code> , and <code>enum</code> tags have a single name space.
Labels	Labels are in their own name space.
Members	Each <code>struct</code> or <code>union</code> has its own name space for its members.
Ordinary identifiers	Ordinary identifiers, including function and object names as well as user-defined type names, are placed in the last name space.

4.1.2.1 Name Space Discrepancies Between Traditional and ANSI C

The definition of name spaces causes discrepancies between traditional and ANSI C in a few situations:

- Structure members in traditional C were nothing more than offsets, allowing the use of a member with a structure to which it does not belong. This is illegal under ANSI C.
- Enumeration constants were special identifiers in traditional C prior to IRIX Release 3.3. In later releases of traditional C, as in ANSI C, these constants are simply integer constants that can be used anywhere they are appropriate.
- Labels reside in the same name space as ordinary identifiers in traditional C. Thus, the following example is legal in ANSI C but not in traditional C:

```
func() {  
  int lab;  
    if (lab) goto lab;  
    func1() ;  
lab:  
  return;  
}
```

4.1.3 Linkage of Identifiers

Two instances of the same identifier appearing in different scopes may, in fact, refer to the same entity. For example, the references to a variable, `counter`, is declared with file scope in the following example:

```
extern int counter;
```

In this example, two separate files refer to the same `int` object. The association between the references to an identifier occurring in distinct scopes and the underlying objects are determined by the identifier's linkage.

The three kinds of linkage are as follows:

Internal linkage	Within a file, all declarations of the same identifier with internal linkage denote the same object.
External linkage	Within an entire program, all declarations of an identifier with external linkage denote the same object.
No linkage	A unique entity, accessible only in its own scope, has no linkage.

An identifier's linkage is determined by whether it appears inside or outside a function, whether it appears in a declaration of a function (as opposed to an object), its storage-class specifier, and the linkage of any previous declarations of the same identifier that have file scope. An identifier's linkage is determined as follows:

1. If an identifier is declared with file scope and the storage-class specifier `static`, it has internal linkage.
2. If the identifier is declared with the storage-class specifier `extern`, or is an explicit or implicit function declaration with block scope, the identifier has the same linkage as any previous declaration of the same identifier with file scope. If no previous declaration exists, the identifier has external linkage.

3. If an identifier for an object is declared with file scope and no storage-class specifier, it has external linkage. (See Section 2.2.3, page 12.)
4. All other identifiers have no linkage. This includes all identifiers that do not denote an object or function, all objects with block scope declared without the storage-class specifier `extern`, and all identifiers that are not members of the ordinary variables name space.

Two declarations of the same identifier in a single file that have the same linkage, either internal or external, refer to the same object. The same identifier cannot appear in a file with both internal and external linkage.

This code gives an example where the linkage of each declaration is the same in both traditional and ANSI C:

```
static int pete;
extern int bert;
int mom;
int func0() {
    extern int mom;
    extern int pete;
    static int dad;
    int bert;
    ...
}
int func1() {
    static int mom;
    extern int dad;
    extern int bert;
    ...
}
```

The declaration of `pete` with file scope has internal linkage by rule 1 above. This means that the declaration of `pete` in `func0()` also has internal linkage by rule 2 and refers to the same object.

By rule 2, the declaration of `bert` with file scope has external linkage, because there is no previous declaration of `bert` with file scope. Thus, the declaration of `bert` in `func1()` also has external linkage (again by rule 2) and refers to the same (external) object. By rule 4, however, the declaration of `bert` in `func0()` has no linkage, and refers to a unique object.

The declaration of `mom` with file scope has external linkage by rule 3, and, by rule 2, so does the declaration of `mom` in `func0()`. (Again, two declarations of the same identifier in a single file that both have either internal or external

linkage refer to the same object.) The declaration of `mom` in `func1()`, however, has no linkage by rule 4 and thus refers to a unique object.

Last, the declarations of `dad` in `func0()` and `func1()` refer to different objects, as the former has no linkage and the latter, by rule 2, has external linkage.

4.1.4 Linkage Discrepancies Between Traditional and ANSI C

Traditional and ANSI C differ on the concept of linkage in the following important ways:

- In traditional C, a function can be declared with block scope and the storage-class specifier `static`. The declaration is given internal linkage. Only the storage class `extern` can be specified in function declarations with block scope in ANSI C.
- In traditional C, if an object is declared with block scope and the storage-class specifier `static`, and a declaration for the object with file scope and internal linkage exists, the block scope declaration has internal linkage. In ANSI C, an object declared with block scope and the storage-class specifier `static` has no linkage.

Traditional and ANSI C handle the concepts of reference and definition differently. For example:

```
extern int mytime;
static int yourtime;
```

In the preceding example, both `mytime` and `yourtime` have file scope. As discussed previously, `mytime` has external linkage, while `yourtime` has internal linkage.

However, there is an implicit difference, which exists in both ANSI and traditional C, between the declarations of `mytime` and `yourtime` in the preceding example. The declaration of `yourtime` allocates storage for the object, whereas the declaration of `mytime` merely references it. If `mytime` had been initialized, as in the following example, it would also have allocated storage:

```
int mytime=0;
```

A declaration that allocates storage is referred to as a definition.

In traditional C, neither of the two declarations below is a definition:

```
extern int bert;  
int bert;
```

In effect, the second declaration includes an implicit `extern` specification. ANSI C does not include such an implicit specification.

Note: In ANSI C, objects with external linkage that are not specified as `extern` at the end of the compilation unit are considered definitions, and, in effect, initialized to zero. (If multiple declarations of the object occur in the compilation unit, only one need have the `extern` specification.)

If two modules contain definitions of the same identifier, the linker complains of “multiple definitions,” even though neither is explicitly initialized.

The ANSI C linker issues a warning when it finds redundant definitions, indicating which modules produced the conflict. However, the linker cannot determine if the initialization of the object is explicit. This may result in incorrectly initialized objects if another module fails to tag the object with `extern`.

Thus, consider the following example:

```
module1.c:  
    int ernie;  
module2.c:  
    int ernie = 5;
```

ANSI C implicitly initializes `ernie` in `module1.c` to zero. To the linker, `ernie` is initialized in two different modules. The linker warns you of this situation, and chooses the first such module it encountered as the true definition of `ernie`. This module may or may not be the one containing the explicitly initialized copy.

4.1.5 Storage Duration

Storage duration denotes the lifetime of an object. Storage duration is of two types: static and automatic.

Objects declared with external or internal linkage, or with the storage-class specifier `static`, have static storage duration. If these objects are initialized, the initialization occurs once, prior to any reference.

Other objects have automatic storage duration. Storage is newly allocated for these objects each time the block that contains their declaration is entered, unless the object has a variable length array type. If the object is variably

modified, and the block is entered by a jump to a labeled statement, then the behavior is undefined.

If an object with automatic storage duration is initialized, the initialization occurs each time the block is entered at the top. This is not guaranteed to occur if the block is entered by a jump to a labeled statement.

4.2 Object Types

The C language supports three fundamental types of objects: character, integer, and floating point.

4.2.1 Character Types

Objects declared as characters (`char`) are large enough to store any member of the implementation's character set. If a genuine character from that character set is stored in a `char` variable, its value is equivalent to the integer code for that character. Other quantities may be stored into character variables, but the implementation is machine dependent. In this implementation, `char` is unsigned by default.

The ANSI C standard has added multibyte and wide character types. In the initial Silicon Graphics release of ANSI C, wide characters are of type `unsigned char`, and multibyte characters are of length one. (See the header files `stddef.h` and `limits.h` for more information.)

4.2.2 Integer and Floating Point Types

Up to five sizes of integral types (signed and unsigned) are available: `char`, `short`, `int`, `long`, and `long long`. Up to three sizes of floating point types are available. The sizes are shown in Table 7, page 39. (The values in the table apply to both ANSI and traditional C, with the exceptions noted in the subsequent discussion.)

Table 7. Storage Class Sizes

Type	Size in Bits (-32)	Size in Bits (-64)	Size in Bits (-64)
char	8	8	8
short	16	16	16
int	32	32	32
long	32	32	64
long long	64	64	64
float	32	32	32
double	64	64	64
long double	64	128	128

Although Silicon Graphics supports `long double` as a type in `-cckr` mode, this is viewed as an extension to traditional C and is ignored in subsequent discussions pertinent only to traditional C.

Differences exist between `-o32` mode, `-n32` mode, and `-64` mode compilations. Types `long` and `int` have different sizes (and ranges) in 64-bit mode; type `long` always has the same size as a pointer value. A pointer (or address) has a 64-bit representation in 64-bit mode and a 32-bit representation in both 32-bit modes. Therefore, an `int` object has a smaller size than a pointer object in 64-bit mode.

The `long long` type is not a valid ANSI C type, so a warning is elicited for every occurrence of `long long` in the source program text in `-ansi` and `-ansiposix` modes.

The `long double` type has equal range in old 32-bit, new 32-bit, and 64-bit mode, but it has increased precision in `-n32` and `-64` modes.

Characteristics of integer and floating point types are defined in the standard header files `<limits.h>` and `<float.h>`. The range of a signed integral type of size n is $[(-2^{n-1})... (2^{n-1} - 1)]$. The range of an unsigned version of the type is $[0... (2^n - 1)]$.

Enumeration constants were special identifiers under various versions of traditional C, before IRIX Release 3.3. In ANSI C, these constants are simply integer constants that may be used anywhere. Similarly, ANSI C allows the assignment of other integer variables to variables of enumeration type, with no error.

You can find additional information on integers, floating points, and structures in the following tables:

- For integer types and ranges, see Table 13, page 159
- For floating point types and ranges, see Table 14, page 161
- For structure alignment, see Table 15, page 163

4.2.3 Derived Types

Because objects of the types mentioned in Section 4.2.2, page 38, can be interpreted usefully as numbers, this manual refers to them as arithmetic types. The types `char`, `enum`, and `int` of all sizes (whether unsigned or not) are collectively called integral types. The `float` and `double` types are collectively called floating types. Arithmetic types and pointers are collectively called scalar types.

The fundamental arithmetic types can be used to construct a conceptually infinite class of derived types, such as the following:

- Arrays of objects of most types
- Functions that return objects of a given type
- Pointers to objects of a given type
- Structures that contain a sequence of objects of various types
- Unions capable of containing any one of several objects of various types

In general, these constructed objects can be used as building blocks for other constructed objects.

4.2.4 `void` Type

The `void` type specifies an empty set of values. It is used as the type returned by functions that generate no value. The `void` type never refers to an object and therefore, is not included in any reference to object types.

4.3 Objects and lvalues

An object is a manipulatable region of storage. An lvalue is an expression referring to an object. An obvious example of an lvalue expression is an

identifier. Some operators yield lvalues. For example, if E is an expression of pointer type, then $*E$ is an lvalue expression referring to the object to which E points. The term lvalue comes from the term “left value.” In the assignment expression $E1 = E2$, the left operand $E1$ must be an lvalue expression.

Most lvalues are modifiable, meaning that the lvalue may be used to modify the object to which it refers. Examples of lvalues that are not modifiable include array names, lvalues with incomplete type, and lvalues that refer to an object, part or all of which is qualified with `const` (see Section 7.6, page 77). Whether an lvalue appearing in an expression must be modifiable is usually obvious. For example, in the assignment expression $E1 = E2$, $E1$ must be modifiable. This document makes the distinction between modifiable and unmodifiable lvalues only when it is not obvious.

Operator Conversions [5]

A number of operators can, depending on the types of their operands, cause an implicit conversion of some operands from one type to another. The following discussion explains the results you can expect from these conversions. The conversions demanded by most operators are summarized in Section 5.6, page 45. When necessary, a discussion of the individual operators supplements the summary.

5.1 Conversions of Characters and Integers

You can use a character or a short integer wherever you can use an integer. Characters are unsigned by default. In all cases, the value is converted to an integer. Conversion of a shorter integer to a longer integer preserves the sign. Traditional C uses “unsigned preserving integer promotion” (unsigned short to unsigned int), while ANSI C uses “value preserving integer promotion” (unsigned short to int).

A longer integer is truncated on the left when converted to a shorter integer or to a char. Excess bits are discarded.

5.2 Conversions of Float and Double

Historically in C, expressions containing floating point operands (either `float` or `double`) were calculated using double precision. This is also true of calculations in traditional C, unless you have specified the compiler option `-float`. With the `-float` option, calculations involving floating point operands and no `double` or `long double` operands take place in single precision. The `-float` option has no effect on argument promotion rules at function calls or on function prototypes.

ANSI C performs calculations involving floating point in the same precision as if `-float` had been specified in traditional C, except when floating point constants are involved.

In traditional C, specifying the `-float` option coerces floating point constants into type `float` if all the other subexpressions are of type `float`. This is not the case in ANSI C. ANSI C considers all floating point constants to be implicitly double precision, and operations involving such constants therefore take place in double precision. To force single precision arithmetic in ANSI C, use the `f` or

`F` suffix on floating point constants. To force long double precision on constants, use the `l` or `L` suffix. For example, `3.14l` is long double precision, `3.14` is double precision, and `3.14f` is single precision in ANSI C.

For a complete discussion with examples, see Section 2.3.2, page 15.

5.3 Conversion of Floating and Integral Types

Conversions between floating and integral values are machine-dependent. Silicon Graphics uses IEEE floating point, in which the default rounding mode is to nearest, or in case of a tie, to even. Floating point rounding modes can be controlled using the facilities of `fp.c`. Floating point exception conditions are discussed in the introductory paragraph of Chapter 6, page 49.

When a floating value is converted to an integral value, the rounded value is preserved as long as it does not overflow. When an integral value is converted to a floating value, the value is preserved unless a value of more than six significant digits is being converted to single precision, or fifteen significant digits is being converted to double precision.

5.4 Conversion of Pointers and Integers

An expression of integral type can be added to or subtracted from an object pointer. In such a case, the integer expression is converted as specified in the discussion of the addition operator in Section 6.7, page 59. Two pointers to objects of the same type can be subtracted. In this case, the result is converted to an integer as specified in the discussion of the subtraction operator, in Section 6.7, page 59.

5.5 Conversion of unsigned Integers

When an unsigned integer is converted to a longer unsigned or signed integer, the value of the result is preserved. Thus, the conversion amounts to padding with zeros on the left.

When an unsigned integer is converted to a shorter signed or unsigned integer, the value is truncated on the left. If the result is signed, this truncation may produce a negative value.

5.6 Arithmetic Conversions

Many types of operations in C require two operands to be converted to a common type. Two sets of conversion rules are applied to accomplish this conversion. The first, referred to as the integral promotions, defines how integral types are promoted to one of several integral types that are at least as large as `int`. The second, called the usual arithmetic conversions, derives a common type in which the operation is performed.

ANSI C and traditional C follow different sets of these rules.

5.6.1 Integral Promotions

The difference between the ANSI C and traditional versions of the conversion rules is that the traditional C rules emphasize preservation of the (un)signedness of a quantity, while ANSI C rules emphasize preservation of its value.

In traditional C, operands of types `char`, `unsigned char`, and `unsigned short` are converted to `unsigned int`. Operands of types `signed char` and `short` are converted to `int`.

ANSI C converts all `char` and `short` operands, whether signed or unsigned, to `int`. Only operands of type `unsigned int`, `unsigned long`, and `unsigned long long` may remain unsigned.

5.6.2 Usual Arithmetic Conversions

Besides differing in emphasis on signedness and value preservation, the usual arithmetic conversion rules of ANSI C and traditional C also differ in the precision of the chosen floating point type.

The following subsections describe two sets of conversion rules, one for traditional C, and the other for ANSI C. Each set is ordered in decreasing precedence. In any particular case, the rule that applies is the first whose conditions are met.

Each rule specifies a type, referred to as the result type. Once a rule has been chosen, each operand is converted to the result type, the operation is performed in that type, and the result is of that type.

5.6.2.1 Traditional C Conversion Rules

The traditional C conversion rules are as follows:

- If any operand is of type `double`, the result type is `double`.
- If any operand is of type `float`, the result type is `float` if you have specified the `[-float]` switch. Otherwise, the result type is `double`.
- The integral promotions are performed on each operand as follows:

If one of the operands is of type:	The result is of type:
<code>unsigned long long</code>	<code>unsigned long long</code>
<code>long long</code>	<code>long long</code>
<code>unsigned long</code>	<code>unsigned long</code>
<code>long</code>	<code>long</code>
<code>unsigned int</code>	<code>unsigned int</code>
<code>otherwise</code>	<code>int</code>

5.6.2.2 ANSI C Conversion Rules

The ANSI C rules are as follows:

- If any operand is of type `long double`, the result type is `long double`.
- If any operand is of type `double`, the result type is `double`.
- If any operand is of type `float`, the result type is `float`.
- The integral promotions are performed on each operand as follows:

If one of the operands is of type:	The result is of type:
<code>unsigned long long</code>	<code>unsigned long long</code>
<code>long long</code>	<code>long long</code>
<code>unsigned long</code>	<code>unsigned long</code>
<code>long</code>	<code>long</code>
<code>unsigned int</code>	<code>unsigned int</code>
<code>otherwise</code>	<code>int</code>

5.7 Conversion of Other Operands

The following three sections discuss conversion of lvalues, function designators, void objects, and pointers.

5.7.1 Conversion of lvalues and Function Designators

Except as noted, if an lvalue that has type “array of <type>” appears as an operand, it is converted to an expression of the type “pointer to <type>.” The resultant pointer points to the initial element of the array. In this case, the resultant pointer ceases to be an lvalue. (For a discussion of lvalues, see Section 4.3, page 40.)

A function designator is an expression that has function type. Except as noted, a function designator appearing as an operand is converted to an expression of type “pointer to function.”

5.7.2 Conversion of void Objects

The (nonexistent) value of a void object cannot be used in any way, and neither explicit nor implicit conversion can be applied. Because a void expression denotes a nonexistent value, such an expression can be used only as an expression statement (see Section 8.1, page 93), or as the left operand of a comma expression (see Section 6.18, page 65).

An expression can be converted to type void by use of a cast. For example, this makes explicit the discarding of the value of a function call used as an expression statement.

5.7.3 Conversion of Pointers

A pointer to void can be converted to a pointer to any object type and back without change in the underlying value.

The NULL pointer constant can be specified either as the integral value zero, or the value zero cast to a pointer to void. If a NULL pointer constant is assigned or compared to a pointer to any type, it is appropriately converted.

Expressions and Operators [6]

This chapter discusses the various expressions and operators available in C. The sections describing expressions and operators are presented roughly in order of precedence.

6.1 Precedence and Associativity Rules in C

Operators in C have rules of precedence and associativity that determine how expressions are evaluated. Table 9, page 50, lists the operators and indicates the precedence and associativity of each. Within each row, the operators have the same precedence. Parentheses can be used to override these rules.

Table 8, page 49, shows some simple examples of precedence and associativity.

Table 8. Precedence and Associativity Examples

Expression	Results	Comments
<code>3 + 2 * 5</code>	13	Multiplication is done before addition.
<code>3 + (2 * 5)</code>	13	Parentheses follow the precedence rules, but clarify the expression for the reader.
<code>(3 + 2) * 5</code>	25	Parentheses override the precedence rules.
<code>TRUE TRUE && FALSE</code>	1 (true)	Logical AND has higher priority than logical OR.
<code>TRUE (TRUE && FALSE)</code>	1 (true)	Parentheses follow the precedence rules, but clarify the expression for the reader.
<code>(TRUE TRUE) && FALSE</code>	0 (false)	Parentheses override the precedence rules.

Except as indicated by the syntax or specified explicitly in this chapter, the order of evaluation of expressions, as well as the order in which side-effects take place, is unspecified. The compiler can arbitrarily rearrange expressions involving a commutative and associative operator (*, +, &, |, ^).

Table 9, page 50, lists the precedence and associativity of all operators.

Table 9. Operator Precedence and Associativity

Tokens (From High to Low Priority)	Operators	Class	Associativity
Identifiers, constants, string literal, parenthesized expression	Primary expression	Primary	
() [] -> .	Function calls, subscripting, indirect selection, direct selection	Postfix	L-R
++ --	Increment, decrement (postfix)	Postfix	L-R
++ --	Increment, decrement (prefix)	Prefix	R-L
! ~ + - & sizeof *	Logical and bitwise NOT, unary plus and minus, address, size, indirection	Unary	R-L
(type)	Cast	Unary	R-L
* / %	Multiplicative	Binary	L-R
+ -	Additive	Binary	L-R
<< >>	Left shift, right shift	Binary	L-R
< <= > >=	Relational comparisons	Binary	L-R
== !=	Equality comparisons	Binary	L-R
&	Bitwise and	Binary	L-R
^	Bitwise exclusive or	Binary	L-R
	Bitwise inclusive or	Binary	L-R
&&	Logical and	Binary	L-R
	Logical or	Binary	L-R
? :	conditional	Ternary	R-L
= += -= *= /= %= ^= &= = <<= >>=	Assignment	Binary	R-L
,	Comma	Binary	L-R

6.2 Primary Expressions

The following are all considered “primary expressions:”

Identifiers	An identifier referring to an object is an lvalue. An identifier referring to a function is a function designator. lvalues and function designators are discussed in “Conversion of lvalues and Function Designators” on page 59.
Constants	A constant’s type is determined by its form and value, as described in Section 3.4, page 24.
String literals	A string literal’s type is “array of char,” subject to modification, as described in Section 5.1, page 43.
Parenthesized expressions	A parenthesized expression’s type and value are identical to those of the unparenthesized expression. The presence of parentheses does not affect whether the expression is an lvalue, rvalue, or function designator. For information on expressions, see “Constant Expressions” on page 79.

6.3 Postfix Expressions

Postfix expressions involving `.`, `->`, subscripting, and function calls associate left to right. The syntax for these expressions is as follows:

```

postfix-expression:      primary-expression
                        postfix-expression [expression]
                        postfix-expression (argument-expression-list opt)
                        postfix-expression. identifier
                        postfix-expression -> identifier
                        postfix-expression ++
                        postfix-expression --

```

argument-expression-list: *argument-expression*
 argument-expression-list, argument-expression

6.3.1 Subscripts

A postfix expression followed by an expression in square brackets is a subscript. Usually, the postfix expression has type “pointer to *<type>*”, the expression within the square brackets has type `int`, and the type of the result is *<type>*. However, it is equally valid if the types of the postfix expression and the expression in brackets are reversed. This is because the expression `E1[E2]` is identical (by definition) to `*((E1)+(E2))`. Because addition is commutative, `E1` and `E2` can be interchanged.

You can find more information on this notation in the discussions on identifiers and in the discussion of the `*` and `+` operators (in Section 6.4, page 55, and Section 6.7, page 59), respectively.

6.3.2 Function Calls

The syntax of function call postfix expressions is as follows:

postfix-expression (argument-expression-list_{opt})
argument-expression-list: *argument-expression*
 argument-expression-list, argument-expression

A function call is a postfix expression followed by parentheses containing a (possibly empty) comma-separated list of expressions that are the arguments to the function. The postfix expression must be of type “function returning *<type>*.” The result of the function call is of type *<type>*, and is not an lvalue.

The behavior of function calls is as follows:

- If the function call consists solely of a previously unseen identifier `foo`, the call produces an implicit declaration as if, in the innermost block containing the call, the following declaration had appeared:

```
extern int foo();
```

- If a corresponding function prototype that specifies a type for the argument being evaluated is in force, an attempt is made to convert the argument to that type.

- If the number of arguments does not agree with the number of parameters specified in the prototype, the behavior is undefined.
- If the type returned by the function as specified in the prototype does not agree with the type derived from the expression containing the called function, the behavior is undefined. Such a scenario may occur for an external function declared with conflicting prototypes in different files.
- If no corresponding prototype is in scope or if the argument is in the variable argument section of a prototype that ends in ellipses (...), the argument is converted according to the following default argument promotions:
 - Type `float` is converted to `double`.
 - Array and function names are converted to corresponding pointers.
 - When using traditional C, types `unsigned short` and `unsigned char` are converted to `unsigned int`, and types `signed short` and `signed char` are converted to `signed int`.
 - When using ANSI C, types `short` and `char`, whether `signed` or `unsigned`, are converted to `int`.
- In preparing for the call to a function, a copy is made of each actual argument. Thus, all argument passing in C is strictly by value. A function can change the values of its parameters, but these changes cannot affect the values of the actual arguments. It is possible to pass a pointer on the understanding that the function can change the value of the object to which the pointer points. (Arguments that are array names can be changed as well, because these arguments are converted to pointer expressions.)
- Because the order of evaluation of arguments is unspecified, side effects may be delayed until the next sequence point, which occurs at the point of the actual call and after all arguments have been evaluated. (For example, in the function call `func(foo++)`, the incrementation of `foo` may be delayed.)
- Recursive calls to any function are permitted.

Silicon Graphics recommends consistent use of prototypes for function declarations and definitions. Do not mix prototyped and nonprototyped function declarations and definitions. Even though the language allows it, never call functions before you declare them. This results in an implicit nonprototyped declaration that may be incompatible with the function definition.

6.3.3 Structure and Union References

A postfix expression followed by a dot followed by an identifier denotes a structure or union reference. The syntax is as follows:

```
postfix-expression . identifier
```

The postfix expression must be a structure or a union, and the identifier must name a member of the structure or union. The value is the value of the named member of the structure or union, and is an lvalue if the first expression is an lvalue. The result has the type of the indicated member and the qualifiers of the structure or union.

6.3.4 Indirect Structure and Union References

A postfix-expression followed by an arrow (built from `-` and `>`) followed by an identifier is an indirect structure or union reference. The syntax is as follows:

```
postfix-expression -> identifier
```

The postfix expression must be a pointer to a structure or a union, and the identifier must name a member of that structure or union. The result is an lvalue referring to the named member of the structure or union to which the postfix expression points. The result has the type of the selected member, and the qualifiers of the structure or union to which the postfix expression points. Thus, the expression `E1->MOS` is the same as `(*E1).MOS`.

Structures and unions are discussed in Section 7.3, page 72.

6.3.5 postfix ++ and postfix --

The syntax of postfix `++` and postfix `--` is as follows:

```
postfix-expression ++
```

```
postfix-expression --
```

When postfix `++` is applied to a modifiable lvalue, the result is the value of the object referred to by the lvalue. After the result is noted, the object is

incremented by 1 (one). See the discussions in Section 6.7, page 59, and Section 6.17, page 64, for information on conversions. The type of the result is the same as the type of the lvalue expression. The result is not an lvalue.

When `postfix --` is applied to a modifiable lvalue, the result is the value of the object referred to by the lvalue. After the result is noted, the object is decremented by 1 (one). See the discussions in Section 6.7, page 59, and Section 6.17, page 64, for information on conversions. The type of the result is the same as the type of the lvalue expression. The result is not an lvalue.

For both `postfix ++` and `postfix --` operators, updating the stored value of the operand may be delayed until the next sequence point.

6.4 Unary Operators

Expressions with unary operators associate from right to left. The syntax for unary operators is as follows:

unary-expression: *postfix-expression*
 `++ unary-expression`
 `-- unary-expression`
 unary-operator cast-expression
 `sizeof unary-expression`
 `sizeof (type-name)`

unary-operator: one of `* & - ! ~ +`

Except as noted, the operand of a unary operator must have arithmetic type.

6.4.1 Address-of and Indirection Operators

The unary `*` operator means “indirection”; the cast expression must be a pointer, and the result is either an lvalue referring to the object to which the expression points, or a function designator. If the type of the expression is “pointer to *<type>*”, the type of the result is *<type>*.

The operand of the unary `&` operator can be either a function designator or an lvalue that designates an object. If it is an lvalue, the object it designates cannot be a bitfield, and it cannot be declared with the storage class register. The result of the unary `&` operator is a pointer to the object or function referred to by the

lvalue or function designator. If the type of the lvalue is *<type>*, the type of the result is “pointer to *<type>*”.

6.4.2 Unary + and - Operators

The result of the unary - operator is the negative of its operand. The integral promotions are performed on the operand, and the result has the promoted type and the value of the negative of the operand. Negation of unsigned quantities is analogous to subtracting the value from 2^n , where n is the number of bits in the promoted type.

The unary + operator exists only in ANSI C. The integral promotions are used to convert the operand. The result has the promoted type and the value of the operand.

6.4.3 Unary ! and ~ Operators

The result of the logical negation operator ! is 1 if the value of its operand is zero, and 0 if the value of its operand is nonzero. The type of the result is *int*. The logical negation operator is applicable to any arithmetic type and to pointers.

The ~ operator (bitwise not) yields the one's complement of its operand. The usual arithmetic conversions are performed. The type of the operand must be integral.

6.4.4 Prefix ++ and -- Operators

The prefix operators ++ and -- increment and decrement their operands. Their syntax is as follows:

<code>++ unary-expression</code>

<code>-- unary-expression</code>

The object referred to by the modifiable lvalue operand of prefix ++ is incremented. The expression value is the new value of the operand but is not an lvalue. The expression ++x is equivalent to x += 1. See the discussions in Section 6.7, page 59, and Section 6.17, page 64, for information on conversions.

The prefix `--` decrements its lvalue operand in the same way that prefix `++` increments it.

6.4.5 `sizeof` Unary Operator

The `sizeof` operator yields the size in bytes of its operand. The size of a `char` is 1 (one). Its major use is in communication with routines such as storage allocators and I/O systems. The syntax of the `sizeof` operator is as follows:

```
sizeof unary-expression
```

```
sizeof (type-name)
```

The operand of `sizeof` cannot have function or incomplete type, or be an lvalue that denotes a bitfield. It can be an object or a parenthesized type name. In traditional C, the type of the result is `unsigned`. In ANSI C, the type of the result is `size_t`, which is defined in `<stddef.h>` as `unsigned int` (in `-o32` and `-n32` modes) or as `unsigned long` (in `-64` mode). The result is a constant and can be used anywhere a constant is required.

When applied to an array, `sizeof` returns the total number of bytes in the array. The size is determined from the declaration of the object in the unary expression. For variable length array types, the result is not a constant expression and is computed at run time.

The `sizeof` operator can also be applied to a parenthesized type name. In that case, it yields the size in bytes of an object of the indicated type.

When `sizeof` is applied to an aggregate, the result includes space used for padding, if any.

6.5 Cast Operators

A cast expression preceded by a parenthesized type name causes the value of the expression to convert to the indicated type. This construction is called a cast. Type names are discussed in Section 7.8, page 85. The syntax of a cast expression is as follows:

cast-expression: *unary-expression*
 (type-name) cast-expression

The type name specifies a scalar type or `void`, and the operand has scalar type. Because a cast does not yield an lvalue, the effect of qualifiers attached to the type name is inconsequential.

When an arithmetic value is cast to a pointer, and vice versa, the appropriate number of bits are simply copied unchanged from one type of value to the other. Be aware of the possible truncation of pointer values in 64-bit mode compilation, when a pointer value is converted to an `(unsigned) int`.

6.6 Multiplicative Operators

The multiplicative operators `*`, `/`, and `%` group from left to right. The usual arithmetic conversions are performed. The following is the syntax for the multiplicative operators:

multiplicative expression: *cast-expression*
 *multiplicative-expression * cast-expression*
 multiplicative-expression / cast-expression
 multiplicative-expression % cast-expression

Operands of `*` and `/` must have arithmetic type. Operands of `%` must have integral type.

The binary `*` operator indicates multiplication, and its result is the product of the operands.

The binary `/` operator indicates division of the first operator (dividend) by the second (divisor). If the operands are integral and the value of the divisor is 0, SIGTRAP is signalled. Integral division results in the integer quotient whose magnitude is less than or equal to that of the true quotient, and with the same sign.

The binary `%` operator yields the remainder from the division of the first expression (dividend) by the second (divisor). The operands must be integral. The remainder has the same sign as the dividend, so that the equality below is true when the divisor is nonzero:

```
(dividend / divisor) * divisor + dividend % divisor == dividend
```


If the value of the divisor is 0, SIGTRAP is signalled.

6.7 Additive Operators

The additive operators + and - associate from left to right. The usual arithmetic conversions are performed. The syntax for the additive operators is as follows:

```

additive-expression:      multiplicative-expression
                           additive-expression + multiplicative-expression
                           additive-expression - multiplicative-expression

```

In addition to arithmetic types, the following type combinations are acceptable for additive expressions:

- For addition, one operand is a pointer to an object type and the other operand is an integral type.
- For subtraction,
 - Both operands are pointers to qualified or unqualified versions of compatible object types.
 - The left operand is a pointer to an object type, and the right operand has integral type.

The result of the + operator is the sum of the operands. The result of the - operator is the difference of the operands.

When an operand of integral type is added to or subtracted from a pointer to an object type, the integral operand is first converted to an address offset by multiplying it by the length of the object to which the pointer points. The result is a pointer of the same type as the original pointer.

For instance, suppose *a* has type “array of *<object>*”, and *p* has type “pointer to *<object>*” and points to the initial element of *a*. Then the result of *p* + *n*, where *n* is an integral operand, is the same as *&a[n]*.

If two pointers to objects of the same type are subtracted, the result is converted (by division by the length of the object) to an integral quantity representing the number of objects separating them. Unless the pointers point to objects in the same array, the result is undefined. The actual type of the result is *int* in traditional C, and *ptrdiff_t* (defined in *<stddef.h>* as *int* in *-o32* and *-n32* modes and as *long* in *-64* mode) in ANSI C.

6.8 Shift Operators

The shift operators `<<` and `>>` associate from left to right. Each operand must be an integral type. The integral promotions are performed on each operand. The syntax is as follows:

```
shift-expression:           additive-expression  
                             shift-expression << additive-expression  
                             shift-expression >> additive-expression
```

The type of the result is that of the promoted left operand. If the right operand is negative, greater than, or equal to the length in bits of the promoted left operand, the result is undefined.

The value of `E1 << E2` is `E1` (interpreted as a bit pattern) left-shifted `E2` bits. Vacated bits are filled with zeros.

The value of `E1 >> E2` is `E1` right-shifted `E2` bit positions. If `E1` is unsigned, or if it is signed and its value is nonnegative, vacated bits are filled with zeros. If `E1` is signed and its value is negative, vacated bits are filled with ones.

6.9 Relational Operators

The relational operators associate from left to right. The syntax is as follows:

```
relational-expression:     shift-expression  
                             relational-expression < shift-expression  
                             relational-expression > shift-expression  
                             relational-expression <= shift-expression  
                             relational-expression >= shift-expression
```

The operators `<` (less than), `>` (greater than), `<=` (less than or equal to), and `>=` (greater than or equal to) all yield a result of type `int` with the value 0 if the specified relation is false and 1 if it is true.

The operands must be one of the following:

- Both arithmetic, in which case the usual arithmetic conversions are performed on them
- Both pointers to qualified or unqualified versions of compatible object types

- Both pointers to qualified or unqualified versions of compatible incomplete types

When two pointers are compared, the result depends on the relative locations in the address space of the pointed-to objects. Pointer comparison is portable only when the pointers point to objects in the same aggregate. In particular, no correlation is guaranteed between the order in which objects are declared and their resulting addresses.

6.10 Equality Operators

The `==` (equal to) and the `!=` (not equal to) operators are exactly analogous to the relational operators except for their lower precedence. (For example, `a < b == c < d` is 1 whenever `a < b` and `c < d` have the same truth value.) The syntax of the equality operators is as follows:

```
equality-expression:           relational-expression
                                equality-expression == relational-expression
                                equality-expression != relational-expression
```

The operands must be one of the following:

- Both arithmetic, in which case the usual arithmetic conversions are performed on them
- Both pointers to qualified or unqualified versions of compatible types
- A pointer to an object or incomplete type, and a pointer to qualified or unqualified void type
- A pointer and a null pointer constant

The semantics detailed in Section 6.9, page 60, apply if the operands have types suitable for those operators. Combinations of other operands have the following behavior:

- Two null pointers to object or incomplete types are equal. If two pointers to such types are equal, they either are null, point to the same object, or point to one object beyond the end of an array of such objects.
- Two pointers to the same function are equal, as are two null function pointers. Two function pointers that are equal are either both null or both point to the same function.

6.11 Bitwise AND Operator

Each operand of the bitwise AND operator must have integral type. The usual arithmetic conversions are performed. The syntax is as follows:

AND-expression: *equality-expression*
AND-expression & equality-expression

The result is the bitwise AND function of the operands, that is, each bit in the result is 0 unless the corresponding bit in each of the two operands is 1.

6.12 Bitwise Exclusive OR Operator

Each operand of the bitwise exclusive OR operator must have integral type. The usual arithmetic conversions are performed. The syntax is as follows:

exclusive-OR-expression: *AND-expression*
exclusive-OR-expression ^ AND-expression

The result has type `int`, `long`, or `long long`, and the value is the bitwise exclusive OR function of the operands. That is, each bit in the result is 0 unless the corresponding bit in one of the operands is 1, and the corresponding bit in the other operand is 0.

6.13 Bitwise Inclusive OR Operator

Each operand of the bitwise inclusive OR operator must have integral type. The usual arithmetic conversions are performed. The syntax is as follows:

inclusive-OR-expression: *exclusive-OR-expression*
inclusive-OR-expression | exclusive-OR-expression

The result has type `int`, `long`, or `long long`, and the value is the bitwise inclusive OR function of the operands. That is, each bit in the result is 0 unless the corresponding bit in at least one of the operands is 1.

6.14 Logical AND Operator

Each of the operands of the logical AND operator must have scalar type. The `&&` operator associates left to right. The syntax is as follows:

logical-AND-expression: *inclusive-OR-expression*
logical-AND-expression &&inclusive-OR-expression

The result has type `int`. If neither of the operands evaluates to 0, the result has a value of 1. Otherwise it has a value of 0.

Unlike `&`, `&&` guarantees left-to-right evaluation; moreover, the second operand is not evaluated if the first operand evaluates to zero. There is a sequence point after the evaluation of the first operand.

6.15 Logical OR Operator

Each of the operands of the logical OR operator must have scalar type. The `||` operator associates left to right. The syntax is as follows:

logical-OR-expression: *logical-AND-expression*
logical-OR-expression || logical-AND-expression

The result has type `int`. If either of the operands evaluates to one, the result has a value of 1. Otherwise it has a value of 0.

Unlike `|`, `||` guarantees left to right evaluation; moreover, the second operand is not evaluated unless the first operand evaluates to zero. A sequence point occurs after the evaluation of the first operand.

6.16 Conditional Operator

Conditional expressions associate from right to left. The syntax is as follows:

conditional-expression: *logical-OR-expression*
logical-OR-expression ? expression :
conditional-expression

The type of the first operand must be `scalar`. Only certain combinations of types are allowed for the second and third operands. These combinations are listed below, along with the type of result that the combination yields:

- Both can be arithmetic types. In this case, the usual arithmetic conversions are performed on them to derive a common type, which is the type of the result.

- Both are compatible structure or union objects. The result has the same type as the operands.
- Both are `void`. The type of the result is `void`.
- One is a pointer, and the other a null pointer constant. The type of the result is the type of the nonconstant pointer.
- One is a pointer to `void`, and the other is a pointer to an object or incomplete type. The second operand is converted to a pointer to `void`. The result is also a pointer to `void`.
- Both are pointers to qualified or unqualified versions of compatible types. The result has a type compatible with each, qualified with all the qualifiers of the types pointed to by both operands.

Evaluation of the conditional operator proceeds as follows:

- The first expression is evaluated, after which a sequence point occurs.
- If the value of the first expression is nonzero, the result is the value of the second operand.
- If the value of the first expression is zero, the result is the value of the third operand.
- Only one of the second and third operands is evaluated.

6.17 Assignment Operators

All assignment operators associate from right to left. The syntax is as follows:

assignment-expression: *conditional-expression*

 unary-expression assignment-operator
 assignment-expression

assignment operator: one = *= /= %= += -= <<= >>= &= ^= |=
of

Assignment operators require a modifiable lvalue as their left operand. The type of an assignment expression is that of its unqualified left operand. The result is not an lvalue. Its value is the value stored in the left operand after the assignment, but the actual update of the stored value may be delayed until the next sequence point.

The order of evaluation of the operands is unspecified.

6.17.1 Assignment Using = (Simple Assignment)

The operands permissible in simple assignment must obey one of the following:

- Both have arithmetic type or are compatible structure or union types.
- Both are pointers, and the type pointed to by the left has all of the qualifiers of the type pointed to by the right.
- One is a pointer to an object or incomplete type, and the other is a pointer to void. The type pointed to by the left must have all of the qualifiers of the type pointed to by the right.
- The left operand is a pointer, and the right is a null pointer constant.

In simple assignment, the value of the right operand is converted to the type of the assignment expression and replaces the value of the object referred to by the left operand. If the value being stored is accessed by another object that overlaps it, the behavior is undefined unless the overlap is exact and the types of the two objects are compatible.

6.17.2 Compound Assignment

For the operators += and -=, either both operators must have arithmetic types, or the left operand must be a pointer and the right an operand integral. In the latter case, the right operand is converted as explained in Section 6.7, page 59. For all other operators, each operand must have arithmetic type consistent with those allowed for the corresponding binary operator.

The expression $E1 \text{ op } = E2$ is equivalent to the expression $E1 = E1 \text{ op } E2$, except that in the former, $E1$ is evaluated only once.

6.18 Comma Operator

A pair of expressions separated by a comma is evaluated left to right, and the value of the left expression is discarded. This operator associates left to right. The syntax of the comma operator is as follows:

expression: *assignment-expression*
expression, assignment-expression

The type and value of the result are the type and value of the right operand. In contexts where the comma is given a special meaning, the comma operator as described in this section can appear only in parentheses. Two such contexts are

lists of actual arguments to functions (described in Section 6.2, page 51) and lists of initializers (see Section 7.11, page 88). For example, the following code has three arguments, the second of which has the value 5:

```
f(a, (t=3, t+2), c)
```

6.19 Constant Expressions

A constant expression can be used any place a constant can be used. The syntax is as follows:

constant-expression: *conditional-expression*

A constant expression cannot contain assignment, increment, decrement, function-call, or comma operators. It must evaluate to a constant that is in the range of representable values for its type. Otherwise, the semantic rules for the evaluation of nonconstant expressions apply.

Constant expressions are separated into three classes:

- An integral constant expression has integral type and is restricted to operands that are integral constants, `sizeof` expressions (whose operands do not have variable length array type or a parenthesized name of such a type), and floating constants that are the immediate operands of integral casts.
- An arithmetic constant expression has arithmetic type and is restricted to operands that are arithmetic constants, and `sizeof` expressions (whose operands do not have variable length array type or a parenthesized name of such a type). Cast expressions in arithmetic constant expressions can convert only between arithmetic types.
- An address constant is a pointer to an lvalue designating an object of static storage duration, or a pointer to a function designator. It can be created explicitly or implicitly, as long as no attempt is made to access an object value.

Either address or arithmetic constant expressions can be used in initializers. In addition, initializers can contain null pointer constants and address constants (for object types), and plus or minus integral constant expressions.

6.20 Integer and Floating Point Exceptions

The following are a few points to keep in mind about integer and floating point exceptions:

- Integer divide-by-zero results in a trap. Other integer exception conditions are ignored.
- Silicon Graphics' floating point conforms to the IEEE standard. Floating point exceptions are ignored by default, yielding the default IEEE results of infinity for divide-by-zero and overflow, not-a-number for invalid operations, and zero for underflow.
- You can gain control over these exceptions and their results most easily by `<_newline>` using the Silicon Graphics IEEE floating point exception handler package (see the `handle_sigfpes(3C)` reference page).
- You can also control these exceptions by implementing your own handler and appropriately initializing the floating point unit (see the `fpc(3C)` reference page).

Declarations [7]

A declaration specifies the interpretation given to a set of identifiers. Declarations have the following form:

declaration: *declaration-specifiers init-declarator-list_{opt}*;

The *init-declarator-list* is a comma-separated sequence of declarators, each of which can have an initializer.

In ANSI C, the *init-declarator-list* can also contain additional type information:

init-declarator-list: *init-declarator*

init-declarator-list , *init-declarator*

init-declarator: *declarator*

declarator = *initializer*

The declarators in the *init-declarator list* contain the identifiers being declared. The declaration specifiers consist of a sequence of specifiers that determine the linkage, storage duration, and part of the type of the identifiers indicated by the declarator. Declaration specifiers have the following form:

declaration-specifiers: *storage-class-specifier declaration-specifiers_{opt}*

type-specifier declaration-specifiers_{opt}

type-qualifier declaration-specifiers_{opt}

If an identifier that is not a tag has no linkage (see Section 4.1, page 31), at most one declaration of the identifier can appear in the same scope and name space. The type of an object that has no linkage must be complete by the end of its declarator or initializer. Multiple declarations of tags and ordinary identifiers with external or internal linkage can appear in the same scope so long as they specify compatible types.

If a sequence of specifiers in a declarator contains a variable length array type, the type specified by the declarator is said to be “variably modified.” All declarations of variably modified types must be declared at either block or function prototype scope. File scope identifiers cannot be declared with a variably modified type.

In traditional C, at most one declaration of an identifier with internal linkage can appear in the same scope and name space, unless it is a tag.

In ANSI C, a declaration must declare at least one of the following:

- A declarator
- A tag
- The members of an enumeration

A declaration may reserve storage for the entities specified in the declarators. Such a declaration is called a definition. (Function definitions have a different syntax and are discussed in Section 7.7.4, page 82, and Chapter 9, page 101.)

7.1 Storage Class Specifiers

The storage class specifier indicates linkage and storage duration. These attributes are discussed in Section 4.1, page 31. Storage class specifiers have the following form:

```
storage-class-specifier:      auto
                               static
                               extern
                               register
                               typedef
```

The `typedef` specifier does not reserve storage and is called a storage-class specifier only for syntactic convenience. See Section 7.10, page 87, for more information.

The following rules apply to the use of storage class specifiers:

- A declaration can have at most one storage class specifier. If the storage class specifier is missing from a declaration, it is assumed to be `extern` unless the declaration is of an object and occurs inside a function, in which case it is assumed to be `auto`. (See Section 2.2, page 10.)
- Identifiers declared within a function with the storage class `extern` must have an external definition (see Chapter 9, page 101) somewhere outside the function in which they are declared.
- Identifiers declared with the storage class `static` have static storage duration, and either internal linkage (if declared outside a function) or no linkage (if declared inside a function). If the identifiers are initialized, the initialization is performed once before the beginning of execution. If no

explicit initialization is performed, static objects are implicitly initialized to zero.

- A register declaration is an auto declaration, with a hint to the compiler that the objects declared will be heavily used. Whether the object is actually placed in fast storage is implementation defined. You cannot take the address of any part of an object declared with the register specifier.

7.2 Type Specifiers

Type specifiers are listed below. The syntax is as follows:

type-specifier: struct-or-union-specifier
 typedef-name
 enum-specifier
 char
 short
 int
 long
 signed
 unsigned
 float
 double
 void

The following is a list of all valid combinations of type specifiers. These combinations are organized into sets. The type specifiers in each set are equivalent in all implementations. The arrangement of the type specifiers appearing in any set can be altered without effect.

- void
- char
- signed char
- unsigned char

- `short`, `signed short`, `short int`, or `signed short int`
- `unsigned short`, or `unsigned short int`
- `int`, `signed`, `signed int`, or no type specifiers
- `unsigned`, or `unsigned int`
- `long`, `signed long`, `long int`, or `signed long int`
- `unsigned long`, or `unsigned long int`
- `long long`, `signed long long`, `long long int`, or `signed long long int`
- `unsigned long long`, or `unsigned long long int`
- `float`
- `double`
- `long double`

In traditional C, the type `long float` is allowed and is equivalent to `double`; its use is not recommended. It elicits a warning if you are not in `-cckr` mode. Use of the type `long double` is not recommended in traditional C.

`long long` is not a valid ANSI C type, so a warning appears for every occurrence of it in the source program text in `-ansi` and `-ansiposix` modes.

Specifiers for structures, unions, and enumerations are discussed in Section 7.3, page 72, and Section 7.5, page 76. Declarations with typedef names are discussed in Section 7.10, page 87.

7.3 Structure and Union Declarations

A structure is an object consisting of a sequence of named members. Each member can have any type. A union is an object that can, at a given time, contain any one of several members. Structure and union specifiers have the same form. The syntax is as follows:

```
struct-or-union-specifier:      struct-or-union {struct-decl-list}  
struct-or-union-specifier:      struct-or-union identifier {struct-decl-list}  
struct-or-union-specifier:      struct-or-union identifier
```

struct-or-union: `struct`
 `union`

The *struct-decl-list* is a sequence of declarations for the members of the structure or union. The syntax, in three possible forms, is as follows:

struct-decl-list: *struct-declaration*
 struct-decl-list struct-declaration

struct-declaration: *specifier-qualifier-list struct-declarator-list;*

struct-declarator-list: *struct-declarator*
 struct-declarator-list , struct-declarator

In the usual case, a *struct-declarator* is just a declarator for a member of a structure or union. A structure member can also consist of a specified number of bits. Such a member is also called a bitfield. Its length, a non-negative constant expression, is separated from the field name by a colon. Section 7.4, page 75, are discussed at the end of this section.

The syntax for *struct-declarator* is as follows:

struct-declarator: *declarator*
 declarator : constant-expression
 : constant-expression

A `struct` or `union` cannot contain any of the following:

- A member with incomplete or function type.
- A member that is an instance of itself. It can, however, contain a member that is a pointer to an instance of itself.
- A member that has a variable length array type.
- A member that is a pointer to a variable length array type.

Within a structure, the objects declared have addresses that increase as the declarations are read left to right. Each non-field member of a structure begins on an addressing boundary appropriate to its type; therefore, there may be unnamed holes in a structure.

A union can be thought of as a structure whose members all begin at offset 0 and whose size is sufficient to contain any of its members. At most, one of the members can be stored in a union at any time.

A structure or union specifier of the second form declares the identifier to be the structure tag (or union tag) of the structure specified by the list. This type of specifier is one of the following:

```
struct identifier {struct-decl-list}  
union identifier {struct-decl-list}
```

A subsequent declaration can use the third form of specifier, one of the following:

```
struct identifier  
union identifier
```

Structure tags allow definition of self-referential structures. Structure tags also permit the long part of the declaration to be given once and used several times.

The third form of a structure or union specifier can be used before a declaration that gives the complete specification of the structure or union in situations in which the size of the structure or union is unnecessary. The size is unnecessary in two situations: when a pointer to a structure or union is being declared and when a `typedef` name is declared to be a synonym for a structure or union. This, for example, allows the declaration of a pair of structures that contain pointers to each other.

The names of members of each `struct` or `union` have their own name space, and do not conflict with each other or with ordinary variables. A particular member name cannot be used twice in the same structure, but it can be used in several different structures in the same scope.

Names that are used for tags reside in a single name space. They do not conflict with other names or with names used for tags in an enclosing scope. This tag name space, however, consists of tag names used for all `struct`, `union`, or `enum` declarations. Therefore, the tag name of an `enum` may conflict with the tag name of a `struct` in the same scope. (See Section 4.1, page 31.)

A simple but important example of a structure declaration is the following binary tree structure:

```
struct tnode {  
    char tword[20];  
    int count;  
    struct tnode *left;  
    struct tnode *right;  
};  
struct tnode s, *sp;
```


This structure contains an array of 20 characters, an integer, and two pointers to instances of itself. Once this structure has been declared, the next line declares a structure of type `struct tnode (s)` and a pointer to a structure of type `struct tnode (sp)`.

With these declarations,

- The expression `sp->count` refers to the count field of the structure to which `sp` points.
- The expression `s.left` refers to the left subtree pointer of the structure `s`.
- The expression `s.right->tword[0]` refers to the first character of the `tword` member of the right subtree of `s`.

7.4 Bitfields

A structure member can consist of a specified number of bits, called a bitfield. In strict ANSI C mode, bitfields should be of type `int`, `signed int`, or `unsigned int`. Silicon Graphics C allows bitfields of any integral type, but warns for non-`int` types in `-ansi` and `-ansiposix` modes.

The default type of field members is `int`, but whether it is `signed` or `unsigned int` is defined by the implementation. Therefore, you should specify the signedness of bitfields when they are declared. In this implementation, the default type of a bitfield is `signed`.

The constant expression that denotes the width of the bitfield must have a value no greater than the width, in bits, of the type of the bitfield. An implementation can allocate any addressable storage unit (referred to in this discussion as simply a “unit”) large enough to hold a bitfield. If an adjacent bitfield will not fit into the remainder of the unit, the implementation defines whether bitfields are allowed to span units or whether another unit is allocated for the second bitfield. The ordering of the bits within a unit is also implementation-defined.

A bitfield with no declarator, just a colon and a width, indicates an unnamed field useful for padding. As a special case, a field with a width of zero (which cannot have a declarator) specifies alignment of the next field at the next unit boundary.

These implementation-defined characteristics make the use of bitfields inherently nonportable, particularly if they are used in situations where the underlying object may be accessed by another data type (in a union, for example).

In the Silicon Graphics implementation of C, the first bitfield encountered in a `struct` is not necessarily allocated on a unit boundary and is packed into the current unit, if possible. A bitfield cannot span a unit boundary. Bits for bitfields are allocated from left (most significant) to right.

There are no arrays of bitfields. Because the address-of operator, `&`, cannot be applied to bitfields, there are also no pointers to bitfields.

7.5 Enumeration Declarations

Enumeration variables and constants have integral type. The syntax is as follows:

```
enum-specifier:           enum {enum-list}
                          enum identifier {enum-list}
                          enum identifier

enum-list:              enumerator
                          enum-list , enumerator

enumerator:            identifier
                          identifier = constant-expression
```

The identifiers in an *enum-list* are declared as `int` constants and can appear wherever such constants are allowed. If no enumerators with `=` appear, then the values of the corresponding constants begin at 0 and increase by 1 as the declaration is read from left to right. An enumerator with `=` gives the associated identifier the value indicated; subsequent identifiers continue the progression from the assigned value. Note that the use of `=` may result in multiple enumeration constants having the same integral value, even though they are declared in the same enumeration declaration.

Enumerators are in the ordinary identifiers name space (see Section 4.1.2, page 33). Thus, an identifier used as an enumerator may conflict with identifiers used for objects, functions, and user-defined types in the same scope.

The role of the identifier in the *enum-specifier* is entirely analogous to that of the structure tag in a *struct-specifier*; it names a particular enumeration. For example,

```
enum color { chartreuse, burgundy, claret = 20, winedark };
...
```

```
enum color *cp, col;
...
col = claret;
cp = &col;
...
if (*cp == burgundy) ...
```

This example makes *color* the enumeration-tag of a type describing various colors, and then declares *cp* as a pointer to an object of that type, *col*. The possible values are drawn from the set {0,1,20,21}. The tags of enumeration declarations are members of the single tag name space, and thus must be distinct from tags of `struct` and `union` declarations.

7.6 Type Qualifiers

Type qualifiers have the following syntax:

```
type-qualifier:          const
                          volatile
                          __restrict
```

The same type qualifier cannot appear more than once in the same specifier list either directly or indirectly (through `typedefs`).

The value of an object declared with the `const` type qualifier is constant. It cannot be modified, although it can be initialized following the same rules as the initialization of any other object. (See the discussion in Section 7.11, page 88.) Implementations are free to allocate `const` objects, that are not also declared `volatile`, in read-only storage.

An object declared with the `volatile` type qualifier may be accessed in unknown ways or have unknown side effects. For example, a `volatile` object may be a special hardware register. Expressions referring to objects qualified as `volatile` must be evaluated strictly according to the semantics. When `volatile` objects are involved, an implementation is not free to perform optimizations that would otherwise be valid. At each sequence point, the value of all `volatile` objects must agree with that specified by the semantics.

The `__restrict` qualifier applies only to pointers and is discussed in Section 7.7.2.1, page 79.

If an array is specified with type qualifiers, the qualifiers are applied to the elements of the array. If a `struct` or `union` is qualified, the qualification applies to each member.

Two qualified types are compatible if they are identically qualified versions of compatible types. The order of qualifiers in a list has no effect on their semantics.

The syntax of pointers allows the specification of qualifiers that affect either the pointer itself or the underlying object. Qualified pointers are covered in Section 7.7.2, page 79.

7.7 Declarators

Declarators have the syntax shown below:

<i>declarator:</i>	<i>pointer</i> _{opt} <i>direct-declarator</i>
<i>direct-declarator:</i>	<i>identifier</i>
	(<i>declarator</i>)
	<i>direct-declarator</i> (<i>parameter-type-list</i> _{opt})
	<i>direct-declarator</i> (<i>identifier-list</i> _{opt})
	<i>direct-declarator</i> [<i>constant-expression</i> _{opt}]

The grouping is the same as in expressions.

7.7.1 Meaning of Declarators

Each declarator is an assertion that when a construction of the same form as the declarator appears in an expression, it designates a function or object with the scope, storage duration, and type indicated by the declaration.

Each declarator contains exactly one identifier; it is this identifier that is declared. If, in the declaration “`T D1;`” `D1` is simply an identifier, then the type of the identifier is `T`. A declarator in parentheses is identical to the unparenthesized declarator. The binding of complex declarators can, however, be altered by parentheses.

7.7.2 Pointer Declarators

Pointer declarators have the form:

```
pointer:           * type-qualifier-listopt
                   * type-qualifier-listopt pointer
```

The following is an example of a declaration:

```
T D1
```

In this declaration, the identifier has type `.. T`, where the `..` is empty if `D1` is just a plain identifier (so that the type of `x` in `int x` is just `int`). Then, if `D1` has the form `*type-qualifier-listopt D`, the type of the contained identifier is `".. (possibly-qualified) pointer to T."`

7.7.2.1 Qualifiers and Pointers

It is important to be aware of the distinction between a qualified pointer to a type and a pointer to a qualified type. In the declarations below, `ptr_to_const` is a pointer to `const long`:

```
const long *ptr_to_const;
long * const const_ptr;
volatile int * const const_ptr_to_volatile;
```

The `long` pointed to by `ptr_to_const` in the first declaration, cannot be modified by the pointer. The pointer itself, however, can be altered. In the second declaration, `const_ptr` can be used to modify the `long` that it points to, but the pointer itself cannot be modified. In the last declaration, `const_ptr_to_volatile` is a constant pointer to a volatile `int` and can be used to modify it. The pointer itself, however, cannot be modified.

The `__restrict` qualifier tells the compiler to assume that dereferencing the qualified pointer is the only way the program can access the memory pointed to by that pointer. Therefore, loads and stores through such a pointer are assumed not to alias with any other loads and stores in the program, except other loads and stores through the same pointer variable.

The following example illustrates the use of the `__restrict` qualifier:

```
float x[ARRAY_SIZE];
float *c = x;

void f4_opt(int n, float * __restrict a, float * __restrict b)
```

```
{
  int i;
  /* No data dependence across iterations because of __restrict */
  for (i = 0; i < n; i++)
    a[i] = b[i] + c[i];
}
```

7.7.2.2 Pointer-related Command Options

The Silicon Graphics C compiler supports the following two alias-related command-line switches that can be useful for improving performance:

-
OPT:alias=restrict Implements the following semantics: memory operations dereferencing different named pointers in the program are assumed not to alias with each other, nor with any named scalar in the program.

For example, if *p* and *q* are pointers, this option means that **p* does not alias with **q*, with *p*, or with any named scalar variable.

-
OPT:alias=disjoint Implements the following semantics: memory operations dereferencing different named pointers in the program are assumed not to alias with each other, and in addition, different dereferencing depths of the same named pointer are assumed not to alias with each other.

For example, if *p* and *q* are of type pointer to pointer, **p* does not alias with **q*, with ***p*, or with ***q*.

Note: With either switch enabled, programs violating the corresponding aliasing assumptions may be compiled incorrectly.

7.7.3 Array Declarators

If in the declaration *T D1*, *D1* has the form *D[expression_{opt}]* or *D[*]*, then the contained identifier has type “array of *T*.” Starting with version 7.2, the Silicon Graphics C compiler now supports variable length arrays as well as fixed length arrays. A variable length array is an array that has a size (at least one dimension) that is determined at run time. The ability to use variable length arrays enhances the compiler’s range of use for numerical programming.

The following rules apply to array declarations:

- If the array is a fixed length array, the expression enclosed in square brackets, if it exists, must be an integral constant expression whose value is greater than zero. (See Section 6.2, page 51.)
- When several “array of” specifications are adjacent, a multi-dimensional array is created; the constant expressions that specify the bounds of the arrays can be missing only for the first member of the sequence.
- The absence of the first array dimension is allowed if the array is external and the actual definition (which allocates storage) is given elsewhere, or if the declarator is followed by initialization. In the latter case, the size is calculated from the number of elements supplied.
- If * is used instead of a size expression, the array is of “variable length array” type with unspecified size. This can only be used in declarations with function prototype scope.
- The array type is “fixed length array” if the size expression is an integer constant expression, and the element type has a fixed size. Otherwise the type is variable length array.
- The size of a variable length array type does not change until the execution of the block containing the declaration has finished.
- Array objects declared with either static or extern storage class specifiers cannot be declared with a variable length array type. However, block scope pointers declared with the static storage class specifier can be declared as pointers to variable length array types.
- In order for two array types to be compatible, their element types must be compatible. In addition, if both of their size specifications are present and are integer constant expressions, they must have the same value. If either size specifier is variable, the two sizes must evaluate to the same value at run time.
- An array can be constructed from one of the basic types, from a pointer, from a structure or union, or from another array (to generate a multi-dimensional array).

The example below declares an array of float numbers and an array of pointers to float numbers:

```
float fa[17], *afp[17];
```

The following example declares a static three-dimensional array of integers, with rank $3 \times 5 \times 7$.

```
static int x3d[3][5][7];
```

In the above example, `x3d` is an array of three items; each item is an array of five items, each of which is an array of seven integers. Any of the expressions `x3d`, `x3d[i]`, `x3d[i][j]`, `x3d[i][j][k]` can reasonably appear in an expression. The first three have type array and the last has type `int`.

7.7.4 Function Declarators and Prototypes

The syntax for function declarators is shown below:

```
direct-declarator (parameter-type-listopt)
direct-declarator (identifier-listopt)
parameter-type-list:
parameter-list
parameter-list , ...
parameter-list:
parameter-declaration
parameter-list , parameter-declaration
parameter-declaration:
declaration-specifiers declarator
declaration-specifiers abstract-declaratoropt
identifier-list:
identifier
identifier-list , identifier
```

Function declarators cannot specify a function or array type as the return type. In addition, the only storage class specifier that can be used in a parameter declaration is `register`. For example, in the declaration `T D1`, `D1` has one of the following forms:

- `D(parameter-type-listopt)`
- `D(identifier-listopt)`

The contained identifier has the type “.. function returning T,” and is possibly a prototype, as discussed later in this section.

A parameter type list declares the types of, and can declare identifiers for, the formal parameters of a function. A declared parameter that is a member of the parameter type list that is not part of a function definition may use the `[*]` notation in its sequence of declarator specifiers to specify a variable length array type.

The absence of a parameter type list indicates that no typing information is given for the function. A parameter type list consisting only of the keyword `void` indicates that the function takes zero parameters. If the parameter type list ends in ellipses (...), the function can have one or more additional arguments of variable or unknown type. (See `<stdarg.h>`.)

The semantics of a function declarator are determined by its form and context. The possible combinations are as follows:

- The declarator is not part of the function definition. The function is defined elsewhere. In this case, the declarator cannot have an identifier list.
 - If the parameter type list is absent, the declarator is an old-style function declaration. Only the return type is significant.
 - If the parameter type list is present, the declarator is a function prototype.
- The declarator is part of the function definition:
 - If the declarator has an identifier list, the declarator is an old-style function definition. Only the return type is significant.
 - If the declarator has a parameter type list, the definition is in prototype form. If no previous declaration for this function has been encountered, a function prototype is created for it that has file scope.

If two declarations (one of which can be a definition) of the same function in the same scope are encountered, they must match, both in type of return value and in parameter type list. If one and only one of the declarations has a parameter type list, the behavior varies between ANSI C and Traditional C.

In traditional C, most combinations pass without any diagnostic messages. However, an error message is emitted for cases where an incompatibility is likely to lead to a run-time failure. For example, a `float` type in a parameter type list of a function prototype is incompatible with any old-style declaration for the same function; therefore, Silicon Graphics considers such redeclarations erroneous.

In ANSI C, if the type of any parameter declared in the parameter type list is other than that which would be derived using the default argument promotions, an error is posted. Otherwise, a warning is posted and the function prototype remains in scope.

In all cases, the type of the return value of duplicate declarations of the same function must match, as must the use of ellipses.

When a function is invoked for which a function prototype is in scope, an attempt is made to convert each actual parameter to the type of the corresponding formal parameter specified in the function prototype, superseding the default argument promotions. In particular, `floats` specified in the type list are not converted to `double` before the call. If the list terminates with an ellipsis (...), only the parameters specified in the prototype have their types checked; additional parameters are converted according to the default argument promotions (discussed in Section 7.6, page 77). Otherwise, the number of parameters appearing in the parameter list at the point of call must agree in number with those in the function prototype.

The following are two examples of function prototypes:

```
double foo(int *first, float second, ... );
int *fip(int a, long l, int (*ff)(float));
```

The first prototype declares a function `foo()` which returns a `double` and has (at least) two parameters: a pointer to an `int` and a `float`. Further parameters can appear in a call of the function and are unspecified. The default argument promotions are applied to any unspecified arguments. The second prototype declares a function `fip()`, which returns a pointer to an `int`. The function `fip()` has three parameters: an `int`, a `long`, and a pointer to a function returning an `int` that has a single (`float`) argument.

7.7.4.1 Prototyped Functions Summarized

When a function call occurs, each argument is converted using the default argument promotions unless that argument has a type specified in a corresponding in-scope prototype for the function being called. It is easy to envision situations that could prove disastrous if some calls to a function are made with a prototype in-scope and some are not. Unexpected results can also occur if a function is called with different prototypes in scope. Therefore, if a function is prototyped, it is extremely important to make sure that all invocations of the function use the prototype.

In addition to adding a new syntax for external declarations of functions, prototypes have added a new syntax for external definitions of functions. This syntax is termed “function prototype form.” It is highly important to define prototyped functions using a parameter type list rather than a simple identifier list if the parameters are to be received as intended.

In ANSI C, unless the function definition has a parameter type list, it is assumed that arguments have been promoted according to the default argument

promotions. Specifically, an in-scope prototype for the function at the point of its definition has no effect on the type of the arguments that the function expects.

The compilers issue error diagnostics when argument-type mismatches are likely to result in faulty run-time behavior.

7.7.5 Restrictions on Declarators

Not all the possibilities allowed by the syntax of declarators are permitted. The following restrictions apply:

- Functions cannot return arrays or functions although they can return pointers.
- No arrays of functions exist although arrays of pointers to functions can exist.
- A structure or union cannot contain a function, but it can contain a pointer to a function.

As an example, the following declaration declares an integer `i`; a pointer to an integer, `ip`; a function returning an integer, `f()`; a function returning a pointer to an integer, `fip()`; and a pointer to a function that returns an integer, `pfi`:

```
int i, *ip, f(), *fip(), (*pfi)();
```

It is especially useful to compare the last two. The binding of `*fip()` is `*(fip())`. The declaration suggests, and the same construction in an expression requires, the calling of a function `fip()`, and then using indirection through the (pointer) result to yield an integer. In the declarator `*pfi()`, the extra parentheses are necessary, because they are also in an expression, to indicate that indirection through a pointer to a function yields a function, which is then called and returns an integer.

7.8 Type Names

In several contexts (for example, to specify type conversions explicitly by means of a cast, in a function prototype, and as an argument of `sizeof`), it is best to supply the name of a data type. This naming is accomplished using a “type name,” whose syntax is a declaration for an object of that type without the identifier.

The syntax for type names is as follows:

type-name: *specifier-qualifier-list abstract-declarator*_{*opt*}

abstract-declarator: *pointer*
*pointer*_{opt} *direct-abstract-declarator*
direct-abstract-declarator: (*abstract-declarator*)
*direct-abstract-declarator*_{opt} [*constant-expression*_{opt}]
*direct-abstract-declarator*_{opt} (*parameter-type-list*_{opt})

The type name created can be used as a synonym for the type of the omitted identifier. The syntax indicates that a set of empty parentheses in a type name is interpreted as function with no parameter information rather than as redundant parentheses surrounding the (omitted) identifier.

Examples of type names are shown in Table 10, page 86.

Table 10. Examples of Type Names

Type	Description
<code>int</code>	Integer
<code>int *</code>	Pointer to integer
<code>int *[3]</code>	Array of three pointers to integers
<code>int (*)[3]</code>	Pointer to an array of three integers
<code>int *(void)</code>	Function with zero arguments returning pointer to integer
<code>int (*)(float, ...)</code>	Pointer to function returning an integer, that has a variable number of arguments the first of which is a <code>float</code>
<code>int (*[3])()</code>	Array of three pointers to functions returning an integer for which no parameter type information is given

7.9 Implicit Declarations

It is not always necessary to specify both the storage class and the type of identifiers in a declaration. The storage class is supplied by the context in external definitions, and in declarations of formal parameters and structure members. Missing storage class specifiers appearing in declarations outside of functions are assumed to be `extern` (see Section 2.5, page 19, for details). Missing type specifiers in this context are assumed to be `int`. In a declaration

inside a function, if a type but no storage class is indicated, the identifier is assumed to be `auto`. An exception to the latter rule is made for functions because `auto` functions do not exist. If the type of an identifier is “function returning `<type>`”, it is implicitly declared to be `extern`.

In an expression, an identifier followed by a left parenthesis (indicating a function call) that is not already declared is implicitly declared to be of type function returning `int`.

7.10 typedef

Declarations with the storage class specifier `typedef` do not define storage. A `typedef` has the following syntax:

```
typedef-name:           identifier
```

An identifier appearing in a `typedef` declaration becomes a synonym for the type rather than becoming an object with the given type. For example, if the `int` type specifier in the following example were preceded with `typedef`, the identifier declared as an object would instead be declared as a synonym for the array type:

```
int intarray[10];
```

This can appear as shown below:

```
typedef int intarray[10];
```

`intarray` could then be used as if it were a basic type, as in the following:

```
intarray ia;
```

In the following example, the last three declarations are legal. The type of `distance` is `int`, that of `metricp` is pointer to `int`, and that of `z` is the specified structure. The `zp` is a pointer to such a structure:

```
typedef int MILES, *KCLICKSP;
typedef struct {
double re, im;
}
complex;

MILES distance;
extern KCLICKSP metricp;
complex z, *zp;
```

The `typedef` does not introduce brand-new types, only synonyms for types that could be specified in another way. For instance, in the preceding example, `distance` is considered to have the same type as any other `int` object.

`typedef` declarations that specify a variably modified type have block scope. The array size specified by the variable length array type is evaluated at the time the type definition is declared and not at the time it is used as a type specifier in an actual declarator.

7.11 Initialization

A declaration of an object or of an array of unknown size can specify an initial value for the identifier being declared. The initializer is preceded by `=` and consists of an expression or a list of values enclosed in nested braces:

```
initializer:           assignment-expression
                       {initializer-list}
                       {initializer-list}
initializer-list:    initializer
                       initializer-list , initializer
```

There cannot be more initializers than there are objects to be initialized. All the expressions in an initializer for an object of static storage duration must be constant expressions (see Section 6.2, page 51). Objects with automatic storage duration can be initialized by arbitrary expressions involving constants and previously declared variables and functions, except for aggregate initialization, which can include only constant expressions.

Identifiers declared with block scope and either external or internal linkage (that is, objects declared in a function with the storage class specifier `extern`) cannot be initialized.

Variables of static storage duration that are not explicitly initialized are implicitly initialized to zero. The value of automatic and register variables that are not explicitly initialized is undefined.

When an initializer applies to a scalar (a pointer or an object of arithmetic type; see Section 4.2.3, page 40), it consists of a single expression, perhaps in braces. The initial value of the object is taken from the expression. With the exception of type qualifiers associated with the scalar, which are ignored during the initialization, the same conversions as for assignment are performed.

7.11.1 Initialization of Aggregates

In traditional C, it is illegal to initialize a union. It is also illegal to initialize a `struct` of automatic storage duration.

In ANSI C, objects that are `struct` or `union` types can be initialized, even if they have automatic storage duration. `unions` are initialized using the type of the first named element in their declaration. The initializers used for a `struct` or `union` of automatic storage duration must be constant expressions if they are in an initializer list. If the `struct` or `union` is initialized using an assignment expression, the expression need not be constant.

When the declared variable is a `struct` or array, the initializer consists of a brace-enclosed, comma-separated list of initializers for the members of the aggregate written in increasing subscript or member order. If the aggregate contains subaggregates, this rule applies recursively to the members of the aggregate.

If the initializer of a subaggregate or union begins with a left brace, its initializers consist of all the initializers found between the left brace and the matching right brace. If, however, the initializer does not begin with a left brace, then only enough elements from the list are taken to account for the members of the subaggregate; any remaining members are left to initialize the next member of the aggregate of which the current subaggregate is a part.

Within any brace-enclosed list, there should not be more initializers than members. If there are fewer initializers in the list than there are members of the aggregate, then the aggregate is padded with zeros.

Unnamed `struct` or `union` members are ignored during initialization.

In ANSI C, if the variable is a union, the initializer consists of a brace-enclosed initializer for the first member of the union. Initialization of `struct` or `union` objects with automatic storage duration can be abbreviated as a simple assignment of a compatible `struct` or `union` object.

A final abbreviation allows a `char` array to be initialized by a string literal. In this case, successive characters of the string literal initialize the members of the array.

In ANSI C, an array of wide characters (that is, whose element type is compatible with `wchar_t`) can be initialized with a wide string literal (see Section 3.5, page 28).

7.11.2 Examples of Initialization

The following example declares and initializes `x` as a one-dimensional array that has three members, because no size was specified and there are three initializers:

```
int x[] = { 1, 3, 5 };
```

The next example shows a completely bracketed initialization: 1, 3, and 5 initialize the first row of the array `y[0]`, namely `y[0][0]`, `y[0][1]`, and `y[0][2]`. Likewise, the next two lines initialize `y[1]` and `y[2]`. The initializer ends early, and therefore, `y[3]` is initialized with 0:

```
float y[4][3] =
{
    { 1, 3, 5 },
    { 2, 4, 6 },
    { 3, 5, 7 },
};
```

The next example achieves precisely the same effect. The initializer for `y` begins with a left brace but that for `y[0]` does not; therefore, three elements from the list are used. Likewise, the next three are taken successively for `y[1]` and `y[2]`:

```
float y[4][3] =
{
    1, 3, 5, 2, 4, 6, 3, 5, 7
};
```

The next example initializes the first column of `y` (regarded as a two-dimensional array) and leaves the rest 0:

```
float y[4][3] = {
    { 1 }, { 2 }, { 3 }, { 4 }
};
```

The following example demonstrates the ANSI C rules. A union object, `dc_u`, is initialized by using the first element only:

```
union dc_u {
    double d;
    char *cptr;
};

union dc_u dc0 = { 4.0 };
```


The final example shows a character array whose members are initialized with a string literal. The length of the string (or size of the array) includes the terminating NULL character, \0:

```
char msg[] = "Syntax error on line %s\n";
```


Statements [8]

A statement is a complete instruction to the computer. Except as indicated, statements are executed in sequence. Statements have the following form:

statement: *expression-statement*
 compound-statement
 selection-statement
 iteration-statement
 jump-statement
 labeled-statement

8.1 Expression Statement

Most statements are expression statements, which have the following form:

expression-statement: *expression*_{opt};

Usually expression statements are expressions evaluated for their side effects, such as assignments or function calls. A special case is the null statement, which consists of only a semicolon.

8.2 Compound Statement or Block

A compound statement (or block) groups a set of statements into a syntactic unit. The set can have its own declarations and initializers, and has the following form:

compound-statement: {*declaration-list*_{opt} *statement-list*_{opt}}

declaration-list: *declaration*
 declaration-list declaration

statement-list: *statement*
 statement-list statement

Declarations within compound statements have block scope. If any of the identifiers in the declaration list were previously declared, the outer declaration is hidden for the duration of the block, after which it resumes its force. In traditional C, however, function declarations always have file scope whenever they appear.

Initialization of identifiers declared within the block is restricted to those that have no linkage. Thus, the initialization of an identifier declared within the block using the `extern` specifier is not allowed. These initializations are performed only once, prior to the first entry into the block, for identifiers with static storage duration. For identifiers with automatic storage duration, it is performed each time the block is entered at the top. It is currently possible (but a bad practice) to transfer into a block; in that case, no initializations are performed.

8.3 Selection Statements

Selection statements include `if` and `switch` statements and have the following form:

```
selection-statement:           if (expression) statement  
                                if (expression) statement else statement  
                                switch (expression) statement
```

Selection statements choose one of a set of statements to execute, based on the evaluation of the expression. The expression is referred to as the controlling expression.

8.3.1 `if` Statement

The controlling expression of an `if` statement must have scalar type.

For both forms of the `if` statement, the first statement is executed if the controlling expression evaluates to nonzero. For the second form, the second statement is executed if the controlling expression evaluates to zero. An `else` clause that follows multiple sequential `else-less if` statements is associated with the most recent `if` statement in the same block (that is, not in an enclosed block).

8.3.2 switch Statement

The controlling expression of a `switch` statement must have integral type. The statement is typically a compound statement, some of whose constituent statements are labeled `case` statements (see Section 8.6, page 99). In addition, at most one labeled default statement can occur in a `switch`. The expression on each `case` label must be an integral constant expression. No two expressions on `case` labels in the same `switch` can evaluate to the same constant.

A compound statement attached to a `switch` can include declarations. Due to the flow of control in a `switch`, however, initialization of identifiers so declared are not performed if these initializers have automatic storage duration.

The integral promotions are performed on the controlling expression, and the constant expression of each `case` statement is converted to the promoted type. Control is transferred to the labeled `case` statement whose expression value matches the value of the controlling expression. If no such match occurs, control is transferred either past the end of the `switch` or to the labeled default statement, if one exists in the `switch`. Execution continues sequentially once control has been transferred. In particular, the flow of control is not altered upon encountering another `case` label. The `switch` statement is exited, however, upon encountering a `break` or `continue` statement (see Section 8.5.3, page 98, and Section 8.5.2, page 98, respectively).

The following is a simple example of a complete `switch` statement:

```
switch (c) {
    case 'o':
        oflag = TRUE;
        break;
    case 'p':
        pflag = TRUE;
        break;
    case 'r':
        rflag = TRUE;
        break;
    default :
        (void) fprintf(stderr,
            "Unknown option\n");
        exit(2);
}
```

8.4 Iteration Statements

Iteration statements execute the attached statement (called the body) repeatedly until the controlling expression evaluates to zero. In the `for` statement, the second expression is the controlling expression. The format is as follows:

```
iteration-statement:           while (expression) statement
                                do statement while (expression) ;
                                for (expressionopt ; expressionopt ; expressionopt)
                                  statement
```

The controlling expression must have scalar type.

The flow of control in an iteration statement can be altered by a `jump` statement (see Section 8.5, page 97).

8.4.1 `while` Statement

The controlling expression of a `while` statement is evaluated before each execution of the body.

8.4.2 `do` Statement

The controlling expression of a `do` statement is evaluated after each execution of the body.

8.4.3 `for` Statement

The `for` statement has the following form:

```
for (expressionopt ;           statement
     expressionopt ;
     expressionopt)
```

The first expression specifies initialization for the loop. The second expression is the controlling expression, which is evaluated before each iteration. The third expression often specifies incrementation. It is evaluated after each iteration.

This statement is equivalent to the following:

```
expression-1;           while (expression-2)
                        {
                        statement
                        expression-3;
                        }
```

One exception exists, however. If a `continue` statement (see Section 8.5.2, page 98) is encountered, *expression-3* of the `for` statement is executed prior to the next iteration.

Any or all of the expressions can be omitted. A missing *expression-2* makes the implied `while` clause equivalent to `while`. Other missing expressions are simply dropped from the previous expansion.

8.5 Jump Statements

Jump statements cause unconditional transfer of control. The syntax is as follows:

```
jump-statement:      goto identifier;
                       continue;
                       break;
                       return expressionopt;
```

8.5.1 goto Statement

Control can be transferred unconditionally by means of a `goto` statement:

```
goto identifier;
```

The identifier must name a label located in the enclosing function. If the label has not yet appeared, it is implicitly declared. (See Section 8.6, page 99, for more information.)

8.5.2 `continue` Statement

The `continue` statement can appear only in the body of an iteration statement. It causes control to pass to the loop-continuation portion of the smallest enclosing `while`, `do`, or `for` statement; that is, to the end of the loop. Consider each of the following statements:

```
while (...)
{
  ..
  contin: ;
}

do {
  ...
  contin: ;
} while (...) ;

for (...) {
  ...
  contin: ;
}
```

A `continue` is equivalent to `gotocontin`. Following the `contin:` is a null statement.

A `goto` statement must not cause a block to be entered by a jump from outside the block to a labeled statement in the block (or an enclosed block) if that block contains the declaration of a variably modified object or variably modified typedef name.

8.5.3 `break` Statement

The `break` statement can appear only in the body of an iteration statement or code attached to a `switch` statement. It transfers control to the statement immediately following the smallest enclosing iteration or `switch` statement, terminating its execution.

8.5.4 `return` Statement

A function returns to its caller by means of the `return` statement. The value of the expression is returned to the caller (after conversion to the declared type of

the function), as the value of the function call expression. The `return` statement cannot have an expression if the type of the current function is `void`.

If the end of a function is reached before the execution of an explicit `return`, an implicit `return` (with no expression) is executed. If the value of the function call expression is used when none is returned, the behavior is undefined.

8.6 Labeled Statements

Labeled statements have the following syntax:

```
labeled-statement:           identifier : statement
                               case constant-expression : statement
                               default : statement
```

A `case` or `default` label can appear only on statements that are part of a `switch`.

Any statement can have a label attached as a simple identifier. The scope of such a label is the current function. Thus, labels must be unique within a function. In traditional C, identifiers used as labels and in object declarations share a name space. Thus, use of an identifier as a label hides any declaration of that identifier in an enclosing scope. In ANSI C, identifiers used as labels are placed in a different name space from all other identifiers and do not conflict. Therefore, the following code fragment is legal in ANSI C but not in traditional C:

```
{
    int foo;
    foo = 1;
    ...
    goto foo;
    ...
    foo: ;
}
```


External Definitions [9]

A C program consists of a sequence of external definitions. An external declaration becomes an external definition when it reserves storage for the object or function indicated. Within the entire program, all external declarations of the same identifier with external linkage refer to the same object or function. Within a particular translation unit, all external declarations of the same identifier with internal linkage refer to the same object or function. External declarations have the following syntax:

external declaration: *function-definition*
 declaration

The syntax for external definitions that are not functions is the same as the syntax for the corresponding external declarations. The syntax for the corresponding external function definition differs from that of the declaration, because the definition includes the code for the function itself.

9.1 External Function Definitions

Function definitions have the following form:

function-definition: *declaration-specifiers*_{opt} *declarator* *declaration-list*_{opt}
 compound statement

The form of a declarator used for a function definition can be as follows:

<pre>pointer_{opt} direct-declarator (parameter-type-list_{opt}) pointer_{opt} direct-declarator (identifier-list_{opt})</pre>
--

In this syntax, the simplest instance of a direct-declarator is an identifier. (For the exact syntax, see Section 7.7, page 78.)

The only storage-class specifiers allowed in a function definition are `extern` and `static`.

If the function declarator has a parameter type list (see Section 7.7, page 78), it is in function prototype form (as discussed in Section 7.7.4, page 82), and the function definition cannot have a declaration list. Otherwise, the function declarator has a possibly empty identifier list, and the declaration list declares

the types of the formal parameters. `register` is the only storage class specifier permitted in declarations that are in the declaration list. Any identifiers in the identifier list of the function declarator that do not have their types specified in the declaration list are assumed to have type `int`.

Each parameter has block scope and automatic storage duration. ANSI C and traditional C place parameters in different blocks. See Section 4.1.1, page 31, for details. Each parameter is also an lvalue, but because function calls in C are by value, the modification of a parameter of arithmetic type cannot affect the corresponding argument. Pointer parameters, while unmodifiable for this reason, can be used to modify the objects to which they point.

Argument promotion rules are discussed in Section 6.3.2, page 52.

The type of a function must be either `void` or an object type that is not an array.

9.2 External Object Definitions

A declaration of an object with file scope that has either an initializer or static linkage is an external object definition.

In ANSI C, a file-scope object declaration with external linkage that is declared without the storage-class specifier `extern`, and also without an initializer, results in a definition of the object at the end of the translation unit. See the discussion in Section 2.1, page 7, for more information.

Multiprocessing Directives [10]

In addition to the usual interpretation performed by a C/C++ compiler, the multiprocessing C/C++ compiler can process explicit multiprocessing directives to produce code that can run concurrently on multiple processors.

Table 11, page 103, lists the multiprocessing `#pragma` directives to use when processing code in parallel regions. The multiprocessing compiler does not know whether an automatic parallelization tool, you the user, or a combination of the two put the directives in the code. The multiprocessing C/C++ compiler does not check for or warn against data dependencies or other restrictions that have been violated.

See the *MIPSpro C and C++ Pragmas*.

Table 11. Multiprocessing C/C++ Compiler Directives

<code>#pragma</code>	Description
<code>#pragma copyin</code>	Copies the value from the master thread's version of an <code>-xlocal</code> -linked global variable into the slave thread's version.
<code>#pragma critical</code>	Protects access to critical statements.
<code>#pragma enter gate</code>	Indicates the point that all threads must clear before any threads are allowed to pass the corresponding exit gate.
<code>#pragma exit gate</code>	Stops threads from passing this point until all threads have cleared the corresponding enter gate.
<code>#pragma independent</code>	Starts an independent code section that executes in parallel with other code in the parallel region.
<code>#pragma local</code>	Tells the compiler the names of all the variables that must be local to each thread.
<code>#pragma no side effects</code>	Tells the compiler to assume that all of the named functions are safe to execute concurrently.
<code>#pragma one processor</code>	Causes the next statement to be executed on only one processor.

#pragma	Description
#pragma parallel	Marks the start of a parallel region.
#pragma pfor	Marks a for loop to run in parallel.
#pragma set chunksize	Tells the compiler which values to use for chunksize.
#pragma set numthreads	Tells the compiler which values to use for numthreads.
#pragma set schedtype	Tells the compiler which values to use for schedtype.
#pragma shared	Tells the compiler the names of all the variables that the threads must share.
#pragma synchronize	Stops threads until all threads reach here.

10.1 OpenMP C/C++ API Multiprocessing Directives

The MIPSpro C and C++ compilers support OpenMP multiprocessing directives. These directives are based on the OpenMP C/C++ Application Program Interface (API) standard. Programs that use these directives are portable and can be compiled by other compilers that support the OpenMP standard.

To enable recognition of the OpenMP directives, specify `-mp` on the `cc` or `CC` command line.

For more information on how to use these directives, see the *MIPSpro C and C++ Pragmas* manual.

10.2 Using Parallel Regions

To understand how most of the multiprocessing C/C++ compiler directives work, consider the concept of a parallel region. On some systems, a parallel region is a single loop that runs in parallel. However, with the multi-processing C/C++ compiler, a parallel region can include several loops and/or independent code segments that execute in parallel.

A simple parallel region consists of only one work-sharing construct, usually a loop. (A parallel region consisting of only a serial section or independent code is a waste of processing resources.)

A parallel region of code can contain sections that execute sequentially as well as sections that execute concurrently. A single large parallel region has a number of advantages over a series of isolated parallel regions: each isolated region executes a single loop in parallel. At the very least, the single large parallel region can help reduce the overhead associated with moving from serial execution to parallel execution.

Large mixed parallel regions avoid the forced synchronization that occurs at the end of each parallel region. The large mixed parallel region also allows you to use `#pragma` directives that execute independent code sections that run concurrently.

Thus, if a thread finishes its work early, it can go on to execute the next section of code—provided that the next section of code is not dependent on the completion of the previous section. However, when you create parallel regions, you need more sophisticated synchronization methods than you need for isolated parallel loops.

10.2.1 Coding Rules of `#pragma` Directives

The `#pragma` directives are modeled after the Parallel Computing Forum (PCF) directives for parallel FORTRAN. The PCF directives define a broad range of parallel execution modes and provide a framework for defining corresponding C/C++ `#pragma` directives.

The following changes have been made to make the `#pragma` directives more C-like:

- Each `#pragma` directive starts with `#pragma` and follows the conventions of ANSI C for compiler directives. You can use white space before and after the `#`, and you must sometimes use white space to separate the words in a `#pragma` directive, as with C syntax. A line that contains a `#pragma` directive can contain nothing else (code or comments).
- `#pragma` directives apply to only one succeeding statement. If a directive applies to more than one statement, you must make a compound statement. C/C++ syntax lets you use curly braces, `{ }`, to do this. Because of the differences between this syntax and FORTRAN, C/C++ can omit the PCF directives that indicate the end of a range (for example, `END PSECTIONS`).

- The `#pragma pfor` directive replaces the `PARALLEL DO` directive because the `for` statement in C is more loosely defined than the `FORTRAN DO` statement.

To make it easier to use `#pragma` directives, you can put several keywords on a single `#pragma` directive line, or spread the keywords over several lines. In either case, you must put the keywords in the correct order, and each directive must contain an initial keyword. For example, the following two code samples do the same thing:

Example 1:

```
#pragma parallel shared(a,b,c, n) local(i) pfor
for (i=0; i<n; i++) a[i]=b[i]+c[i];
```

Example 2:

```
#pragma parallel
#pragma shared( a )
#pragma shared( b, c, n )
#pragma local( i )
#pragma pfor
    for (i=0; i<n; i++) a[i]=b[i]+c[i];
```

10.3 Parallel Regions

A parallel region consists of a number of work-sharing constructs. The C/C++ compiler supports the following work-sharing constructs:

- A loop executed in parallel
- “Local” code run (identically) by all threads
- An independent code section executed in parallel with the rest of the code in the parallel region
- Code executed by only one thread
- Code run in “protected mode” by all threads

In addition, the C/C++ compiler supports three types of explicit synchronization. To account for data dependencies, it is sometimes necessary for threads to wait for all other threads to complete executing an earlier section of code. Three sets of directives implement this coordination:

```
#pragma critical, #pragma synchronize, and #pragma enter gate
and #pragma exit gate.
```


The parallel region should have a single entry at the top and a single exit at the bottom.

To start a parallel region, use the `#pragma parallel` directive. To mark a `for` loop to run in parallel, use the `#pragma pfor` directive. To start an independent code section that executes in parallel with the rest of the code in the parallel region, use the `#pragma independent`.

When you execute a program, nothing actually runs in parallel until it reaches a parallel region. Then multiple threads begin (or continue, if this is not the first parallel region), and the program runs in parallel mode. When the program exits a parallel region, only a single thread continues (sequential mode) until the program again enters a parallel region and the process repeats.

Figure 1, page 108, shows the execution of a typical parallel program with parts running in sequential mode and other parts running in parallel mode.




Figure 1. Program Execution

10.4 Parallel Reduction Operations in C and C++

A reduction operation applies to an array of values and reduces (combines) the array values into a single value.

Consider the following example:

```
int a[10000];
int i;
int sum = 0;
for (i = 0; i < 10000; i++)
    sum = sum + a[i];
```

The loop computes the cumulative sum of the elements of the array. Because the value of a sum computed in one iteration is used in the next iteration, the loop as written cannot be executed in parallel directly on multiprocessors.

However, you can rewrite the above loop to compute the local sum on each processor by introducing a local variable. This breaks the iteration dependency of sum and the loop can be executed in parallel on multiprocessors. This loop computes the local sum of the elements on each processor, which can subsequently be serially added to yield the final sum.

The multiprocessing C/C++ compiler provides a `reduction` clause as a modifier for a `pfor` statement. Using this clause, the above loop can be parallelized as follows:

```
int a[10000];
int i;
int sum = 0
#pragma parallel shared(a, sum) local(i)
#pragma pfor reduction(sum)
for i=0; i<10000; i++)
    sum = sum + a[i];
```

10.4.1 Restrictions on the Reduction Clause

The following restrictions are imposed on the `reduction` clause:

- You can specify only variables of integer types (`int`, `short`, and so forth) or of floating point types (`float`, `double`, and so forth).
- You can use the `reduction` clause only with the primitive operations plus (+), and times (*), which satisfy the associativity property as illustrated in the following example:

```
a op (b op c) == (a op b) op c.
```

The preceding example that uses a reduction clause has the same semantics as the following code that uses local variables and explicit synchronization. In this code, because *sum* is shared, the computation of the final sum has to be done in a critical region to allow each processor exclusive access to *sum*:

```
int a[10000];
int i;
int sum, localsum;
sum = 0;
#pragma parallel shared(a,sum) local(i,localsum)
{
    localsum = 0;
#pragma pfor iterate(;;)
    for (i = 0; i < 10000; i++) localsum +=a[i];
#pragma critical
    sum = sum + localsum;
}
```

The general case of reduction of a binary operation, *op*, on an array $a_1, a_2, a_3, \dots, a_n$ involves the following computation:

```
a1 op a2 op a3 op ... op an
```

When the various operations are performed in parallel, they can be invoked in any order. In order for the reduction to produce a unique result, the binary operation, **op**, must therefore satisfy the associativity property, as shown below:

```
a op (b op c) == (a op b) op c
```

10.4.2 Performance Considerations

The reduction example in Section 10.4.1, page 109, has the drawback that when the number of processors increases, there is more contention for the lock in the critical region.

The following example uses a shared array to record the result on individual processors. The array entries are *CacheLine* apart to prevent write contention on the cache line (128 bytes in this example). The array permits recording results for up to *NumProcs* processors. Both these variables *CacheLine* and *NumProcs* can be tuned for a specific platform:

```
#define CacheLine 128
int a[10000];
```

```

int i, sum;
int *localsum = malloc(NumProcs * CacheLine);

for (i = 0; i < NumProcs; i++)
    localsum [i] = 0;

sum = 0;
#pragma parallel shared (a, sum, localsum) local (i) local (myid)
{
    myid = mp_my_threadnum();

    #pragma pfor
    for (i = 0; i < 10000; i++)
        localsum [myid] += a [i];
    }
    for (i = 0; i < numprocs; i++)
        sum += localsum[i];

```

The only operation in the critical region is the computation of the final result from the local results on individual processors.

In the case when the reduction applies to an array of integers, the reduction function can be specialized by using an intrinsic operation `__fetch_and_<op>` rather than the more expensive critical region. (See Section 11.4, page 125.)

For example, to add an array of integers, the critical region can be replaced by the following call:

```
__fetch_and_add(&sum, localsum);
```

The intrinsic `__fetch_and_<op>` is defined only for the following operations: `add`, `sub`, `or`, `xor`, `nand`, `mpy`, `min`, and `max`; and for the type integers together with their size and signed variants. Therefore, it cannot be used in the general case.

10.4.3 Reduction on User-Defined Types in C++

In C++ a generalized reduction function can be written for any user-defined binary operator `op` that satisfies the associative property.

10.4.3.1 Reduction Example

The following generic function performs reduction on an array of elements of type `ElemType`, with array indices of type `IndexType`, and a binary operation

`op` that takes two arguments of type `ElemType` and produces a result of type `ElemType`. The type `IndexType` is assumed to have operators `<`, `-`, and `++` defined on it. The use of a function object `plus` is in keeping with the spirit of generic programming as in the Standard Template Library (STL). A function object is preferred over a function pointer because it permits inlining:

```
template <class ElemType, class IndexType, class BinaryOp>
ElemType reduction(IndexType first, IndexType last,
ElemType zero, ElemType ar[],
BinaryOp op) {

ElemType result = zero;
IndexType i;
#pragma parallel shared (result, ar) local (i) byvalue(zero, first, last)
{
    ElemType localresult = zero;

#pragma pfor
{
for (i = first; i < last - first; i++)
localresult = op(localresult,ar[i]);
}

#pragma critical
    result = op(result,localresult);
}

return result;
}
```

With the preceding definition of `reduction`, you can perform the following reduction:

```
adsum = reduction(0,size,0,ad,plus<double>());
acsum = reduction(0,size,czero,ac,plus<Complex>());
```

10.5 Restrictions for the C++ Compiler

This section summarizes some restrictions that are relevant only for the C++ compiler. It also lists some restrictions that result from the interaction between pragmas and C++ semantics.

10.5.1 Restrictions on `pfor`

If you are writing a `pfor` loop for the multiprocessing C++ compiler, the index variable `i` can be declared within the `for` statement using the following:

```
int i = 0;
```

The ANSI C++ standard states that the scope of the index variable declared in a `for` statement extends to the end of the `for` statement, as in this example:

```
#pragma pfor
for (int i = 0, ...) { ... }
```

The MIPSpro 7.2 C++ compiler does not enforce this rule. By default, the scope extends to the end of the enclosing block. The default behavior can be changed by using the command line option `-LANG:ansi-for-init-scope=on` which enforces the ANSI C++ standard.

To avoid future problems, write `for` loops in accordance with the ANSI standard, so a subsequent change in the compiler implementation of the default scope rules does not break your code.

10.5.2 Restrictions on Exception Handling

The following restrictions apply to exception handling by the multiprocessing C++ compiler:

- A `throw` cannot cross an multiprocessing parallel region boundary; it must be caught within the multiprocessing region.

A thread that throws an exception must catch the exception as well. For example, the following program is valid. Each thread throws and catches an exception:

```
extern ``C`` printf(char *,...);
extern ``C`` int mp_my_threadnum();
main() {
    int localmax,n;

    #pragma parallel local (localmax,n)
    {
        localmax = 0;

        try {
            throw 10;
```

```
    }
    /* .... */
    catch (int) {
        printf(`!!!exception caught in process \n`);
        printf(`My thread number is %d\n`,mp_my_threadnum());
    } /* end of try block */
} /* end of parallel region */
}
```

- An attempt to throw (propagate) an exception past the end of a parallel program region results in a runtime abort. All other threads abort.

For example, if the following program is executed, all threads abort:

```
extern `C` printf(char *,...);
void ehfn() {
try {
    throw 10;
}
catch (double) // not a handler for throw 10
{
    printf(`exception caught in process \n`);
}
}

main() {
#pragma parallel
{
    ehfn();
}
}
```

The program aborts even if a handler is present in `main()`, as in the following example:

```
main() {
#pragma parallel
{
    try {
        ehfn();
    }
    catch (...) {};
}
}
```


The reason this program aborts is that the throw propagates past the multiprocessing region.

10.5.3 Scoping Restrictions

The following default scope rules apply for the C++ multiprocessing compiler.

- Class objects or structures that have constructors [that is, non-pods (plain old data structures)] cannot be placed on the local list of `#pragma parallel`.

The following is invalid:

```
class C {
    ....
};

main() {

    C c;
    #pragma parallel local (c) // Class object c cannot be in local list
    {
        ....
    }
}
```

Instead, declaring such objects within the parallel region allows the default rule to be used to indicate that they are local (as the following example illustrates):

```
main() {
    #pragma parallel
    {
        C c;
        ....
    }
}
```

- Structure fields and class object members cannot be placed on the local list. Instead, the entire class object must be made local.
- Values of variables in the local list are not copied into each processor's local variables; instead, initialize locals within the parallel program text. For example,

```
main() {
int i;

i = 0;
#pragma parallel local(i)
{
// Here i is not 0.
// Explicit initialization of i within the parallel region
// is necessary
}
}
```

Multiprocessing Advanced Features [11]

A number of features are provided so that you can override the multiprocessing defaults and customize the parallelism to your particular applications. The following sections provide brief explanations of these features.

11.1 Run-time Library Routines

The Silicon Graphics multiprocessing C and C++ compiler provides the following routines for customizing your program.

11.1.1 `mp_block` and `mp_unblock`

The `mp_block` routine puts the slave threads into a blocked state using the `blockproc` system call. The slave threads stay blocked until a call is made to the `mp_unblock` routine. These routines are useful if the job has bursts of parallelism separated by long stretches of single processing, as with an interactive program. You can block the slave processes so they consume CPU cycles only as needed, thus freeing the machine for other users. The system automatically unblocks the slaves on entering a parallel region if you neglect to do so.

11.1.2 `mp_setup`, `mp_create`, and `mp_destroy`

The `mp_setup`, `mp_create`, and `mp_destroy` subroutine calls create and destroy threads of execution. This can be useful if the job has only one parallel portion or if the parallel parts are widely scattered. When you destroy the extra execution threads, they cannot consume system resources; they must be recreated when needed. Use of these routines is discouraged because they degrade performance; the `mp_block` and `mp_unblock` routines should be used in almost all cases.

`mp_setup` takes no arguments. It creates the default number of processes as defined by previous calls to `mp_set_numthreads`, by the `MP_SET_NUMTHREADS` environment variable, or by the number of CPUs on the current hardware platform. `mp_setup` is called automatically when the first parallel loop is entered to initialize the slave threads.

`mp_create` takes a single integer argument, the total number of execution threads desired. Note that the total number of threads includes the master

thread. Thus, `mp_create(n)` creates one thread less than the value of its argument. `mp_destroy` takes no arguments; it destroys all the slave execution threads, leaving the master untouched.

When the slave threads die, they generate a `SIGCLD` signal. If your program has changed the signal handler to catch `SIGCLD`, it must be prepared to deal with this signal when `mp_destroy` is executed. This signal also occurs when the program exits; `mp_destroy` is called as part of normal cleanup when a parallel job terminates.

11.1.3 `mp_blocktime`

The slave threads spin wait until there is work to do. This makes them immediately available when a parallel region is reached. However, this consumes CPU resources. After enough wait time has passed, the slaves block themselves through `blockproc`. Once the slaves are blocked, it requires a system call to `unblockproc` to activate the slaves again (refer to the `unblockproc(2)` man page for details). This makes the response time much longer when starting up a parallel region.

This trade-off between response time and CPU usage can be adjusted with the `mp_blocktime` call. The `mp_blocktime` routine takes a single integer argument that specifies the number of times to spin before blocking. By default, it is set to 10,000,000; this takes roughly one second. If called with an argument of 0, the slave threads will not block themselves no matter how much time has passed. Explicit calls to `mp_block`, however, will still block the threads.

This automatic blocking is transparent to the user's program; blocked threads are automatically unblocked when a parallel region is reached.

11.1.4 `mp_numthreads`, `mp_suggested_numthreads`, `mp_set_numthreads`

Occasionally, you may want to know how many execution threads are available. The `mp_numthreads` routine is a zero-argument integer function that returns the total number of execution threads for this job. The count includes the master thread. In addition, this routine has the side effect of freezing (for eternity) the number of threads to the returned value, so this routine should be used sparingly. To determine the number of threads without this freeze property, use `mp_suggested_numthreads`.

`mp_suggested_numthreads` takes an unsigned integer and uses the supplied value as a hint about how many threads to use in subsequent parallel regions. It returns the previous value of the number of threads to be employed in

parallel regions. It does not affect currently executing parallel regions, if any. The implementation may ignore this hint depending on factors such as overall system load. This routine may also be called with the value 0, in which case it simply returns the number of threads to be employed in parallel regions.

`mp_set_numthreads` takes a single integer argument. It changes the default number of threads to the specified value. A subsequent call to `mp_setup` will use the specified value rather than the original defaults. If the slave threads have already been created, this call will not change their number. It has an effect only when `mp_setup` is called.

11.1.5 `mp_my_threadnum`

The `mp_my_threadnum` routine is a zero-argument function that allows a thread to differentiate itself while in a parallel region. If there are n execution threads, the function call returns a value between zero and $n - 1$. The master thread is always thread zero. This function can be useful when parallelizing certain kinds of loops. Most of the time the loop index variable can be used for the same purpose. Occasionally, the loop index may not be accessible, as, for example, when an external routine is called from within the parallel loop. This routine provides a mechanism for those cases.

11.1.6 `mp_setlock`, `mp_unsetlock`, `mp_barrier`

The `mp_setlock`, `mp_unsetlock`, and `mp_barrier` zero-argument subroutines provide convenient (although limited) access to the locking and barrier functions provided by `ussetlock`, `usunsetlock`, and `barrier`. These subroutines are convenient because you do not need to initialize them; calls such as `usconfig` and `usinit` are done automatically. The limitation is that there is only one lock and one barrier. For most programs, this amount is sufficient. If your program requires more complex or flexible locking facilities, use the `ussetlock` family of subroutines directly.

11.1.7 `mp_set_slave_stacksize`

The `mp_set_slave_stacksize` routine sets the stack size (in bytes) to be used by the slave processes when they are created (using `sprobsp`). The default size is 16 MB. Slave processes only allocate their local data onto their stack, shared data (even if allocated on the master's stack) is not counted.

11.2 Run-time Environment Variables

The Silicon Graphics multiprocessing C and C++ compiler provides the following environment variables that you can use to customize your program.

11.2.1 `MP_SET_NUMTHREADS`, `MP_BLOCKTIME`, `MP_SETUP`

The `MP_SET_NUMTHREADS`, `MP_BLOCKTIME`, and `MP_SETUP` environment variables act as an implicit call to the corresponding routine(s) of the same name at program start-up time.

For example, the following `cs` command causes the program to create two threads regardless of the number of CPUs actually on the machine, as does the source statement below it:

`cs` command:

```
% setenv MP_SET_NUMTHREADS 2
```

Source statement:

```
mp_set_numthreads (2)
```

Similarly, the following `sh` commands prevent the slave threads from autoblocking, as does the source statement:

`sh` commands:

```
% set MP_BLOCKTIME 0
% export MP_BLOCKTIME
```

Source statement:

```
mp_blocktime (0);
```

For compatibility with older releases, the environment variable `NUM_THREADS` is supported as a synonym for `MP_SET_NUMTHREADS`.

To help support networks with several multiprocessors and several CPUs, the environment variable `MP_SET_NUMTHREADS` also accepts an expression involving integers `+`, `-`, `min`, `max`, and the special symbol `"all"`, which stands for the number of CPUs on the current machine. For example, the following command selects the number of threads to be two fewer than the total number of CPUs (but always at least one):

```
% setenv MP_SET_NUMTHREADS max(1,all-2)
```

11.2.2 MP_SUGNUMTHD, MP_SUGNUMTHD_MIN, MP_SUGNUMTHD_MAX, MP_SUGNUMTHD_VERBOSE

In an environment with long running jobs and varying workloads, it may be preferable to vary the number of threads during execution of some jobs.

Setting `MP_SUGNUMTHD` causes the run-time library to create an additional, asynchronous process that periodically wakes up and monitors the system load. When idle processors exist, this process increases the number of threads, up to a maximum of `MP_SET_NUMTHREADS`. When the system load increases, it decreases the number of threads, possibly to as few as 1. When `MP_SUGNUMTHD` has no value, this feature is disabled and multithreading works as before.

The environment variables `MP_SUGNUMTHD_MIN` and `MP_SUGNUMTHD_MAX` are used to limit this feature as desired. When `MP_SUGNUMTHD_MIN` is set to an integer value between 1 and `MP_SET_NUMTHREADS`, the process will not decrease the number of threads below that value.

When `MP_SUGNUMTHD_MAX` is set to an integer value between the minimum number of threads and `MP_SET_NUMTHREADS`, the process will not increase the number of threads above that value.

If you set any value in the environment variable `MP_SUGNUMTHD_VERBOSE`, informational messages are written to `stderr` whenever the process changes the number of threads in use.

Calls to `mp_numthreads` and `mp_set_numthreads` are taken as a sign that the application depends on the number of threads in use. The number in use is frozen upon either of these calls; and if `MP_SUGNUMTHD_VERBOSE` is set, a message to that effect is written to `stderr`.

11.2.3 MP_SCHEDTYPE, CHUNK

These environment variables specify the type of scheduling to use on `for` loops that have their scheduling type set to `RUNTIME`. For example, the following `cs`h commands cause loops with the `RUNTIME` scheduling type to be executed as interleaved loops with a chunk size of 4:

```
% setenv MP_SCHEDTYPE INTERLEAVE
% setenv CHUNK 4
```

The defaults are the same as on the `#pragma pfor` directive; if neither variable is set, `SIMPLE` scheduling is assumed. If `MP_SCHEDTYPE` is set, but `CHUNK` is not set, a `CHUNK` of 1 is assumed. If `CHUNK` is set, but `MP_SCHEDTYPE` is not, `DYNAMIC` scheduling is assumed.

11.2.4 MP_SLAVE_STACKSIZE

The stack size of slave processes can be controlled through the environment variable `MP_SLAVE_STACKSIZE`, which may be set to the desired stacksize in bytes. The default value is 16 MB (4 MB for more than 64 threads).

11.2.5 MPC_GANG

`MPC_GANG` specifies gang scheduling. Set `MPC_GANG` to `ON` to enable gang scheduling. To disable gang scheduling, set `MPC_GANG` to `OFF`.

11.3 Communicating Between Threads Through Thread Local Data

The routines described in this section allow you to perform explicit communication between threads within their multiprocessing C program. These communication mechanisms are similar to message-passing, one-sided-communication, or `shmem`, and may be desirable for reasons of performance and/or style.

The operations allow a thread to fetch from (`get`) or send to (`put`) data belonging to other threads. Therefore, these operations can be performed only on data that has been declared to be `-Xlocal` (that is, each thread has its own private copy of that data; see the `ld(1)` man page for details on `Xlocal`). A `get` operation requires that the *source* parameter point to `Xlocal` data, while a `put` operation requires that the *target* parameter point to `Xlocal` data.

The following routines are available as part of the Message Passing Toolkit (MPT) and are similar to the original `shmem` routines (see the `shmem` reference page), but are prefixed by `mp_`:

```
void mp_shmem_get32 (int *target,  
int *source,  
int length,  
int source_thread)
```

```
void mp_shmem_put32 (int *target,  
int *source,  
int length,  
int target_thread)
```

```
void mp_shmem_iget32 (int *target,  
int *source,  
int target_inc,
```



```
int source_inc ,
int length ,
int source_thread )

void mp_shmem_iput32 (int *target ,
int *source ,
int target_inc ,
int source_inc ,
int length ,
int target_thread )

void mp_shmem_get64(long long *target ,
long long *source ,
int length ,
int source_thread )

void mp_shmem_put64 (long long *target ,
long long *source ,
int length ,
int target_thread )

void mp_shmem_iget64 (long long *target ,
long long *source ,
int target_inc ,
int source_inc ,
int length ,
int source_thread )

void mp_shmem_iput64 (long long *target ,
long long *source ,
int target_inc ,
int source_inc ,
int length ,
int target_thread )
```

The following rules apply to the preceding listed routines:

- Both *source* and *target* are pointers to 32-bit quantities for the 32-bit versions, and to 64-bit quantities for the 64-bit versions of the calls. The actual type of the data is not important, because the routines perform a bit-wise copy.
- For a *put* operation, the *target* must be *xlocal*. For a *get* operation, the *source* must be *xlocal*.

- *length* specifies the number of elements to be copied, in units of 32 or 64-bit elements, as appropriate.
- *source_thread* and *target_thread* specify the thread-number of the remote processing element (PE).
- A get operation copies from the remote PE. A put operation copies to the remote PE.
- *target_inc* and *source_inc* are specified for the strided *iget* and *iput* operations. They specify the increment (in units of 32-bit or 64-bit elements) for source and target when performing the data transfer. The number of elements copied during a strided *put* or *get* operation is still determined by *length*.

Note: Call these routines only after the threads have been created (typically, the first `pthread_for`/parallel region). Performing these operations while the program is still serial leads to a run-time error because each thread's copy has not yet been created.

In the example below, compiling with `-Wl,-Xlocal`, `myvars` ensures that each thread has a private copy of `x` and `y`.

```
struct {
    int x;
    double y[100];
} myvars;
```

The following example copies the value of `x` on thread 3 into the private copy of `x` for the current thread.

```
mp_shmem_get32 (&x, &x, 1, 3)
```

The next example copies the value of `localvar` into the thread 5 copy of `x`.

```
mp_shmem_put32 (&x, &localvar, 1, 5)
```

The example below fetches values from the thread 7 copy of array `y` into `localarray`.

```
mp_shmem_get64 (&localarray, &y, 100, 7)
```

The next example copies the value of every other element of `localarray` into the thread 9 copy of `y`.

```
mp_shmem_iput64 (&y, &localarray, 2, 2, 50, 9)
```

11.4 Synchronization Intrinsic

The intrinsics described in this section provide a variety of primitive synchronization operations. Besides performing the particular synchronization operation, each of these intrinsics has two key properties:

- The function performed is guaranteed to be atomic (typically achieved by implementing the operation using a sequence of load-linked and/or store-conditional instructions in a loop).
- Associated with each intrinsic are certain memory barrier properties that restrict the movement of memory references to visible data across the intrinsic operation (by either the compiler or the processor).

A visible memory reference is a reference to a data object potentially accessible by another thread executing in the same shared address space. A visible data object can be one of the following:

- C/C++ global data
- Data declared `extern`
- Volatile data
- Static data (either file-scope or function-scope)
- Data accessible via function parameters
- Automatic data (local-scope) that has had its address taken and assigned to some visible object (recursively)

The memory barrier semantics of an intrinsic can be one of the following three types:

acquire barrier	Disallows the movement of memory references to visible data from after the intrinsic (in program order) to before the intrinsic. (This behavior is desirable at lock-acquire operations.)
release barrier	Disallows the movement of memory references to visible data from before the intrinsic (in program order) to after the intrinsic. (This behavior is desirable at lock-release operations.)
full barrier	Disallows the movement of memory references to visible data past the intrinsic (in

either direction), and is thus both an acquire and a release barrier. A barrier restricts only the movement of memory references to visible data across the intrinsic operation: between synchronization operations (or in their absence), memory references to visible data may be freely reordered subject to the usual data-dependence constraints.

By default, it is assumed that a memory barrier applies to all visible data. If you know the precise set of data objects that must be restricted by the memory barrier, you can specify the set of data objects as additional arguments to the intrinsic. In this case, the memory barrier restricts the movement of memory references to the specified list of data objects only, possibly resulting in better performance. The specified data objects must be simple variables and cannot be expressions (for example, `&p` and `*p` are disallowed).



Caution: Conditional execution of a synchronization intrinsic (such as within an `if` or a `while` statement) does not prevent the movement of memory references to visible data past the overall `if` or `while` construct.

11.4.1 Atomic fetch-and-op Operations

The fetch-and-op operations are as follows:

```
<type> __fetch_and_add (<type>* ptr, <type> value, ...)  
<type> __fetch_and_sub (<type>* ptr, <type> value, ...)  
<type> __fetch_and_or (<type>* ptr, <type> value, ...)  
<type> __fetch_and_and (<type>* ptr, <type> value, ...)  
<type> __fetch_and_xor (<type>* ptr, <type> value, ...)  
<type> __fetch_and_nand(<type>* ptr, <type> value, ...)  
<type> __fetch_and_mpy (<type>* ptr, <type> value, ...)  
<type> __fetch_and_min (<type>* ptr, <type> value, ...)  
<type> __fetch_and_max (<type>* ptr, <type> value, ...)
```

`<type>` can be any of the following:

```
int  
long  
long long  
unsigned int  
unsigned long  
unsigned long long
```

The ellipses (. . .) refer to an optional list of variables protected by the memory barrier.

Each of these operations behaves as follows:

- Atomically performs the specified operation with the given value on **ptr*, and returns the old value of **ptr*.

```
{tmp = *ptr; *ptr <op>= value; return tmp;}
```

- Full barrier

11.4.2 Atomic op-and-fetch Operations

The op-and-fetch operations are as follows:

```
<type> __add_and_fetch (<type>* ptr, <type> value, ...)
<type> __sub_and_fetch (<type>* ptr, <type> value, ...)
<type> __or_and_fetch (<type>* ptr, <type> value, ...)
<type> __and_and_fetch (<type>* ptr, <type> value, ...)
<type> __xor_and_fetch (<type>* ptr, <type> value, ...)
<type> __nand_and_fetch(<type>* ptr, <type> value, ...)
<type> __mpy_and_fetch (<type>* ptr, <type> value, ...)
<type> __min_and_fetch (<type>* ptr, <type> value, ...)
<type> __max_and_fetch (<type>* ptr, <type> value, ...)
```

<type> can be any of the following:

```
int
long
long long
unsigned int
unsigned long
unsigned long long
```

Each of these operations behaves as follows:

- Atomically performs the specified operation with the given value on **ptr*, and returns the new value of **ptr*.

```
{*ptr <op>= value; return *ptr;}
```

- Full barrier

11.4.3 Atomic compare-and-swap Operation

The compare-and-swap operation is as follows:

```
int __compare_and_swap (<type>* ptr, <type> oldvalue, <type> newvalue, ...)
```

<type> can be one of the following:

```
int
long
long long
unsigned int
unsigned long
unsigned long long
```

This operation behaves as follows:

- Atomically compares **ptr* to *oldvalue*. If equal, it stores the new value and returns 1, otherwise it returns 0.

```
if (*ptr != oldvalue) return 0;
else {
    *ptr = newvalue;
    return 1;
}
```

- Full barrier

11.4.4 Atomic synchronize Operation

The synchronize operation is as follows:

```
__synchronize (...)
```

The ellipses (...) refer to an optional list of variables protected by the memory barrier.

This operation behaves as follows:

- Issues a sync operation
- Full barrier

11.4.5 Atomic lock and unlock Operations

11.4.5.1 Atomic lock-test-and-set Operation

The lock-test-and-set operation is as follows:

```
<type> __lock_test_and_set (<type>* ptr, <type> value, ...)
```

<type> can be any of the following:

```
int
long
long long
unsigned int
unsigned long
unsigned long long
```

This operation behaves as follows:

- Atomically stores the supplied value in **ptr* and returns the old value of **ptr*

```
{tmp = *ptr; *ptr = value; return tmp;}
```
- Acquire barrier

11.4.5.2 Atomic lock-release Operation

The lock_release operation is as follows:

```
void __lock_release (<type>* ptr, ...)
```

<type> can be one of the following:

```
int
long
long long
unsigned int
unsigned long
unsigned long long
```

This operation behaves as follows:

- Issues sync then sets **ptr* to 0 and flushes it from the register

```
{*ptr = 0}
```
- Release barrier

11.4.6 Example of Implementing a Pure Spin-Wait Lock

The following example shows implementation of a spin-wait lock:

```
int lockvar = 0;
while (__lock_test_and_set (&lockvar, 1) != 0); /* acquire the lock */
    ... read and update shared variables ...
__lock_release (&lockvar); /* release the lock */
```

The memory barrier semantics of the intrinsics guarantee that no memory reference to visible data is moved out of the above critical section, either ahead of the lock-acquire or past the lock-release.

Note: Pure spin-wait locks can perform poorly under heavy contention.

If the data structures protected by the lock are known precisely (for example, *x*, *y*, and *z* in the example below), then those data structures can be precisely identified as follows:

```
int lockvar = 0;
while (__lock_test_and_set (&lockvar, 1, x, y, z) != 0);
    ... read/modify the variables x, y, and z ...
__lock_release (&lockvar, x, y, z);
```


Parallel Programming on Origin Servers [12]


This chapter describes the support provided for writing parallel programs on the Origin200 and Origin2000 servers. It assumes that you are familiar with basic parallel constructs.

You can find additional information on parallel programming in *Topics in IRIX Programming*.

Note: The multiprocessing features described in this chapter require support from the MP run-time library (`libmp`). IRIX operating system versions 6.3 (and above) are automatically shipped with this new library. If you want to access these features on a system running IRIX 6.2, then contact your local Silicon Graphics service provider or Silicon Graphics Customer Support (1-800-800-4744) for `libmp`.

12.1 Performance Tuning of Parallel Programs on an Origin2000 Server

An Origin2000 server provides cache-coherent, shared memory in the hardware. Memory is physically distributed across processors. Consequently, references to locations in the remote memory of another processor take substantially longer (by a factor of two or more) to complete than references to locations in local memory. This can severely affect the performance of programs that suffer from a large number of cache misses. Figure 2, page 132, shows a simplified version of the Origin2000 memory hierarchy.



a11999

Figure 2. Origin2000 Memory Hierarchy

12.1.1 Improving Program Performance

To obtain good performance in such programs, it is important to schedule computation and distribute data across the underlying processors and memory modules, so that most cache misses are satisfied from local rather than remote memory. The primary goal of the programming support, therefore, is to enable user control over data placement and computation scheduling.

Cache behavior continues to be the largest single factor affecting performance, and programs with good cache behavior usually have little need for explicit data placement. In programs with high cache misses, if the misses correspond to true data communication between processors, then data placement is unlikely to help. In these cases, it may be necessary to redesign the algorithm to reduce inter-processor communication. Figure 3, page 134, shows this scenario.

If the misses are to data referenced primarily by a single processor, then data placement may be able to convert remote references to local references, thereby reducing the latency of the miss. The possible options for data placement are automatic page migration or explicit data distribution, either regular or reshaped. The differences between these choices are shown in Figure 3, page 134.

Automatic page migration requires no user intervention and is based on the run-time cache miss behavior of the program. It can therefore adjust to dynamic changes in the reference patterns. However, the page migration heuristics are deliberately conservative, and may be slow to react to changes in the references patterns. They are also limited to performing page-level allocation of data.

Regular data distribution (performing just page-level placement of the array) is also limited to page-level allocation, but is useful when the page migration heuristics are slow and the desired distribution is known to the programmer.

Finally, reshaped data distribution changes the layout of the array, thereby overcoming the page-level allocation constraints; however, it is useful only if a data structure has the same (static) distribution for the duration of the program. Given these differences, it may be necessary to use each of these options for different data structures in the same program.

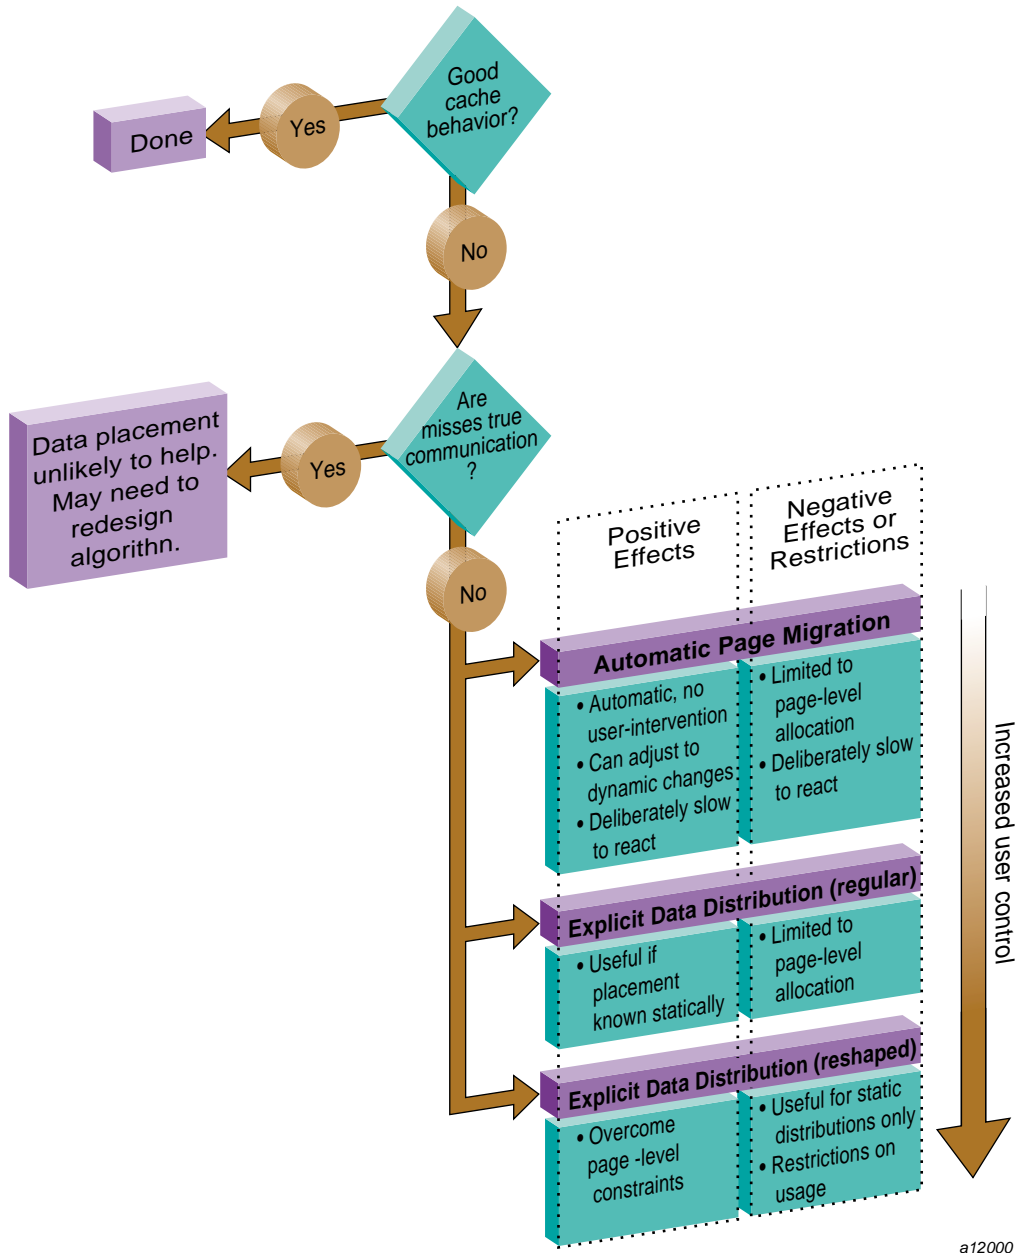


Figure 3. Cache Behavior and Solutions

12.2 Types of Data Distribution

This section describes the two types of data distribution: regular and reshaped.

12.2.1 Regular Data Distribution

The regular data distribution directives try to achieve the desired distribution solely by influencing the mapping of virtual addresses to physical pages without affecting the layout of the data structure. Because the granularity of data allocation is a physical page (at least 16 KB), the achieved distribution is limited by the underlying page granularity. However, the advantages are that regular data distribution directives can be added to an existing program without any restrictions, and can be used for affinity scheduling (see Section 12.3, page 140).

12.2.2 Data Distribution With Reshaping

Similar to regular data distribution, the `#pragma distribute_reshape` directive specifies the desired distribution of an array. In addition, however, the `#pragma distribute_reshape` directive declares that the program makes no assumptions about the storage layout of that array. The compiler performs aggressive optimizations for reshaped arrays that violate standard layout assumptions but guarantee the desired data distribution for that array.

12.2.3 Implementation of Reshaped Arrays

The compiler transforms a reshaped array into a pointer to a “processor array.” The processor array has one element per processor, with the element pointing to the portion of the array local to the corresponding processor.

Figure 4, page 136, shows the effect of the `#pragma distribute_reshape` directive with a `block` distribution on a one-dimensional array. N is the size of the array dimension, P is the number of processors, and B is the block size on each processor.

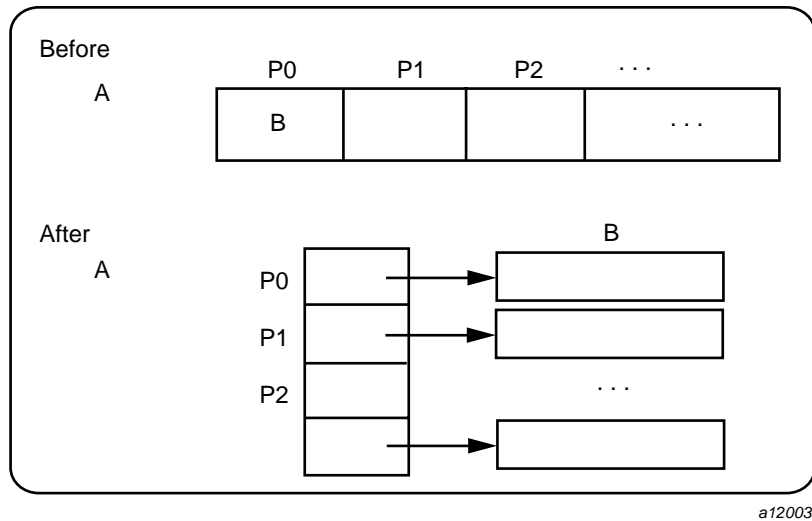


Figure 4. Implementation of block Distribution

With this implementation, an array reference $A[i]$ is transformed into a two-dimensional reference $A[i/B][i\%B]$, where B , the size of each block, is equal to N/P , rounded up to the nearest integer value ($\text{ceiling}(N/P)$). $A[i/B]$, therefore, points to a processor's local portion of the array, and $A[i/B][i\%B]$ refers to a specific element within the local processor's portion.

A cyclic distribution with a chunk size of 1 is implemented as shown in Figure 5, page 137.

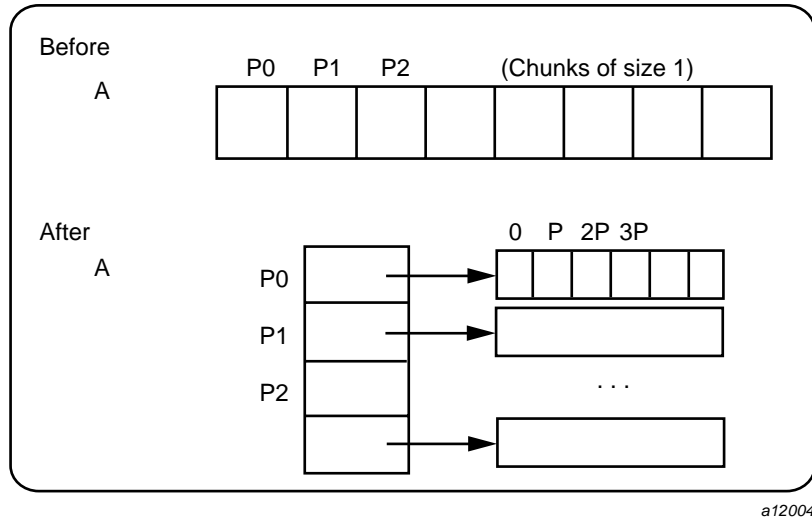


Figure 5. Implementation of cyclic(1) Distribution

An array reference, $A[i]$, is transformed to $A[i \% P][i / P]$, where P is the number of threads in that distributed dimension.

Finally, a `cyclic` distribution with a chunk size that is either a constant greater than 1 or a run-time value (also called `block-cyclic`) is implemented as shown in Figure 6, page 138.

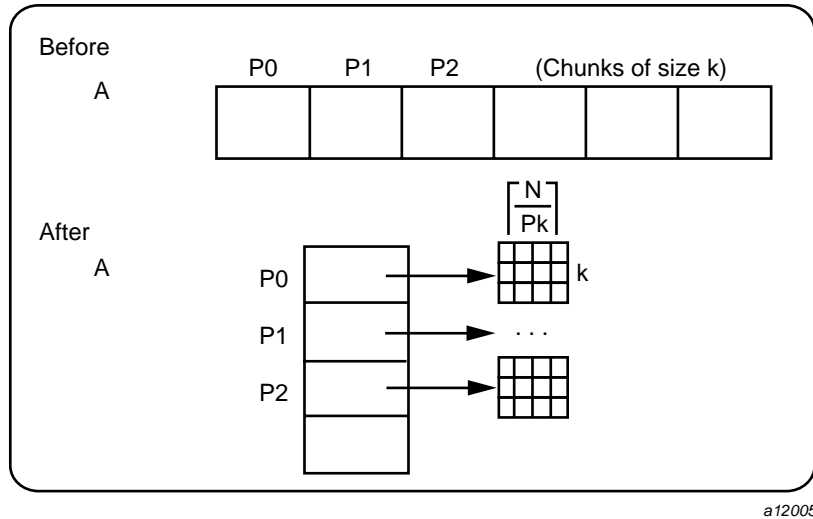


Figure 6. Implementation of block-cyclic Distribution

An array reference, $A[i]$, is transformed to the three-dimensional reference $A[(i/k)\%P][i/(P*k)][i\%k]$, where P is the total number of threads in that dimension, and k is the chunk size.

The compiler tries to optimize these divide and modulo operations out of inner loops through aggressive loop transformations such as blocking and peeling.

12.2.4 Regular Versus Reshaped Data Distribution

In summary, consider the differences between regular and reshaped data distribution. The advantage of regular distributions is that they do not impose any restrictions on the distributed arrays and can be freely applied in existing codes. Furthermore, they work well for distributions where page granularity is not a problem. For example, consider a `block` distribution of the columns of a two-dimensional array of size $A[r][c]$ (row-major layout) and distribution $[block][*]$. If the size of each processor's portion, $\text{ceiling}(r/P) \times c \times \text{element_size}$, is significantly greater than the page size (16KB on the Origin2000 server), then regular data distribution should be effective in placing the data in the desired fashion.

However, regular data distribution is limited by page-granularity considerations. For instance, consider a $[block][block]$ distribution of a two-dimensional array where the size of a row is much smaller than a page.

Each physical page is likely to contain data belonging to multiple processors, making the data-distribution quite ineffective. (Data distribution may still be desirable from affinity-scheduling considerations.)

Reshaped data distribution addresses the problems of regular distributions by changing the layout of the array in memory so as to guarantee the desired distribution. However, because the array no longer conforms to standard storage layout, there are restrictions on the usage of reshaped arrays.

Given both types of data distribution, you can choose between the two based on the characteristics of the particular array in an application.

12.2.5 Choosing Between Multiple Options

For a given data structure in the program, you can choose a data distribution option based on the following criteria:

- If the program repeatedly references the data structure and benefits from reuse in the cache, then data placement is not needed.
- If the program incurs a large number of cache misses on the data structure, then you should identify the desired distribution in the array dimensions (such as `block` or `cyclic`) based on the desired parallelism in the program. For example, the following code suggests a distribution of `A[block][*]`:

```
#pragma pfor
for(i = 2; i <= n; i++)
for(j = 2; j <= n; j++)
A[i][j] = 3*i + 4*j + A[i][j-1];
```

Whereas the next code segment suggests a distribution of `A[*][block]`:

```
for(i = 2; i <= n; i++)
#pragma pfor
for(j = 2; j <= n; j++)
A[i][j] = 3*i + 4*j + A[i][j-1];
```

- Having identified the desired distribution, you can select either regular or reshaped distribution based on the size of an individual processor's portion of the distributed array. Regular distribution is useful only if each processor's portion is substantially larger than the page-size in the underlying system (16 KB on the Origin2000 server). Otherwise, regular distribution is unlikely to be useful, and you should use `distribute_reshape`, where the compiler changes the layout of the array to overcome page-level constraints.

For example, consider the following code:

```
double A[m][n];  
#pragma distribute A[*][block]
```

In this example, each processor's portion is approximately $(m*n/P)$ elements ($8*(m*n/P)$ bytes), where P is the number of processors. Due to the row-major layout of the array, however, each contiguous piece is only $8*(n/P)$ bytes. If n is 1000,000 then each contiguous piece is likely to exceed a page and regular distribution is sufficient. If instead n is small, say 10,000, then `distribute_reshape` is required to obtain the desired distribution.

In contrast, consider the following distribution:

```
#pragma distribute A[block][*]
```

In this example, the size of each processor's portion is a single contiguous piece of $(m*n)/P$ elements ($8*(m*n)/P$ bytes). So if m is 100, for instance, regular distribution may be sufficient even if n is only 10,000.

As this example illustrates, distributing the outer dimensions of an array increases the size of an individual processor's portion (favoring regular distribution), whereas distributing the inner dimensions is more likely to require reshaped distribution.

Finally, the IRIX operating system on Origin2000 follows a default "first-touch" page-allocation policy; that is, each page is allocated from the local memory of the processor that incurs a page-fault on that page. Therefore, in programs where the array is initialized (and consequently first referenced) in parallel, even a regular distribution directive may not be necessary, because the underlying pages are allocated from the desired memory location automatically due to the first-touch policy.

12.2.6 Explicit Placement of Data

For irregular data structures, you can explicitly place data in the physical memory of a particular processor using the `#pragma page_place` directive.

12.3 Affinity Scheduling

The goal of affinity scheduling is to control the mapping of iterations of a parallel loop for execution onto the underlying threads. Specify affinity scheduling with an additional clause to a `#pragma pfor` directive. An affinity clause, if supplied, overrides the `schedtype` clause.

12.3.1 Data and Thread Affinity

12.3.1.1 Data Affinity

The syntax of the `#pragma pfor` directive with the `affinity` clause is as follows:

```
#pragma pfor affinity(idx) = data(array(expr))
```

idx is the loop-index variable; *array* is the distributed array; and *expr* indicates an element owned by the processor on which you want this iteration executed.

The following code shows an example of data affinity:

```
#pragma distribute A[block]
#pragma parallel shared (A, a, b) local (i)
#pragma pfor affinity(i) = data(A[a*i + b])
for (i = 0; i < n; i++)
    A[a*i + b] = 0;
```

The multiplier for the loop index variable (*a*) and the constant term (*b*) must both be literal constants, with *a* greater than zero.

The effect of this clause is to distribute the iterations of the parallel loop to match the data distribution specified for the array *A*, such that iteration *i* is executed on the processor that owns element $A[a*i + b]$, based on the distribution for *A*. The iterations are scheduled based on the specified distribution, and are not affected by the actual underlying data-distribution (which may differ at page boundaries, for example).

In the case of a multi-dimensional array, affinity is provided for the dimension that contains the loop-index variable. The loop-index variable cannot appear in more than one dimension in an affinity directive. For example,

```
#pragma distribute A[block][cyclic(1)]
#pragma parallel shared (A, n) local (i, j)
#pragma pfor
#pragma affinity (i) = data(A[i + 3, j])
for (i = 0; i < n; i++)
for (j = 0; j < n; j++)
    A[i + 3, j] = A[i + 3, j-1];
```

In this example, the loop is scheduled based on the block distribution of the first dimension. See Chapter 12, page 131, for more information about distribution directives.

Data affinity for loops with non-unit stride can sometimes result in non-linear affinity expressions. In such situations the compiler issues a warning, ignores the affinity clause, and defaults to simple scheduling.

12.3.2 Data Affinity for Redistributed Arrays

By default, the compiler assumes that a distributed array is not dynamically redistributed, and directly schedules a parallel loop for the specified data affinity. In contrast, a redistributed array can have multiple possible distributions, and data affinity for a redistributed array must be implemented in the run-time system based on the particular distribution.

However, the compiler does not know whether or not an array is redistributed, because the array may be redistributed in another function (possibly even in another file). Therefore, you must explicitly specify the `#pragma dynamic` declaration for redistributed arrays. This directive is required only in those functions that contain a `parallel` loop with data affinity for that array. This informs the compiler that the array can be dynamically redistributed. Data affinity for such arrays is implemented through a run-time lookup.

12.3.2.1 Data Affinity for a Formal Parameter

You can supply a `#pragma distribute` directive on a formal parameter, thereby specifying the distribution on the incoming actual parameter. If different calls to the subroutine have parameters with different distributions, then you can omit the `#pragma distribute` directive on the formal parameter; data affinity loops in that subroutine are automatically implemented through a run-time lookup of the distribution. (This is permissible only for regular data distribution. For reshaped array parameters, the distribution must be fully specified on the formal parameter.)

12.3.3 Data Affinity and the `#pragma parallel nest` Clause

This section discusses how the `nest` clause interacts with the `affinity` clause when the program has reshaped arrays.

When you combine a `nest` clause and an `affinity` clause, the default scheduling is `simple`, except when the program has reshaped arrays and is compiled `-O3`. In that case, the default is to use data affinity scheduling for the most frequently accessed reshaped array in the loop (chosen heuristically by the compiler). To obtain `simple` scheduling even at `-O3`, you can explicitly specify the `schedtype` on the parallel loop.

The following example illustrates a nested `pfor` with an `affinity` clause:

```
#pfor nest(i, j) affinity(i, j) = data(A[i][j])
for (i = 2; i < n; i++)
for (j = 2; j < m; j++)
A[i][j] = A[i][j] + i * j;
```

12.4 Directives for Performance Tuning on Origin2000

The programming support consists of extensions to the existing C `#pragma` directives. Table 12, page 143, summarizes the new directives. Like the other C directives, these new directives are ignored except under multiprocessor compilation. For detailed information about these directives, see the *MIPSpro C and C++ Pragmas*

Table 12. Loop Nest Optimization Directives Specific to the Origin2000 Server

<code>#pragma</code>	Short Description
<code>#pragma distribute</code>	Specifies data distribution.
<code>#pragma redistribute</code>	Specifies dynamic redistribution of data.
<code>#pragma distribute_reshape</code>	Specifies data distribution with reshaping.
<code>#pragma dynamic</code>	Tells the compiler that the specified array may be redistributed in the program.
<code>#pragma page_place</code>	Allows the explicit placement of data.
<code>#pragma pfor</code>	<code>affinity</code> clause allows data-affinity or thread-affinity scheduling; <code>nest</code> clause exploits nested concurrency.

12.5 Query Intrinsic for Distributed Arrays

You can use the following set of intrinsics to obtain information about an individual dimension of a distributed array. C array dimensions are numbered

starting at 0. All routines work with 64-bit integers as shown below, and return -1 in case of an error (except `dsm_this_startingindex`, where -1 may be a legal return value).

- `dsm_numthreads`: Called with a distributed array and a dimension number, and returns the number of threads in that dimension:

```
extern long long dsm_numthreads (void* array, long long dim)
```

- `dsm_chunksize`: Returns the chunk size (ignoring partial chunks) in the given dimension for each of `block`, `cyclic(..)`, and `star` distributions:

```
extern long long dsm_chunksize (void* array, long long dim)
```

- `dsm_this_chunksize`: Returns the chunk size for the chunk containing the given index value for each of `block`, `cyclic(..)`, and `star`. This value may be different from `dsm_chunksize` due to edge effects that may lead to a partial chunk.

```
extern long long dsm_this_chunksize (void* array, long long dim,  
long long index)
```

- `dsm_rem_chunksize`: Returns the remaining chunk size from index to the end of the current chunk, inclusive of each end point. Essentially it is the distance from index to the end of that contiguous block, inclusive.

```
extern long long dsm_rem_chunksize (void* array, long long dim,  
long long index)
```

- `dsm_this_startingindex`: Returns the starting index value of the chunk containing the supplied index:

```
extern long long dsm_this_startingindex (void* array, long long dim,  
long long index)
```

- `dsm_numchunks`: Returns the number of chunks (including partial chunks) in given dim for each of `block`, `cyclic(..)`, and `star` distributions:

```
extern long long dsm_numchunks (void* array, long long dim)
```

- `dsm_this_threadnum`: Returns the thread number for the chunk containing the given index value for each of `block`, `cyclic(..)`, and `star` distributions:

```
extern long long dsm_this_threadnum(void* array, long long dim,  
long long index)
```

- `dsm_distribution_block`, `dsm_distribution_cyclic`, and `dsm_distribution_star`: Boolean routines to query the distribution of a given dimension:

```
extern long long dsm_distribution_block (void* array, long long dim)
extern long long dsm_distribution_cyclic (void* array, long long dim)
extern long long dsm_distribution_star (void* array, long long dim)
```

- `dsm_isreshaped`: Boolean routine to query whether reshaped or not:

```
extern long long dsm_isreshaped (void* array)
```

- `dsm_isdistributed`: Boolean routine to query whether distributed (regular or reshaped) or not:

```
extern long long dsm_isdistributed (void* array)
```

12.6 Optional Environment Variables and Compile-Time Options

You can control various run-time features through the following optional environment variables and options. This section describes:

- Section 12.6.1, page 145
- Section 12.6.2, page 146

12.6.1 Multiprocessing Environment Variables

Details about the following multiprocessing environment variables can be found on the `pe_environ(5)` man page.

- `_DSM_OFF`
- `_DSM_BARRIER`
- `_DSM_PPM`
- `_DSM_PLACEMENT`
- `_DSM_MIGRATION`
- `_DSM_MIGRATION_LEVEL`
- `_DSM_WAIT`
- `_DSM_VERBOSE`
- `MP_SIMPLE_SCHED`

- `MP_SUGNUMTHD`
- `PAGESIZE_STACK`, `PAGESIZE_DATA`, and `PAGESIZE_TEXT`

12.6.2 Compile-Time Options

Useful compile-time options include:

`-MP:dsm={on, off}` (default on)

All the data-distribution and scheduling features described in this chapter are enabled by default under `-mp` compilation. To disable all the DSM-specific directives (for example, distribution and affinity scheduling), compile with `-MP:dsm=off`.

Note: Under `-mp` compilation, the compiler silently generates some book-keeping information under the directory `rii_files`. This information is used to implement data distribution directives, as well as perform consistency checks of these directives across multiple source files. To disable the processing of the data distribution directives and not generate the `rii_files`, compile with the `-MP:dsm=off` option.

`-MP:clone={on, off}` (default on)

The compiler automatically clones procedures that are called with reshaped arrays as parameters for the incoming distribution. However, if you have explicitly specified the distribution on all relevant formal parameters, then you can disable auto-cloning with `-MP:clone=off`. The consistency checking of the distribution between actual and formal parameters is **not** affected by this flag, and is always enabled.

`-MP:check_reshape={on, off}` (default off)

Enables generation of the run-time consistency checks across procedure boundaries when passing reshaped arrays (or portions thereof) as parameters.

12.7 Examples

The examples in this section include the following:

- Section 12.7.1, page 147
- Section 12.7.2, page 148
- Section 12.7.3, page 149
- Section 12.7.4, page 150
- Section 12.7.5, page 151

12.7.1 Distributing a Matrix

The example below distributes sequentially the rows of a matrix. Such a distribution places data effectively only if the size of an individual row (n) exceeds that of a page.

```
double A[n][n];

/* Distribute columns in cyclic fashion */
#pragma distribute A [cyclic(1)][*]

/* Perform Gaussian elimination across rows
The affinity clause distributes the loop iterations based
on the row distribution of A */
for (i = 0; i < n; i++)
#pragma pfor affinity(j) = data(A[i][j])
for (j = i+1; j < n; j++)
... reduce row j by row i ...
```

If the rows are smaller than a page, then it may be beneficial to reshape the array. This is easily specified by changing the keyword from `distribute` to `distribute_reshape`.

In addition to overcoming size constraints as shown above, the `#pragma distribute_reshape` directive is useful when the desired distribution is contrary to the layout of the array. For instance, suppose you want to distribute the columns of a two-dimensional matrix. In the following example, the `#pragma distribute_reshape` directive overcomes the storage layout constraints to provide the desired distribution:

```
double A[n][n];

/* Distribute columns in block fashion */
#pragma distribute_reshape A [*][block]

double sum[n];
#pragma distribute sum[block]

/* Perform sum-reduction on the elements of each column */
#pragma pfor local(j) affinity(i) = data(A[j][i])
for (i = 0; i < n; i++)
for (j = 0; j < n; j++)
    sum[i] = sum[i] + A[j][i];
```

12.7.2 Using Data Distribution and Data Affinity Scheduling

This example demonstrates regular data distribution and data affinity. This example, run on a four-processor Origin2000 server, uses simple block scheduling. Processor 0 will calculate the results of the first 25,000 elements of *A*, processor 1 will calculate the second 25,000 elements of *A*, and so on. Arrays *B* and *C* are initialized using one processor; therefore, all of the memory pages are touched by the master processor (processor 0) and are placed in processor 0's local memory.

Using data distribution changes the placement of memory pages for arrays *A*, *B*, and *C* to match the data reference pattern. Thus, the code runs 33% faster on a four-processor Origin2000 than it would run using SMP directives without data distribution.

Without Data Distribution

```
double a[1000000], b[1000000];
double c[1000000];
int i;

#pragma parallel shared(a, b, c) local(i)
#pragma pfor
for (i = 0; i < 100000; i++)
a[i] = b[i] + c[i];
```

```

With Data Distribution

double a[1000000], b[1000000];
double c[1000000];
int i;
#pragma distribute a[block], b[block], c[block]

#pragma parallel shared(a, b, c) local(i)
#pragma pfor affinity(i) = data(a[i])
for (i = 0; i < 100000; i++)
    a[i] = b[i] + c[i];

```

12.7.3 Parameter Passing

A distributed array can be passed as a parameter to a subroutine that has a matching declaration on the formal parameter:

```

double A [m][n];
#pragma distribute_reshape A [block][*]
foo (m, n, A);

foo (int p, int q, double A[p][q]){
#pragma distribute_reshape A[block][*]
#pragma pfor affinity (i) = data (A[i][j])
for (i = 0; i < P; i++)
...
}

```

Because the array is reshaped, it is required that the `#pragma distribute_reshape` directive in the caller and the callee match exactly. Furthermore, all calls to function `foo()` must pass in an array with the exact same distribution.

If the array was only distributed (that is, not reshaped) in the above example, then the subroutine `foo()` can be called from different places with different incoming distributions. In this case, you can omit the distribution directive on the formal parameter, thereby ensuring that any data affinity within the loop is based on the distribution (at run time) of the incoming actual parameter.

```

double A[m][n], B[p][q];
double A [block][*];
double B [cyclic(1)][*];
foo (m, n, A);
foo (p, q, B);

```

```
...
foo (int s, int t, double X[s][t]) {
#pragma pfor affinity (i) = data (X[i+2][j])
for (i = ...)
}
```

12.7.4 Redistributed Arrays

This example shows how an array is redistributed at run time:

```
bar(int n, double X[n][n]) {
...
#pragma redistribute X [cyclic(<expr>)]
...
}

foo() {
double LocalArray [1000][1000];
#pragma distribute LocalArray [block]
/* the call to bar() may redistribute LocalArray */
#pragma dynamic LocalArray
...
bar(1000, LocalArray);
/* The distribution for the following pfor */
is not known statically */
#pragma pfor affinity (i) = data (LocalArray[i][j])
...
}
```

The next example illustrates a situation where the `#pragma dynamic` directive can be optimized away. The main routine contains a local array *A* that is both distributed and dynamically redistributed. This array is passed as a parameter to `foo()` before being redistributed, and to `bar()` after being (possibly) redistributed. The incoming distribution for `foo()` is statically known; you can specify a `#pragma distribute` directive on the formal parameter, thereby obtaining more efficient static scheduling for the affinity `pfor`. The subroutine `bar()`, however, can be called with multiple distributions, requiring run-time scheduling of the affinity `pfor`.

```
main () {
double A[m][n];
#pragma distribute A [block][*]
#pragma dynamic A
```

```

foo (A);
if (...) {
#pragma redistribute A [cyclic(x)][*]
}
bar (A);
}

void foo (double A[m][n]) {
/* Incoming distribution is known to the user */
#pragma distribute A[block][*]
#pragma pfor affinity (i) = data (A[i][j])
    ...
}

void bar (double A[m][n]) {
/* Incoming distribution is not known statically */
#pragma dynamic A
#pragma pfor affinity (i) = data (A[i][j])
    ...
}

```

12.7.5 Irregular Distributions and Thread Affinity

The example below consists of a large array that is conceptually partitioned into unequal portions, one for each processor. This array is indexed through an index array *idx*, which stores the starting index value and the size of each processor's portion.

```

double A[N];
/* idx ---> index array containing start index into A [idx[p][0]]
and size [idx[p][1]] for each processor */
int idx [P][2];
#pragma page_place A[idx[0][0], idx[0][1]*8, 0)
#pragma page_place A[idx[1][0], idx[1][1]*8, 1)
#pragma page_place A[idx[2][0], idx[2][1]*8, 2)
...
#pragma pfor affinity (i) = thread(i)
for (i = 0; i < P-1; i++)
    ... process elements on processor i ...
    ... A[idx[i][0]] to A[idx[i][0]+idx[i][1]] ...

```


The Auto-Parallelizing Option (APO) [13]

Note: APO is licensed and sold separately from the MIPSpro C/C++ compilers. APO features in your code are ignored unless you are licensed for this product. For sales and licensing information, contact your sales representative.

The Auto-Parallelizing Option (APO) enables the MIPSpro C/C++ compilers to optimize parallel codes and enhances performance on multiprocessor systems. APO is controlled with command line options and source directives.

APO is integrated into the compiler; it is not a source-to-source preprocessor. Although run-time performance suffers slightly on single-processor systems, parallelized programs can be created and debugged with APO enabled.

Parallelization is the process of analyzing sequential programs for parallelism and restructuring them to run efficiently on multiprocessor systems. The goal is to minimize the overall computation time by distributing the computational workload among the available processors. Parallelization can be automatic or manual.

During *automatic parallelization*, the compiler analyzes and restructures the program with little or no intervention by you. With APO, the compiler automatically generates code that splits the processing of loops among multiple processors. An alternative is *manual parallelization*, in which you perform the parallelization using compiler directives and other programming techniques.

APO integrates automatic parallelization with other compiler optimizations, such as interprocedural analysis (IPA), optimizations for single processors, and loop nest optimization (LNO). In addition, run-time and compile-time performance is improved.

For details on using APO command line options and source directives, see the *MIPSpro C and C++ Pragmas* manual.

Implementation-Defined Behavior [A]

The sections in this appendix describe implementation-defined behavior. Each section is keyed to the ANSI C Standard (ANSI X3.159-1989), Appendix F, and each point is keyed to the section number of the ANSI C Standard. The bold lines, usually marked with bullets, are items from Appendix F of the ANSI C Standard. Text following the italic lines describes the Silicon Graphics implementation.

A.1 Translation (F.3.1)

- **Whether each nonempty sequence of white-space characters other than newline is retained or replaced by one space character (2.1.1.2).**

A nonempty sequence of white-space characters (other than newline) is retained.

- **How a diagnostic is identified (2.1.1.3).**

Successful compilations are silent. Diagnostics are, in general, emitted to standard error. Diagnostic messages have the general pattern of **file-name,line-number:severity(number): message** in `-n32` and `-64` modes. Diagnostics have a slightly different pattern in `-o32` mode. Also, the range of numbers in `-o32` mode is disjointed from the range in `-n32` and `-64` modes.

For example, typical messages from the ANSI C compiler front end in `-n32` and `-64` mode look like this:

```
"t4.c", line 4: error(1020):identifier "x" is undefined
"t4.c", line 5: warning(1551):variable "y" is used before its
value is set
```

Messages can also be issued by other internal compiler passes.

- **Classes of diagnostic messages, their return codes and control over them.**

Three classes of messages exist: warning, error, and remark. Warning messages include the notation "warning" (which can be capitalized), and allow the compilation to continue (return code 0). Error messages cause the compilation to fail (return code 1).

Remark messages appear in `-n32` and `-64` modes only. Typically, remarks are issued only if the `-fullwarn` option appears on the command line. More control is available with the `-diag_warning`, `-diag_remark`, and `-diag_error` options. (See the `cc` reference page for more information.)

Warning messages from the compiler front end have a unique diagnostic number. You can suppress these messages individually by putting the number in the numberlist of a `-woff numberlist` switch to the `cc` command. *numberlist* is a comma-separated list of warning numbers and ranges of warning numbers. For example, to suppress the warning message in the previous example, enter

```
-woff 1551
```

To suppress warning messages numbered 1642, 1643, 1644, and 1759, enter

```
-woff 1642-1644,1759
```

A.2 Environment (F.3.2)

- **Support of freestanding environments.**

No support is provided for a freestanding environment.

- **The semantics of the arguments to `main` (2.1.2.2.1).**

`main` is defined to have the two required parameters *argc* and *argv*. A third parameter, *envp*, is provided as an extension. That is, `main` would have the equivalent of the following prototype:

```
int main(int argc, char *argv[], char *envp[])
```

The parameters have the following semantics:

- *argc* is the number of arguments on the command line.
 - *argv*[0..*argc*-1] are pointers to the command-line arguments (strings).
 - *argv*[0] is the program name, as it appeared on the command line.
 - *argv*[*argc*] is a null pointer.
 - *envp* is an array of pointers to strings of the form *NAME=value*, where *NAME* is the name of an environment variable and *value* is its value. The array is terminated by a null pointer.
- **What constitutes an interactive device (2.1.2.3).**

Asynchronous terminals, including windows, are interactive devices and are, by default, line buffered. In addition, the standard error device, `stderr`, is unbuffered by default.

A.3 Identifiers (F.3.3)

- **The number of significant initial characters (beyond 31) in an identifier without external linkage (3.1.2).**

All characters are significant.

- **The number of significant initial characters (beyond 6) in an identifier with external linkage (3.1.2).**

All characters are significant.

- **Whether case distinctions are significant in an identifier with external linkage (3.1.2).**

Case distinctions are always significant.

A.4 Characters (F.3.4)

- **The members of the source and execution character sets, except as explicitly specified in the standard (2.2.1).**

Only the mandated characters are present. The source character set includes all printable ASCII characters, hexadecimal 0x20 through 0x7e, and 0x7 through 0xc (the standard escape sequences).

- **The values to which the standard escape sequences are translated (2.2.2).**

The escape sequences are translated as specified for standard ASCII: `\a = 0x7`, `\b = 0x8`, `\f = 0xc`, `\n = 0xa`, `\r = 0xd`, `\t = 0x9`, `\v=0xb`

- **The shift states used for the encoding of multibyte characters (2.2.1.2).**

The multibyte character set is identical to the source and execution character sets. There are no shift states.

- **The number of bits in a character in the execution character set (2.2.4.2.1).**

There are eight bits per character.

- **The mapping of members of the source character set (in character constants and string literals) to members of the execution character set (3.1.3.4).**

The mapping is the identity mapping.

- **The value of an integer character constant that contains a character or escape sequence not represented in the basic execution character set or in the extended character set for a wide character constant (3.1.3.4).**

With the exception of newline (0xa), backslash ('\'), and 0xff (end-of-file), eight-bit values appearing in an integer character constant are placed in the resultant integer in the same fashion as are characters that are members of the execution character set (see below). A backslash, newline, or 0xff can be placed in a character constant by preceding it with a backslash (that is, “escaping” it).

- **The value of an integer character constant that contains more than one character or a wide character constant that contains more than one multibyte character (3.1.3.4).**

You can assign up to four characters to an `int` using a character constant, as the following example illustrates:

```
int t = 'a'; /* integer value 0x61 */
int t2 = 'ab'; /* integer value 0x6162 */
int t4 = 'abcd'; /* integer value 0x61626364 */
int t4 = 'abcde'; /* error: too many characters for */
/* character constant */
```

The encoding of multiple characters in an integer consists of the assignment of the corresponding character values of the *n* characters in the constant to the least-significant *n* bytes of the integer, filling any unused bytes with zeros. The most significant byte assigned contains the value of the lexically first character in the constant.

Because the multibyte character set is identical to the source and execution character sets, the above discussion applies to the assignment of more than one multibyte character to a wide character constant.

- **The current locale used to convert multibyte characters into corresponding wide character (codes) for a wide character constant (3.1.3.4).**

The mapping is the identity mapping to the standard ASCII character set. The C locale is used.

- Whether a “plain” `char` has the same range of values as `signed char` or `unsigned char`.

Plain `char` is the same as `unsigned char` by default. Use the `-signed` option to `cc` to switch the range to be that of `signed char`.

A.5 Integers (F.3.5)

- The representations and sets of values of the various types of integers (3.1.2.5).

Integers are two’s complement binary. Table 13, page 159, lists the sizes and ranges of the various types of integer. The use of `long long` results in a warning in `-ansi` and `-ansiposix` modes.

In `-o32` and `-n32` mode implementations, to take full advantage of the support for 64-bit integral values in `-ansi` and `-ansiposix` modes, you can define the macro `_LONGLONG` on the `cc` command line when using the types `__uint64_t`, `__int64_t`, or library routines that are prototyped in terms of these types.

Table 13. Integer Types and Ranges

Type	Range: Low	High	Size (bits)
<code>signed char</code>	-128	127	8
<code>char</code> , <code>unsigned char</code>	0	255	8
<code>short</code> , <code>signed short</code>	-32768	32767	16
<code>unsigned short int</code>	0	65535	16
<code>int</code> , <code>signed int</code>	-2147483648	2147483647	32
<code>unsigned int</code>	0	4294967295	32
<code>long</code> , <code>signed long int</code>	-2147483648 (-32 and -n32 modes) -9223372036854775808 (-64 mode)	2147483647 (-32 and -n32 modes) 9223372036854775807 (-64 mode)	32 64

Type	Range: Low	High	Size (bits)
unsigned long int	0	4294967295 (-32 and -n32 modes) 18446744073709551615 (-64 mode)	32 64
long long signed long long int	-9223372036854775808	9223372036854775807	64
unsigned long long int	0	18446744073709551615	64

- **The result of converting an integer to a shorter signed integer, or the result of converting an unsigned integer to a signed integer of equal length, if the value cannot be represented (3.2.1.2).**

The least significant n bits (n being the length of the result integer) of the source are copied to the result.

- **The results of bitwise operations on signed integers (3.3).**

With the exception of right-shift of a negative signed integer (defined below), operations on signed and unsigned integers produce the same bitwise results.

- **The sign of the remainder on integer division (3.3.5).**

The sign of the remainder is that of the numerator.

- **The result of a right shift of a negative-valued signed integral type (3.3.7).**

The sign bit is propagated, so the result value is still negative.

A.6 Floating Point (F.3.6)

- **The representations and sets of values of the various types of floating-point numbers (3.1.2.5).**

The representation is IEEE:

- Single (for `float` values)
- Double (for `double` values and for `long double` values in `-o32` mode)
- Quad precision (for `long double` values in `-n32` and `-64` mode).

See ANSI/IEEE Standard 754-1985 and IEEE Standard for Binary Floating-Point Arithmetic. Table 14, page 161, lists ranges of floating point types.

Table 14. Ranges of floating point Types

Type	Range: Min	Max	Size (Bits)
float	1.1755e-38	3.4028e+38	32
double	2.225e-308	1.7977e+308	64
long double	2.225e-308	1.7977e+308	128 (-n32 and -64 modes)

- **The type of rounding or truncation used when representing a floating-point constant which is within its range.**

Per IEEE, the rounding is round-to-nearest (IEEE Standard 754, sections 4.1 and 5.5). If the two values are equally near, then the one with the least significant bit zero is chosen.

- **The direction of truncation when an integral number is converted to a floating-point number that cannot exactly represent the original value (3.2.1.3).**

Conversion of an integral type to a float type, if the integral value is too large to be exactly represented, gives the next higher value.

- **The direction of truncation or rounding when a floating-point number is converted to a narrower floating-point number.**

Per IEEE, the rounding is round-to-nearest (IEEE Standard 754, Section 4.1 and 5.5). If the two values are equally near, then the one with the least significant bit zero is chosen.

A.7 Arrays and Pointers (F.3.7)

- **The type of integer required to hold the maximum size of an array— that is, the type of the sizeof operator, `size_t`(3.3.3.4, 4.1.1).**

An unsigned `long` holds the maximum array size.

- **The size of integer required for a pointer to be converted to an integer type (3.3.4).**

`long ints` are large enough to hold pointers in `-n32` and `-o32` mode. Both are 32 bits wide.

`long ints` are large enough to hold pointers in `-64` mode. Both are 64 bits wide.

- **The result of casting a pointer to an integer or vice versa (3.3.4).**

The result is bitwise exact provided the integer type is large enough to hold a pointer.

- **The type of integer required to hold the difference between two pointers to elements of the same array, `ptrdiff_t` (3.3.6, 4.1.1).**

An `int` is large enough to hold the difference between two pointers to elements of the same array in `-o32` and `-n32` modes.

A `long int` is large enough to hold the difference between two pointers to elements of the same array in `-n32`, `-o32`, and `-64` modes.

A.8 Registers (F.3.8)

- **The extent to which objects can actually be placed in registers by use of the register storage-class specifier (3.5.1).**

The compilation system can use up to eight of the register storage-class specifiers for nonoptimized code in `-32` mode, and it ignores register specifiers for formal parameters. Use of register specifiers is not recommended.

The register storage-class specifier is always ignored and the compilation system makes its own decision about what should be in registers for optimized code (`-O2` and above).

A.9 Structures, Unions, Enumerations, and Bitfields (F.3.9)

- **What is the result if a member of a union object is accessed using a member of a different type (3.3.2.3).**

The bits of the accessed member are interpreted according to the type used to access the member. For integral types, the `N` bits of the type are simply

accessed. For floating types, the access might cause a trap if the bits are not a legal floating point value. For pointer types, the 32 bits (64 bits if in `-64` mode) of the pointer are picked up. The usability of the pointer depends on whether it points to a valid object or function, and whether it is used appropriately. For example, a pointer whose least-significant bit is set can point to a character, but not to an integer.

- **The padding and alignment of members of structures (3.5.2.1).**

This should present no problem unless binary data written by one implementation are read by another.

Members of structures are on the same boundaries as the base data type alignments anywhere else. A word is 32 bits and is aligned on an address, which is a multiple of 4. `unsigned` and `signed` versions of a basic type use identical alignment. Type alignments are given in Table 15, page 163, .

Table 15. Alignment of Structure Members

Type	Alignment
<code>long double</code>	Double- word boundary (<code>-32</code> mode) Quad-word boundary (<code>-n32</code> and <code>-64</code> modes)
<code>double</code>	Double-word boundary
<code>float</code>	Word boundary
<code>long long</code>	Double-word boundary
<code>long</code>	Word boundary (<code>-n32</code> and <code>-32</code> modes) double-word boundary (<code>-64</code> mode)
<code>int</code>	Word boundary
<code>pointer</code>	Word boundary
<code>short</code>	Half-word boundary
<code>char</code>	Byte boundary

- **Whether a “plain” `int` bit-field is treated as a `signed int` bit-field or as an `unsigned int` bit-field (3.5.2.1).**

A “plain” `int` bit-field is treated as a `signed int` bit-field.

- **The order of allocation of bitfields within a unit (3.5.2.1).**

Bits in a bitfield are allocated with the most-significant bit first within a unit.

- **Whether a bitfield can straddle a storage-unit boundary (3.5.2.1).**

Bitfields cannot straddle storage unit boundaries (relative to the beginning of the `struct` or `union`), where a storage unit can be of size 8, 16, 32, or 64 bits.

- **The integer type chosen to represent the values of an enumeration type (3.5.2.2).**

The `int` type is always used.

Note: `long` or `long long` enumerations are not supported.

A.10 Qualifiers (F.3.10)

- **What constitutes an access to an object that has volatile-qualified type (3.5.3).**

Objects of volatile-qualified type are accessed only as specified by the abstract semantics, and as would be expected on a RISC architecture, no complex instructions exist (for example, read-modify-write). Volatile objects appearing on the left side of an assignment expression are accessed once for the write. If the assignment is not simple, an additional read access is performed. Volatile objects appearing in other contexts are accessed once per instance. Incrementation and decrementation require both a read and a write access.

Volatile objects that are memory-mapped are accessed only as specified. If such an object is of size `char`, for example, adjacent bytes are not accessed. If the object is a bitfield, a read may access the entire storage unit containing the field. A write of an unaligned field necessitates a read and write of the storage unit that contains it.

A.11 Declarators (F.3.11)

- **The maximum number of declarators that can modify an arithmetic, structure, or union type (3.5.4).**

There is no limit.

A.12 Statements (F.3.12)

- **The maximum number of case values in a switch statement (3.6.4.2).**

There is no limit.

A.13 Preprocessing Directives (F.3.13)

- **Whether the value of a single-character character constant in a constant expression that controls conditional inclusion matches the value of the same character constant in the execution character set. Whether such a character constant can have a negative value (3.8.1).**

The preprocessing and execution phases use exactly the same meanings for character constants.

A single-character character constant is always positive.

- **The method for locating includable source files (3.8.2).**

For filenames surrounded by `< >`, the includable source files are searched for in `/usr/include`.

The default search list includes `/usr/include`. You can change this list with various compiler options. See the `cc(1)` reference page and the `-I` and `-nostdinc` options.

- **The support of quoted names for includable source files (3.8.2).**

Quoted names are supported for includable source files. For filenames surrounded by ```''`, the includable source files are searched for in the directory of the current include file, then in `/usr/include`.

The default search list includes `/usr/include`. You can change this list with various compiler options. See the `cc(1)` reference page and the `-I` and `-nostdinc` options.

- **The mapping of source file character sequences (3.8.2).**

The mapping is the identity mapping.

- **The behavior on each recognized `#pragma` directive.**

See *MIPSpro C and C++ Pragmas* on the Silicon Graphics Tech Pubs Library (<http://techpubs.sgi.com/library>) for details on all supported `#pragma` directives.

- **The definitions for `__DATE__` and `__TIME__` when, respectively, the date and time of translation are not available.**

The date and time of translation are always available in this implementation.

- **What is the maximum nesting depth of include files (3.8.2).**

The maximum nesting depth of include files is 200.

A.14 Library Functions (F.3.14)

- **The null pointer constant to which the macro `NULL` expands (4.1.5).**

The `NULL` pointer constant expands to an `int` with value zero. That is,

```
#define NULL 0
```

- **The diagnostic printed by and the termination behavior of the `assert` function (4.2).**

If an assertion given by `assert(EX)` fails, the following message is printed on `stderr` using `_write` to its underlying `fileno`:

```
Assertion failed: EX, file <filename>, line <linenumber>
```

This is followed by a call to `abort` (which exits with a `SIGABRT`).

- **The sets of characters tested for by the `isalnum`, `isalpha`, `iscntrl`, `islower`, `isprint`, and `isupper` functions (4.3.1).**

The statements in the following list are true when operating in the C locale. The C locale is in effect at program start up for programs compiled for pure ANSI C (that is, `-ansi`), or by invoking `setlocale(LC_ALL, 'C')`. The C locale can be overridden at start up for any program that does not explicitly invoke `setlocale` by setting the value of the `CHRCLASS` environment variable. (See the `ctype(3C)` reference page.)

- `isalnum` is nonzero for the 26 letters `a-z`, the 26 letters `A-Z`, and the digits `0-9`.
- `isalpha` is nonzero for the 26 letters `a-z` and the 26 letters `A-Z`.
- `islower` is nonzero for the 26 letters `a-z`.
- `isupper` is nonzero for the 26 letters `A-Z`.

- `isprint` is nonzero for the ASCII characters space through tilde (~) (0x20 through 0x7e).
- `iscntrl` is nonzero for the ASCII characters NUL through US (0x0 through 0x1f).
- **The values returned by the mathematics functions on domain errors (4.5.1).**

The value returned by the math functions on domain errors is the default IEEE Quiet NaN in all cases except the following:

- The functions `pow` and `powf` return `-HUGE_VAL` when the first argument is zero and the second argument is negative. When both arguments are zero, `pow()` and `powf()` return 1.0.
- The functions `atan2` and `atan2f` return zero when both arguments are zero.
- **Whether mathematics functions set the integer expression `errno` to the value of the macro `RANGE` on underflow range errors (4.5.1).**

Yes, except intrinsic functions that have been inlined. Note that `fabs`, `fabsf`, `sqrt`, `sqrtf`, `hypotf`, `fhypot`, `pow`, and `powf` are intrinsic by default in `-xansi` and `-cckr` modes and can be made intrinsic in `-ansi` mode by using the `D__INLINE_INTRINSICS` compiler option.

- **Whether a domain error occurs or zero is returned when the `fmod` function has a second argument of zero (4.5.6.4).**

`fmod(x, 0)` gives a domain error and returns the default IEEE Quiet NaN.

A.14.1 Signals

- **The set of signals for the signal function (4.7.1.1).**

The signal set is listed in Table 16, page 168, which is from the `signal(2)` reference page. The set of signals conforms to the SVR4 ABI. Note that some of the signals are not defined in `-ansiposix` mode. References in square brackets beside the signal numbers are described under “Signal Notes” in the discussion of signal semantics.

Table 16. Signals

Signal	Number[Note]	Meaning
SIGHUP	01	Hangup
SIGINT	02	Interrupt
SIGQUIT	03[1]	Quit
SIGILL	04[1]	Illegal instruction (not reset when caught)
SIGTRAP	05[1][5]	Race trap (not reset when caught)
SIGIOT	06	IOT instruction
SIGABRT	06[1]	Abort
SIGEMT	07[1][4]	MT instruction
SIGFPE	08[1]	Floating point exception
SIGKILL	09	Kill (cannot be caught or ignored)
SIGBUS	10[1]	Bus error
SIGSEGV	11[1]	Segmentation violation
SIGSYS	12[1]	Bad argument to system call
SIGPIPE	13	Write on a pipe with no one to read it
SIGALRM	14	Alarm clock
SIGTERM	15	Software termination signal
SIGUSR1	16	User-defined signal 1
SIGUSR2	17	User-defined signal 2
SIGCLD	18[2]	Termination of a child process
SIGGHLD	18	4.3 BSD and POSIX [®] name
SIGPWR	19[2]	Power fail (not reset when caught)
SIGWINCH	20[2]	Window size changes
SIGURG	21[2]	Urgent condition on I/O channel
SIGIO	22[2]	Input/output possible
SIGPOLL	22[3]	Selectable event pending

Signal	Number[Note]	Meaning
SIGSTOP	23[6]	Stop (cannot be caught or ignored)
SIGTSTP	24[6]	Stop signal generated from keyboard
SIGCONT	25[6]	Continue after stop (cannot be ignored)
SIGTTIN	26[6]	Background read from control terminal
SIGTTOU	27[6]	Background write to control terminal
SIGVTALRM	28	Virtual time alarm
SIGPROF	29	Profiling alarm
SIGXCPU	30	CPU time limit exceeded [see <code>setrlimit(2)</code>]
SIGXFSZ	31	File size limit exceeded [see <code>setrlimit(2)</code>]
SIG32	32	Reserved for kernel usage

- **The semantics for each signal recognized by the signal function (4.7.1.1).**

In the signal invocation `signal(sig, func)`, `func` can be the address of a signal handler, `handler`, or one of the two constant values (defined in `<sys/signal.h>`) `SIG_DFL` or `SIG_IGN`. The semantics of these values are as follows:

<code>SIG_DFL</code>	Terminate process upon receipt of signal <code>sig</code> . (This is the default if no call to <code>signal</code> for signal <code>sig</code> occurs.) Upon receipt of the signal <code>sig</code> , the receiving process is to be terminated with all of the consequences outlined in the <code>exit(2)</code> reference page. See note 1 under Section A.14.1.1, page 171.
<code>SIG_IGN</code>	Ignore signal. The signal <code>sig</code> is to be ignored.
<code>handler</code>	Catch signal. <code>func</code> is the address of function <code>handler</code> .

Note: The signals `SIGKILL`, `SIGSTOP`, and `SIGCONT` cannot be ignored.

If *func* is the address of handler, upon receipt of the signal *sig*, the receiving process is to invoke handler as follows:

```
handler (int sig, int code, struct sigcontext *sc);
```

The remaining arguments are supplied as extensions and are optional. The value of the second argument *code* is meaningful only in the cases shown in Table 17, page 170.

Table 17. Valid Codes in a Signal-Catching Function

Condition	Signal	Code
User breakpoint	SIGTRAP	BRK_USERBP
User breakpoint	SIGTRAP	BRK_SSTEPBP
Integer overflow	SIGTRAP	BRK_OVERFLOW
Divide by zero	SIGTRAP	BRK_DIVZERO
Multiply overflow	SIGTRAP	BRK_MULOVF
Invalid virtual address	SIGSEGV	EFAULT
Read-only address	SIGSEGV	EACCESS
Read beyond mapped object	SIGSEGV	ENXIO

The third argument, *sc*, is a pointer to a `struct sigcontext` (defined in `<sys/signal.h>`) that contains the processor context at the time of the signal. Upon return from *handler*, the receiving process resumes execution at the point where it was interrupted.

Before entering the signal-catching function, the value of *func* for the caught signal is set to `SIG_DFL`, unless the signal is `SIGILL`, `SIGTRAP`, or `SIGPWR`. This means that before exiting the handler, a call to `signal` is necessary to catch future signals.

Suppose a signal that is to be caught occurs during one of the following routines:

- A `read`, `write`, or `open`
- An `ioctl` system call on a slow device (like a terminal, but not a file)
- A `pause` (system call)

- A `wait` system call that does not return immediately due to the existence of a previously stopped or zombie process

The signal catching function is executed and then the interrupted system call returns a `-1` to the calling process with `errno` set to `EINTR`.

Note: The signals `SIGKILL` and `SIGSTOP` cannot be caught.

A.14.1.1 Signal Notes

1. If `SIG_DFL` is assigned for `SIGQUIT`, `SIGILL`, `SIGTRAP`, `SIGABRT`, `SIGEMT`, `SIGFPE`, `SIGBUS`, `SIGSEGV`, or `SIGSYS`, in addition to the process being terminated, a “core image” is constructed in the current working directory of the process, if the following two conditions are met:
 - a. The effective user ID and the real user ID of the receiving process are equal.
 - b. An ordinary file named `core` exists and is writable or can be created.

If the file must be created, it has the following properties:

- A mode of `0666` modified by the file creation mask (see the `umask(2)` reference page)
- A file owner ID that is the same as the effective user ID of the receiving process
- A file group ID that is the same as the effective group ID of the receiving process

Note: The core file can be truncated if the resultant file size would exceed either `ulimit` (see the `ulimit(2)` reference page) or the process’s maximum core file size (see the `setrlimit(2)` reference page).

2. For the signals `SIGCLD`, `SIGWINCH`, `SIGPWR`, `SIGURG`, and `SIGIO`, the actions associated with each of the three possible values for `func` are as follows:

<code>SIG_DFL</code>	Ignore signal. The signal is to be ignored.
<code>SIG_IGN</code>	Ignore signal. The signal is to be ignored. Also, if <code>sig</code> is <code>SIGCLD</code> , the calling process’s child processes do not create zombie processes when they terminate (see the <code>exit(2)</code> reference page).

`handler` Catch signal. If the signal is `SIGPWR`, `SIGURG`, `SIGIO`, or `SIGWINCH`, the action to be taken is the same as that previously described when *func* is the address of a function. The same is true if the signal is `SIGCLD` with one exception: while the process is executing the signal-catching function, all terminating child processes are queued. The `wait` system call removes the first entry of the queue. If the signal system call is used to catch `SIGCLD`, the signal handler must be reattached when exiting the handler, and at that time—if the queue is not empty—`SIGCLD` is raised again before `signal` returns. (See the `wait(2)` reference page.)

In addition, `SIGCLD` affects the `wait` and `exit` system calls as follows:

`wait` If the handler parameter of `SIGCLD` is set to `SIG_IGN` and a `wait` is executed, the `wait` blocks until all of the calling process's child processes terminate; it then returns a value of `-1` with `errno` set to `ECHILD`.

`exit` If, in the exiting process's parent process, the handler parameter of `SIGCLD` is set to `SIG_IGN`, the exiting process does not create a zombie process.

When processing a pipeline, the shell makes the last process in the pipeline the parent of the preceding processes. Do not set `SIGCLD` to be caught for a process that can be piped into in this manner (and thus become the parent of other processes).

3. `SIGPOLL` is issued when a file descriptor corresponding to a `STREAMS` (see `intro(2)`) file has a "selectable" event pending. A process must specifically request that this signal be sent using the `I_SETSIG` `ioctl` call. Otherwise, the process never receives `SIGPOLL`.
4. `SIGEMT` is never generated on an IRIS 4D system.
5. `SIGTRAP` is generated for breakpoint instructions, overflows, divide by zeros, range errors, and multiply overflows. The second argument code gives specific details of the cause of the signal. Possible values are described in `<sys/signal.h>`.

6. The signals `SIGSTOP`, `SIGTSTP`, `SIGTTIN`, `SIGTTOU`, and `SIGCONT` are used by command interpreters like the C shell (see the `csh(1)` reference page) to provide job control. The first four signals listed stop the receiving process unless the signal is caught or ignored. `SIGCONT` resumes a stopped process. `SIGTSTP` is sent from the terminal driver in response to the `SWTCH` character being entered from the keyboard (see the `termio(7)` reference page). `SIGTTIN` is sent from the terminal driver when a background process attempts to read from its controlling terminal. If `SIGTTIN` is ignored by the process, then the read returns `EIO`. `SIGTTOU` is sent from the terminal driver when a background process attempts to write to its controlling terminal when the terminal is in `TOSTOP` mode. If `SIGTTOU` is ignored by the process, then the write succeeds, regardless of the state of the controlling terminal.

`signal` does not catch an invalid function argument, *func*, and results are undefined when an attempt is made to execute the function at the bad address.

`SIGKILL` immediately terminates a process, regardless of its state.

Processes stopped via job control (typically `CTRL+Z`) do not act upon any delivered signals other than `SIGKILL` until the job is restarted. Processes blocked via a `blockproc` system call unblock if they receive a signal that is fatal (that is, a non-job-control signal that they are not catching). These processes remained stopped, however, if the job they are a part of is stopped. Only upon restart do they die. Any non-fatal signals received by a blocked process do not cause the process to be unblocked. An `unblockproc` or `unblockprocall` system call is necessary.

If an instance of signal *sig* is pending when `signal(sig, func)` is executed, the pending signal is cancelled unless it is `SIGKILL`.

`signal` fails if *sig* is an illegal signal number, including `SIGKILL` and `SIGSTOP`, or if an illegal operation is requested (such as ignoring `SIGCONT`, which is ignored by default). In these cases, `signal` returns `SIG_ERR` and sets `errno` to `EINVAL`.

After a `fork`, the child inherits all handlers and signal masks. If any signals are pending for the parent, they are not inherited by the child.

The `exec` routines reset all caught signals to the default action; ignored signals remain ignored; the blocked signal mask is unchanged and pending signals remain pending.

The following reference pages contain other relevant information: `intro(2)`, `blockproc(2)`, `kill(2)`, `pause(2)`, `ptrace(2)`, `sigaction(2)`, `sigset(2)`, `wait(2)`, `setjmp(3C)`, `sigvec`, and `kill(1)`.

A.14.1.2 Diagnostics

Upon successful completion, `signal` returns the previous value of `func` for the specified signal `sig`. Otherwise, a value of `SIG_ERR` is returned and `errno` is set to indicate the error. `SIG_ERR` is defined in the `<sys/signal.h>` header file.



Caution: Signals raised by the instruction stream, `SIGILL`, `SIGEMT`, `SIGBUS`, and `SIGSEGV`, will cause infinite loops if their handler returns, or the action is set to `SIG_IGN`. The POSIX signal routines (`sigaction`, `sigpending`, `sigprocmask`, `sigsuspend`, `sigsetjmp`), and the BSD 4.3 signal routines (`sigvec`, `signal`, `sigblock`, `sigpause`, `sigsetmask`) must never be used with `signal` or `sigset`.

Before entering the signal-catching function, the value of `func` for the caught signal is set to `SIG_DFL`, unless the signal is `SIGILL`, `SIGTRAP`, or `SIGPWR`. This means that before exiting the handler, a `signal` call is necessary to again set the disposition to catch the signal.

Note that handlers installed by `signal` execute with no signals blocked, not even the one that invoked the handler.

- **The default handling and the handling at program startup for each signal recognized by the signal function (4.7.1.1).**

Each signal is set to `SIG_DFL` at program start up.

- **If the equivalent of `signal(sig, SIG_DFL)`; is not executed prior to the call of a signal handler, the blocking of the signal that is performed(4.7.1.1).**

The equivalent of `signal(sig, SIG_DFL)` is executed prior to the call of a signal handler unless the signal is `SIGILL`, `SIGTRAP`, or `SIGPWR`. See the `signal` reference page for information on the support for the BSD 4.3 signal facilities.

- **Whether the default handling is reset if the `SIGILL` signal is received by a handler specified to the signal function (4.7.1.1).**

No.

A.14.2 Streams and Files

- **Whether the last line of a text stream requires a terminating newline character (4.9.2).**

There is no requirement that the last line of a text stream have a terminating newline: the output is flushed when the program terminates, if not earlier (as a result of `fflush` call). However, subsequent processes or programs reading the text stream or file might expect the newline to be present; it customarily is in IRIX text files.

- **Whether space characters that are written out to a text stream immediately before a newline character appear when read in (4.9.2).**

All text characters (including spaces before a newline character) written out to a text stream appear exactly as written when read back in.

- **The number of null characters that can be appended to data written to a binary stream (4.9.2).**

The library never appends nulls to data written to a binary stream. Only the characters written by the application are written to the output stream, whether binary or text. Text and binary streams are identical: there is no distinction.

- **Whether the file position indicator of an append mode stream is initially positioned at the beginning or end of the file (4.9.2).**

The file position indicator of an append stream is initially positioned at the end of the file.

- **Whether a write on a text stream causes the associated file to be truncated beyond that point (4.9.3).**

A write on a text stream does not cause the associated file to be truncated.

- **The characteristics of file buffering (4.9.3).**

Files are fully buffered, as described in paragraph 3, section 4.9.3, of ANSI X3.159-1989.

- **Whether a zero-length file actually exists (4.9.3).**

Zero-length files exist, but have no data, so a read on such a file returns an immediate EOF.

- **The rules for composing valid file names (4.9.3).**

Filenames consist of 1 to `FILENAME_MAX` characters. These characters can be selected from the set of all character values excluding `\0` (null) and the ASCII code for `/` (slash).

It is generally unwise to use `*`, `?`, `[`, or `]` as part of filenames because of the special meaning attached to these characters by the shell (see the `sh(1)` reference page). Although permitted, the use of unprintable characters should be avoided.

- **Whether the same file can be opened multiple times (4.9.3).**

A file can be open any number of times.

- **The effect of the remove function on an open file (4.9.4.1).**

For local disk files, a `remove(1)` removes a directory entry pointing to the file but has no effect on the file or the program with the file open. For files remotely mounted via NFS software, the effect is unpredictable (the file might be removed making further I/O impossible through open streams, or it might behave like a local disk file) and might depend on the version(s) of NFS involved.

- **The effect if a file with the new name exists prior to a call to the rename function (4.9.4.2).**

If the new name exists, the file with that new name is removed (See the `rm(1)` reference page) before the rename is done.

- **The output for `%p` conversion in the `fprintf` function (4.9.6.1).**

`%p` is treated the same as `%x`.

- **The input for `%p` conversion in the `fscanf` function (4.9.6.2).**

`%p` is treated the same as `%x`.

- **The interpretation of a `-` character that is neither the first nor the last character in the scanlist for `%[` conversion in the `fscanf` function (4.9.6.2).**

A `-` character that does not fit the pattern mentioned above is used as a shorthand for ranges of characters. For example, `[abcdefgh]` and `[xa-h]` mean that characters a through h and the character x are in the range (called a scanset in 4.9.6.2).

A.14.3 Temporary Files

- **Whether a temporary file is removed if a program terminates abnormally (4.9.4.3).**

Temporary files are removed if a program terminates abnormally.

A.14.4 `errno` and `perror`

- **The value to which the macro `errno` is set by the `fgetpos` or `ftell` function on failure (4.9.9.1, 4.9.9.4).**

`errno` is set to `EBADF` (9) by the `fgetpos` or `ftell` function on failure.

- **The messages generated by the `perror` function (4.9.10.4).**

The message generated is simply a string. The content of the message given for each legal value of `errno` is given in the list below, which is of the format `errno_value:message`.

- 1: No permission match (-o32 mode) 1: Not privileged (-n32 and -64 modes)
- 2: No such file or directory
- 3: No such process
- 4: Interrupted system call
- 5: I/O error
- 6: No such device or address
- 7: Arg list too long
- 8: Exec format error
- 9: Bad file number
- 10: No child processes
- 11: Resource temporarily unavailable
- 12: Not enough space
- 13: Permission denied
- 14: Bad address

- 15: Block device required
- 16: Device or resource busy (-o32 mode) 16: Device busy (-n32 and -64 modes)
- 17: File exists
- 18: Cross-device link
- 19: No such device
- 20: Not a directory
- 21: Is a directory
- 22: Invalid argument
- 23: Too many open files in system (-o32 mode) 23: File table overflow (-n32 and -64 modes)
- 24: Too many open files in a process (-o32 mode) 24: Too many open files (-n32 and -64 modes)
- 25: Inappropriate IOCTL operation (-o32 mode) 25: Not a typewriter (-n32 and -64 modes)
- 26: Text file busy
- 27: File too large
- 28: No space left on device
- 29: Illegal seek
- 30: Read-only filesystem
- 31: Too many links
- 32: Broken pipe
- 33: Argument out of domain
- 34: Result too large
- 35: No message of desired type
- 36: Identifier removed
- 37: Channel number out of range

- 38: Level 2 not synchronized
- 39: Level 3 halted
- 40: Level 3 reset
- 41: Link number out of range
- 42: Protocol driver not attached
- 43: No CSI structure available
- 44: Level 2 halted
- 45: Deadlock situation detected/avoided
- 46: No record locks available
- 47: Error 47
- 48: Error 48
- 49: Error 49
- 50: Bad exchange descriptor
- 51: Bad request descriptor
- 52: Message tables full
- 53: Anode table overflow
- 54: Bad request code
- 55: Invalid slot
- 56: File locking deadlock
- 57: Bad font file format
- 58: Error 58
- 59: Error 59
- 60: Not a stream device
- 61: No data available
- 62: Timer expired
- 63: Out of stream resources

- 64: Machine is not on the network
- 65: Package not installed
- 66: Object is remote
- 67: Link has been severed
- 68: Advertise error
- 69: Srmount error
- 70: Communication error on send
- 71: Protocol error
- 72: Error 72
- 73: Error 73
- 74: Multihop attempted
- 75: Error 75
- 76: Error 76
- 77: Not a data message
- 78: Error 78 (-o32 mode) 78: Filename too long (-n32 and -64 modes)
- 79: Error 79 (-o32 mode) 79: Value too large for defined data type (-n32 and -64 modes)
- 80: Name not unique on network
- 81: File descriptor in bad state
- 82: Remote address changed
- 83: Cannot access a needed shared library
- 84: Accessing a corrupted shared library
- 85: .lib section in a.out corrupted
- 86: Attempting to link in more shared libraries than system limit
- 87: Cannot exec a shared library directly
- 88: Invalid System Call (-o32 mode) 88: Illegal byte sequence (-n32 and -64 modes)

- 89: Error 89 (-o32 mode) 89: Operation not applicable (-n32 and -64 modes)
- 90: Error 90 (-o32 mode) 90: Too many symbolic links in pathname traversal (-n32 and -64 modes)
- 91: Error 91 (-o32 mode) 91: Restartable system call (-n32 and -64 modes)
- 92: Error 92 (-o32 mode) 92: If pipe/FIFO, don't sleep in stream head (-n32 and -64 modes)
- 93: Error 93 (-o32 mode) 93: Directory not empty (-n32 and -64 modes)
- 94: Error 94 (-o32 mode) 94: Too many users (-n32 and -64 modes)
- 95: Error 95 (-o32 mode) 95: Socket operation on non-socket (-n32 and -64 modes)
- 96: Error 96 (-o32 mode) 96: Destination address required (-n32 and -64 modes)
- 97: Error 97 (-o32 mode) 97: Message too long (-n32 and -64 modes)
- 98: Error 98 (-o32 mode) 98: Protocol wrong type for socket (-n32 and -64 modes)
- 99: Error 99 (-o32 mode) 99: Option not supported by protocol (-n32 and -64 modes)
- 100: Error 100
- 101: Operation would block (-o32 mode) 101: Error 101 (-n32 and -64 modes)
- 102: Operation now in progress (-o32 mode) 102: Error 102 (-n32 and -64 modes)
- 103: Operation already in progress (-o32 mode) 103: Error 103 (-n32 and -64 modes)
- 104: Socket operation on non-socket (-o32 mode) 104: Error 104 (-n32 and -64 modes)
- 105: Destination address required (-o32 mode) 105: Error 105 (-n32 and -64 modes)
- 106: Message too long (-o32 mode) 106: Error 106 (-n32 and -64 modes)

107: Protocol wrong type for socket (-o32 mode) 107: Error 107 (-n32 and -64 modes)

108: Option not supported by protocol (-o32 mode) 108: Error 108 (-n32 and -64 modes)

109: Protocol not supported (-o32 mode) 109: Error 109 (-n32 and -64 modes)

110: Socket type not supported (-o32 mode) 110: Error 110 (-n32 and -64 modes)

111: Operation not supported on socket (-o32 mode) 111: Error 111 (-n32 and -64 modes)

112: Protocol family not supported (-o32 mode) 112: Error 112 (-n32 and -64 modes)

113: Address family not supported by protocol family (-o32 mode) 113: Error 113 (-n32 and -64 modes)

114: Address already in use (-o32 mode) 114: Error 114 (-n32 and -64 modes)

115: Can't assign requested address (-o32 mode) 115: Error 115 (-n32 and -64 modes)

116: Network is down (-o32 mode) 116: Error 116 (-n32 and -64 modes)

117: Network is unreachable (-o32 mode) 117: Error 117 (-n32 and -64 modes)

118: Network dropped connection on reset (-o32 mode) 118: Error 118 (-n32 and -64 modes)

119: Software caused connection abort (-o32 mode) 119: Error 119 (-n32 and -64 modes)

120: Connection reset by peer (-o32 mode) 120: Protocol not supported (-n32 and -64 modes)

121: No buffer space available (-o32 mode) 121: Socket type not supported (-n32 and -64 modes)

122: Socket is already connected (-o32 mode) 122: Operation not supported on transport endpoint (-n32 and -64 modes)

123: Socket is not connected (-o32 mode) 123: Protocol family not supported (-n32 and -64 modes)

124: Can't send after socket shutdown (-o32 mode) 124: Address family not supported by protocol family (-n32 and -64 modes)

125: Too many references: can't splice (-o32 mode) 125: Address already in use (-n32 and -64 modes)

126: Connection timed out (-o32 mode) 126: Cannot assign requested address (-n32 and -64 modes)

127: Connection refused (-o32 mode) 127: Network is down (-n32 and -64 modes)

128: Host is down (-o32 mode) 128: Network is unreachable (-n32 and -64 modes)

129: Host is unreachable (-o32 mode) 129: Network dropped connection because of reset (-n32 and -64 modes)

130: Too many levels of symbolic links (-o32 mode) 130: Software caused connection abort (-n32 and -64 modes)

131: Filename too long (-o32 mode) 131: Connection reset by peer (-n32 and -64 modes)

132: Directory not empty (-o32 mode) 132: No buffer space available (-n32 and -64 modes)

133: Disk quota exceeded (-o32 mode) 133: Transport endpoint is already connected (-n32 and -64 modes)

134: Stale NFS $\frac{1}{4}$ [®] file handle (-o32 mode) 134: Transport endpoint is not connected (-n32 and -64 modes)

135: Structure needs cleaning (-n32 and -64 modes)

136: Error 136 (-n32 and -64 modes)

137: Not a name file (-n32 and -64 modes)

138: Not available (-n32 and -64 modes)

139: Is a name file (-n32 and -64 modes)

140: Remote I/O error (-n32 and -64 modes)

141: Reserved for future use (-n32 and -64 modes)

- 142: Error 142 (-n32 and -64 modes)
 - 143: Cannot send after socket shutdown (-n32 and -64 modes)
 - 144: Too many references: cannot splice (-n32 and -64 modes)
 - 145: Connection timed out (-n32 and -64 modes)
 - 146: Connection refused (-n32 and -64 modes)
 - 147: Host is down (-n32 and -64 modes)
 - 148: No route to host (-n32 and -64 modes)
 - 149: Operation already in progress (-n32 and -64 modes)
 - 150: Operation now in progress (-n32 and -64 modes)
 - 151: Stale NFS file handle (-n32 and -64 modes)
- See the `perror(3C)` reference page for further information.

A.14.5 Memory Allocation

- **The behavior of the `calloc`, `malloc`, or `realloc` function if the size requested is zero (4.10.3).**

The `malloc` in `libc.a` returns a pointer to a zero-length space if a size of zero is requested. Successive calls to `malloc` return different zero-length pointers. If the library `libmalloc.a` is used, `malloc` returns 0 (the NULL pointer).

A.14.6 `abort` Function

- **The behavior of the `abort` function with regard to open and temporary files (4.10.4.1).**

Open files are not flushed, but are closed. Temporary files are removed.

A.14.7 `exit` Function

- **The status returned by the `exit` function if the value of the argument is other than zero, `EXIT_SUCCESS` or `EXIT_FAILURE` (4.10.4.3).**

The status returned to the environment is the least significant eight bits of the value passed to `exit`.

A.14.8 `getenv` Function

- **The set of environment names and the method for altering the environment list used by the `getenv` function (4.10.4.4).**

Any string can be used as the name of an environment variable, and any string can be used for its value. The function `putenv` alters the environment list of the application. For example,

```
putenv( ``MYNAME=foo`` )
```

This sets the value of the environment variable `MYNAME` to “foo.” If the environment variable `MYNAME` already existed, its value is changed. If it did not exist, it is added. The string passed to `putenv` actually becomes part of the environment, and changing it later alters the environment. Further, the string should not be space that was automatically allocated (for example, an auto array); rather, it should be space that is either global or `malloced`. For more information, see the `putenv(3C)` reference page.

It is not wise to alter the value of well-known environment variables. For the current list, see the `environ(5)` reference page.

A.14.9 `system` Function

- **The contents and mode of execution of the string passed to the `system` function (4.10.4.5).**

The contents of the string should be a command string, as if typed to a normal IRIX shell, such as `sh(1)`. A shell (`sh`) is forked, and the string is passed to it. The current process waits until the shell has completed and returns the exit status of the shell as the return value.

A.14.10 `strerror` Function

- **The contents of the error message strings returned by the `strerror` function (4.11.6.2).**

The string is exactly the same as the string output by `perror`, which is documented in Section A.14.4, page 177.

A.14.11 Time Zones and the `clock` Function

- **The local time zone and daylight saving time (4.12.1).**

Local time and daylight saving time are determined by the value of the TZ environment variable. TZ is set by `init` to the default value indicated in the file `/etc/TIMEZONE`, and this value is inherited in the environment of all processes. If TZ is unset, the local time zone defaults to GMT (Greenwich mean time, or coordinated universal time), and daylight saving time is not in effect. See the reference pages `ctime(3C)`, `time(2)`, `timezone(4)`, `environ(5)`, `getenv(3)`, and other related reference pages for the format of TZ.

- **The era for the clock function (4.12.2.1).**

`clock` counts seconds from 00:00:00: GMT, January 1, 1970. What was once known as Greenwich mean time (GMT) is now known as coordinated universal time, though the reference pages do not reflect this change yet. See the `ctime(3C)` reference page for further information.

A.15 Locale-Specific Behavior (F.4)

For information on locale-specific behavior, see the chapter titled “Internationalizing Your Application” in *Topics in IRIX Programming*. That chapter covers some locale-specific topics to consider when internationalizing an application. Topics include

- Overview of Locale-Specific Behavior
- Native Language Support and the NLS Database
- Using Regular Expressions
- Cultural Data

Also, that chapter describes setting a locale, location of locale-specific data, cultural items to consider, and GUI concerns.

For additional information on locale-specific behavior, refer to the *X/Open Portability Guide, Volume 3, “XSI Supplementary Definitions,”* published by Prentice Hall, Englewood Cliffs, New Jersey 07632, ISBN 0-13-685-850-3.

A.16 Common Extensions (F.5)

The following extensions are widely used in many systems, but are not portable to all implementations. The inclusion of any extension that can cause a strictly conforming program to become invalid renders an implementation nonconforming. Examples of such extensions are new keywords, or library

functions declared in standard headers or predefined macros with names that do not begin with an underscore. The Standard's description of each extension is followed by a definition of any Silicon Graphics support/nonsupport of each common extension.

A.16.1 Environment Arguments (F.5.1)

- In a hosted environment, the main function receives a third argument, `char *envp[]`, that points to a null-terminated array of pointers to `char`. Each of these pointers points to a string that provides information about the environment for this execution of the process (2.1.2.1.1).

This extension is supported.

A.16.2 Specialized Identifiers

- Characters other than the underscore `_`, letters, and digits, that are not defined in the required source character set (such as dollar sign `$`, or characters in national character sets) can appear in an identifier.

If the `-dollar` option is given to `cc`, then the dollar sign (`$`) is allowed in identifiers.

A.16.3 Lengths and Cases of Identifiers

- All characters in identifiers (with or without external linkage) are significant and case distinctions are observed (3.1.2).

All characters are significant. Case distinctions are observed.

A.16.4 Scopes of Identifiers (F.5.4)

- A function identifier, or the identifier of an object (the declaration of which contains the keyword `extern`) has file scope.

This is true of the compiler when invoked with `cc -cckr` (that is, when requesting traditional C). When compiling in ANSI mode (by default or with one of the ANSI options) function identifiers (and all other identifiers) have block scope when declared at block level.

A.16.5 Writable String Literals (F.5.5)

- **String literals are modifiable. Identical string literals shall be distinct (3.1.4).**

All string literals are distinct and writable when the `-use_readwrite_const` option is in effect. Otherwise, string literals may not be writable.

A.16.6 Other Arithmetic Types (F.5.6)

- **Other arithmetic types, such as `long long int` and their appropriate conversions, are defined (3.2.2.1).**

Yes.

A.16.7 Function Pointer Casts (F.5.7)

- **A pointer to an object or to void can be cast to a pointer to a function, allowing data to be invoked as a function (3.3.4). A pointer to a function can be cast to a pointer to an object, or to void, allowing a function to be inspected or modified (for example, by a debugger) (3.3.4).**

Function pointers can be cast to a pointer to an object, or to void, and vice versa.

Data can be invoked as a function.

Casting a pointer to a function to a pointer to an object or void does allow a function to be inspected. Normally, functions cannot be written to, because text space is read-only. Dynamically loaded functions are loaded (by a user program) into data space and can be written to.

A.16.8 Non-`int` Bit-Field Types (F.5.8)

- **Types other than `int`, `unsigned int`, and `signed int` can be declared as bitfields, with appropriate maximum widths (3.5.2.1).**

A bitfield can be any integral type in `-xansi` and `-cckr` modes. However, bitfields of types other than `int`, `signed int`, and `unsigned int` result in a warning diagnostic in `-ansi` mode.

A.16.9 `fortran` Keyword (F.5.9)

- The `fortran` declaration specifier can be used in a function declaration to indicate that calls suitable for Fortran should be generated, or that different representations for external names are to be generated (3.5.4.3).

The `fortran` keyword is not supported in this ANSI C. With `cc -cckr`, that keyword is accepted but ignored.

A.16.10 `asm` Keyword (F.5.10)

- The `asm` keyword can be used to insert assembly language code directly into the translator output. The most common implementation is via statement of the form `asm (character-string-literal)` (3.6).

The `asm` keyword is not supported.

A.16.11 Multiple External Definitions (F.5.11)

- There can be more than one external definition for the identifier of an object, with or without the explicit use of the keyword `extern`. If the definitions disagree, or more than one is initialized, the behavior is undefined (3.7.2).

With ANSI C, only one external definition of the object is permitted. If more than one is present, the linker (`ld(1)`) gives a warning message. The Strict Ref/Def model is followed (ANSI C Rationale, 3.1.2.2, page 23).

With `cc -cckr`, the Relaxed Ref/Def model is followed (ANSI C Rationale, 3.1.2.2, page 23): multiple definitions of the same identifier of an object in different files are accepted and all but one of the definitions are treated (silently) as if they had the `extern` keyword.

If the definitions in different source units disagree, the mismatch is not currently detected by the linker (`ld`), and the resulting program will probably not work correctly.

A.16.12 Empty Macro Arguments (F.5.12)

- A macro argument can consist of no preprocessing tokens (3.8.3).

This extension is supported. For example, one could define a macro such as

```
#define notokargs() macrovalue
```

A.16.13 Predefined Macro Names (F.5.13)

- **Macro names that do not begin with an underscore, describing the translation and execution environments, may be defined by the implementation before translation begins (3.8.8).**

This is **not** true for `cc -ansi`, which defines ANSI C. Only macro names beginning with two underscores or a single underscore followed by a capital letter are predefined by the implementation before translation begins. The name space is not polluted.

With `cc -cckr` (traditional C), a C preprocessor is used with a full set of the predefined symbols. For example, `sgi` is predefined.

With `cc -xansi` (which is the default for `cc`), an ANSI C preprocessor and compiler are used and a full set of predefined symbols is defined (including `sgi`, for example).

A.16.14 Extra Arguments for Signal Handlers (F.5.14)

- **Handlers for specific signals can be called with extra arguments in addition to the signal number.**

Silicon Graphics supports System V, POSIX, and BSD signal handlers. Extra arguments to the handler are available for your use. See the `signal` reference page.

A.16.15 Additional Stream Types and File-Opening Modes (F.5.15)

- **Additional mappings from files to streams may be supported (4.9.2), and additional file-opening modes may be specified by characters appended to the mode argument of the `fopen` function (4.9.5.3).**

There are no additional modes supported. There are no additional mappings. The UNIX approach is used, as mentioned in the ANSI C Rationale, Section 4.9.2, page 90.

A.16.16 Defined File Position Indicator (F.5.16)

- **The file position indicator is decremented by each successful call to the `ungetc` function for a text stream, except if its value was zero before a call (4.9.7.11).**

The Silicon Graphics C compiler supports only the one character of pushback guaranteed by the standard.

lint-style Comments [B]

Table 18, page 193, lists the lint-style comments available with the Silicon Graphics C compiler, along with a short description. See the lint(1) reference page for more details.

Table 18. lint-style Comments

Comment	Short Description
<code>/*PRINTFLIKEn*/</code>	Applies lint-style check to the first ($n-1$) arguments as usual. The n th argument is interpreted as a <code>printf</code> format string that is used to check the remaining arguments.
<code>/*SCANFLIKEn*/</code>	Applies lint-style check to the first ($n-1$) arguments as usual. The n th argument is interpreted as a <code>scanf</code> format string that is used to check the remaining arguments.
<code>/*ARGSUSEDn*/</code>	Applies lint-style check to only the first n arguments for usage; a missing n is taken to be 0 (this option acts like the <code>-v</code> option for the next function).
<code>/*VARARGSn*/</code>	Suppresses the usual checking for variable numbers of arguments in the following function declaration. The data types of the first n arguments are checked; a missing n is taken to be 0. The use of the ellipsis terminator (...) in the definition is suggested in new or updated code.
<code>/*NOTREACHED*/</code>	Stops comments about unreachable code when placed at appropriate points. (This comment is typically placed just after calls to functions like <code>exit</code>).
<code>/*REFERENCED*/</code>	Tells the compiler that the variable defined after comment is referenced.

Built-in Functions [C]

Table 19, page 195, list the built-in functions available in the Silicon Graphics C compiler, along with a short description.

Table 19. Built-in Functions

Intrinsic	Short Description
<code>void *__builtin_alloca(unsigned size)</code>	Returns a pointer to a specified number of bytes of uninitialized local stack space.
<code>float __builtin_fsqrt(float x)</code>	Computes the non-negative square root of a floating point argument.
<code>double __builtin_sqrt(double x)</code>	Computes the non-negative square root of a double argument.
<code>float __builtin_fabs(float x)</code>	Computes the absolute value of a float argument.
<code>double __builtin_dabs(double x)</code>	Computes the absolute value of a double argument.
<code>int __builtin_cast_f2i(float x)</code>	Treats float as int.
<code>float __builtin_cast_i2f(int x)</code>	Treats int as float.
<code>long long __builtin_cast_d2ll(double x)</code>	Treats double as long long.
<code>double __builtin_cast_ll2d(long long x)</code>	Treats long long as double.
<code>int __builtin_copy_dhi2i(double x)</code>	Copies high part of double to int.
<code>double __builtin_copy_i2dhi(int x)</code>	Copies int to high part of double.

Intrinsic	Short Description
<code>int __builtin_copy_dlo2i(double x)</code>	Copies low part of double to int.
<code>double __builtin_copy_i2dlo(int x, double y)</code>	Copies int to low part of double.
<code><type> __high_multiply (<type>, <type>)</code>	Multiplies two parameters as 32 (or 64) bit integers and returns the upper 32 (or 64) bits of a 64 (or 128) bit result. <type> can be signed or unsigned, int, long, or long long.

- ! operator, 56
 - != operator, 61
 - % operator, 58
 - & operator, 55, 62
 - fields and, 76
 - && operator, 63
 - * operator, 58
 - + operator, 56, 59
 - ++ operator, 56
 - += operator, 65
 - - operator, 54
 - character
 - in fscanf function, 177
 - operator, 56, 59
 - operator, 65
 - / operator, 58
 - < operator, 60
 - << operator, 60
 - <= operator, 60
 - = operator, 65
 - == operator, 61
 - > operator, 60
 - >= operator, 60
 - >> operator, 60
 - ? operator, 63
 - ^ operator, 62
 - | operator, 62
 - || operator, 63
 - ~ operator, 56
 - 32-bit mode
 - type differences, 39
 - 64-bit mode, 58
- A**
- ABI
 - N32 APO, 153
 - N64 APO, 153
 - abort function
 - effect on temporary files, 184
 - acpp
 - changes, 7
 - Additive operators
 - pointers and, 59
 - Address constant, 66
 - Address-of operator, 55
 - fields and, 76
 - affinity scheduling, 140
 - AND operator
 - bitwise, 62
 - logical, 63
 - ANSI C
 - allocating storage, 13
 - conversion rules, 45, 46
 - disambiguating identifiers, 10
 - floating point, 43
 - function prototype scope, 11
 - function prototype error, 18
 - guidelines, 2
 - identifiers, 31
 - libraries, 2
 - linkage discrepancies, 36
 - linker
 - warnings, 13
 - lint, 3
 - macro replacement, 8
 - name space, 12
 - name space discrepancies, 33
 - name spaces, 1, 33
 - preprocessor, 7
 - scoping differences, 10
 - strictly conforming programs, 1
 - string literals, 8
 - switches, 2
 - trigraph sequences, 26
 - value preserving integer promotion, 43

- warnings, 2
- ANSI C standard header files, 22
- ansi compiler option
 - external names and, 21
 - macros, 8
 - string literals, 8
 - tokens, 9
- APO, 153
 - licensing, 153
- Append mode stream
 - initial file position, 175
- Application Program Interface
 - See "API", 104
- Argument promotions, 53
- Argument type promotions
 - changes, 17
- Arguments
 - passing, 53
- Arithmetic constant expressions, 66
- Arithmetic conversions, 45
- Arithmetic expressions, 14
- Arithmetic types, 40
- Arithmetic value
 - 64-bit mode, 58
- Array
 - type required to hold maximum size, 162
- Array declarators, 80
- Arrays
 - and data affinity, 142
 - index
 - example, 151
 - processor arrays, 135
 - redistributed, 142
 - example, 150
 - reshaped
 - affinity clause, 142
 - nest clause, 142
 - reshaping, 135
 - variable length, 80
- arrays
 - query dimensions, 144
- asm keyword, 189
- Assert, 166

- diagnostic, 166
- Assignment operators, 64
 - +=, 65
 - =, 65
 - =, 65
- Associativity
 - examples, 49
- atan2, 167
- atan2f, 167
- Atomic compare-and-swap operation, 128
- Atomic fetch-and-op operations, 126
- Atomic lock and unlock operations, 129
- Atomic lock-release operation, 129
- Atomic lock-test-and-set operation, 129
- Atomic op-and-fetch operations, 127
- Atomic synchronize operation, 128
- auto, 70
- Auto keyword, 71
- Auto storage class, 71
- Auto-Parallelizing Option
 - See "APO", 153
- Autoblocking, 120
- Automatic page migration, 133
- Automatic parallelization
 - definition, 153
- Automatic storage duration, 37

B

- barrier, 119
- Barrier function, 119
- Behavior
 - locale-specific, , 186
- Binary streams
 - null characters in, 175
- Bitfield
 - diagnostics, 188
 - integral type, 188
- Bitfields, , 75, 162
 - integer types, 4
 - order of allocation, 164

signedness of, 163
 spanning unit boundary, 76
 straddling int boundaries, 164

Bits

bitfields, 76
 Bits per character, 158
 Bitwise and operator, 62
 Bitwise not operator, 56
 Bitwise operations
 signed integers, 160
 Bitwise OR operator
 inclusive, 62
 Bitwise or operator
 exclusive, 62

Blanks, 23

block distribution, , 136

Block scope
 definition, 32

Block statements, 93

block-cyclic distribution, 138

Blocking

 automatic, 118

Blocking slave threads, 117

blockproc, 118

break statements, 95, 98

Built-in functions, 195

 __builtin_alloca, 195
 __builtin_cast_d2ll, 195
 __builtin_cast_dhi2i, 195
 __builtin_cast_dlo2i, 196
 __builtin_cast_f2i, 195
 __builtin_cast_i2dhi, 195
 __builtin_cast_i2dlo, 196
 __builtin_cast_i2f, 195
 __builtin_cast_ll2d, 195
 __builtin_dabs, 195
 __builtin_fabs, 195
 __builtin_fsqrt, 195
 __builtin_sqrt, 195
 __high_multiply, 196
 __builtin_alloca, 195
 __builtin_cast_d2ll, 195
 __builtin_cast_dhi2i, 195

 __builtin_cast_dlo2i, 196

 __builtin_cast_f2i, 195

 __builtin_cast_i2dhi, 195

 __builtin_cast_i2dlo, 196

 __builtin_cast_i2f, 195

 __builtin_cast_ll2d, 195

 __builtin_dabs, 195

 __builtin_fabs, 195

 __builtin_fsqrt, 195

 __builtin_sqrt, 195

C

Cache

 misses, 131

Cache behavior, 132

calloc, 184

Case distinctions

 in identifiers, 157

case label, 95

case labels, 99

Case values

 maximum number of, 165

Cast operators, 57

Casting

 pointer to a function, 188

-cckr compiler option, 8

 external names and, 21

 tokens, 9

cConversions

 void, 47

char, 38

 default sign, 159

 unsigned vs. "plain", 159

Character

 white space, 155

Character constant, 165

Character constants, 25

 wide, 25

Character set, 157

Character types, 38

- Characters, 157
 - conversions to integer, 43
 - integer constants, 158
 - multibyte, 38, 157, 159
 - nongraphic, 25
 - number of bits, 158
 - shift states, 157
 - source set vs. execution set, 158
 - special, 25
 - wide, 159
 - initialization, 89
- CHRCLASS environment variable, 166
- CHUNK, 121
- clock function, 186
- Coding rules
 - #pragma directives, 105
- Comma operator, 65
- Comments, 23
- common compiler option, 13
- Communication
 - between processors, 122
- compare-and-swap operation, 128
- Compatible types, 16
- Compilation, 2
- Compilation mode
 - effect on names, 21
- Compile-time options
 - multiprocessing, 146
 - parallel programming, 146
- Compiler options
 - MP
 - check_reshape, 146
 - clone, 146
 - dsm, 146
- Compiler restrictions, 113
 - exception handling, 113
 - on #pragma pfor, 113
 - scoping, 115
- Compound assignment, 65
- Compound statements, 93
 - scope of declarations, 94
- Computation scheduling
 - user control, 132
- Conditional operator, 63
- const object, 4
- const type qualifier
 - qualifiers
 - const, 77
- Constant expression, 165
 - arithmetic, 66
- Constant expressions, 24, 66
 - address constant, 66
 - integral, 66
- Constants, 51
 - character, 25
 - enumeration, 27
 - floating, 27
 - integer, 24
 - long double precision, 43
 - types of, 24
 - wide character, 25
- Continue statements, 95
- continue statements, 97, 98
- Controlling expression
 - definition, 94
- Conversions, 43
 - arithmetic, 45
 - character, 43
 - floating-point, 43
 - function designators, 47
 - integer, 44
 - promotions, 45
 - lvalues, 47
 - pointer, 44
 - pointers, 47
 - rules
 - ANSI C, 46
 - Traditional C, 46
- cpp
 - changes, 7
- cyclic distribution, 137

D

Data

- explicit placement, 140
- placement in memory, 140
- placement of, 140

Data affinity, , 141

- example, 148
- formal parameter, 142
- redistributed arrays, 142

Data area names changes, 20

Data distribution, 135

- criteria, 139
- differences, 138
- example, 148
- regular, 135
- reshaping, 135
- rii_files, 146

Data placement

- user control, 132

Data structures

- irregular, 140

Date

- availability, 166
- __DATE__, 166

Daylight saving time, 186

Declarations

- as definitions, 70
- enumerations, 76
- implicit, 86
- multiple, 69
- structure, 72
- union, 72

Declarators

- array, 80
- definition, 78
- maximum number of, 165
- pointer, 79
- restrictions, 85
- syntax, 78

Decrement operator, 56

Default argument promotions, 53

Default labels, 95, 99

Definition

- declaration, 70

Definitions

- external, 101

Denoting a bitfield, 57

Derived types, 40

Device

- interactive, 157

Diagnostics

- classes, 156
- control, 156
- identification errors, 155
- return codes, 155, 156

dimensions

- arrays, 144
- of arrays, 144

Directives

- multiprocessing, 104, 143
- OpenMP, 104
- #pragma
 - coding rules, 105
- #pragma critical, 106
- #pragma distribute, 135
- #pragma distribute_reshape, 135
 - example, 147
- #pragma enter gate, 106
- #pragma exit gate, 106
- #pragma independent, 107
- #pragma page_place, 140
- #pragma parallel, 107
- #pragma pfor, 107, 140
- #pragma reshape, 135
- #pragma synchronize, 106
- preprocessing, 165
 - regular data distribution, 135

Disambiguating identifiers, 10

Disambiguating names, 31

distributed shared memory, 145

Distribution

- matrix
 - example, 147

Division

- integer, 58
 - sign of remainder, 160

- Division by zero, 58, 67

- do statements, 96

- Domain errors

 - return values, 167

- Double, 160

 - representation of, 160

- double, 39

- Double precision, 43

- `_DSM_BARRIER`, 145

- `_DSM_MIGRATION`, 145

- `_DSM_MIGRATION_LEVEL`, 145

- `_DSM_OFF` environment variable, 145

- `_DSM_PLACEMENT`, 145

- `_DSM_PPM`, 145

- `_DSM_VERBOSE`, 145

- `_DSM_WAIT`, 145

E

- else statements, 94

- `-ansi` switch, 2

- `-xansi` switch, 2

- enum, 74

 - changes, 12

- Enumeration constants, 27, 40, 76

 - changes, 12

- Enumeration types

 - type of int used, 164

- Enumeration variables, 76

- Enumerations, , 162

- Environment

 - altering, 185

 - names, 185

 - variables, 185

- Environment variables

 - `CHRCLASS`, 166

 - `CHUNK`, 121

 - `_DSM_BARRIER`, 145

 - `_DSM_MIGRATION`, 145

 - `_DSM_MIGRATION_LEVEL`, 145

 - `_DSM_OFF`, 145

 - `_DSM_PLACEMENT`, 145

 - `_DSM_PPM`, 145

 - `_DSM_VERBOSE`, 145

 - `_DSM_WAIT`, 145

 - gang scheduling, 122

 - `MP_BLOCKTIME`, 120

 - `MP_SCHEDTYPE`, 121

 - `MP_SET_NUMTHREADS`, 120

 - `MP_SETUP`, 120

 - `MP_SIMPLE_SCHED`, 146

 - `MP_SLAVE_STACKSIZE`, 122

 - `MP_SUGNUMTHD`, 121, 146

 - `MP_SUGNUMTHD_MAX`, 121

 - `MP_SUGNUMTHD_MIN`, 121

 - `MP_SUGNUMTHD_VERBOSE`, 121

 - `MPC_GANG`, 122

 - `PAGESIZE_DATA`, 146

 - `PAGESIZE_STACK`, 146

 - `PAGESIZE_TEXT`, 146

- environment variables

 - parallel programming, 145

- Environments, , 156

 - freestanding, 156

- Equality operators, 61

- `ERANGE` macro, 167

- `errno`, 167

- `errno` macro, 177

- Escape sequences, 157

 - hexadecimal, 26

- Examples

 - data affinity, 148

 - data distribution, 148

 - distributing a matrix, 147

 - irregular distributions, 151

 - multiprocessing, 147

 - parameter passing, 149

 - redistributed arrays, 150

 - thread affinity, 151

- Exception handling, 67

- Exception handling restrictions, 113

- Exclusive or operator, 62

-
- exit function, 184
 - Expression statements, 93
 - Expressions
 - ++, 54
 - , 54
 - constant, 66
 - postfix, 51
 - function calls, 52
 - structure references, 54
 - subscripts, 52
 - union references, 54
 - primary, 51
 - Extensions, , 2, 186
 - arithmetic types, 188
 - asm keyword, 189
 - bitfield types, 188
 - defined file position indicator, 191
 - empty macro arguments, 189
 - environment arguments, 187
 - external definitions, 189
 - fortran keyword, 189
 - function pointer casts, 188
 - identifiers, 187
 - predefined macro names, 190
 - scopes, 187
 - signal handlers, 190
 - specialized identifiers, 187
 - stream types, 190
 - writable string literals, 188
 - extern, 70
 - definitions, 13
 - function definitions, 101
 - External definitions, 101
 - External function definitions, 101
 - External linkage, 34
 - External names, 19
 - compiler options and, 21
 - External object definitions, 102
- F**
- fetch-and-op operations
 - __fetch_and_add, 126
 - __fetch_and_and, 126
 - __fetch_and_max, 126
 - __fetch_and_min, 126
 - __fetch_and_mpy, 126
 - __fetch_and_nand, 126
 - __fetch_and_or, 126
 - __fetch_and_sub, 126
 - __fetch_and_xor, 126
 - fgetpos function
 - errno on failure, 177
 - File buffering, 175
 - File position indicator
 - initial position, 175
 - filenames, 176
 - Files
 - renaming, 176
 - rri_files, 146
 - temporary, 184, 177
 - zero-length, 175
 - files
 - opening multiple times, 176
 - remove on an open file, 176
 - valid names, 176
 - Float
 - representation of, 160
 - float, 39
 - float variables, 14
 - float compiler option, 43
 - effect on conversions, 43
 - type promotions, 14
 - floating constants
 - definition, 27
 - Floating point constants, 15
 - Floating point conversions, , 16
 - Floating types, 40
 - Floating-point
 - conversions, 43
 - exception handling, 67
 - sizes, 39
 - types, 160
 - fmod, 167

- for statements, 96
- Formal parameter
 - and data affinity, 142
- fortran keyword, 189
- fprintf, 176
- fscanf, 176
- ftell function
 - errno on failure, 177
- fullwarn compiler option
 - scope, 11
- Function definitions, 101
- Function names
 - changes, 20
- Function pointer casts, 188
- Function prototype scope, 11, 32
- Function prototypes, 18
 - incompatible types, 19
 - inconsistent, 19
- Function scope, 11, 32
- Functions
 - built-in
 - See "Built-in functions", 195
 - calls, 52
 - declarators, 82
 - designators
 - conversions, 47
 - external
 - definition, 101
 - mixed use, 18
 - non-void, 4
 - nonprototyped, 54
 - prototyped, 53
 - prototypes, 82, 84
 - storage-class specifiers, 101
 - type, 102

G

- Gang scheduling, 122
- getenv function, 185
- goto statements, 97

H

- Header files
 - changes, 22
- Headers
 - standard, 21
- Hexadecimal escape sequences, 26
- __high_multiply, 196
- Hints, 3

I

- I compiler option, 165
- Identifiers, 23, 157
 - case distinctions, 157
 - definition, 23
 - disambiguating, 10
 - linkage, 12, 34
 - scope, 10
- IEEE
 - floating point, 44
- if statements, 94
- iIntegers
 - sizes, 159
- Implicit declarations, 86
- Include files, 165
 - maximum nesting depth, 166
 - quoted names, 165
- Inclusive or operator, 62
- Incompatible types
 - function prototypes and, 19
- Increment operator, 56
- Index array
 - example, 151
- Indirect references, 54
- Indirection operator, 55
- init-declarator-list
 - definition, 69
- Initialization, 88
 - and storage duration, 70
 - examples, 90

- structs, 89
 - unions, 89
 - initialization
 - aggregates, 89
 - int, 39
 - pointer conversion, 58
 - Integer
 - conversions to character, 43
 - divide-by-zero, 67
 - sizes, 39
 - Integer character constants, 158
 - Integer constants, 24
 - Integer division, 58
 - sign of remainder, 160
 - Integers
 - bitwise operations, 160
 - conversions, 160
 - exception conditions, 67
 - pointers, 44
 - ranges, 159
 - representations, 159
 - unsigned
 - conversions, 44
 - Integral constant expressions, 66
 - Integral promotions, 45, 46
 - Integral types, 40
 - Interactive device, 157
 - Internal linkage, 34
 - Intrinsics
 - example, 130
 - memory barrier, 125
 - synchronization, 125
 - IPA
 - automatic parallelization, 153
 - Irregular data structures, 140
 - Irregular distributions
 - example, 151
 - isalnum, 166
 - isalpha, 166
 - iscntrl, 166
 - islower, 166
 - isprint, 166
 - isupper, 166
 - Iteration statements, 96
 - controlling expression, 96
 - flow of control, 96
- J**
- Jump statements, 97
- K**
- Keywords
 - list of, 23
- L**
- Labeled statements, 99
 - Labels
 - case, 95
 - default, 95
 - name spaces, 99
 - libmalloc.a, 184
 - Libraries
 - C, 2
 - Math, 2
 - shared, 3
 - Library
 - multiprocessing library, libmp, 131
 - Library functions, , 166
 - prototypes, 21
 - Linkage, 70
 - determination of, 34
 - discrepancies, 36
 - external, 34
 - identifiers, 34
 - internal, 34
 - none, 34
 - Linker-defined names, 20
 - lint, 3
 - lint-style comments, 193

- `/*ARGUSED*/`, 193
- `/*NOTREACHED*/`, 193
- `/*PRINTFLIKE*/`, 193
- `/*REFERENCED*/`, 193
- `/*SCANF like*/`, 193
- `/*VARARGS*/`, 193

Literals, 28

LNO

- automatic parallelization, 153

Local time, 186

Locale-specific behavior, , 186

lock and unlock operations, 129

Lock example, 130

lock-release operation, 129

lock-test-and-set operation, 129

Logical operators

- AND, 63

- OR, 63

long, 39

long double, 27, 39, 160

long double precision, 43

long long, 39

LONGLONG macro, 159

lvalue

- conversions, 47

- definition, 40

lvalues, 57

M

Macros

- in -ansi mode, 8

- in -cckr mode, 8

- LONGLONG, 159

main

- arguments to, 156

malloc, 184

Manual parallelization, 153

Mapping

- character sequences, 165

Mathematics functions

- domain errors, 167

- underflow range errors, 167

Matrix

- example, 147

Memory

- acquire barrier, 125

- allocation, 184

- full barrier, 125

- placement of data, 140

- release barrier, 125

memory

- distributed shared, 145

Memory barrier, 125

Message passing, 122

Message Passing Toolkit (MPT), 122

Messages

- diagnostic, 155

- error, 155

- multiple definition, 13

minus, 56

- prefix, 56

MP

- compiler options, 146

- libmp library, 131

-MP

- check_reshape compiler option, 146

- clone compiler option, 146

- dsm compiler option, 146

mp_barrier, 119

mp_block, 117

MP_BLOCKTIME, 120

mp_blocktime, 118

mp_create, 117

mp_destroy, 117

mp_my_threadnum, 119

mp_numthreads, 118

MP_SCHEDTYPE, 121

MP_SET_NUMTHREADS, 120

mp_set_numthreads, 117, 118

mp_set_slave_stacksize, 119

mp_setlock, 119

MP_SETUP, 120

mp_setup, 117, 119

mp_shmem, 122
 MP_SIMPLE_SCHED, 146
 MP_SLAVE_STACKSIZE, 122
 mp_suggested_numthreads, 118
 MP_SUGNUMTHD, 121, 146
 MP_SUGNUMTHD_MAX, 121
 MP_SUGNUMTHD_MIN, 121
 MP_SUGNUMTHD_VERBOSE, 121
 mp_unblock, 117
 mp_unsetlock, 119
 MPC_GANG, 122
 Multibyte characters, 38, 157, 159
 Multiple definition messages, 13
 Multiplicative operators, 58
 Multiprocessing

- control of, 131
- data distribution, 135
- libmp library, 131
- options, 146
- rii_files directory, 146

 multiprocessing

- environment variables, 145

 Multiprocessing directives, 104, 143

N

Name

- definition, 23

 Name spaces, 1, 33

- changes, 12
- discrepancies, 33
- labels, 99

 Names

- compilation mode effect on, 21
- data area, 20
- external, 19
- functions
 - changes, 20
- linker-defined, 20

 Negation, 56
 Negative integers

- right shift on, 160

newline

- in text streams, 175

 Newlines, 23
 Non-ANSI C

- compilation mode, 3

 Non-void function, 4
 Nongraphic characters, 25
 Nonprototyped function declarations, 53
 -nostdinc compiler option, 165
 NUL character, 26
 Null, 26
 Null characters

- in binary streams, 175

 NULL pointer, 166
 NULL pointer constant, 47
 Null statement, 93
 NUM_THREADS, 120

O

Object

- definition, 40

 Objects

- definitions
 - external, 102
- external, 102
- types, 38

 offsetof() macro, 3
 op-and-fetch operations

- __add_and_fetch, 127
- __and_and_fetch, 127
- __max_and_fetch, 127
- __min_and_fetch, 127
- __mpy_and_fetch, 127
- __nand_and_fetch, 127
- __or_and_fetch, 127
- __sub_and_fetch, 127
- __xor_and_fetch, 127

 OpemMP

- multiprocessing directives, 104

 Operator

- bitwise not, 56
- operator, 56
- ++operator, 55
- Operators
 - !, 56
 - %, 58
 - &, 62
 - *, 58
 - +, 59
 - unary, 56
 - ++
 - prefix, 56
 - , 59
 - unary, 56
 - /, 58
 - <<, 60
 - >>, 60
 - ~, 56
 - additive, 59
 - address-of, 55
 - AND, 62
 - assignment, 64
 - +=, 65
 - =, 65
 - =, 65
 - associativity, 49
 - bitwise
 - AND, 62
 - cast, 57
 - comma, 65
 - conditional, 63
 - conversions, 43
 - equality, 61
 - evaluation, 49
 - exclusive or, 62
 - grouping, 49
 - inclusive OR, 62
 - indirection, 55
 - list of, 28
 - logical
 - AND, 63
 - minus, 56
 - multiplicative, 58

- OR
 - exclusive, 62
 - inclusive, 62
 - logical, 63
- order of evaluation, 49
- precedence, 49
- relational, 60
- shift, 60
- sizeof, 57
- unary, 55
- OPT
 - alias=disjoint, 80
 - alias=restrict, 80
- Optimization
 - APO, 153
- Options
 - control multiprocessing, 146
 - parallel programming, 146
- OR operator
 - exclusive, 62
 - inclusive, 62
 - logical, 63
- Order of evaluation
 - operators, 49
- Origin200
 - parallel programming, 131
- Origin2000
 - memory model, 131
 - parallel programming, 131
 - performance tuning, 131
 - programming examples, 147
- Overflow handling, 67

P

- %p conversion
 - in fprintf function, 176
 - in fscanf function, 176
- PAGESIZE_DATA, 146
- PAGESIZE_STACK, 146
- PAGESIZE_TEXT, 146

- Parallel computing forum (PCF), 105
- Parallel Fortran
 - communication between threads, 122
- Parallel programming
 - examples, 147
 - options, 146
- parallel programming
 - environment variables, 145
- Parallel programming on Origin200, 131
- Parallel programming on Origin2000, 131
- Parallel programs
 - improving performance, 131
 - tuning, 131
- Parallel reduction operations, 109
- Parallel regions, 105
 - work-sharing constructs, , 106
- Parallelization
 - automatic, 153
 - definition, 153
 - manual, 153
- Parameter list, 83
- Parameter passing
 - example, 149
- Parameter, formal
 - data affinity, 142
- Parenthesized expressions, 51
- Passing arguments, 53
- Performance
 - and cache behavior, 132
 - directives, 143
 - improving Origin2000, 131
- Performance tuning
 - examples, 147
- perror function, 177
- Pointer
 - convert to int, 58
 - truncation of value, 58
- Pointer constant
 - NULL, 47
- Pointer declarators, 79
- Pointers
 - additive operators on, 59
 - command options, 80
 - conversion to int, 162
 - conversions, 47
 - differences of, 162
 - integer additions and, 44
 - qualifiers, 79
 - restricted, 79
 - to qualified types, 79
 - to void, 47
- pointers
 - casting to int, 162
- Postfix expressions, 51
 - ++, 54
 - , 54
 - function calls, 52
 - indirect references, 54
 - structure references, 54
 - subscripts, 52
 - union references, 54
- pow, 167
- powf, 167
- #pragma critical, 106
- #pragma directives
 - changes from Fortran directives, 105
 - coding rules, 105
- #pragma distribute_reshape, 147
- #pragma enter gate, 106
- #pragma exit gate, 106
- #pragma independent, 107
- #pragma page_place, 140
- #pragma parallel, 107
- #pragma pfor, 107, 140
 - restrictions, 113
- #pragma synchronize, 106
- Precedence
 - examples, 49
- Precedence of operators, 49
- Precision, 27
- Preprocessing directives, 165
- Preprocessor
 - changes, 7
- Primary expressions, 51
- Processor

- arrays, 135
- Programming hints, 3
- Promotions
 - arguments, 53
 - arithmetic expressions, 14
 - floating-point, 14
 - integral, 14, 45
- Prototyped function declarations, 53
- Prototyped functions, 84
- Prototypes, 82
 - function, 18
 - incompatible types, 19
 - inconsistent, 19
- ptrdiff_t, 162
- Punctuators
 - definition, 29
 - list of, 29
- putenv function, 185

Q

- Quad precision, 160
- Qualified objects, 4
- Qualifiers, 77
 - access to volatile, 164
 - volatile, 77
- query intrinsics
 - distributed arrays, 144

R

- Ranges
 - floating points, 161
 - integers, 159
- realloc, 184
- Recommendations
 - coding practices, 3
 - things to avoid, 4
- Redistributed arrays
 - example, 150
- Reduction

- example, 112
 - on user-defined type in C++, 111
- reduction clause, 109
- Reduction operations, 109
- Register, 82
 - 32 mode, 162
 - optimized code, 162
- register
 - function declaration lists, 101
- Register keyword, 71
- Register storage-class specifier, 162
- Registers, , 162
- Regular data distribution, 133
 - vs. reshaped distribution, 138
- Relational operators, 60
- Remainder
 - sign of, 160
- remove function
 - on an open file, 176
- rename function, 176
- Reserved keywords, 23
- Reshaped arrays, 135
- Reshaped data distribution, 133
 - vs. regular distribution, 138
- __restrict type qualifier, 79
 - example, 79
- Result type
 - definition, 45
- return statements, 98
- Right shift
 - on negative integers, 160
- rii_files directory, 146
- Rounding
 - type used, 161
- Routines, 119
 - barrier, 119
 - mp_barrier, 119
 - mp_block, 117
 - mp_blocktime, 118
 - mp_create, 117
 - mp_destroy, 117
 - mp_my_threadnum, 119

- mp_numthreads, 118
- mp_set_numthreads, 117, 118
- mp_set_slave_stacksize, 119
- mp_setlock, 119
- mp_setup, 117
- mp_suggested_numthreads, 118
- mp_unblock, 117
- mp_unsetlock, 119
- ussetlock, 119
- usunsetlock, 119

Run-time library routines

- See "Routines", 117

S

- Scalar types, 40
- Scheduling, 121
 - simple, 142
- Scheduling methods
 - between processors, 122
 - gang, 122
- Scope
 - block, 32
 - changes, 11
 - definition, 31
 - file, 32
 - function, 32
 - function prototype, 32
- Scoping
 - changes, 10
- Scoping restrictions, 115
- Selection statements, 94
- setlocale, 166
- shared memory
 - distributed, 145
- Shift operators, 60
- Shift states, 157
- shmem
 - See "mp_shmem", 122
- shmem routines
 - rules, 123
- short, 39

- SIGCLD, 118
- Signal-catching functions
 - valid codes, 171
- Signals
 - semantics, 169
 - set of, 167
 - SIGCLD, 118
- signals
 - default handling, 174
- Simple assignment, 65
- Single precision, 43
- size
 - arrays, 144
 - size_t, 57, 162
 - sizeof, 57, 85, 162
 - type of result, 57
- Sizes
 - floating points, 161
 - integers, 159
- Slave processes
 - stack size, 122
- Slave threads
 - blocking, 117, 118
- &space, 78
- special characters, 25
- Spin-wait lock example, 130
- sprocksp, 119
- Stack size, 120, 122
- Standard headers, 21
- Standards
 - ANSI C, 1
- Statements
 - block, 93
 - break, 95, 98
 - compound, 93
 - scope of declarations, 94
 - continue, 97, 98
 - do, 96
 - else, 94
 - expression, 93
 - for, 96
 - goto, 97

- if, 94
- iteration, 96
- jump, 97
- labeled, 99
- null, 93
- return, 98
- selection, 94
- switch, 94, 95
- while, 96
- static
 - function definitions, 101
- Static keyword, 70
- Static storage duration, 37, 70
- stdarg, 4, 83
- stderr, 157
- Storage class sizes, 39
- storage class sizes, 39
- Storage class specifiers, 70
- Storage duration, 70
 - auto, 71
 - automatic, 37
 - static, 37, 70
- strerror function, 185
- String literals, 5, 28, 51, 188
 - wide, 28
 - wide characters, 89
- Struct
 - namespace
 - changes, 12
- struct, 72
 - initialization, 89
 - members
 - restrictions, 73
- Structs
 - alignment, 163
- Structure
 - declaration, 72
 - indirect references, 54
 - members
 - restrictions, 73
 - references, 54
- Structure designators, 4
- Structures, , 162
 - alignment, 163
 - padding, 163
- structures
 - initialization, 89
- Subroutines
 - See "Routines", 119
- Subscripts
 - in postfix expressions, 52
- Switch statements
 - maximum number of case values, 165
- switch statements, 94, 95
 - labels, 99
- Switches, 2
 - ansi, 2
 - xansi, 2
- Synchronization intrinsics, 111, 125
- Synchronize operation
 - __synchronize, 128
- system function, 185

T

- Tabs, 23
- Temporary files, 184, 177
- Text stream
 - last line, 175
 - newline, 175
- Text streams
 - writes on, 175
- Thread
 - master, 119
 - slave, 119
- Thread affinity, , 141
 - example, 151
- Threads
 - and processors, 122
- Time
 - availability, 166
 - clock function, 186
 - daylight savings, 186
 - local, 186

- Time zone, 186
 - __TIME__, 166
 - Token
 - definition, 23
 - Token concatenation, 9
 - Tokens
 - classes of, 23
 - in -ansi mode, 9
 - in -cckr mode, 9
 - Traditional C
 - allocating storage, 13
 - conversion rules, 45, 46
 - enumeration constants, 12
 - floating point, 43
 - function prototype error, 18
 - identifiers, 31
 - incompatibilities, 83
 - linkage discrepancies, 36
 - name space discrepancies, 33
 - scope, 11
 - scoping differences, 10
 - unsigned preserving integer promotion, 43
 - Traditional c
 - compiler option, 3
 - Translation, , 155
 - Trigraph sequences, 26
 - Truncation
 - direction of, 161
 - pointer value, 58
 - type used, 161
 - tType qualifiers
 - __restrict, 79
 - Type, 15
 - Type names, 85
 - Type qualifiers, 77
 - Type specifiers
 - list of, 71
 - typedef, 70, 72, 74, 77, 87
 - Types, 38
 - 32-bit mode, 39
 - 64-bit mode, 39
 - arithmetic, 40
 - changes, 14
 - character, 38
 - compatibility, 14, 17
 - derived, 40
 - differences, 39
 - float, 43
 - floating, 44
 - floating-point, 39
 - int, 59
 - integer, 39
 - integral, 44
 - long double, 27
 - multibyte characters, 38
 - promotion in arithmetic expressions, 14
 - promotion rules, 14
 - promotions
 - arguments, 17
 - floating-point, 14
 - integral, 14
 - sizes, 39
 - unsigned char, 38
 - variably modified, 69
 - void, 40
 - types
 - integral, 40
 - TZ environment variable, 186
- ## U
- Unary operators, 55
 - unblockproc, 118
 - Underflow handling, 67
 - Underflow range errors
 - math functions, 167
 - Union
 - indirect references, 54
 - namespace
 - changes, 12
 - references, 54
 - union, 72
 - declaration, 72
 - initialization, 89

- members
 - restrictions, 73
- Unions, , 162
 - 32-bit mode, 163
 - 64-bit mode, 163
 - accessing members, 163
- unlock operation, 129
- unsigned char, 38
 - default, 159
- unsigned integers
 - conversions, 44
- usconfig, 119
- use_readwrite_const option, 188
- User name space, 1
- usinit, 119
- ussetlock, 119
- usunsetlock, 119

V

- valid filenames, 176
- Variable length array
 - as specifier for type declarator, 69
- Variable length arrays, 80
- Variables
 - float, 14
- variables
 - to control multiprocessing, 145
- void, 40, 83
 - conversions, 47
 - pointers to, 47

- return statements, 98
- Volatile, 77
- volatile object, 4
- Volatile-qualified types
 - access to, 164

W

- Warnings, 2
- while statements, 96
- White space, 23, 155
- Wide characters, 159
- Wide string literals, 28
- Words
 - alignment, 163
 - size, 163
- Work-sharing constructs, 106
- write(), 20

X

- xansi compiler option
 - external names and, 21
- Xlocal, 122

Z

- Zero-length files, 175