

HD I/O Option Board Owner's Guide

Document Number 007-3968-002

CONTRIBUTORS

Written by Carolyn Curtis

Illustrated by Cheri Brown, Dan Young, Dany Galgani, and Carolyn Curtis

Production by Carlos Miqueo

Engineering contributions by Bill Warner, Ted Marsh, Greg Sadowski, Ed Miszkiewicz, Kirk Knapp, Michael Poimboeuf, and Scott Pritchett

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

© 1999, Silicon Graphics, Inc.— All Rights Reserved

The contents of this document may not be copied or duplicated in any form, in whole or in part, without the prior 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 Unpublished rights reserved under the Copyright Laws of the United States.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.

FCC Warning

This equipment has been tested and found compliant with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

Attention

This product requires the use of external shielded cables in order to maintain compliance pursuant to Part 15 of the FCC Rules.

European Union Statement

This device complies with the European Directives listed on the “Declaration of Conformity” which is included with each product. The CE mark insignia displayed on the device is an indication of conformity to the aforementioned European requirements.



International Special Committee on Radio Interference (CISPR)

This equipment has been tested to and is in compliance with the Class A limits per CISPR publication 22, Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment; Germany’s BZT Class A limits for Information Technology Equipment; and Japan’s VCCI Class 1 limits.

Canadian Department of Communications Statement

This digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus as set out in the Radio Interference Regulations of the Canadian Department of Communications.

Attention

Cet appareil numérique n’émet pas de perturbations radioélectriques dépassant les normes applicables aux appareils numériques de Classe A prescrites dans le Règlement sur les interférences radioélectriques établi par le Ministère des Communications du Canada.

VCCI Class 1 Statement for Japan

この装置は、情報処理装置等電波障害自主規制協議会 (VCCI) の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Chinese Class A Warning

警告使用者：

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Silicon Graphics, SGI, the Silicon Graphics logo, OpenGL, Origin, and IRIS are registered trademarks and IRIX, XIO, Onyx, Onyx2, Origin200, Origin2000, Graphics Library, REACT, RealityEngine, XFS, and Sirius Video are trademarks of Silicon Graphics, Inc.

Philips is a registered trademark and Spirit Datacine is a trademark of Philips Electronics, N.V. D19.

QuickTime is a registered trademark of Apple Computer, Inc.

Contents

List of Figures ix

List of Tables xi

About This Guide xiii

Audience xiii

Structure of This Guide xiii

Other Documents xiv

Conventions Used in This Guide xv

- 1. Features and Capabilities** 1
 - HD I/O Features 1
 - Supported Video Formats 1
 - Genlock and Timing Features 2
 - Other Features 2
 - HD I/O Panel and Cable 5
- 2. Programming the HD I/O Option** 7
 - VL Basics for the HD I/O Option Board 7
 - VL Concepts 8
 - VL Syntax Elements 9
 - VL Object Classes 9
 - VL Nodes for the HD I/O Option 10
 - VL Data Transfer Functions 12
 - HD I/O Data Flow 13

- HD I/O Controls 14
 - Setting Controls 15
 - HD I/O Control Summary 16
 - VL_TIMING 19
 - VL_FORMAT 20
 - VL_PACKING 21
 - VL_COLORSPACE 22
 - VL_CAP_TYPE 27
 - VL_SIZE and VL_OFFSET 28
 - VL_ZOOM 28
- Field Dominance 29
- Automatically Correcting for Output Underflow 30
- Capturing Graphics to Video 31
- HD I/O Events 31
- Reporting 32
- Examples 32
 - Capture to Memory for Disk Recording 33
 - Playback From Memory for Disk Playback 33
 - Capture to Memory for Graphics 34
- 3. Synchronizing Data Streams and Signals 35**
 - Using UST, MSC, and Buffered Media Streams for Synchronization 35
 - Media Library Interfaces for UST and MSC 38
- A. HD I/O Option Board Specifications 39**
 - Cable Connectors 40
 - GPI Interface 44
 - GPI Connector 44
 - GPI Transmitter 46
 - GPI Receiver 47
 - Genlock 48

- B. Setting Up the HD I/O Board for Your Video Hardware 49**
 - Setting Up Digital Source Video 50
 - Setting Up the Output (Drain) 52
 - Setting Up Sync 54
 - Setting Up Internal Sync 54
 - Setting Up External Sync 55
 - Saving Settings 56
- C. Pixel Packings and Color Spaces 57**
 - HD I/O Pixel Packings 57
 - Packings and Color Spaces 57
 - Packing Diagram Conventions 58
 - Packings and Library Tokens 60
 - Packing Naming Conventions 60
 - 16-Bit Pixel Packings 62
 - 20-Bit Pixel Packings 63
 - 24-Bit Pixel Packings 64
 - 32-Bit Pixel Packings 65
 - Sampling Patterns 74
 - 4:4:4 and 4:4:4:4 Sampling 74
 - 4:2:2 and 4:2:2:4 Sampling 75
- D. Programming Methods for Real-Time Digital Media Recording and Playback 77**
 - Digital Media Buffers 78
 - Direct I/O 78
 - Multiprocessing 80
 - Asynchronous I/O 81
- E. Installing HD I/O Software on a New Disk 83**
 - Index 85

List of Figures

Figure 1-1	HD I/O Option Board	3
Figure 1-2	HD I/O Board Diagram, Simplified	4
Figure 1-3	HD I/O Board Connectors	5
Figure 1-4	HD I/O Cables	5
Figure 2-1	Simple VL Path	8
Figure 2-2	Data Flow, Input	13
Figure 2-3	Data Flow, Output	13
Figure 2-4	Control Flow	14
Figure 2-5	Color-Space Conversion Example	26
Figure 2-6	Fields and Frames for SMPTE 274M	29
Figure A-1	HD I/O Cables	40
Figure A-2	GPI Connector	44
Figure A-3	GPI Pinouts	44
Figure A-4	GPI Pins and HD I/O Video Pipes	45
Figure A-5	GPI Transmitter Electrical Specifications	46
Figure A-6	GPI Receiver Electrical Specifications	47
Figure A-7	Genlock BNCs	48
Figure B-1	HD I/O Ports	49
Figure B-2	HD I/O Cables	50
Figure B-3	Selecting Digital Input Video Format in vcp	51
Figure B-4	Selecting Video Drain Format	53
Figure B-5	Setting Standalone or Genlock Sync	54
Figure C-1	VL_PACKING_444_8	58
Figure C-2	VL_PACKING_242_8	62
Figure C-3	VL_PACKING_R242_8	62
Figure C-4	VL_PACKING_242_10	63
Figure C-5	VL_PACKING_R242_10	63

Figure C-6	VL_PACKING_444_8	64
Figure C-7	VL_PACKING_R444_8	65
Figure C-8	VL_PACKING_4444_8	66
Figure C-9	VL_PACKING_R4444_8	67
Figure C-10	VL_PACKING_R0444_8	68
Figure C-11	VL_PACKING_0444_8	69
Figure C-12	VL_PACKING_4444_10_10_10_2	70
Figure C-13	VL_PACKING_R4444_10_10_10_2	70
Figure C-14	VL_PACKING_2424_10_10_10_2Z	71
Figure C-15	VL_PACKING_R2424_10_10_10_2Z	71
Figure C-16	VL_PACKING_242_10_in_16_L	72
Figure C-17	VL_PACKING_242_10_in_16_R	72
Figure C-18	VL_PACKING_R242_10_in_16_L	73
Figure C-19	VL_PACKING_R242_10_in_16_R	73
Figure C-20	4:4:4 Sampling	74
Figure C-21	4:2:2 Sampling	75

List of Tables

Table 1-1	50-Pin Cable Connectors	6
Table 2-1	HD I/O Node Controls	16
Table 2-2	Controls for the HD I/O Option	17
Table 2-3	Values for VL_TIMING	19
Table 2-4	Genlock Sync Source and Timing	20
Table 2-5	VL_FORMAT and VL_COLORSPACE Combinations Supported	23
Table 2-6	Color-Space Values	24
Table 2-7	HD I/O Events	32
Table A-1	Panasonic 50-Pin Connector Pinout (HD-D5)	41
Table A-2	Philips 50-Pin Connector Pinout (Spirit DataCine)	42
Table A-3	LINK A and LINK B Usage in 4:2:2:4 Mode	42
Table A-4	LINK A and LINK B Usage in RGBA Mode	43
Table A-5	LINK A and LINK B Usage in 4:4:4:4 Mode	43
Table A-6	GPI Pinouts	45
Table A-7	GPI Transmitter Electrical Specifications	46
Table A-8	GPI Receiver Input Optoisolator Electrical Specifications	48
Table C-1	HD I/O Packings	61

About This Guide

The SGI HD I/O option board allows the Silicon Graphics Onyx2 supercomputing workstation to generate and receive uncompressed high-definition television (HDTV) signals in real time. The board is also supported in SGI Origin 2000 servers.

Note: This option board requires IRIX 6.5.4 or later; earlier versions of IRIX do not recognize the board.

Features of this option are controlled with the Video Library (VL). VL device-independent calls and controls are explained in the *Digital Media Programming Guide* (007-1799-060 or later; online only).

Audience

This guide was written for the sophisticated video user in a professional or research environment. You should be familiar with video standards, the operation of the Silicon Graphics workstation or server, and the VL information in the *Digital Media Programming Guide*.

Structure of This Guide

This guide includes the following chapters and appendices:

- Chapter 1, “Features and Capabilities,” outlines the main components of the HD I/O option.
- Chapter 2, “Programming the HD I/O Option,” describes using the VL to accomplish common specific tasks.
- Chapter 3, “Synchronizing Data Streams and Signals,” explains how to use unadjusted system time (UST) and media stream count (MSC) for the HD I/O board.

- Appendix A, "HD I/O Option Board Specifications," summarizes technical specifications for the option board.
- Appendix B, "Setting Up the HD I/O Board for Your Video Hardware," describes connecting video equipment to HD I/O board connectors and using the control panel *vcp* to configure the board for the equipment.
- Appendix C, "Pixel Packings and Color Spaces," sets forth all packing formats used by the HD I/O hardware.
- Appendix D, "Programming Methods for Real-Time Digital Media Recording and Playback," explains programming concepts, such as real-time disk I/O, and gives examples.
- Appendix E, "Installing HD I/O Software on a New Disk," gives the steps for installing the software, should that become necessary in the case of a new disk.

An index completes this guide.

Other Documents

Besides this guide, *Digital Media Connections* (007-3525-003 or later) is shipped with the HD I/O option board. The *Digital Media Programming Guide* (007-1799-060) is available with the IRIX digital media development environment software (*dmedia_dev*); the online version of this manual is included with IRIX 6.5.4 and later.

It is also a good idea to have your system owner's guide available. If you do not have these guides handy, the information is also online in the following locations:

- IRIS InSight Library: from the Toolchest, choose Help > Online Books > SGI EndUser or SGI Admin, and select the applicable guide.
- Technical Publications Library: if you have access to the Internet, enter the following URL in your Web browser location window:
<http://techpubs.sgi.com/library/>

Once you are in the library, choose Catalogs > Hardware Catalog > and look under the Owner's Guides for the applicable owner's guide. For software guides, look on the bookshelf for the applicable IRIX version.

Conventions Used in This Guide

In command syntax descriptions and examples, square brackets ([]) surrounding an argument indicate an optional argument. Variable parameters are in *italics*. Replace these variables with the appropriate string or value.

In text descriptions, IRIX filenames are in *italics*.

Helvetica Bold font is used for labels on hardware, such as for ports and LEDs on the I/O panel.

Messages and prompts that appear on-screen are shown in `typewriter` font. Entries that are to be typed exactly as shown are in **boldface typewriter font**.

In each chapter or appendix in which the *Digital Media Programming Guide* is referenced, it is referred to by its full title at the first occurrence and thereafter as the DMPG.

Features and Capabilities

The Silicon Graphics HD I/O option board is an XIO card for the Silicon Graphics Onyx2 supercomputing workstation for interfacing to American Television Standards Committee high-definition television video formats, so that it can generate and receive uncompressed HDTV signals in real time.

This chapter consists of the following sections:

- “HD I/O Features” on page 1
- “HD I/O Panel and Cable” on page 5

HD I/O Features

Features include the following:

- “Supported Video Formats” on page 1
- “Genlock and Timing Features” on page 2
- “Other Features” on page 2

Supported Video Formats

The HD I/O board supports all video formats defined by the Advanced Television Standards Committee (ATSC) that have a pixel clock of up to 74.25 MHz:

- 16 x 9 aspect ratio with 1920, 1280 and 720 active pixels in line, depending on video format
- Support for 24/1.001, 24, 25, 30/1.001, 50, and 60/1.001 Hz vertical rates

Examples of supported formats are SMPTE 274M (interlaced), SMPTE 296 (progressive), SMPTE 295M (interlaced), SMPTE 260M, and SMPTE 293M.

HD I/O formats supported for this release are as follows:

- 1920x1080i@59.94Hz 4:2:2 (8,10-bit)
- 1280x720p@59.94Hz 4:2:2 (8,10-bit)

Both formats are used for content creation and telecine output, and support serial-parallel conversion. Subsequent releases will support additional formats.

Genlock and Timing Features

Genlock and timing features are all designed to SMPTE 274M timing specifications:

- Genlock to external analog and internal digital sync reference inputs:
 - Trilevel analog sync at high-definition (HD) rates
 - Bilevel analog sync at standard-definition (SD) rates (with and without output format timing conversion of SD to HD rates)
- Genlock to digital input: locking digital output to jitter-attenuated digital input clock
- Standalone free-run timing:

Uncalibrated free-run mode provides 25 ppm-accurate free-run frequency. After calibration, free-run mode provides higher than 10 ppm-accurate free-run frequency, which can be recalibrated for local references and conditions to as high as 1 ppm accuracy. For information on calibration, see the HD I/O release notes.

Other Features

Other salient features of the HD I/O option board are:

- Bit-parallel ECL differential interface (special connectors, cable, adapters)
- Automatic adjustment of input phase between Link A and Link B (up to eight input clocks)
- YCrCb 10 or 8 bits per component (4:2:2 or 4:4:4 sampling rates)
- Alpha channel support: 10 or 8 bits (full sampling rate; special connector/cable required)
- Video interface support for RGB 10 or 8 bits (special connector/cable required)

- Support for up to 13 bits RGB in memory
- Real-time transparent color-space conversion and key scaling
- User-programmable horizontal and vertical phase adjustment of the output video
- UST support on input and output
- General-Purpose Interface (GPI) connector, two channels in and two channels out
- Gamma correction support through 13-bit-wide lookup tables
- Board internal loopback mode

Figure 1-1 shows the board.




Figure 1-1 HD I/O Option Board

Figure 1-2 is a simplified top-level diagram of HD I/O option board.

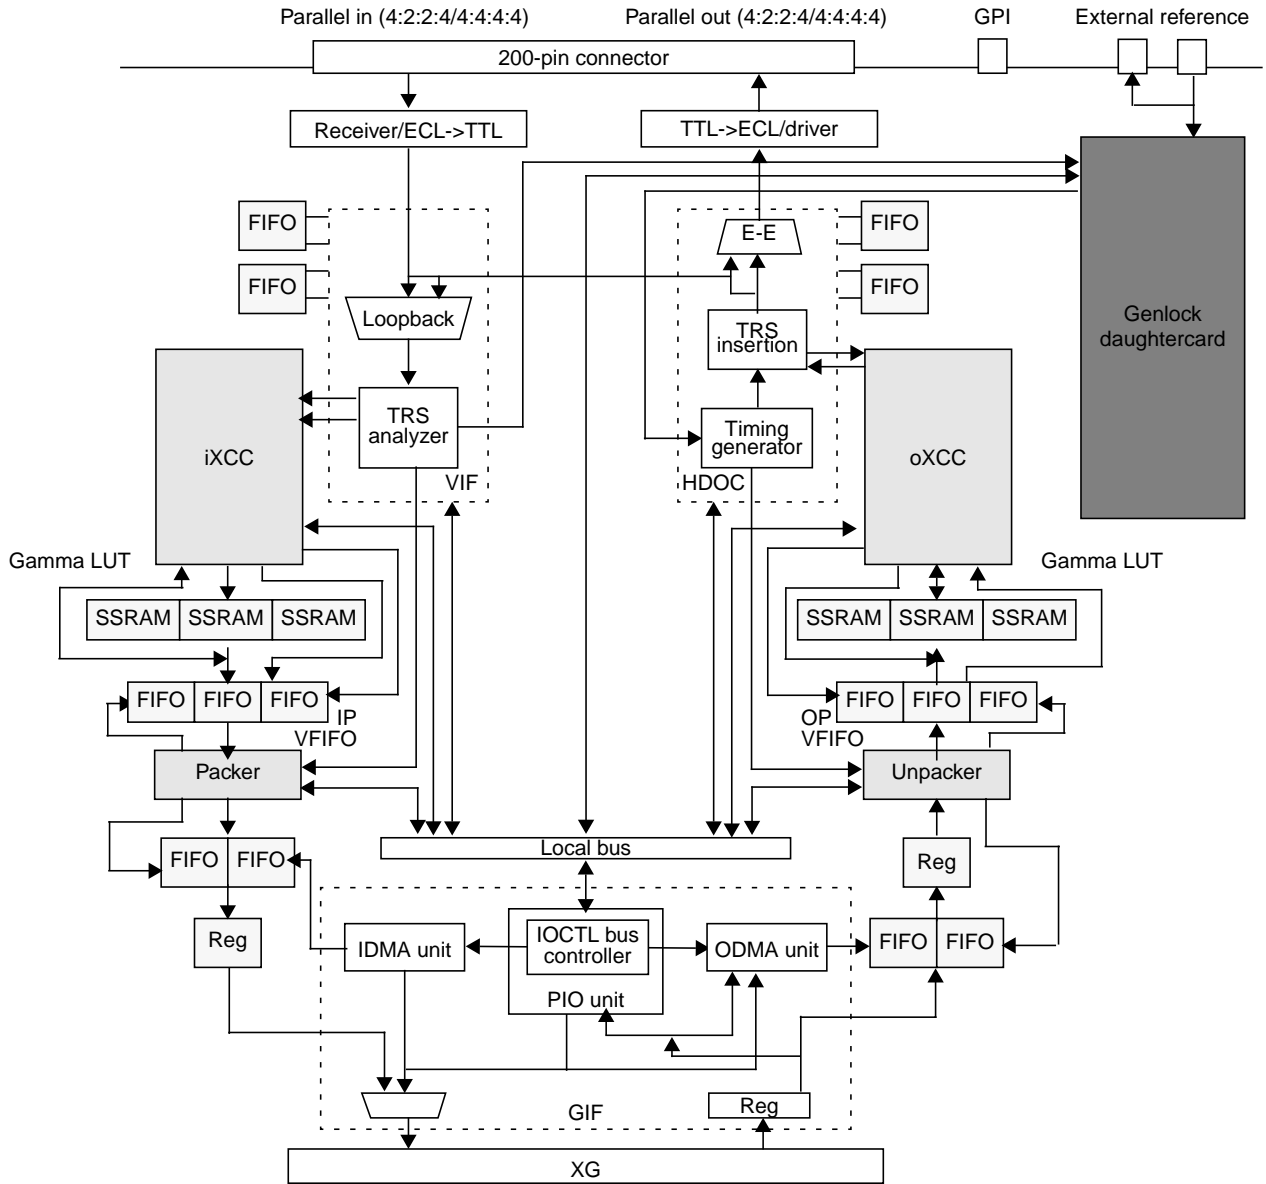


Figure 1-2 HD I/O Board Diagram, Simplified

HD I/O Panel and Cable

Figure 1-3 shows connectors on the HD I/O front panel.




Figure 1-3 HD I/O Board Connectors

Figure 1-3 shows the two multiheaded cables included with the board; each has four 50-pin connectors for link A input, link B input, link A output, and link B output. The 50-pin connectors differ for each type of cable, following the Panasonic and Philips 50-pin video equipment interface standard.

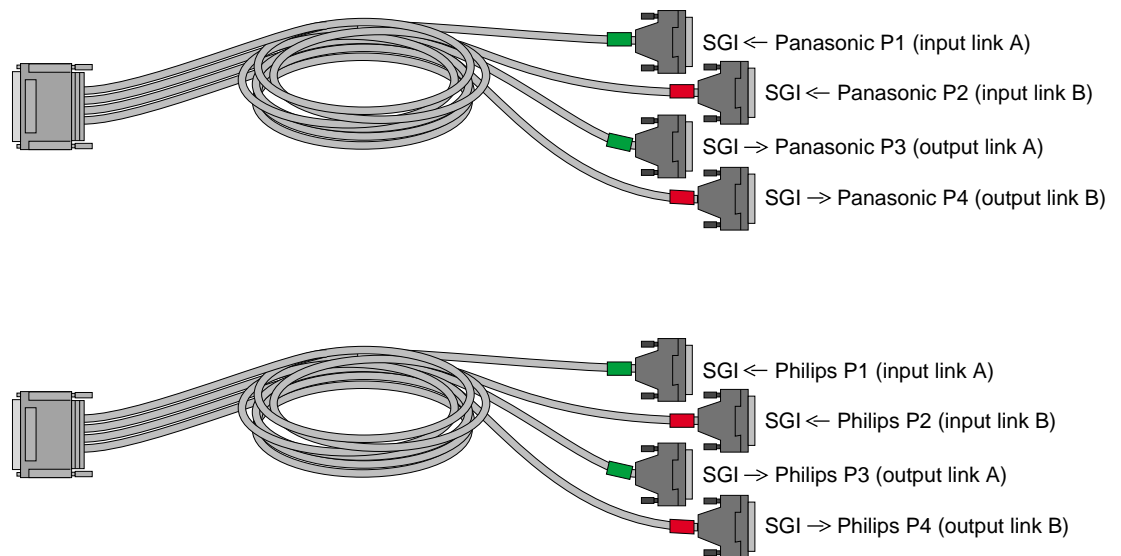


Figure 1-4 HD I/O Cables

The cable 50-pin connectors are used as indicated in Table 1-1.

Table 1-1 50-Pin Cable Connectors

Label Color	Label: Panasonic Cable	Label: Philips Cable	Use
Green	SGI <--- PANASONIC P1	SGI <--- PHILIPS P1	Input Link A
Red	SGI <--- PANASONIC P2	SGI <--- PHILIPS P2	Input Link B
Green	SGI ---> PANASONIC P3	SGI ---> PHILIPS P3	Output Link A
Red	SGI ---> PANASONIC P4	SGI ---> PHILIPS P4	Output Link B

For pinouts, see “Cable Connectors” on page 40 in Appendix A.

The four 50-pin connectors can be used for 4:4:4:4 or 4:2:2:4 in dual-link mode, or 4:2:2 in single-link mode where alpha is ignored:

- In 4:4:4:4 mode, Link A carries Y plus Cr and Cb from even-numbered sample points; Link B carries alpha plus Cr and Cb from odd-numbered sample points.
- In 4:2:2:4 mode, Link A carries Y plus Cr and Cb; Link B carries alpha only.
- In RGBA mode, Link A carries B and G; Link B carries R and alpha.

The video format selected determines Link A and Link B usage. For more information, see the following standards, which contain provisions for video signals:

- SMPTE 240M
 SMPTE 240M corresponds to the early 1035i HDTV format; it is being replaced by SMPTE 274M. The HD I/O option supports this model for compatibility with some equipment still using this standard. (This standard also defines color spaces, which must be set with a control, as explained in “VL_COLORSPACE” on page 22 in Chapter 2.)
- SMPTE 274M (subset, up to 74.25 MHz)
- SMPTE 296M (progressive)
- Recommendation 709 (ITU-R BT.709-2), which defines color primaries used by SMPTE 274M

Programming the HD I/O Option

The HD I/O option board supports the Video Library (VL). This API is described in the *Digital Media Programming Guide* (007-1799-060 or later; hereafter referred to as the DMPG).

This chapter consists of these sections:

- “VL Basics for the HD I/O Option Board” on page 7
- “HD I/O Controls” on page 14
- “Field Dominance” on page 29
- “HD I/O Events” on page 31
- “Capturing Graphics to Video” on page 31
- “Reporting” on page 32
- “Examples” on page 32

VL Basics for the HD I/O Option Board

To build programs that run under VL, you must

- install the *dmedia_dev* and *dmedia_eoe* options
- link with *libvl*
- include *dmedia/vl.h* and *dmedia/vl_xthd.h* for device-dependent functionality

The client library for VL is */usr/lib32/libvl.so*. The header files for the VL are in */usr/include/dmedia*; the main file is *vl.h*. This file contains the main definition of the VL API and controls that are common across all hardware. Several useful digital media programming examples are in */usr/share/src/dmedia/video/XTHD*.

Note: When building a VL-based program, you must add *-lvl* to the linking command.

For more information on the Video Library and the API usage, see the latest version of the DMPG.

This section explains

- “VL Concepts” on page 8
- “VL Syntax Elements” on page 9
- “VL Object Classes” on page 9
- “VL Nodes for the HD I/O Option” on page 10
- “VL Data Transfer Functions” on page 12
- “HD I/O Data Flow” on page 13

VL Concepts

The Video Library defines a basic set of primitives and mechanisms to specify interconnections and controls to achieve the desired setup. The two central concepts for VL are

- *path*: an abstraction for a way of moving data around
- *node*: an endpoint of the path

The basic nodes are a *source* (such as a VTR) and a *drain* (such as memory). Figure 2-1 diagrams the simplest VL path, with one of each of these two nodes.




Figure 2-1 Simple VL Path

The HD I/O board has a video source node (the video input), a video drain node (the video output), a memory source node (for output from application), and a memory drain node (for input to application). For transfers, each path must contain exactly one video node and one memory node. HD I/O nodes are further discussed in “VL Nodes for the HD I/O Option” on page 10.

VL Syntax Elements

VL syntax elements are as follows:

- VL types and constants begin with uppercase VL; for example, VLServer
- VL functions begin with lowercase vl; for example, `vlOpenVideo()`

VL Object Classes

The VL recognizes these classes of objects:

- *devices*, each including sets of nodes
- *nodes*, which are sources, drains, and internal nodes (as discussed in the preceding section)
- *paths*, connecting sources and drains (as discussed in the preceding section)
- *buffers*, for sending and receiving field/frame data to and from host memory

The HD I/O option requires the use of DMbuffers (digital media buffers).

DMbuffers, an abstraction of main memory, allow efficient and API-independent interchange of data between the different digital media libraries. For example, video fields can be captured into DMbuffers via VL and then displayed in graphics using OpenGL. They can also be passed between two processes without the data having to be copied explicitly. Refer to Chapter 5, “Digital Media Buffers,” in the DMPG for details.

- *events*, for monitoring video I/O status
- *controls*, or parameters that modify how data flows through nodes; for example:
 - video device parameters, such as sync source
 - video data parameters such as packing, size, and color space

VL controls fall into two categories:

- *device-global* or *device-independent* (prefix VL_), which can be used by several Silicon Graphics video products

For details of the device-independent controls, refer to the DMPG.

- *device-dependent* (prefix VL_XTHD_), specific to a particular video device, in this case, the HD I/O option

Both types of VL controls are explained in this chapter with respect to their usage with The HD I/O option.

VL Nodes for the HD I/O Option

Use **vlGetNode()** to specify nodes. This call returns the node's handle, which is used when setting controls or setting up paths. Its function prototype is:

```
VLNode vlGetNode(VLServer svr, int type, int kind, int number)
```

In this prototype, variables are as follows:

svr Names the server (as returned by **vlOpenVideo()**).

type Specifies the type of node:

- VL_SRC: source, such as a digital tapedeck connected to an input port

Note: The HD I/O option has only one input.

- VL_DRN: drain, such as system memory
- VL_DEVICE: global control, such as a default source; Table 2-1 summarizes the values for this type

Note: If you are using VL_DEVICE, the *VLNode* should be set to 0.

kind Specifies the kind of node:

- VL_VIDEO: connection to a video device equipment; for example, a video tapedeck or camera
- VL_MEM: workstation memory

number Number of the node in cases of two or more identical nodes, such as two video source nodes. The default value for all *kinds* is 0.

VL_ANY can also be used as a value for *number* to reference the first available node of the specified *type* and *kind*.

In general, a path for the HD I/O option has a memory node and a video node. The following fragment creates a digital video input source node and a memory drain node, and creates the path:

```

VLServer svr;
VLPath path;
VLNode src;
VLNode drn;
                                /*Set up video source node */
src = vlGetNode(svr, VL_SRC, VL_VIDEO, VL_ANY);
                                /*Set up memory drain node */
drn = vlGetNode(svr, VL_DRN, VL_MEM, VL_ANY);
                                /* Create source-to-drain path */
if((path = vlCreatePath(svr, VL_ANY, src, drn)) < 0){
    fprintf(stderr, "%s\n", vlStrError(vlGetErrno()));
    exit(1);
}
                                /* Set up path with shared src and drn node */
vlSetupPaths(svr, (VLPathList)&path, 1, VL_SHARE, VL_SHARE);

```

After nodes are specified, use **vlSetControl()** to specify parameters:

- memory nodes: you must set packing, color space, size, and capture type
- video nodes: if desired, set video timing, format (for example digital component), and color space; otherwise, the default values displayed in the video control panel (*vcp*) are applied

Controls for each node are defined in “HD I/O Controls” on page 14 in this chapter, and are summarized in Table 2-2. You can set controls in any order.

VL Data Transfer Functions

This section summarizes VL data transfer categories, and gives the basic steps of creating an application. For the HD I/O option, VL data transfers always involve memory (video to memory, memory to video) and require setting up a DMbuffer pool.

In the VL programming model, the process of creating a VL application consists of these steps:

1. Open a connection to the video daemon (**vlOpenVideo()**).
2. Specify nodes on the data path (**vlGetNode()**).
3. Create the path (**vlCreatePath()**).
4. Optional step: add more nodes to a path (**vlAddNode()**).
5. Set up the hardware for the path (**vlSetupPaths()**).
6. Specify path-related events to be captured (**vlSelectEvents()**, **vlAddCallback()**).
7. Set input and output parameters (controls) for the nodes on the path (**vlSetControl()**). Video format and timing set in the *vcp* are persistent and default to reasonable values.
8. Create a dmBuffer pool to hold data for memory transfers (**vlDMGetParams()**, **dmBufferSetPoolDefaults()**, **dmBufferCreatePool()**, **vlGetTransferSize()**).
9. Register the buffer (**vlDMPoolRegister()**, **vlDMPoolDeregister()**).
10. Start the data transfer (**vlBeginTransfer()**).
11. Get the data (**vlDMBufferGetValid()**, **vlDMBufferPutValid()**, **dmBufferAllocate()**, **dmBufferAllocateSize()**, **dmBufferGetPoolState()**, **dmBufferGetPoolFD()**, **dmBufferSetPoolSelectSize()**, **dmBufferMapData()**, **dmBufferFree()**) to manipulate frame data.
12. Handle data stream events (**vlSelectEvents()**, **vlNextEvent()**, **vlPending()**).
13. Clean up (**vlEndTransfer()**, **vlDMPoolDeregister()**, **vlDestroyPath()**, **vlCloseVideo()**).

Note: Error handling (**vlPerror()**) is accomplished throughout.

HD I/O Data Flow

Figure 2-2 diagrams data flow for input.

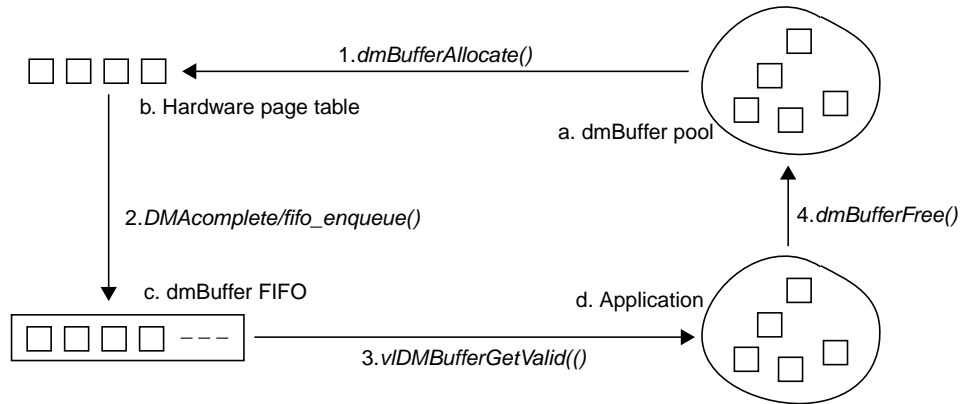


Figure 2-2 Data Flow, Input

Figure 2-2 diagrams data flow for output.

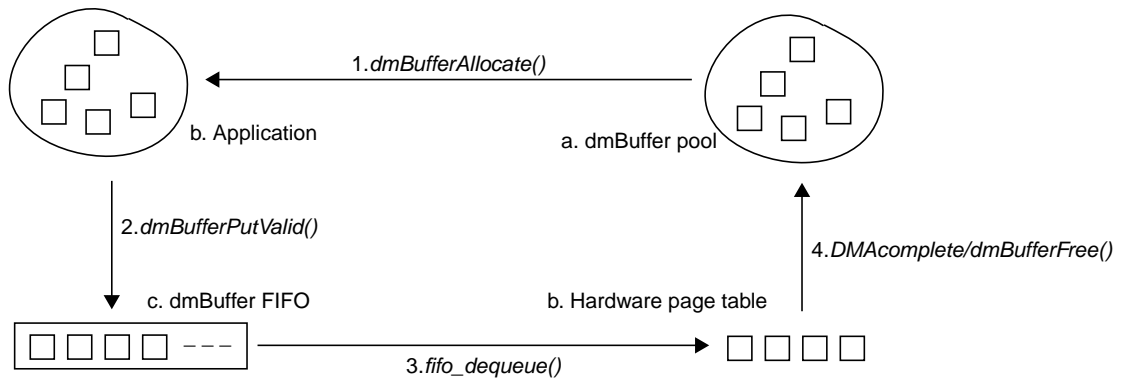


Figure 2-3 Data Flow, Output

Figure 2-4 diagrams control flow.

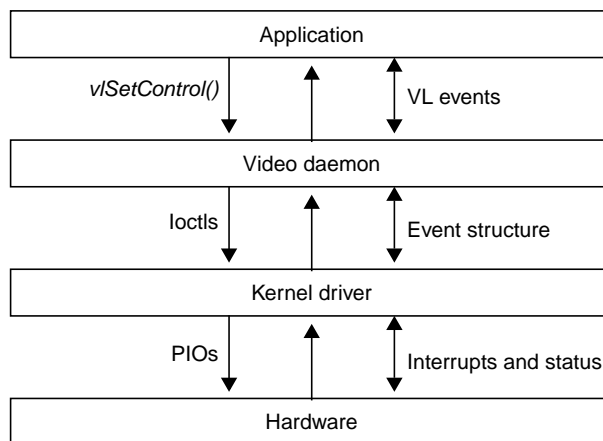


Figure 2-4 Control Flow

HD I/O Controls

This section consists of the following:

- “Setting Controls” on page 15
- “HD I/O Control Summary” on page 16
- “VL_TIMING” on page 19
- “VL_FORMAT” on page 20
- “VL_PACKING” on page 21
- “VL_COLORSPACE” on page 22
- “VL_CAP_TYPE” on page 27
- “VL_SIZE and VL_OFFSET” on page 28
- “VL_ZOOM” on page 28

Setting Controls

After setting up a path, and before starting a transfer, the memory and video nodes must be configured appropriately. Controls on the video node are persistent; that is, they retain their values between path destruction and creation. Memory node controls, however, are reset at each path creation. Because the video node controls are persistent, they need not be set by the application that can deal with arbitrary settings.

For a memory node, the following controls must be set because the default values are likely to be inappropriate:

- VL_PACKING
- VL_COLORSPACE
- VL_CAP_TYPE
- VL_SIZE

The recommended way to set VL_SIZE is to query VL_SIZE on the video node (after setting VL_TIMING appropriately) and use the returned value to set VL_SIZE on the memory node.

On the video node, applications will probably need to set the some controls as well. Most of them can also be set with the video control panel application (*vcp*).

To determine the available devices (that is, video options in the workstation, such as the HD I/O option board) and the nodes available on them, run *vlinfo*. To determine possible controls for each device, enter

```
vlinfo -l
```

Note: VL controls specified as true with **vlSetControl()** are executed immediately. However, they are not guaranteed to happen at a specific time.

To set controls for HD I/O nodes, use **vlSetControl()**. The following example sets video format and timing on a node:

```
timing.intVal = VL_TIMING_1125_1920x1080_5994i;
format.intVal = VL_FORMAT_DIGITAL_COMPONENT;

if (vlSetControl(svr, path, drn, VL_TIMING, &timing) <0)
{
    vlPerror("VlSetControl:TIMING");
    exit(1);
}
```

```

if (vlSetControl(svr, path, drn, VL_FORMAT, &format) <0)
{
    vlPerror("VlSetControl:FORMAT");
    exit(1);
}

```

For details on **vlSetControl()** and **vlGetControl()**, see the latest version of the DMPG.

HD I/O Control Summary

Table 2-1 summarizes node controls for the HD I/O option.

Table 2-1 HD I/O Node Controls

Control	Video Source	Memory Source	Video Drain	Memory Drain
VL_CAP_TYPE		X		X
VL_COLORSPACE	X (YCrCb and RGB_H only)	X	X YCrCb and RGB_H only)	X
VL_FIELD_DOMINANCE	X		X	
VL_FORMAT	X		X	
VL_H_PHASE			X	
VL_OFFSET	X (read only)	X	X (read only)	X
VL_PACKING		X		X
VL_SIZE	X (read only)	X	X (read only)	X
VL_SYNC			X	
VL_SYNC_SOURCE	X		X	
VL_TIMING	X		X	
VL_V_PHASE			X	

Table 2-1 (continued) HD I/O Node Controls

Control	Video Source	Memory Source	Video Drain	Memory Drain
VL_XTHD_EE_MODE		X	X	
VL_XTHD_INTERFACE_PRECISION	X		X	
VL_XTHD_LOOPBACK	X			
VL_XTHD_OUTPUT_REPEAT	X			X
VL_ZOOM		X		X

Note: Except for VL_XTHD_VITC_LINE_OFFSET, VL_H_PHASE, and VL_VPHASE, no controls can be changed during a transfer.

Table 2-2 summarizes the values and uses of controls for the HD I/O option.

Table 2-2 Controls for the HD I/O Option

Control	Values or Range	Use
VL_CAP_TYPE	VL_CAPTURE_INTERLEAVED VL_CAPTURE_NONINTERLEAVED VL_CAPTURE_FIELDS	Selects type of frame(s) or field(s) to capture. See “VL_CAP_TYPE” on page 27 in this chapter.
VL_COLORSPACE	VL_COLORSPACE_REC601_YCRCB VL_COLORSPACE_REC601_YUV VL_COLORSPACE_REC601_RGB_H VL_COLORSPACE_REC601_RGB_F VL_COLORSPACE_240M_YCRCB VL_COLORSPACE_240M_YUV VL_COLORSPACE_240M_RGB_H VL_COLORSPACE_240M_RGB_F VL_COLORSPACE_REC709_YCRCB VL_COLORSPACE_REC709_YUV VL_COLORSPACE_REC709_RGB_H VL_COLORSPACE_REC709_RGB_F	Specifies color space of video data in memory or for input/output. See “VL_COLORSPACE” on page 22 in this chapter.
VL_FIELD_DOMINANCE	VL_F1_IS_DOMINANT VL_F2_IS_DOMINANT Note: Frames that are output are deinterlaced differently depending on the choice of output field dominance. Deinterlacing is specified in the application.	Identifies frame boundaries in a field sequence (interlaced formats); ignored by progressive timings. See “Field Dominance” on page 29 in this chapter.

Table 2-2 (continued) Controls for the HD I/O Option

Control	Values or Range	Use
VL_FORMAT	VL_FORMAT_DIGITAL_COMPONENT_DUAL VL_FORMAT_DIGITAL_COMPONENT	Specifies sampling format of data in or out: Standard YCrCb color space sampled at 4:4:4:4 or 4:4:4, depending on packing. 4:2:2:4 or 4:2:2, depending on packing. See “VL_FORMAT” on page 20 in this chapter.
VL_H_PHASE	[-1000,1000] pixels (increment of 1)	Sets horizontal phase.
VL_OFFSET	Fixed at (0,0).	Sets the position within the video raster to (0,0).
VL_PACKING	See Appendix C, “Pixel Packings and Color Spaces” for values.	Sets packing format for memory source or drain node. See “VL_PACKING” on page 21 in this chapter and Appendix C, “Pixel Packings and Color Spaces.”
VL_SIZE	Raster size.	Memory: Sets size of the desired region of the video raster, which must be the same as the default active region. Video: Returns size of video raster (frame); dependent on timing.
VL_SYNC	VL_SYNC_INTERNAL VL_SYNC_GENLOCK	Sets sync mode for video output.
VL_SYNC_SOURCE	VL_SYNC_HOUSE VL_SYNC_DIGITAL_INPUT_LINK_A VL_SYNC_DIGITAL_INPUT_LINK_B VL_XTHD_SYNC_DIGITAL_INPUT_LINK_A VL_XTHD_SYNC_DIGITAL_INPUT_LINK_B	Selects the genlock source if VL_SYNC_GENLOCK is used. VL_SYNC_HOUSE is analog reference. VL_XTHD_SYNC_DIGITAL_INPUT_LINK_[A B] on video source node only.
VL_TIMING	VL_TIMING_1125_1920x1080_5994i VL_TIMING_750_1280x720_5994p	Sets or gets timing; see “VL_TIMING” on page 19 in this chapter.
VL_V_PHASE	[-600,600] lines (increment of 1)	Sets vertical phase.
VL_XTHD_EE_MODE	VL_XTHD_EE_MODE_OFF VL_XTHD_EE_MODE_ON	Causes digital output to transmit a copy of digital input data; output must be genlocked.
VL_XTHD_INTERFACE_PRECISION	VL_XTHD_INTERFACE_PRECISION_8 VL_XTHD_INTERFACE_PRECISION_10	Specifies whether external video interface is 8 bits or 10 bits wide.

Table 2-2 (continued) Controls for the HD I/O Option

Control	Values or Range	Use
VL_XTHD_LOOPBACK	TRUE FALSE	For video source only, sets video input to the output pipe rather than the input jack (TRUE).
VL_XTHD_OUTPUT_REPEAT	VL_XTHD_OUTPUT_REPEAT_DISABLED VL_XTHD_OUTPUT_REPEAT_LAST_FIELD VL_XTHD_OUTPUT_REPEAT_LAST_FRAME	Controls whether system repeats DMbuffers output is underflowing; see “Automatically Correcting for Output Underflow” on page 30.
VL_ZOOM	1, -1	Sets position of top video line in the buffer; see “VL_ZOOM” on page 28 in this chapter.

VL_TIMING

The VL_TIMING control sets timing type, which expresses the timing of video presented to a source or drain.

Note: For the HD I/O option, VL_TIMING is specified on the video node only, and not also on the memory mode, as with other Silicon Graphics video options.

Table 2-3 summarizes VL_TIMING values for the HD I/O option. Future releases will support additional values.

Table 2-3 Values for VL_TIMING

Timing	Sampling Rate (MHz)	Standard
VL_TIMING_1125_1920x1080_5994i	~74.18 (74.25/1.001)	SMPTE 274M
VL_TIMING_750_1280x720_5994p	~74.18 (74.25/1.001)	SMPTE 296M

Each value for VL_TIMING indicates the raster configuration of a particular SMPTE specification, such as SMPTE 274M-1995 system 3. The values are named according to the raster format:

- The first field is the number of total lines, such as 1125, 750, 525, or 625.
- The second field is the size of the active region, in pixels by lines.

- The third field is the vertical refresh rate and the scanning format; the scanning format is
 - i: interlaced
 - p: progressive (noninterlaced)

For example, VL_TIMING_1125_1920x1080_5994i specifies 1125 total lines, an active region of 1920 pixels by 1080 lines, 59.94 fields per second, and 2:1interlacing.

Note: Although the VL defines timings with sampling rates up through 148.5 MHz, the HD I/O option supports only those through 74.25 MHz.

Timing rates and sync source are interrelated for HDTV. Table 2-4 summarizes the relationships.

Table 2-4 Genlock Sync Source and Timing

Connector	Sync Source	Output Timing	Output Timing	Calibration Timing	Calibration Timing
GEN IN	Analog 1920x1080 5994i trilevel	1920x1080 59.94i	1280x720 59.94p	1920x1080 59.94i (74.18 MHz)	1920x1080 59.94i (74.18 MHz)
	Analog NTSC color black bilevel	1920x1080 59.94i	1280x720 59.94p	1920x1080 59.94i (74.18 MHz)	525 (108 MHz)
	Analog PAL black bilevel	N/A	N/A	1920x1080 59.94i (74.25 MHz)	625 (108 MHz)
LINK A or B	1920x1080 59.94i	1920x1080 59.94i	1280x720 59.94p	N/A	N/A
	1280x720 59.94p	1920x1080 59.94i	1280x720 59.94p	N/A	N/A

VL_FORMAT

The VL_FORMAT control is used on video nodes only. It specifies the sampling format of data on the wire, specifically, dual-link (4:4:4 or 4:4:4:4) or single-link (4:2:2 or 4:2:2:4):

- VL_FORMAT_DIGITAL_COMPONENT_DUAL: either 4:4:4 or 4:4:4:4, depending on the specified packing
- VL_FORMAT_DIGITAL_COMPONENT: Standard YCrCb color space sampled at 4:2:2 or 4:2:2:4, depending on the specified packing

The memory packing mode VL_PACKING determines how components are actually sampled.

VL_FORMAT does not imply color space, nor does it imply whether the second link is used. The second link is used whenever alpha/key channel is used, which depends on the VL_PACKING setting, or when color is sampled at 4:4:4.

VL_PACKING

A video packing describes how a video signal is stored in memory, in contrast to a video format, which describes the characteristics of the video signal. For example, the memory source node accepts packed video from a DMbuffer and outputs video in a given format. Packings are specified through the VL_PACKING control on the memory nodes.

Note: Because of HDTV's multiple color spaces, "old style" packings, such as VL_PACKING_Y_8_P, are ambiguous and therefore no longer supported for the HD I/O board. You must specify the packing and color space explicitly.

The HD I/O option supports common packings up through 32 bits per pixel, including 4:4:4, 4:2:2, 4:4:4:4, and 4:2:2:4, at 8, 10, and 12 bits per component. Specifically, it supports packings compatible with OpenGL and IRIS GL, and those in common use from other video products. Appendix C, "Pixel Packings and Color Spaces," shows the layout of each packing for the HD I/O option. It also gives the corresponding names for these packings that are used by other libraries.

An application must set both VL_PACKING and VL_COLORSPACE. Note that changes in one parameter may change the values of other parameters set earlier; for example, clipped size may change if VL_PACKING is set after VL_SIZE. For example:

```
VLControlValue val;  
  
val.intVal = VL_PACKING_444_8;  
vlSetControl(vlSvr, path, memdrn, VL_PACKING, &val);
```

Note: Changing this control at the beginning of data transfer takes several seconds to go into effect.

VL_COLORSPACE

The VL_COLORSPACE control specifies color space of video data in memory or for input and output. A color space is a color component encoding format, for example, RGB and YUV. Because video equipment uses more than one color space, the HD I/O video nodes, in addition to the memory nodes, support the VL_COLORSPACE control.

Each component of an image has

- a color that it represents
- a canonical minimum value
- a canonical maximum value

Normally, a component stays within the minimum and maximum values. For example, for a luma signal such as Y, you can think of these limits as the black level and the peak white level, respectively. For an unsigned component with n bits, there are two possibilities for [minimum value, maximum value]:

- full range: $[0, (2^n)-1]$, which provides the maximum resolution for each component
- compressed (headroom) range, which provides numerical headroom, which is often useful when processing video images:
 - Cr and Cb: $[(2^n)/16, 15*(2^n)/16]$
 - Y, A, R, G, and B: $[(2^n)/16, 235*(2^n)/256]$

Color Spaces and Color Models

Various HDTV specifications define color models differently from those defined in ITU-R BT.601, which is used by most standard-definition digital video equipment. For HDTV, the VL defines three color models:

- SMPTE 240M
- SMPTE 274M (Recommendation 709, which is ITU-R BT.709-2)
- Recommendation 601

Within each color model, four different color spaces exist:

- YCrCb: headroom range
 Headroom range means that black is at, for example, code 64 rather than 0, and white is at, for example, code 940 rather than 1023. Headroom-range color spaces can accommodate overshoot (superwhite) and undershoot (superblack) colors. Full-range color spaces clamp these out-of-range colors to black and white.
- YUV: full range
- RGB_H: headroom range
- RGB_F: full range

For memory nodes, these four color spaces are defined for each of three color models, resulting in 12 color spaces. Note that all 12 are supported on memory nodes, but only YCrCb and RGB_H color spaces are supported on video nodes.

Table 2-5 summarizes currently supported combinations of VL_FORMAT and VL_COLORSPACE; future releases will support more combinations.

Table 2-5 VL_FORMAT and VL_COLORSPACE Combinations Supported

Video Node	Memory Node
4:2:2 YCrCb	4:4:4 RGB_F
4:2:2 YCrCb	4:2:2 YCrCb

Color-space conversion is performed within a color model if the color spaces are different on the memory and video nodes. Conversion between the color models is not supported.

Table 2-3 summarizes VL_COLORSPACE values for the HD I/O option.

Table 2-6 Color-Space Values

Value	Standard	Nodes	Definition
VL_COLORSPACE_REC601_YCRCB	ITU-R BT.601	All	Same as VL_COLORSPACE_CCIR601
VL_COLORSPACE_REC601_YUV	ITU-R BT.601	Memory	Full-range YUV (same as VL_COLORSPACE_YUV)
VL_COLORSPACE_REC601_RGB_H	ITU-R BT.601	All	Headroom-range RGB (VL_COLORSPACE_RP175)
VL_COLORSPACE_REC601_RGB_F	ITU-R BT.601	Memory	Full-range Rec. 601 RGB (VL_COLORSPACE_RGB)
VL_COLORSPACE_240M_YCRCB	SMPTE 240M	All	SMPTE 240M YCrCb
VL_COLORSPACE_240M_YUV	SMPTE 240M	Memory	Like 240M_YCRCB but full-range (no headroom)
VL_COLORSPACE_240M_RGB_H	SMPTE 240M	All	Headroom-range RGB
VL_COLORSPACE_240M_RGB_F	SMPTE 240M	Memory	Full-range RGB
VL_COLORSPACE_REC709_YCRCB	ITU-R BT.709-2	All	60-Hz YCrCb
VL_COLORSPACE_REC709_YUV	ITU-R BT.709-2	Memory	Like REC709_YCRCB but full-range (no headroom)
VL_COLORSPACE_REC709_RGB_H	ITU-R BT.709-2	All	Headroom-range RGB
VL_COLORSPACE_REC709_RGB_F	ITU-R BT.709-2	Memory	Full-range RGB

Note: Changing this control at the beginning of data transfer takes several seconds to go into effect.

Two sets of colors are commonly used together, RGB (RGBA) and YCrCb/YUV (VYUA). YCrCb (YUV), the most common representation of color from the video world, represents each color by a luma component called Y and two components of chroma, called Cr (or V), and Cb (or U). The luma component is loosely related to brightness or luminance, and the chroma components make up a quantity loosely related to hue. These components are defined rigorously in ITU-R BT.601 (also known as Rec. 601 and CCIR 601), ITU-R BT.709-2, and SMPTE 240M.

The alpha channel is not a real color. For that channel, the canonical minimum value means completely transparent, and the canonical maximum value means completely opaque.

For OpenGL, IRIS GL, and DM:

- the library constant indicates whether the data is RGBA or VYUA
- RGBA data is full-range by default
- VYUA data in DM can be full-range or compressed-range; you must determine this from context

For more information about color spaces, see *A Technical Introduction to Digital Video*, by Charles A. Poynton (New York: Wiley, 1996).

VL_COLORSPACE Control of Blanking

Along with memory node color space, VL_COLORSPACE determines the color-conversion matrix values. In addition, this control affects the type of blanking output by the board during horizontal and vertical blanking, and during active video area when not transferring data. On a video drain node, VL_COLORSPACE affects the type of blanking that the board outputs, in accordance with SMPTE 274M:

- YCrCb: blanking is $Y = 64$, $Cr/Cb = 512$, $A = 64$
- RGB_H: blanking is $R = 64$, $G = 64$, $B = 64$, $A = 64$

VL_COLORSPACE and Lookup Tables

The HD I/O board supports lookup tables (LUTs) on input and output for gamma correction or decorection. If an application is to work with linear components, these can be used to convert between linear and nonlinear spaces.

The HD I/O hardware includes a separate LUT for each RGB color component. Each of the three LUTs is a table of 8192 entries; each entry stores 13 bits. The application programs the entries in each table. LUTs accomplish offsets, if they are required by the memory storage format.

The LUTs perform rounding as follows:

- If the LUT is not explicitly programmed by the application, it is in pass-through mode, and all rounding is performed in the color-space converter.
- If the LUT has been programmed explicitly by the application, the application can control rounding as part of the lookup table function. The packer (hardware that reads the LUT and formats data for the host memory; see Figure 1-2 on page 4) performs a final conversion from 13-bit LUT format to host memory format.

VL_COLORSPACE Example

Figure 2-5 summarizes a complicated conversion between color spaces with for 4:2:2 sampling rates and different primary colors. The example shows conversion of material produced using SMPTE 240M colorimetry to material in SMPTE 274M format (ITU-R BT.709.2 colors).

Note: In Figure 2-5, RGB are values in linear space and R'G'B' are values in nonlinear space after the optoelectric transfer function is applied as specified in ITU-R BT.709. To convert between RGB and R'G'B', the LUTs can be used to apply this function or its inverse.

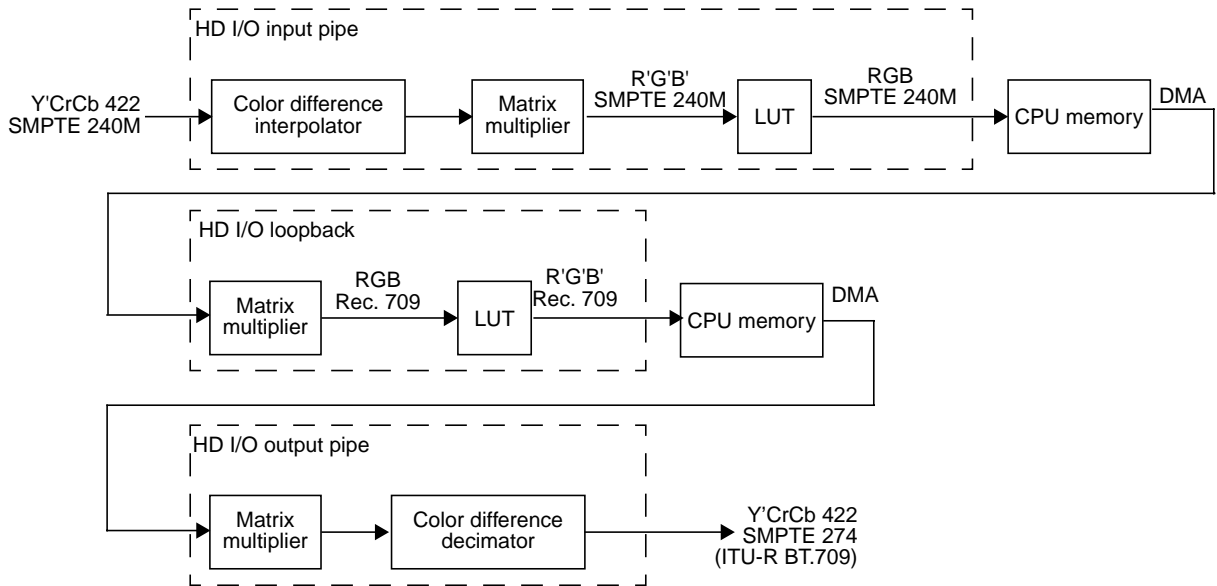


Figure 2-5 Color-Space Conversion Example

VL_CAP_TYPE

An application can request that the HD I/O option capture or play back a video stream in a number of ways. For example, the application can request that each field be placed in its own buffer, or that each buffer contain an interleaved frame. This section explains the capture types that the HD I/O option supports.

Capture types are as follows:

- VL_CAPTURE_NONINTERLEAVED
- VL_CAPTURE_INTERLEAVED

VL_CAPTURE_FIELDS is equivalent to VL_CAPTURE_NONINTERLEAVED. The HD I/O option does not support VL_CAPTURE_EVEN_FIELDS and VL_CAPTURE_ODD_FIELDS.

Note: VL_SIZE refers to the size of each frame, rather than the size of the buffer in memory, so it is independent of VL_CAP_TYPE; the VL_CAP_TYPE setting does not change VL_SIZE.

VL_CAPTURE_NONINTERLEAVED

The VL_CAPTURE_NONINTERLEAVED capture type specifies that frame-size units are captured noninterleaved. Each field is placed in its own buffer, with the dominant field in the first buffer. If one of the fields of a frame is dropped, all fields are dropped. Consequently, an application is guaranteed that the field order is maintained; no special synchronization is necessary to ensure that fields from different frames are mixed.

If VL_CAPTURE_NONINTERLEAVED is specified for playback, similar guarantees apply as for capture. If one field is lost during playback, it is not possible to “take back” the field. The HD I/O option resynchronizes on the next frame boundary, although black or “garbage” video might be present between the erring field and the frame boundary.

VL_CAPTURE_INTERLEAVED

Interleaved capture interleaves the two fields of a frame and places them in a single buffer; the order of the fields depends on the value set for VL_DOMINANCE_FIELD.

The HD I/O option guarantees that the interleaved fields are from the same frame: if one field of a frame is dropped, then both are dropped.

During playback, a frame is deinterleaved and output as two consecutive fields, with the dominant field output first. If one of the fields is lost, the HD I/O option resynchronizes to a frame boundary before playing the next frame. During the resynchronization period, black or “garbage” data may be displayed.

VL_SIZE and VL_OFFSET

VL_SIZE refers to the size of each frame, rather than the size of the buffer in memory, so it is independent of VL_CAP_TYPE; the VL_CAP_TYPE setting does not change VL_SIZE.

For memory nodes, VL_SIZE specifies the desired region of the video raster, which must be the same as the default active region. There is no default; this control must be set. On video nodes, this control is read-only and cannot be set.

For memory nodes, VL_OFFSET is related to VL_SIZE; it specifies the origin of the video region to be captured at (0,0), the standard beginning of active video. The recommended way to set VL_SIZE for the memory node is to query it on the video node (after setting VL_TIMING appropriately) and use the returned value to set VL_SIZE on the memory node.

VL_ZOOM

The VL_ZOOM control specifies the site of the top video line in memory; it does not scale (zoom and decimate) a video image as for some other Silicon Graphics video options. The VL_ZOOM value can be 1 or -1:

- 1: Normal line order; the top video line is at the lowest address in the buffer.
- -1: Inverted line order; as in OpenGL, the top video line is at the highest address in the buffer.

Field Dominance

Field dominance identifies the frame boundaries in a field sequence; that is, it specifies which pair of fields in a field sequence constitute a frame. The control VL_FIELD_DOMINANCE allows you to specify whether an edit occurs on the nominal video field boundary (Field 1, or F1) or on the intervening field boundary (Field 2, or F2).

- F1 dominant: The edit occurs on the nominal video field boundary.
- F2 dominant: The edit occurs on the intervening field boundary.

Whether a field is Field 1 or Field 2 is determined by the setting of bit 9, the F bit, in the XYZ word of the EAV and SAV sequences. The setting means the following:

- For Field 1 (also called the odd field), the F bit is 0.
- For Field 2 (also called the even field), the F bit is 1.

Note: Field dominance has no effect on progressive timings.

Figure 2-6 shows fields and frames as defined for digital 1080-line formats for the HD I/O option.




Figure 2-6 Fields and Frames for SMPTE 274M

Editing is usually on Field 1 boundaries, where Field 1 is defined as the first field in the video standard's two-field output sequence. However, some users may want to edit on F2 boundaries, which fall on the field in between the video standard's frame boundary. To do so, use this control, then program your deck to select F2 edits.

A set of frames to be output must be deinterlaced into fields differently, depending on the choice of output field dominance: for SMPTE 274M, the top line is in F1, as shown in Figure 2-6; for SMPTE 240M, the top line is in F2. For example, when F1 dominance is selected, the field with the topmost line must be the first field to be transferred; when F2 dominance is selected, the field with the topmost line must be the second field to be transferred.

Automatically Correcting for Output Underflow

If the application is not sending buffers fast enough for the receiving equipment's video frame rate, you can set `VL_XTHD_OUTPUT_REPEAT` to repeat DMbuffers automatically. The values for this control vary depending on whether the transfer is progressive or interlaced.

For interlaced noninterleaved transfers, choices are as follows:

- `VL_XTHD_OUTPUT_REPEAT_LAST_FIELD`
This setting repeats the last buffer twice. This setting is spacially imperfect, but does not cause flicker.
- `VL_XTHD_OUTPUT_REPEAT_LAST_FRAME` (the default)
This setting repeats the last pair of buffers. This setting is spacially better than `VL_XTHD_OUTPUT_REPEAT_LAST_FIELD`, but causes flicker.
- `VL_XTHD_OUTPUT_REPEAT_DISABLED`
This setting, which does not resend any data, is the most useful for debugging, since underflow is then quite visible on output.

For progressive or interleaved transfers, choices are as follows:

- `VL_XTHD_OUTPUT_REPEAT_LAST_FRAME` (the default)
This setting repeats the last buffer.
- `VL_XTHD_OUTPUT_REPEAT_DISABLED`

Capturing Graphics to Video

To capture graphics to video, you can use OpenGL to read pixels into memory. However, the coordinate system differs between video and Open GL; under OpenGL, the origin is at the lower left corner and in video, origin is in the upper left corner. To adjust for this difference, set `VL_ZOOM` to -1; see “`VL_ZOOM`” on page 28.

HD I/O Events

The VL provides several ways of handling data stream events, such as completion or failure of data transfer, vertical retrace event, loss of the path to another client, lack of detectable sync, or dropped fields or frames. The method you use depends on the kind of application you are writing:

- For a strictly VL application, use
 - `vlSelectEvents()` to choose the events to which you want the application to respond
 - `vlCallback()` to specify the function called when the event occurs
 - your own event loop or a main loop (`vlMainLoop()`) to dispatch the events
- For an application that also accesses another program or device driver, or if you are adding video capability to an existing X or OpenGL application, set up an event loop in the main part of the application and use the IRIX file descriptor (FD) of the event(s) you want to add.

For more information on these functions, see Chapter 4 in the DMPG.

Table 2-7 summarizes events for the HD I/O option. For these options, this table supersedes the table of events in Chapter 14, “VL Event Handling,” in the DMPG; the HD I/O option supports only the events listed in Table 2-7.

Table 2-7 HD I/O Events

Event	Use
VLSyncLost	Sync is not detected
VLStreamStarted	Stream started delivery
VLStreamStopped	Stream stopped delivery
VLSequenceLost	A field/frame was dropped
VLControlChanged	A control on the path has changed
VLTransferComplete	A field/frame transfer has completed
VLTransferFailed	A transfer has failed and DMA is aborted
VLTransferError	A transfer error was discovered; field may be invalid

Reporting

The DMediaInfo structure reports the Unadjusted System Time (UST).

The HD I/O option makes use of the error events noted in Chapter 4 of the DMPG, as well as VLTransferErrorEvent, which reports nonfatal video transfer errors. The VLTransferComplete and VLSequenceLost events also report the Media Stream Count (MSC) of the field or frame transferred or failed.

Examples

This section contains three examples:

- “Capture to Memory for Disk Recording” on page 33
- “Playback From Memory for Disk Playback” on page 33
- “Capture to Memory for Graphics” on page 34

Capture to Memory for Disk Recording

This example shows an application that simply captures video data and records it to disk, for later playback. The video is left in its original 4:2:2 sampling and color space. Because 10 bits are required, the R242_10 packing mode is chosen. Frames start on F1 boundaries. After creating the VL path, the application sets up the video node, mostly as a result of user input, since it depends on external equipment.

```
videoSource.VL_TIMING = VL_TIMING_1125_1920x1080_5994i
videoSource.VL_COLORSPACE = VL_COLORSPACE_REC709_YCRCB
videoSource.VL_FIELD_DOMINANCE = VL_F1_IS_DOMINANT
videoSource.VL_FORMAT = VL_FORMAT_DIGITAL_COMPONENT
videoSource.VL_XTHD_LOOPBACK = VL_XTHD_LOOPBACK_OFF
videoSource.VL_XTHD_INTERFACE_PRECISION = VL_XTHD_INTERFACE_PRECISION_10
```

Next, application configures memory node:

```
memoryDrain.VL_SIZE = videoSource.VL_SIZE
memoryDrain.VL_COLORSPACE = videoSource.VL_COLORSPACE
memoryDrain.VL_CAP_TYPE = VL_CAPTURE_NONINTERLEAVED
memoryDrain.VL_PACKING = VL_PACKING_R242_10
```

Next, the application calls `vlGetTransferSize()` to get the required buffer size to hold each field. Using that size, it creates and registers a DMS buffer pool and starts the transfer.

Playback From Memory for Disk Playback

This example shows an application that plays back data captured in the previous example.

The application in the previous example has stored various attributes along with the data, including color space, frame rate, dominance, and size. These attributes, along with some user-specified options, are used to derive the control values. After creating the VL path, the application sets up the video node:

```
videoDrain.VL_TIMING = VL_TIMING_1125_1920x1080_5994i
videoDrain.VL_COLORSPACE = VL_COLORSPACE_REC709_YCRCB
videoDrain.VL_FIELD_DOMINANCE = VL_F1_IS_DOMINANT
videoDrain.VL_FORMAT = VL_FORMAT_DIGITAL_COMPONENT
videoDrain.VL_XTHD_INTERFACE_PRECISION = VL_XTHD_INTERFACE_PRECISION_10
videoDrain.VL_SYNC = VL_SYNC_GENLOCK
videoDrain.VL_SYNC_SOURCE = VL_XTHD_SYNC_HOUSE
```

Next, the application configures the memory node:

```
memorySource.VL_SIZE = videoDrain.VL_SIZE
memorySource.VL_COLORSPACE = videoDrain.VL_COLORSPACE
memorySource.VL_CAP_TYPE = VL_CAPTURE_NONINTERLEAVED
memorySource.VL_PACKING = VL_PACKING_R242_10
```

Next, the application calls `vlGetTransferSize()` to get the required buffer size to hold each field. Using that size, it creates a DMS buffer pool, allocates, fills, and prequeues some buffers, and starts the transfer.

Capture to Memory for Graphics

In this example, the application captures video and draws it on the graphics screen using OpenGL. This example resembles the *videoin* application.

The incoming video is converted to 10-bit 4:4:4 full-range RGB in an OpenGL-compatible packing format. The 4444_10_10_10_2 packing is chosen; it is compatible with OpenGL using packed-pixel extension `GL_UNSIGNED_INT_10_10_10_2_EXT` pixel format. Video is interleaved in memory into frames, and is written upside down to be compatible with OpenGL's default coordinate system (`glPixelZoom` of (1.0, 1.0)).

First, video node is configured:

```
videoSource.VL_TIMING = VL_TIMING_1125_1920x1080_5994i
videoSource.VL_COLORSPACE = VL_COLORSPACE_REC709_YCRCB
videoSource.VL_FIELD_DOMINANCE = VL_F1_IS_DOMINANT
videoSource.VL_FORMAT = VL_FORMAT_DIGITAL_COMPONENT
videoSource.VL_XTHD_LOOPBACK = VL_XTHD_LOOPBACK_OFF
videoSource.VL_XTHD_INTERFACE_PRECISION = VL_XTHD_INTERFACE_PRECISION_10
```

Next, application configures memory node:

```
memoryDrain.VL_SIZE = videoSource.VL_SIZE
memoryDrain.VL_COLORSPACE = VL_COLORSPACE_REC709_RGB_F
memoryDrain.VL_CAP_TYPE = VL_CAPTURE_INTERLEAVED
memoryDrain.VL_PACKING = VL_PACKING_4444_10_10_10_2
memoryDrain.VL_ZOOM = -1/1
```

Next, the application calls `vlGetTransferSize()` to get the required buffer size to hold each field. Using that size, it creates and registers a DMS buffer pool and starts the transfer.

Synchronizing Data Streams and Signals

You can use special signals recognized or generated by the HD I/O board—UST (unadjusted system time), MSC (media stream count)—to synchronize data streams. This chapter explains

- “Using UST, MSC, and Buffered Media Streams for Synchronization” on page 35
- “Media Library Interfaces for UST and MSC” on page 38

Using UST, MSC, and Buffered Media Streams for Synchronization

Whenever a VL path is open in continuous mode, the HD I/O board and certain other Silicon Graphics video devices continuously try to dequeue media stream samples from the path’s buffer for input, or to enqueue media stream samples onto the path’s buffer for output. If the buffer between the application and each device never underflows or overflows, then the application can measure and schedule the timing of input and output signals to 100% of the accuracy of the underlying device.

Occasionally, the application is held off and audio, video, or both come out late. Buffer underflow on output and overflow on input can result from the application not keeping the buffer adequately filled for the following reasons:

- The application is busy with other tasks, allowing too much time between putting fields into the buffer.
- Processes are subject to various interruptions (10 to 80 ms for some processes) under IRIX because
 - the process for filling the buffer is running at too low a priority
 - the process cannot get a resource from IRIX that it needs, such as memory pages

To get around this problem, a mechanism built into the VL helps keep track of data flow into and out of buffers by providing accurate timing information for each frame of video that enters or leaves the system. This mechanism, called UST/MSC, produces matched pairs of two numbers:

- unadjusted system time (UST), a time value that is used to state timing measurements to applications
- media stream count (MSC), a count value that identifies a particular media stream sample (a video field or frame)

The device keeps a counter called the device media stream count (device MSC), which increments by one every time the device attempts to enqueue or dequeue a media stream sample, whether or not the enqueue or dequeue attempt is successful. UST/MSC was designed to return timing information in a form that is valid whenever the buffer is not underflowing or overflowing.

The UST/MSC capability and the buffering that goes with it are appropriate for applications and devices such as movie players and digital video editing devices.

UST/MSC affords maximally accurate synchronization when scheduling cannot be guaranteed and some buffering is acceptable. Also, if scheduling becomes reliable at some later point, UST/MSC continues to function the same way with no code changes required; the buffers can be made smaller, and the result is a low-latency application with the same accurate synchronization.

Note that UST/MSC itself

- does not add any latency to an application
The buffer adds latency: it increases the time the application would take to respond to some output event by changing its input (and vice versa). This solution to the synchronization problem is useful for applications in which a small latency can be sacrificed for more accuracy.
- does not require that an application trade off latency for accuracy
- does not require that an application use any particular size buffer
- delivers the full accuracy of the underlying hardware's timing support regardless of the scheduling characteristics of the application
- could be useful for graphics and texture even for low-latency applications

For the HD I/O option, UST/MSC pairs are maintained in software and are valid only during a transfer. Make calls to `vlGetUSTMSCPair()` and `vlGetFrontierMSC()` only during a transfer; these calls block until at least one buffer has been successfully transferred. Note the following:

- For interlaced timings, MSC always increments by 1 per field.
- For progressive timings, MSC always increments by 1 per frame.

The code below is a high-level algorithm to maintain synchronization of two buffered media streams that send data from memory to hardware outputs; a corresponding one is necessary for the other direction:

```

create video buffer between me and the audio output;
create audio buffer between me and the video output;
while (1)
{
    sleep until one of the buffers is getting empty;
    for (video buffer)
    {
        use UST/MSC to determine:
        "at what time (what UST) will the next video data I enqueue
        on the buffer actually go out the jack of the machine?";
    }

    for (audio buffer)
    {
        (exact same thing as above, except for audio)
    }

    From the predicted video and audio USTs, determine
    "what is the synchronization error between the audio and video
    streams?"

    Enqueue more frames to fill up the audio and video buffer queues.
    If there is synchronization error, enqueue new frames to either skip
    frames on the stream that is behind or repeat frames on the stream
    that is ahead.
    }
}

```

The answers to the questions in the pseudocode above are obtained with three VL calls that manipulate UST and MSC and are explained in the next section.

Media Library Interfaces for UST and MSC

UST/MSC calls allow you to associate a UST with a particular piece of data that just left a buffer or is about to enter a buffer. The VL calls for determining the MSC and UST—`vlGetUSTMSCPair(3dm)`, `vlGetFrontierMSC(3dm)`, and `vlGetUSTPerMSC(3dm)`—help synchronize input and output of different data streams in cases where the application is getting data from or putting data into each device via a buffer. The application is at the “frontier” end of this buffer and the devices are at the “device” end of the buffer.

- **vlGetUSTMSCPair()** gets the timing information for each frame or field as it enters or leaves the physical jack of a device.

This call returns an atomic UST/MSC pair for the jack (specified with the `VL_NODE`) for a given path that contains a `VL_MEM` node. The returned MSC is not guaranteed to be the one currently at the jack, nor is it even guaranteed to be the number of any media stream sample currently in the application’s buffer. To relate the returned MSC to a particular item in the application’s buffer, you must use **vlGetFrontierMSC()**.

- **vlGetFrontierMSC()** gets the frontier MSC associated with a particular `VL_MEM` node. The frontier MSC, at the application end of the media stream, is the MSC of the next item that the application removes from or puts into the buffer.
- **vlGetUSTPerMSC()** gets the time spacing of fields or frames in a path (the nominal average UST time elapsed between media stream samples in a given `VLPath` that includes a `VL_MEM` node).

These calls are used for extrapolating a UST/MSC pair as shown in **vlGetFrontierMSC()**. For other types of media streams, a similar mechanism extrapolates the UST/MSC pair; for example, for audio, use equivalent AL calls.

Once you have calculated the extrapolated UST/MSC pairs for both media streams, you can determine the synchronization error. The difference in the audio and video USTs for matching frame numbers is the amount they are out of sync. To resynchronize them, you must enqueue new frames to either skip frames on the stream that is behind or repeat frames on the stream that is ahead. The number of frames to be skipped or repeated is the difference in USTs divided by the frame rate.

To use UST/MSC, the application must have separate handles for each separate piece of data coming in or going out of some kind of buffer. The application can use these handles to specify, for example, a particular frame to output or pixels of a particular field to get.

Note: For complete details, including syntax, code examples, and caveats, see the man pages for these calls.

HD I/O Option Board Specifications

This appendix summarizes hardware specifications for the HD I/O option board and its cable, in these sections:

- “Cable Connectors” on page 40
- “GPI Interface” on page 44
- “Genlock” on page 48

Cable Connectors

Figure A-1 shows the two multiheaded cables included with the board; each has four 50-pin connectors for link A input, link B input, link A output, and link B output. The 50-pin connectors differ for each type of cable, following the Panasonic and Philips 50-pin video equipment interface standard.




Figure A-1 HD I/O Cables

Table A-1 summarizes the Panasonic (HD-D5) pinout for the cable's 50-pin connector.

Table A-1 Panasonic 50-Pin Connector Pinout (HD-D5)

Pin	Signal	Pin	Signal	Pin	Signal
1	CLK+			34	CLK-
2	Data9+	18	Data1+	35	Data9-
3	Data8+	19	Data1-	36	Data8-
4	Data7+	20	Data0+	37	Data7-
5	Data6+	21	Data0-	38	Data6-
6	Data5+	22	GND	39	Data5-
7	Data4+	23	GND	40	Data4-
8	Data3+	24	GND	41	Data3-
9	Data2+	25	GND	42	Data2-
10	Data19+	26	Data11+	43	Data19-
11	Data18+	27	Data11-	44	Data18-
12	Data17+	28	Data10+	45	Data17-
13	Data16+	29	Data10-	46	Data16-
14	Data15+	30	GND	47	Data15-
15	Data14+	31	GND	48	Data14-
16	Data13+	32	GND	49	Data13-
17	Data12+	33	GND	50	Data12-

Table A-2 summarizes the Philips (Spirit DataCine) pinout for the cable’s 50-pin connector.

Table A-2 Philips 50-Pin Connector Pinout (Spirit DataCine)

Pin	Signal	Pin	Signal	Pin	Signal
1	CLK+			34	CLK-
2	Data9+	18	GND	35	Data9-
3	Data8+	19	GND	36	Data8-
4	Data7+	20	Data1+	37	Data7-
5	Data6+	21	Data1-	38	Data6-
6	Data5+	22	Data0+	39	Data5-
7	Data4+	23	Data0-	40	Data4-
8	Data3+	24	GND	41	Data3-
9	Data2+	25	GND	42	Data2-
10	Data19+	26	GND	43	Data19-
11	Data18+	27	GND	44	Data18-
12	Data17+	28	Data11+	45	Data17-
13	Data16+	29	Data11-	46	Data16-
14	Data15+	30	Data10+	47	Data15-
15	Data14+	31	Data10-	48	Data14-
16	Data13+	32	GND	49	Data13-
17	Data12+	33	GND	50	Data12-

Table A-3 summarizes the use of **LINK A** and **LINK B** connectors for 4:2:2:4 mode. If **LINK B** is not used in 4:2:2:4 format, the resulting format is 4:2:2. The **LINK A** connector carries 10-bit wide UYY information; the **LINK B** connector carries 10-bit alpha. The cables included with the HD I/O option each carry two channels in parallel.

Table A-3 LINK A and LINK B Usage in 4:2:2:4 Mode

Sample	LINK A	LINK B
0	Cb0 Y0	x A0
1	Cr1 Y1	x A1
2	Cb2 Y2	x A2
3	Cr3 Y3	x A3

Table A-4 summarizes usage for 10-bit RGBA.

Table A-4 LINK A and LINK B Usage in RGBA Mode

Sample	LINK A	LINK B
0	B0 G0	R0 A0
1	B1 G1	R1 A1
2	B2 G2	R2 A2
3	B3 G3	R3 A3

Table A-5 summarizes the use of **LINK A** and **LINK B** connectors for 4:4:4:4 mode. The **LINK A** connector carries a 4:2:2 sampled portion of 10-bit wide UVY; the **LINK B** connector carries the remaining 10-bit UV samples and 10-bit alpha. Usage is similar for 10-bit RGBA.

Table A-5 LINK A and LINK B Usage in 4:4:4:4 Mode

Sample	LINK A	LINK B
0	C _b 0 Y0	C _r 0 A0
1	C _b 1 Y1	C _r 1 A1
2	C _b 2 Y2	C _r 2 A2
3	C _b 3 Y3	C _r 3 A3

GPI Interface

The General Purpose Interface (GPI) port provides two channels of input and output trigger signal pairs on one connector. This section consists of the following subsections:

- “GPI Connector” on page 44
- “GPI Transmitter” on page 46
- “GPI Receiver” on page 47

GPI Connector

Figure A-2 points out the GPI connector on the HD I/O panel.




Figure A-2 GPI Connector

Figure A-3 shows the pinouts for the GPI connector.




Figure A-3 GPI Pinouts

Table A-6 gives the uses of the pins in Figure A-3.

Table A-6 GPI Pinouts

Pin	Use	For Video Pipe	CCT/CCR
8	In transmit +	In	CCT +
4	In transmit -	In	CCT -
5	Out transmit +	Out	CCT +
2	Out transmit -	Out	CCT -
6	In receive +	In	CCR +
7	In receive -	In	CCR -
3	Out receive +	Out	CCR +
1	Out receive -	Out	CCR -

Figure A-4 diagrams the relationship between the HD I/O board's video pipes and the GPI pins.




Figure A-4 GPI Pins and HD I/O Video Pipes

GPI Transmitter

GPI contact closure transmit (CCT) outputs use an optically coupled solid-state relay (SSR) to provide a means of electrical isolation for destination equipment.

Figure A-5 and Table A-7 show electrical specifications for the GPI transmitter.

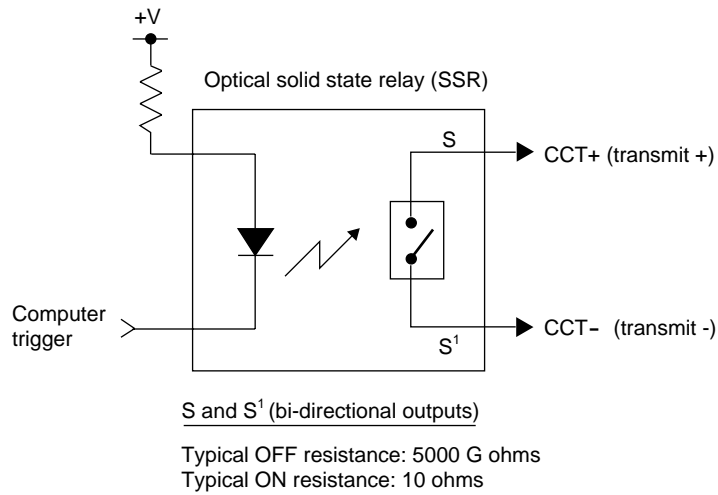


Figure A-5 GPI Transmitter Electrical Specifications

Table A-7 GPI Transmitter Electrical Specifications

Parameter	Value
On resistance	10 ohms typical, 15 ohms maximum
Off resistance	5000 G ohms
Current limit	360 mA typical, 460 mA maximum
Output capacitance	60 pF
Continuous DC load current	180 mA
Output power dissipation	600 mW
Isolation voltage	3750 V rms

The GPI transmitter can be interfaced to the destination equipment by tying the CCT- terminal to GND and using the CCT+ terminal as a current sink. The input device of the destination equipment can consist of a logic device with active pullup, an optoisolator LED with series-limiting resistor, or relay primary with series-limiting resistor.

The GPI transmitter's logic sense can be swapped (inverted) by tying the CCT+ terminal to the logic power supply (VCC) of the destination equipment and using the CCT- terminal to drive the input of the receiving device.

GPI Receiver

GPI receive (CCR) inputs use an optical isolator device to provide a means of electrical isolation from source equipment. The device consists of a bidirectional input LED optically coupled to a bipolar transistor. A voltage pulse applied across the CCR+/- pins causes the LED to become momentarily forward-biased, which produces a GPI trigger to the computer.

Figure A-6 shows electrical specifications for the GPI receiver.

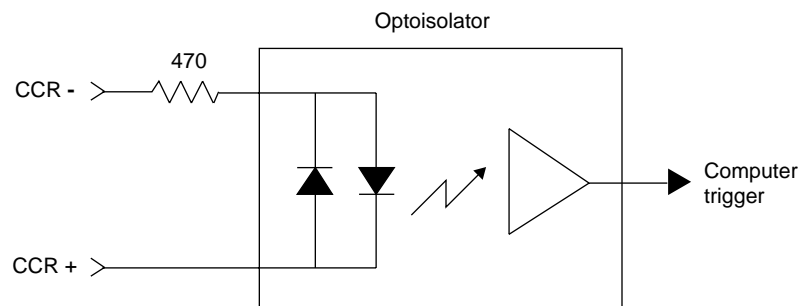


Figure A-6 GPI Receiver Electrical Specifications

Table A-8 summarizes electrical specifications for the GPI receiver optoisolator.

Table A-8 GPI Receiver Input Optoisolator Electrical Specifications

Parameter	Value
Forward voltage (V_F)	1.55 V, 1.2 V typical ($I_F = 10$ mA)
Continuous forward current (I_F)	30 mA
Peak forward current	1000 mA (10 μ s duration, 1% DC)
Reverse current (I_R)	0.1 μ A, 100 μ A maximum ($V_R = 6$ V)
Isolation surge voltage (V_{10})	2500 VAC _{RMS} (t = 1 min)

The GPI receiver can be interfaced to the source equipment by tying either the CCR+ terminal or the CCR- terminal across the output terminals of an optoisolator, solid-state relay, or any device that acts like a single-pole contact switch. The other terminal (CCR- or CCR+) must then be appropriately tied to power or ground. Whenever the logic device is sourcing current (driving a logic high), a GPI trigger is generated.

Note: Proper biasing and current limitations must be observed for the input LED. See Table A-8.

Genlock

The **GEN IN** and **GEN OUT** BNC connectors on the HD I/O front panel make up a passive genlock loopthrough connection. If you attach a cable to one **GEN** connector, you must attach to the other **GEN** connector either a 75-ohm BNC terminator or a cable to other equipment accepting analog sync. If another cable is connected, it must be terminated at the end of the loopthrough chain. Figure A-7 points out the genlock BNCs.




Figure A-7 Genlock BNCs

Setting Up the HD I/O Board for Your Video Hardware

This appendix illustrates how to attach video equipment to connectors on the HD I/O front panel and how to use the video control panel *vcp* to set the option to match your installation.

This appendix explains

- “Setting Up Digital Source Video” on page 50
- “Setting Up the Output (Drain)” on page 52
- “Setting Up Sync” on page 54
- “Saving Settings” on page 56

Figure B-1 shows connectors on the HD I/O front panel.




Figure B-1 HD I/O Ports

Setting Up Digital Source Video

The cables supplied with the HD I/O option each have four 50-in connectors for attaching to equipment that complies with the SMPTE 274M standard. The cables can be used for 4:4:4:4 or 4:2:2:4 dual-link mode; for 4:2:2 single-link mode, ignore the alpha:

- In 4:4:4:4 mode, Link A carries Y plus the U and V from even-numbered sample points; Link B carries alpha plus the U and V from odd-numbered sample points.
- In 4:2:2:4 mode, Link A carries Y plus the U and V from all sample points; Link B carries alpha only.

To set up the option for a digital video source, follow these steps:

1. Using the Panasonic-standard or Philips-standard cable included with the HD I/O option, attach the connector with the green label **SGI ← PANASONIC P1** or **SGI ← PHILIPS P1** to a device outputting data.

Depending on the equipment to be connected to, use the Panasonic or Philips cable. If necessary, see “Cable Connectors” on page 40 in Appendix A for an explanation of the labeling and color coding.




Figure B-2 HD I/O Cables

2. If you are using a device for alpha key data, also attach the connector with the red label **SGI <--- PANASONIC P2** or **SGI <--- PHILIPS P2** to a device outputting alpha.
3. Call up the panel:


```
/usr/sbin/vcp
```
4. In the Inputs(s): HD I/O Digital Video Source section of the control panel *vcp* (pointed out in Figure B-3) for the channel(s) you are using, select the format that matches your equipment: Digital Component 4:2:2:4 (the default) or Digital Component 4:4:4:4.




Figure B-3 Selecting Digital Input Video Format in *vcp*

5. In the Digital Video Source portion of the panel for the channel(s) you are using, select the timing that matches your equipment.

Setting Up the Output (Drain)

To set up the digital video output, follow these steps:

1. Using the Panasonic-standard or Philips-standard cable included with the HD I/O option, attach the connector with the green label **SGI ---> PANASONIC P3** or **SGI ---> PHILIPS P3** to a device to which you want to send data. If you use only one output, you must use this connector.

Depending on the equipment to be connected to, use the Panasonic or Philips cable. If necessary, see "Cable Connectors" on page 40 in Appendix A for an explanation of the labeling and color coding.

2. If you are using a device for alpha key data, also attach the connector with the red label **SGI ---> PANASONIC P4** or **SGI ---> PHILIPS P4** to a device receiving alpha key data.
3. If necessary, call up the panel (`/usr/sbin/vcp`).

4. In the Output(s): Video Drain section of the control panel (pointed out in Figure B-4), select the format that matches your equipment: Digital Component 4:2:2:4 (the default) or Digital Component 4:4:4:4.




Figure B-4 Selecting Video Drain Format

5. Select the timing that matches your equipment.
6. To set field dominance, at the "Input Timing" menu item select "F1 dominant" for the edit to occur on the nominal video field boundary, or "F2 dominant" for the edit to occur on the intervening field boundary. See "Field Dominance" on page 29 in Chapter 2 for more information on field dominance.

Setting Up Sync

This section explains

- “Setting Up Internal Sync” on page 54
- “Setting Up External Sync” on page 55

Setting Up Internal Sync

In the Output(s): Digital Video Drain section of the control panel, select the Sync format that matches your equipment:

- standalone (not synced to another device): select Standalone (internal)
- output sync to an external source connected to the genlock in: select Genlock

These two choices toggle; Figure B-5 shows the default.

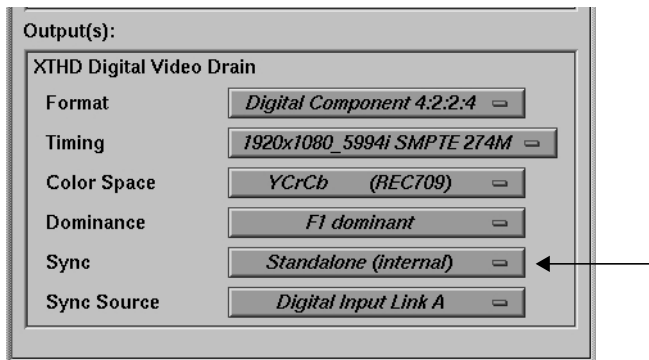


Figure B-5 Setting Standalone or Genlock Sync

Setting Up External Sync

To set up the HD I/O option for an external sync source, follow these steps:

1. If necessary, call up the panel (`/usr/sbin/vcp`).
2. Select the appropriate external sync source in Output(s): Sync Source:
 - For a device connected to **GEN IN**, select External 1920x1080_5994_i or External 525 NTSC depending on your sync source.
 - If you are syncing to the device attached to **SIG ← PANASONIC/PHILIPS P1**, select Digital Input Link A.
 - If you are syncing to the device attached to **SIG ← PANASONIC/PHILIPS P2**, select Digital Input Link B.

Note: See Table 2-4 on page 20 in Chapter 2 for details of the interrelationships of timing and sync source.

If a genlock source has not been cabled to the HD I/O panel yet, any application in progress might generate an error message.

3. Connect the sync source equipment to one of the following connectors:
 - the **GEN IN** BNC on the I/O panel (see Figure B-1 if necessary)or
 - the cable connector **SIG ← PANASONIC/PHILIPS P1** or **SIG ← PANASONIC/PHILIPS P2** (if this cable connector is not already attached to the device supplying sync)
4. If you are using the same signal for other equipment, attach a BNC cable to the **GEN OUT** BNC to loop the signal through the board. Make sure the final element in the chain is terminated.

If the HD I/O board is the last element in the sync chain, attach a terminator to the **GEN OUT** BNC.

5. Check the genlock LEDs:
 - If the yellow **VALID** LED is lit, an acceptable sync source is attached to the genlock input.
 - If the green **LOCKED** LED is lit, genlock is enabled (via `vcp` or an application), the reference is adequate, and the board is ready to use.

Saving Settings

Once you have set values in *vcp* to match your installation, save them; they are written to */usr/etc/video/videod.defaults*. Select "Restore Settings" on the video control panel File menu to load the values in this file to *vcp*.

The last settings saved are automatically loaded every time the system is reinitialized. If the panel is running, current settings are in effect.

Note: You do not need to open the panel to put its settings into effect.

You can also use File menu choices to restore the factory defaults and close the panel.

Pixel Packings and Color Spaces

This appendix explains

- “HD I/O Pixel Packings” on page 57
- “Sampling Patterns” on page 74

HD I/O Pixel Packings

This section presents each packing used by the HD I/O hardware, giving a diagram and its tokens in the pertinent libraries. It explains

- “Packings and Color Spaces” on page 57
- “Packing Diagram Conventions” on page 58
- “Packings and Library Tokens” on page 60
- “Packing Naming Conventions” on page 60
- “16-Bit Pixel Packings” on page 62
- “20-Bit Pixel Packings” on page 63
- “24-Bit Pixel Packings” on page 64
- “32-Bit Pixel Packings” on page 65

Packings and Color Spaces

A packing

- determines which of the four components are sampled, either RGBA or VYUA (more correctly, CrYCbA)
- determines the sampling pattern (for example, 4:4:4:4 or 4:2:2:4), which specifies where and how often each component of the image is sampled

- allocates a certain number of bits to represent the component samples, and positions those samples along with possible padding in memory; each sample is an unsigned number, unless specified otherwise in the description of the packing

A color space

- determines the color in each component by specifying the color set
- specifies a canonical minimum and maximum value for each component, either full range or compressed range (headroom range); see “VL_COLORSPACE” on page 22 in Chapter 2 for an explanation

In most Silicon Graphics libraries, a single token encodes both color space and packing. For example, VL_PACKING_RGBA_8 is a 32-bit packing in the RGBA color space. For the VL of SGI advanced video products, the two parameters are specified separately with different controls: VL_PACKING and VL_COLORSPACE. The color space must be defined with the VL_COLORSPACE control.

Packing Diagram Conventions

In all illustrations, as you move from left to right:

- each byte goes from the most significant bit to the least significant bit
- the bytes increase in memory address by 1
- component samples go from most significant bit to least significant bit

Each illustration shows the smallest repeating spatial pattern of component samples that is a multiple of 8 bits wide. No additional padding or alignment is to be inferred. For example, a 24-bit-per-pixel diagram, such as that for VL_PACKING_444_8 (Figure C-1), indicates 3-byte quantities packed together in memory; the values are not padded out to 32-bit boundaries.

Pixel 1		
Byte 1	Byte 2	Byte 3
r r r r r r r r	g g g g g g g g	b b b b b b b b
v v v v v v v v	y y y y y y y y	u u u u u u u u

Figure C-1 VL_PACKING_444_8

The packing defines a bit layout, but for convenience, as shown in Figure C-1, the component slots are filled with the RGBA or VYUA color set as appropriate. (See “VL_COLORSPACE” on page 22 in Chapter 2 for more information.) For chroma components, Cr and Cb are more accurate terms than V and U; however, this chapter uses the letters V and U in the illustrations of packings for typographical convenience.

Packings that use 4:2:2 sampling also show the location of each component sample: left and right for 4:2:2. The diagrams assume row-major, left-to-right ordering of pixels in memory.

An x (“don’t care”) in a bit means the following:

- Readers may get any garbage in this bit.
- Writers can leave this bit as garbage.

A 0 means the following:

- Readers may assume this bit is zero.
- Writers can leave this bit as garbage.

An s indicates a padding bit that is a sign extension bit. For the HD I/O option, this convention applies only to the more significant bits in 12-bit and 13-bit packings with rightward orientation; that is, VL_PACKING_4444_12_in_16_R and VL_PACKING_4444_13_in_16_R.

A p indicates a padding bit in the least significant bits of a left-justified 10-, 12-, or 13-bit word, such as VL_PACKING_R242_10_in_16_L or VL_PACKING_4444_13_in_16_L:

- Readers can assume that the bits are replicated from the component found in the same word: With bits numbered starting with 0 for the least significant, there are n contiguous p bits to the right of the component. The p bits contain a copy of bits [9,9-n+1] of the component.
- Writers can leave the p bits as garbage.

The HD I/O device can natively transfer data of all the packings shown in this appendix in real time.

Packings and Library Tokens

Following each packing diagram are comments and library tokens for that packing, listing, where applicable, the color set (RGBA or VYUA) and the library (VL, OpenGL, and DM) for each library token.

- DM refers to the tokens in */usr/lib/dmedia/dm_image.h*, which are used by several libraries (*libdmedia* (*dmParams*, *dmIC*, *dmColor*), *libmoviefile*, *libmovieplay*, and others). See “VL_COLORSPACE” on page 22 in Chapter 2 for more information.
- The VL includes the packing control value and a color-space control value; for example, `VL_PACKING_4_8 + VL_COLORSPACE_{CCIR,YUV}`. For the HD I/O option, you set packing and color space separately for memory nodes. These definitions provide a more flexible way to specify memory layout of pixels and their color spaces.

Note: Because of HDTV’s multiple color spaces, “old style” packings, such as `VL_PACKING_Y_8_P`, are ambiguous and therefore not supported longer for the HD I/O board. You must specify the packing and color space explicitly.

All HD I/O packings are also supported by the DIVO and DIVO-DVC boards.

Packing Naming Conventions

In packing tokens, the following applies:

- `_L` and `_R` appended to the end of tokens with padding (0 bits) indicate that the 0 bits are at the left end or the right end of the pattern, respectively; for example, `VL_PACKING_4444_10_in_16_L` and `VL_PACKING_4444_10_in_16_R`.
- `X` before the numerical part of the token at the end of a token indicates a component order other than the standard (RGBA or ABGR, VYUA or AUYV); for example, `DM_IMAGE_PACKING_XBGR`.
- `R` before the numerical part of the token indicates reverse order of the components; for example, `VL_PACKING_242_8` and `VL_PACKING_R242_8` have the same pattern of component bits, but their order is reversed in `VL_PACKING_R242_8`.
- `Z` at the end of the token name means that the packing is padded to the word boundary; for example, the packing in `VL_PACKING_2424_10_10_10_2Z` is 30 bits per pixel, but it is padded to 32 bits per pixel.

Table C-1 lists the HD I/O packings by number of bits in the pattern of component samples—the order in which they are described in the rest of this section.

Table C-1 HD I/O Packings

Packing	Bits	Color Space
VL_PACKING_242_8	16	VYUA
VL_PACKING_R242_8	16	VYUA
VL_PACKING_242_10	20	VYUA
VL_PACKING_R242_10	20	VYUA
VL_PACKING_444_8	24	RGBA/VYUA
VL_PACKING_R444_8	24	RGBA/VYUA
VL_PACKING_4444_8	32	RGBA/VYUA
VL_PACKING_R4444_8	32	RGBA/VYUA
VL_PACKING_R0444_8	32	RGBA/VYUA
VL_PACKING_0444_8	32	RGBA/VYUA
VL_PACKING_4444_10_10_10_2	32	RGBA/VYUA
VL_PACKING_R4444_10_10_10_2	32	RGBA/VYUA
VL_PACKING_2424_10_10_10_2Z	32	VYUA
VL_PACKING_R2424_10_10_10_2Z	32	VYUA
VL_PACKING_242_10_in_16_L	32	VYUA
VL_PACKING_242_10_in_16_R	32	VYUA
VL_PACKING_R242_10_in_16_L	32	VYUA
VL_PACKING_R242_10_in_16_R	32	VYUA

The packings are explained in these categories:

- “16-Bit Pixel Packings” on page 62
- “20-Bit Pixel Packings” on page 63

- “24-Bit Pixel Packings” on page 64
- “32-Bit Pixel Packings” on page 65

16-Bit Pixel Packings

Figure C-2 shows VL_PACKING_242_8, a 16-bit VYUA packing.

Pixels 1-2			
Byte 1	Byte 2	Byte 3	Byte 4
v v v v v v v v	y y y y y y y y	u u u u u u u u	y y y y y y y y
left			right

Figure C-2 VL_PACKING_242_8

Note: Cr and Cb are more accurate terms than V and U; however, this appendix uses the letters V and U in the illustrations of packings for typographical convenience.

This rarely used packing is VL_PACKING_242_8 + VL_COLORSPACE_{CCIR,YUV} in the VL. It samples chroma and luma in a 4:2:2 pattern. See “Sampling Patterns” on page 74.

Figure C-3 shows VL_PACKING_R242_8, a 16-bit 4:2:2 VYUA packing. The most commonly used 4:2:2 packing, it is used by other Silicon Graphics video hardware.

Pixels 1-2			
Byte 1	Byte 2	Byte 3	Byte 4
u u u u u u u u	y y y y y y y y	v v v v v v v v	y y y y y y y y
left			right

Figure C-3 VL_PACKING_R242_8

This packing is

- VL_PACKING_R242_8 + VL_COLORSPACE_{CCIR,YUV}
- GL_YCRCB_422_SGIX GL_UNSIGNED_BYTE in OpenGL
- DM_IMAGE_PACKING_CbYCrY in DM

20-Bit Pixel Packings

Figure C-4 shows VL_PACKING_242_10, a 20-bit VYUA packing.

Pixels 1-2				
Byte 1	Byte 2	Byte 3	Byte 4	Byte 5
v v v v v v v v	v v y y y y y y	y y y y u u u u	u u u u u u y y	y y y y y y y y
left			right	

Figure C-4 VL_PACKING_242_10

This packing is VL_PACKING_242_10 + VL_COLORSPACE {CCIR,YUV}. It uses 4:2:2 sampling.

Figure C-5 shows VL_PACKING_R242_10, a 20-bit VYUA packing.

Pixels 1-2				
Byte 1	Byte 2	Byte 3	Byte 4	Byte 5
u u u u u u u u	u u y y y y y y	y y y y v v v v	v v v v v v y y	y y y y y y y y
left			right	

Figure C-5 VL_PACKING_R242_10

This packing is VL_PACKING_R242_10 + VL_COLORSPACE {CCIR,YUV}. It uses 4:2:2 sampling.

24-Bit Pixel Packings

Figure C-6 shows VL_PACKING_444_8, a 24-bit RGBA/VYUA packing.

Pixel 1		
Byte 1	Byte 2	Byte 3
r r r r r r r r	g g g g g g g g	b b b b b b b b
v v v v v v v v	y y y y y y y y	u u u u u u u u

Figure C-6 VL_PACKING_444_8

This packing is

- **RGBA:**
 - GL_RGB GL_UNSIGNED_BYTE in OpenGL
 - VL_PACKING_444_8 + VL_COLORSPACE_{RGB,RP175}
 - PM_RGB GL_UNSIGNED_BYTE on RealityEngine (IRIS GL)
 - DM_IMAGE_PACKING_RGB in DM
- **VYUA:** VL_PACKING_444_8 + VL_COLORSPACE_{YUV/YCRCB}

Figure C-7 shows VL_PACKING_R444_8, a 24-bit RGBA/VYUA packing.

Pixel 1		
Byte 1	Byte 2	Byte 3
b b b b b b b b	g g g g g g g g	r r r r r r r r
u u u u u u u u	y y y y y y y y	v v v v v v v v

Figure C-7 VL_PACKING_R444_8

This packing is

- **RGBA:**
 - VL_PACKING_R444_8 + VL_COLORSPACE_{RGB,RP175}
 - DM_IMAGE_PACKING_BGR in DM
 - PM_BGR PM_UNSIGNED_BYTE on RealityEngine (IRIS GL)
- **VYUA:**
 - VL_PACKING_R444_8 + VL_COLORSPACE_{CCIR,YUV}
 - DM_IMAGE_PACKING_CbYCr in DM

32-Bit Pixel Packings

This section explains

- “OpenGL-Like 32-Bit Pixel Packing” on page 66
- “IRIS GL-Like 32-Bit Pixel Packings” on page 67
- “32-Bit Pixel Packing for QuickTime” on page 69
- “4:4:4:4 10_10_10_2 32-Bit Pixel Packings” on page 70
- “4:2:2:4 10_10_10_2 32-Bit Pixel Packings” on page 71
- “4:2:2 10_in_16 32-Bit Pixel Packings” on page 72

OpenGL-Like 32-Bit Pixel Packing

Figure C-8 shows VL_PACKING_4444_8, an OpenGL-like 32-bit packing. This packing, supported by many Silicon Graphics video products, is the most commonly used OpenGL packing.

Pixel 1			
Byte 1	Byte 2	Byte 3	Byte 4
r r r r r r r r	g g g g g g g g	b b b b b b b b	a a a a a a a a
v v v v v v v v	y y y y y y y y	u u u u u u u u	a a a a a a a a

Figure C-8 VL_PACKING_4444_8

This packing is

- **RGBA:**
 - VL_PACKING_4444_8 + VL_COLORSPACE_{RGB,RP175}
 - GL_RGBA GL_UNSIGNED_BYTE in OpenGL (the default)
 - DM_IMAGE_PACKING_RGBA in DM
 - PM_RGBA PM_UNSIGNED_BYTE on RealityEngine (IRIS GL)
- **VYUA:** VL_PACKING_4444_8 + VL_COLORSPACE_{CCIR,YUV}

IRIS GL-Like 32-Bit Pixel Packings

Figure C-9 shows VL_PACKING_R4444_8, an IRIS GL-like 32-bit packing. This packing, supported by many Silicon Graphics video products, is the default IRIS GL packing.

Pixel 1			
Byte 1	Byte 2	Byte 3	Byte 4
a a a a a a a a	b b b b b b b b	g g g g g g g g	r r r r r r r r
a a a a a a a a	u u u u u u u u	y y y y y y y y	v v v v v v v v

Figure C-9 VL_PACKING_R4444_8

This packing is

- **RGBA:**
 - VL_PACKING_R4444_8 + VL_COLORSPACE_{RGB,RP175}
 - GL_ABGR_EXT GL_UNSIGNED_BYTE in OpenGL
 - DM_IMAGE_PACKING_ABGR in DM
 - PM_ABGR PM_UNSIGNED_BYTE (default IRIS GL packing)
- **VYUA:** VL_PACKING_R4444_8 + VL_COLORSPACE_{CCIR,YUV}

Figure C-10 shows VL_PACKING_R0444_8, an IRIS GL-like 32-bit packing. This packing is supported by many Silicon Graphics video products.

Pixel 1			
Byte 1	Byte 2	Byte 3	Byte 4
x x x x x x x x	b b b b b b b b	g g g g g g g g	r r r r r r r r
x x x x x x x x	u u u u u u u u	y y y y y y y y	v v v v v v v v

Figure C-10 VL_PACKING_R0444_8

- **RGBA:**
 - VL_PACKING_R0444_8 + VL_COLORSPACE_{RGB,RP175}
 - DM_IMAGE_PACKING_XBGR
 - Use DM_IMAGE_PACKING_ABGR instead of this packing unless you specifically want to inform a piece of software (such as dmColor) not to spend processing time on the alpha channel.
- **VYUA:** VL_PACKING_R0444_8 + VL_COLORSPACE_{CCIR,YUV}

32-Bit Pixel Packing for QuickTime

Figure C-11 shows VL_PACKING_0444_8, a 32-bit packing used for QuickTime files (uncompressed format without alpha).

Pixel 1			
Byte 1	Byte 2	Byte 3	Byte 4
x x x x x x x x	r r r r r r r r	g g g g g g g g	b b b b b b b b
x x x x x x x x	v v v v v v v v	y y y y y y y y	u u u u u u u u

Figure C-11 VL_PACKING_0444_8

This packing is

- RGBA:
 - VL_PACKING_0444_8 + VL_COLORSPACE_{RGB,RP175}
 - DM_IMAGE_PACKING_XRGB in DM
- VYUA: VL_PACKING_0444_8 + VL_COLORSPACE_{CCIR,YUV}

4:4:4:4 10_10_10_2 32-Bit Pixel Packings

Figure C-12 shows VL_PACKING_4444_10_10_10_2, a 32-bit 4:4:4:4 10_10_10_2 packing.

Pixel 1			
Byte 1	Byte 2	Byte 3	Byte 4
r r r r r r r r	r r g g g g g g	g g g g b b b b	b b b b b b a a
v v v v v v v v	v v y y y y y y	y y y y u u u u	u u u u u u a a

Figure C-12 VL_PACKING_4444_10_10_10_2

This packing is

- RGBA:
 - VL_PACKING_4444_10_10_10_2 + VL_COLORSPACE_{RGB,RP175}
 - GL_RGBA GL_UNSIGNED_INT_10_10_10_2_EXT in OpenGL
- VYUA: VL_PACKING_4444_10_10_10_2 + VL_COLORSPACE_{CCIR,YUV}

Figure C-13 shows VL_PACKING_R4444_10_10_10_2, an alternate 32-bit 4:4:4:4 10_10_10_2 packing.

Pixel 1			
Byte 1	Byte 2	Byte 3	Byte 4
a a b b b b b b	b b b b g g g g	g g g g g g r r	r r r r r r r r
a a u u u u u u	u u u u y y y y	y y y y y y v v	v v v v v v v v

Figure C-13 VL_PACKING_R4444_10_10_10_2

This packing is

- RGBA:
 - VL_PACKING_R4444_10_10_10_2 + VL_COLORSPACE_{RGB,RP175}
 - GL_ABGR GL_UNSIGNED_INT_10_10_10_2_EXT in OpenGL
- VYUA: VL_PACKING_R4444_10_10_10_2 + VL_COLORSPACE_{CCIR,YUV}

4:2:2:4 10_10_10_2 32-Bit Pixel Packings

Figure C-14 shows VL_PACKING_2424_10_10_10_2Z, the 4:2:2:4 10_10_10_2 32-bit VYUA packing.

Pixels 1-2							
Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6	Byte 7	Byte 8
v v v v v v v v	v v y y y y y y	y y y y a a a a	a a a a a a 0 0	u u u u u u u u	u u y y y y y y	y y y y a a a a	a a a a a a 0 0
left				0 0	left	right	0 0

Figure C-14 VL_PACKING_2424_10_10_10_2Z

This packing is VL_PACKING_2424_10_10_10_2Z + VL_COLORSPACE_{CCIR,YUV} in the VL.

Figure C-15 shows VL_PACKING_R2424_10_10_10_2Z, an alternate 4:2:2:4 10_10_10_2 32-bit packing.

Pixels 1-2							
Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6	Byte 7	Byte 8
u u u u u u u u	u u y y y y y y	y y y y a a a a	a a a a a a 0 0	v v v v v v v v	v v y y y y y y	y y y y a a a a	a a a a a a 0 0
left				0 0	left	right	0 0

Figure C-15 VL_PACKING_R2424_10_10_10_2Z

This packing is VL_PACKING_R2424_10_10_10_2Z + VL_COLORSPACE_{CCIR,YUV} in the VL.

4:2:2 10_in_16 32-Bit Pixel Packings

The diagrams of packings that use 4:2:2 sampling show the location (left and right) of each component sample. Only DIVO and DIVO-DVC use this packing.

Figure C-16 shows VL_PACKING_242_10_in_16_L, a 4:2:2 10_in_16 32-bit VYUA packing. This packing is supported by several Silicon Graphics video products. For an explanation of the p bit, see "Packing Diagram Conventions" on page 58.

Pixels 1-2							
Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6	Byte 7	Byte 8
v v v v v v v v	v v p p p p p p	y y y y y y y y	y y p p p p p p	u u u u u u u u	u u p p p p p p	y y y y y y y y	y y p p p p p p
left	p p p p p p	left	p p p p p p	left	p p p p p p	right	p p p p p p

Figure C-16 VL_PACKING_242_10_in_16_L

This packing is VL_PACKING_242_10_in_16_L + VL_COLORSPACE_{CCIR,YUV} in the VL.

Figure C-17 shows VL_PACKING_242_10_in_16_R, a 4:2:2 10_in_16 32-bit VYUA packing.

Pixels 1-2							
Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6	Byte 7	Byte 8
0 0 0 0 0 0 v v	v v v v v v v v	0 0 0 0 0 0 y y	y y y y y y y y	0 0 0 0 0 0 u u	u u u u u u u u	0 0 0 0 0 0 y y	y y y y y y y y
0 0 0 0 0 0 left	0 0 0 0 0 0 left	0 0 0 0 0 0 left	0 0 0 0 0 0 left	0 0 0 0 0 0 right	0 0 0 0 0 0 right	0 0 0 0 0 0 right	0 0 0 0 0 0 right

Figure C-17 VL_PACKING_242_10_in_16_R

This packing is VYUA VL_PACKING_242_10_in_16_R + VL_COLORSPACE_{CCIR,YUV} in the VL.

Figure C-18 shows VL_PACKING_R242_10_in_16_L, a 4:2:2 10_in_16 32-bit VYUA packing. This packing is supported by several Silicon Graphics video products. For an explanation of the p bit, see “Packing Diagram Conventions” on page 58.

Pixels 1-2							
Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6	Byte 7	Byte 8
u u u u u u u u	u u p p p p p p	y y y y y y y y	y y p p p p p p	v v v v v v v v	v v p p p p p p	y y y y y y y y	y y p p p p p p
left		left		left		right	

Figure C-18 VL_PACKING_R242_10_in_16_L

This packing is VL_PACKING_R242_10_in_16_L + VL_COLORSPACE_{CCIR,YUV} in the VL.

Figure C-19 shows VL_PACKING_R242_10_in_16_R, a 4:2:2 10_in_16 32-bit VYUA packing.

Pixels 1-2							
Byte 1	Byte 2	Byte 3	Byte 4	Byte 5	Byte 6	Byte 7	Byte 8
0 0 0 0 0 0 u u	u u u u u u u u	0 0 0 0 0 0 y y	y y y y y y y y	0 0 0 0 0 0 v v	v v v v v v v v	0 0 0 0 0 0 y y	y y y y y y y y
0 0 0 0 0 0 left		0 0 0 0 0 0 left		0 0 0 0 0 0 left		0 0 0 0 0 0 right	

Figure C-19 VL_PACKING_R242_10_in_16_R

This packing is VYUA VL_PACKING_R242_10_in_16_R + VL_COLORSPACE_{CCIR,YUV} in the VL.

Sampling Patterns

Sampling patterns are

- “4:4:4 and 4:4:4:4 Sampling” on page 74
- “4:2:2 and 4:2:2:4 Sampling” on page 75

4:4:4 and 4:4:4:4 Sampling

Some of the diagrams in the “HD I/O Pixel Packings” section indicate 4:4:4 or 4:4:4:4 sampling. This video industry terminology means that each of the three or four components is sampled at every pixel. Figure C-20 diagrams this sampling pattern.




Figure C-20 4:4:4 Sampling

4:2:2 and 4:2:2:4 Sampling

The packings shown in diagrams that indicate 4:2:2 sampling make sense only in the VYUA color spaces. For every two pixels, there are two luma samples (two Ys) but only one chroma sample (one sample of Cr and Cb, which together determine the chroma), as shown in Figure C-21.




Figure C-21 4:2:2 Sampling

The chroma samples belong at the same instant in space as the left Y sample (the chrominance samples and the left Y are co-sited). The diagrams for 4:2:2 packings in the “HD I/O Pixel Packings” section of this appendix show the location of each Y, Cr, or Cb component as left or right. The first pixel of each line is a left pixel.

Converting 4:4:4 video to 4:2:2 video is like converting 44.1 kHz audio into 22.05 kHz audio: just dropping every other Cr,Cb sample yields extremely poor results. Video devices that need to convert between 4:4:4 and 4:2:2 use carefully designed filters. The characteristics of the required filter are specified in ITU-R BT.601 (Rec. 601).

4:2:2 sampled packings that also include alpha are called 4:2:2:4. This method has one alpha value per pixel, like the Y value.

Programming Methods for Real-Time Digital Media Recording and Playback

This appendix explains the following real-time disk I/O concepts:

- “Digital Media Buffers” on page 78
- “Direct I/O” on page 78
- “Multiprocessing” on page 80
- “Asynchronous I/O” on page 81

The methods described in this appendix used digital media buffers (DMbuffers), a real-time data transport facility. See the *Digital Media Programming Guide* (document number 007-1799-060 or later, hereafter referred to as the DMPG) for more details. The emphasis here is not on how data is acquired from or transported to the video device, but rather on how data is moved to disk in real time.

The DMPG covers basic digital media programming concepts. The directory `/usr/share/src/dmedia/video/XTHD` contains two simple programming examples:

- `sdhd_memtovid.c` illustrates how video data is copied out of the DMbuffers for the simpler non-real-time case
- `sdhd_vidtogfx.c` illustrates how video data is transferred into memory and from there into graphics

At an abstract level, high-bandwidth throughput is simple; the work is in the details, as explained in this appendix.

Note: For source code examples for developer I/O, contact the SGI Developer program, for example, through the Web site <http://www.sgi.com>.

Digital Media Buffers

In IRIX 6.5.4 and later, you can tune how memory is allocated for digital media buffer pools. This capability can offer a performance advantage for some HD I/O applications.

For example, the parameter `DM_POOL_NODE_MASK` allows you to request memory allocation on particular physical node cards. Each bit in the node mask enables memory allocation on that node. For example, to allocate buffers distributed across nodes 0 and 1, set the node mask to 3:

```
long long mask = 0x3;

if (dmParamsSetLongLong(plist, DM_POOL_NODE_MASK, mask) ) {
    fprintf(stderr, "Error setting dmparam for pool node mask\n");
    exit(-1);
}

if (dmBufferCreatePool(plist, &pool) != DM_SUCCESS) {
    fprintf(stderr, "Error creating dmbuffer pool\n");
    exit(-1);
}
```

Note that the node mask is a long long parameter. To see the effect of setting this parameter, use `osview -a` and examine the free memory on each node. For more information on `dmBuffers`, see the `dmbuffercreatepool(3dm)` man page.

Direct I/O

The most efficient way to move data on and off a disk device is to use the XFS filesystem with direct I/O mode and large data transfer sizes. If large transfer sizes cannot be achieved, you can combine memory pages from noncontiguous locations using `writev(2)` or `readv(2)`. Finally, you can use asynchronous I/O to queue multiple I/O requests to the kernel without waiting for blocked calls to return. Other real-time software features and products, such as REACT, can be used to assure low-latency interrupts and high-priority scheduling, but are not absolutely necessary for digital media applications.

Normally, when a disk file is opened with no status flags specified, a call to `write(2)` for that file returns as soon as the data has been copied to a buffer managed by the device driver (see `open(2)`). The actual disk write may not take place until considerable time has passed. A common pool of disk buffers is used for all disk files.

Disk buffering is integrated with the virtual memory paging mechanism. A daemon executes periodically and initiates output of buffered blocks according to the age of the data and the needs of the system. You can force the writing of all pending output for a file by calling `fsync(2)` or by opening the file and specifying the `O_SYNC` flag. However, the process blocks until the data has been written to disk, and all output data must still be copied from the buffer in the user address space to a buffer in the kernel address space. See Chapter 8, “Optimizing Disk I/O for a Real-Time Program,” in the *REACT Real Time Programmer’s Guide* for details.

If you use the `O_DIRECT` flag, writes to the file take place directly from your program’s buffer, and the data is not copied to a buffer in the kernel first. Because the filesystem cache is bypassed, your application must manage buffer alignment and block size specification. To use `O_DIRECT`, you must transfer data in quantities that are multiples of the filesystem block size. The following code shows how to query the filesystem block size and system DMA transfer size limit.

```
struct dioattr da;
struct stat fileStat;
char *ioFileName = "videodata";
int ioBlockSize, ioMaxXferSize;

ioFileFD = open(ioFileName,O_DIRECT | O_RDWR | O_CREAT | O_TRUNC,0644);
if (ioFileFD < 0)
    return(DM_FAILURE);
if (fcntl(ioFileFD, F_DIOINFO, &da) < 0)
    return(DM_FAILURE);
ioBlockSize = da.d_miniosz;
ioMaxXferSize = da.d_maxiosz;
```

The two important constraints of direct I/O with XFS are memory address alignment and buffer length. Direct I/O requires all memory addresses to be page-aligned. XFS requires buffers to be allocated as a multiple of the filesystem block size, *ioBlockSize*. DMbuffers are guaranteed to be page-aligned, but to ensure that the buffers are properly padded, you must set the buffer size, *bytesPerXfer*, to the size of the image data you will transfer rounded up to the nearest multiple of *ioBlockSize*.

```
VLServer vlServer;
VLPath vlPath;
DMparams * paramsList;
int dmBufferPoolSize = 30; /* 1 second of video */
int vlBytesPerImage = vlGetTransferSize(vlServer, vlPath);
int ioBlocksPerImage = (vlBytesPerImage+ioBlockSize - 1) / ioBlockSize;
int bytesPerXfer = ioBlocksPerImage * ioBlockSize;
```

```
if (dmBufferSetPoolDefaults(paramsList,dmBufferPoolSize,bytesPerXfer,
    DM_TRUE, DM_TRUE) == DM_FAILURE) {
    fprintf(stderr, "error setting pool defaults\n");
    return(DM_FAILURE);
}
```

All Silicon Graphics systems have a configurable maximum DMA transfer size (see `sysctl(1M)`). This value should be compared with the user's I/O request size.

```
if (bytesPerXfer > ioMaxXferSize) {
    fprintf("DMA request size is too small. Reconfigure with
        sysctl()\n");
    return(DM_FAILURE);
}
```

Multiprocessing

Some aspects of digital media programming lend themselves to a multiprocessing programming model. On a multiprocessor system, the various tasks of moving multiple streams of video and audio data on and off disk, serial I/O control of external video equipment and input devices, processing of video data, or the transport of video data in and out of the graphics framebuffer can be assigned to different processors. New processes must be created with all virtual space attributes (shared memory, mapped files, data space) shared. The following fragment illustrates how to create a process to perform video recording.

```
if ((video_recorder_pid = sproc(video_recorder, PR_SADDR|PR_SFDS))<0){
    perror("video_recorder");
    exit(DM_FAILURE);
}
```

If you use multiprocessing, note the following caveats:

- When VL calls are made, VL objects such as `VLServer`, `VLPath`, `VLNode`, and so on, are passed through the kernel to the video driver. However, you cannot create any VL objects without first creating a `VLServer`, from which everything else is instanced.
- In a process share group, only one VL call whose arguments derive from a `VLServer` can execute at a time. This requirement applies even to VL calls that do not explicitly take a `VLServer` as an argument (for example, `vlBufferAdvise(3dm)`).

- You can use objects derived from a given VLServer in any number of threads as long as you use a locking scheme, such as `usnewsema(3P)` or `pthread_mutex_init(3P)`, to make the use in each thread mutually exclusive of a use in any of the other threads.

The VL error state, returned by `vlGetErrno(3dm)`, is global to a share group, not per VLServer. If a VL call using one VLServer in one thread executes simultaneously with a VL call using another VLServer in another thread, both calls try to set the error state returned by `vlGetErrno()`. This call should be global only to the thread, not to the entire process share group.

Asynchronous I/O

Asynchronous I/O allows an application to process multiple read or write requests simultaneously. On Silicon Graphics platforms, asynchronous I/O is available through the *aio* facility. This facility, based on `sproc(2)`'ed processes, provides all of the benefits of multiprocessing for free. Because multiple I/O requests might be outstanding, when you use asynchronous I/O, the round-trip delay between making a request, having it serviced, and issuing another request is removed. Any process-scheduling delay between these steps is also eliminated.

Because asynchronous I/O operations complete out of sequence, the application must keep track of the order in which data appears in the DMbuffers. DMbuffers are contained in a DMbufferPool; the pool itself is unordered and buffers can be obtained and returned to the pool in any order. Ordering is achieved by a first-in-first-out queue and maintained only while the buffers reside in the queue. The application is free to impose any processing order once buffers are dequeued.

Installing HD I/O Software on a New Disk

Use *inst* to install the software. See the *IRIX Admin: Software Installation and Licensing* manual (007-1364-xxx) if you need more detailed information on booting the miniroot from a remote CD-ROM.

Follow these steps:

1. Make sure a CD-ROM drive is attached to the system (or is available on the network).
2. If necessary, upgrade the system software to IRIX 6.5.4 or later, from the CD set included with the HD I/O software.
3. Install the HD I/O software.
4. Reboot the system.
5. Enter **hinv** at the IRIX level to determine whether the system sees all boards installed. A typical output might be as follows:

```
# hinv
2 180 MHZ IP27 Processors
CPU: MIPS R10000 Processor Chip Revision: 2.6
FPU: MIPS R10010 Floating Point Chip Revision: 0.0
Main memory size: 256 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 1 Mbyte
Integral SCSI controller 0: Version QL1040B, single ended
  Disk drive: unit 1 on SCSI controller 0
Integral SCSI controller 1: Version QL1040B, single ended
IOC3 serial port: tty1
IOC3 serial port: tty2
Integral Fast Ethernet: ef0, version 1, module 1, slot io1, pci 2
Origin BASEIO board, module 1 slot 1: Revision 3
XT-HDIO Video: controller 2, unit 0, version 0x0
IOC3 external interrupts: 1
```

You can use *hinv -m* to get the board serial number, revision level, and other statistics. A typical output might be as follows.

```
# hinv -m
      MODULEID Board: barcode K0002031   part          rev
      8P12-MPLN Board: barcode CEM652    part 013-1547-003 rev A
      IP27 Board: barcode DAW914        part 030-1266-001 rev C
      BASEIO Board: barcode DCH915      part 030-1124-002 rev A
      HDTV1 Board: barcode DPX101       part 030-1282-001 rev A
      HDTV_GENLOCK Board: barcode DPD453 part 030-1382-003 rev A
2 180 MHZ IP27 Processors
CPU: MIPS R10000 Processor Chip Revision: 2.6
FPU: MIPS R10010 Floating Point Chip Revision: 0.0
Main memory size: 256 Mbytes
Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 1 Mbyte
Integral SCSI controller 0: Version QL1040B, single ended
  Disk drive: unit 1 on SCSI controller 0
Integral SCSI controller 1: Version QL1040B, single ended
IOC3 serial port: tty1
IOC3 serial port: tty2
Integral Fast Ethernet: ef0, version 1, module 1, slot iol, pci 2
Origin BASEIO board, module 1 slot 1: Revision 3
XT-HDIO Video: controller 2, unit 0, version 0x0
IOC3 external interrupts: 1
```

Index

Numbers

- 0 bit in packing, 59
- 1280x720p@59.94Hz, 2, 19, 20
- 1920x1080i@59.94Hz, 2, 19, 20, 55
- 4:2:2
 - format, 6
 - sampling, 75
 - video, converting, 75
- 4:2:2:4
 - connector usage, 42
 - control for setting, 18
 - format, and Links A and B, 6, 50
 - sampling, 75
- 4:4:4
 - sampling, 74
 - video, converting, 75
- 4:4:4:4
 - connector usage, 43
 - control for setting, 18
 - format, and Links A and B, 6, 50
 - sampling, 74
- 8-bit or 10-bit precision, control, 17, 18

A

- alpha, 2, 6
- aspect ratio, 1
- asynchronous I/O, 81

B

- blanking and VL_COLORSPACE, 25
- buffer, 9
 - and examples, 77
 - and I/O, 78-81
 - and UST/MSC, 35-38
 - pool, 78

C

- cable, 5-6, 50, 55
- calibration, 2
- capture
 - control, 16, 17
 - type, 27-28
- color space, 22-25
 - and blanking, 25
 - and color model, 22
 - and lookup tables, 25
 - compressed-range, 22-25
 - control, 16, 17, 22-25
 - conversion, 23
 - example, 26
 - full-range, 22-25
 - values, 24
- compressed-range color space, 22-25
- control, 14-31
 - determining for device, 15
 - device-dependent, 10
 - device-global, 10
 - device-independent. *See* control, device-global.

HD I/O-specific, 10
order, 11
prefix, 10
setting, 15
values and uses, 17-19
conventions, xv

D

device, 9
determining, 15
Digital Media Programming Guide, xiii
digital video drain, setting up, 52-53
digital video source
setting up, 50-51
timing in panel, 51
direct I/O, 78-80
DMbuffer, 9
See also buffer.
drain node. *See* node, drain.
dual-link mode, 6, 50

E

events, 31-32

F

field dominance, 29-30
and outputting frames, 30
control, 16, 17
in *vcp*, 53
format control, 15, 16, 18
full-range color space, 22-25

G

GEN IN, 48
genlock, 2, 48
LEDs, 55
GEN OUT, 48
GPI, 44-48
pinouts, 44-45
receiver, 47-48
transmitter, 46-47
interfacing, 47
graphics to video, 31

H

HD I/O
board
connectors, 5, 49
diagram, 4
illustration, 3
ports, 5, 49
cable, 5-6
controls for, 9, 14-31
panel, 5, 49
path, 11
setting up for hardware, 49-56
software, installing, 83-84
headroom-range color space. *See* compressed-range
color space.
horizontal phase control, 16, 18

I

installing software, 83-84
interlaced, 20
interleaving, 27-28

I/O

- asynchronous, 81
- direct, 78-80

IRIX version required, xiii

ITU-R BT.601, 22, 24, 75

ITU-R BT.709-2, 6, 22, 24

K

kind, 10

L

Link A, 6, 50

- and sync source, 55
- transfer mode usage, 42, 43

Link B, 6, 50

- and sync source, 55
- transfer mode usage, 42, 43

linking, 7

lookup tables and VL_COLORSPACE, 25

loopback control, 17, 19

loopthrough for genlock input, 48

M

manual

- conventions, xv
- other required, xiv

media stream count. *See* MSC.

MSC, 32, 35-38

multiprocessing, 80-81

N

node, 8, 9, 10-11

- drain, 8, 10

- source, 8, 10

O

offset control, 16, 18, 28

OpenGL to read pixels into memory, 31

origin, different in OpenGL and video, 31

P

packing, 21, 57-75

- 0 bit, 59

- 20-bit, 63

- 24-bit, 64-65

- 32-bit, 65-73

- and sampling pattern, 57, 74-75

- control, 16, 18

- x bit, 59

panel (*vcp*), 51-56

- callup, 51

- Digital Video Drain, 52-53

- Digital Video Source, 50-51

- external sync source, 55

- restoring settings, 56

- saving settings, 56

path, 8, 9, 11

phase

- horizontal, control, 16, 18

- vertical, control, 16, 18

precision control, 17, 18

progressive, 20

R

Recommendation 601, 22, 24
Recommendation 709, 6, 22
reporting, 32
RGB, 24
RGB_F, 23
RGB_H, 23
RGBA, 24

S

sampling pattern, 74-75
 and packing, 57
single-link mode, 50
size control, 16, 18
SMPTE 240M, 6, 22, 24
 field dominance, 30
SMPTE 260M, 1
SMPTE 274M, 1, 6, 22
 field dominance, 29-30
 genlock, 2
SMPTE 293M, 1
SMPTE 295M, 1
SMPTE 296M, 6
software, installing, 83-84
source node. *See* node, source.
specifications, 39-42
sync
 connectors, 48
 control, 16, 18
 setting up, 54-55
 source
 and timing, 20
 control, 16, 18
 setting up, 55
synchronizing data streams, 35-38

T

timing
 and sync source, 20
 control, 16, 18
 free-run, 2
 setting, 19-20
 in *vcp*, 51
type, 10

U

unadjusted system time. *See* UST.
underflow correction, 30
 control, 17, 19, 30
UST, 32, 35-38

V

vcp, 49-56
 callup, 51
 See also panel.
vertical
 phase control, 16, 18
 rates supported, 1
video
 formats supported, 1-2
 pipes and GPI connector, 45
Video Library. *See* VL.
VL
 central concepts, 8
 data transfer functions summarized, 12
 header files, 7
 object classes, 9-10
 path, 8
 requirements for running, 7
VL_ANY, 11

VL_CAP_TYPE, 16, 17, 27-28
VL_COLORSPACE, 16, 17, 22-26
 values, 24
VL_FIELD_DOMINANCE, 16, 17, 29-30
VL_FORMAT, 16, 18, 20-21
 supported combinations with VL_COLORSPACE,
 23
VL_H_PHASE, 16, 18
VL_MEM, 10
VL_OFFSET, 16, 18, 28
VL_PACKING, 16, 18, 21
VL_PACKING_0444_8, 69
VL_PACKING_242_10, 63
VL_PACKING_242_8, 62
VL_PACKING_2424_10_10_10_2Z, 71
VL_PACKING_444_8, 58, 64
VL_PACKING_4444_10_10_10_2, 70
VL_PACKING_4444_8, 58, 66
VL_PACKING_R0444_8, 68
VL_PACKING_R242_10, 63
VL_PACKING_R242_10_in_16_L, 72, 73
VL_PACKING_R242_10_in_16_R, 72, 73
VL_PACKING_R242_8, 62
VL_PACKING_R2424_10_10_10_2Z, 71
VL_PACKING_R444_8, 65
VL_PACKING_R4444_8, 67
VL_SIZE, 16, 18, 28
VL_SYNC, 16, 18
VL_SYNC_SOURCE, 16, 18
VL_TIMING, 16, 18, 19
VL_V_PHASE, 16, 18
VL_VIDEO, 10
VL_XTHD_EE_MODE, 17, 18
VL_XTHD_INTERFACE_PRECISION, 17, 18
VL_XTHD_LOOPBACK, 17, 19
VL_XTHD_OUTPUT_REPEAT, 17, 19, 30
VL_ZOOM, 17, 19, 28, 31
vlGetFrontierMSC(), 38
vlGetNode(), 10
vlGetUSTMSCPair(), 38
vlGetUSTPerMSC(), 38
vlinfo, 15
vlOpenVideo(), 9, 10
vlSetControl(), 11, 15
VYUA, 24

X

x bit in packing, 59

Y

YCrCb, 23, 24
YUV, 23, 24

Z

zoom, 28
 factor control, 17, 19

Tell Us About This Manual

As a user of Silicon Graphics products, you can help us to better understand your needs and to improve the quality of our documentation.

Any information that you provide will be useful. Here is a list of suggested topics:

- General impression of the document
- Omission of material that you expected to find
- Technical errors
- Relevance of the material to the job you had to do
- Quality of the printing and binding

Please send the title and part number of the document with your comments. The part number for this document is 007-3968-002.

Thank you!

Three Ways to Reach Us

- To send your comments by **electronic mail**, use either of these addresses:
 - On the Internet: techpubs@sgi.com
 - For UUCP mail (through any backbone site): *[your_site]!sgi!techpubs*
- To **fax** your comments (or annotated copies of manual pages), use this fax number: 650-932-0801
- To send your comments by **traditional mail**, use this address:

Technical Publications
Silicon Graphics, Inc.
2011 North Shoreline Boulevard, M/S 535
Mountain View, California 94043-1389