Indy Video™ and Indigo^2 Video™ Technical Report

Section 1 Product Overview

Indy Video and Indigo^2 Video are video option cards for the Indy™ and Indigo^2™ desktop workstations respectively. They provide cost-effective, high-quality video input and output in both 525/60 and 625/50 video timings They also provide onboard realtime keying and special effects mixing as well as realtime scan conversion of graphics output to video.

1.1 Product Features

Indy Video and Indigo^2 Video have these features:

- Analog input in composite or S-VHS
- Analog output in composite or S-VHS
- Analog output genlocks to video input or to external black burst (house reference)
- Full-size video windows in 525 (640 x 486) and 625 (768 x 576) formats
- Alpha blending (keying and mixing) of video and graphics in real time
- Video-to-screen conversion with anti-aliasing:
  - Integer zoomed and decimated sizes (1/7 to 1:1 to 7:1)
  - Pan (on pixel boundaries)
  - Selectable de-interlace filtering
- Computer-generated graphics filtered to 525/60 and 625/50 video:
  - Anti-aliasing
  - Selectable flicker reduction
  - Output may be video-sized (1:1) or nearly full-screen scan converted to video size
- Keys generated from chroma/luma of the video signal or from pixels in graphics window
- X-Y pixel-wipes and fades
- Two 24-bit non-overlapping video windows, or one 24-bit and two 12-bit non-overlapping video windows on screen simultaneously
- Frame buffer for synchronizing video signals, storage of video frames, or transfer of still images to the workstation
- Crosspoint switch selects input for major board components
- Video is routed internally to the video option card in YUV 4:2:2 color space (converting to RGB color space is unnecessary)

1.2 Physical Characteristics

Indy Video and Indigo^2 Video are video expansion cards. Indigo^2 Video occupies the space allotted to one slot in the Indigo^2 workstation chassis. Indy Video occupies the space allotted to one GIO slot in the Indy chassis. The boards use Philips television components and two custom SGI gate-arrays.

Section 2 How to Use Indy Video and Indigo^2 Video

With the video expansion card, the Indy and Indigo^2 workstations provides basic video input/output and graphics overlays at a low cost. An application can input analog video, convert any graphic image to video, create video keys and special effects within the application, and output the results as analog video.

2.1 Analog Video

2.1.1 Input

The Indy Video and Indigo^2 Video video expansion cards support composite and Y/C analog input. The video is converted to 24-bit RGB for display on the workstation monitor. The video may be displayed in full size video windows or scaled down; it may be zoomed and panned, as well. The input may also be used for video timing within the Indy and Indigo^2 Video systems. With Indy and Indigo^2 XL Graphics, a video input can be used to frame lock the graphics subsystem.

Video display does not affect workstation performance, since the Indy Video and Indigo^2 Video cards perform all input, scaling and switching between graphics pixels and video pixels in real time.

2.1.2 Output

Indy Video and Indigo^2 Video can output composite and Y/C (S-VHS) with sync in 525/60 and 625/50 timings.

Almost all of the full graphics screen (1280x960) or an NTSC- or PAL-sized portion of it may be output in real time. The output is taken from the graphics subsystem frame buffer and scan converted with flicker reduction and filtering. The outputting and scan converting are performed by the video board, so they place no additional demands on the workstation CPU or graphics subsystem. This filtered scaling is set in the hardware to 2:1 for NTSC output and 5:3 for PAL output.

All output signals are low-pass-filtered and sin x/x-corrected. The board has a notch filter that an application may select as needed. It reduces the cross-chroma artifacts created when dithered graphics images are encoded to composite video.

Applications may use the board as a frame buffer to record single frames to a frame-accurate VTR.

2.2 SGI Digital Video

4:2:2 YUV SGI digital video (only on Indigo^2 Video) may be input and output at the digital expansion port available (see Section 3.4). Applications should use the port for I/O to a compression/decompression board, for reading and writing to digital video devices (such as storage media), and for other workstation I/O. The port has two channels; one is used only for input and one may be configured for input or output.

Silicon Graphics sells the Cosmo Compress JPEG compression board (see the Cosmo Compress Technical Report) which utilizes this port on the Indy Video and Indigo^2 Video boards.

2.3 Video Overlay and Blending

Video and graphics may be combined in simple mixes, such as fades, wipes and keys. An application can generate a chroma or luma key from either the video input or the input from the graphics subsystem and overlay the key over the other signal. Graphics and video can be set to any level of opacity and blended (or mixed) with the other signal.

On the Indy workstation, the standard video inputs can be used as the source of the graphics subsystem input to the keyer/mixer circuitry on the Indy Video option card. Using the standard video input from the Indy Video option card allows keying and mixing of two composite or Y/C video channels.

Section 3 Hardware Features

3.1 System Architecture

Figure 1 is a block diagram of Indy Video and Indigo^2 Video.

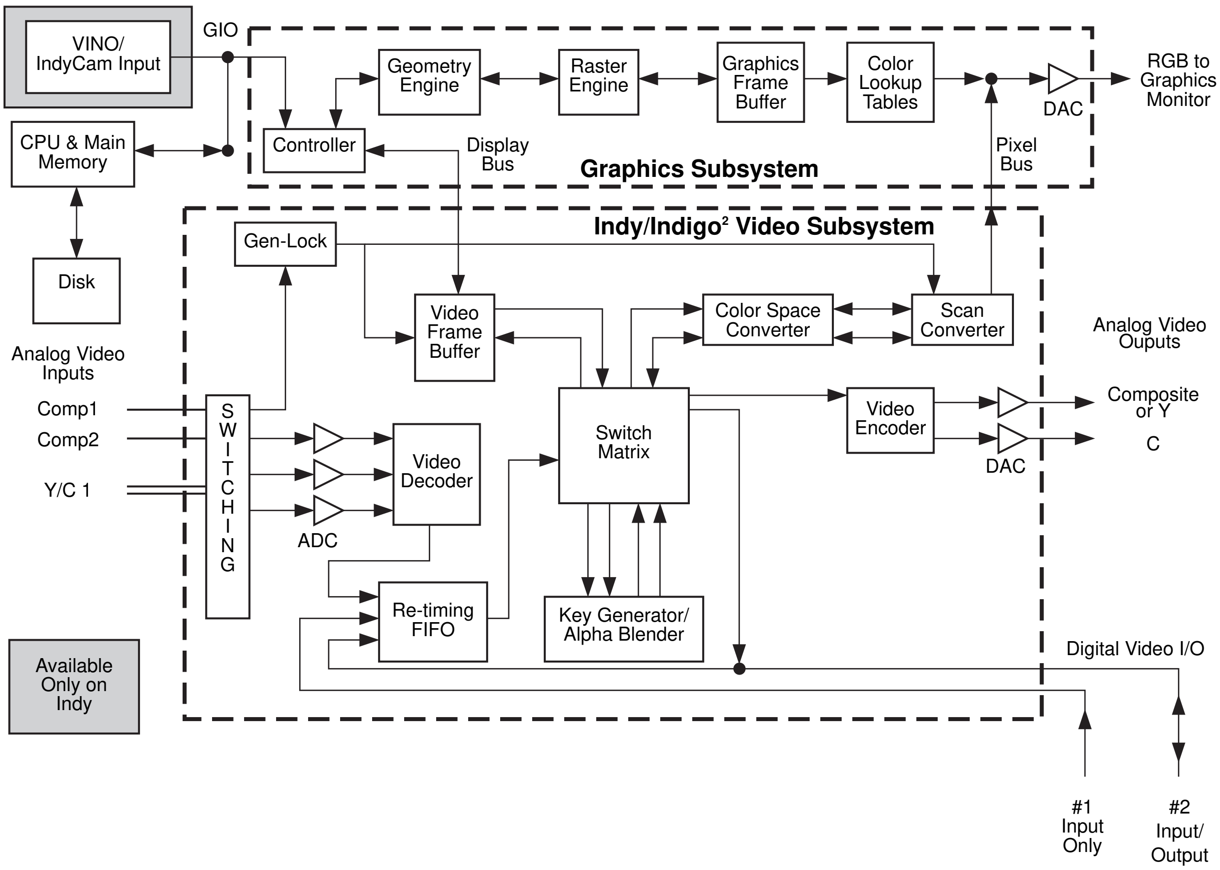

Figure 1:   Indy Video and Indigo2 Video Block Diagram

3.2 Analog Input

Indy Video and Indigo^2 Video have one analog input channel which is decoded by a Philips 7191 decoder. The video is decoded into YUV components at the sampling rates shown in Table 1 below.

YUV Component	525/60 Square Pixel	625/50 Square Pixel
Y	12.2727 MHz	14.75 MHz
U (B-Y)	6.13635 MHz	7.375 MHz
V (R-Y)	6.13635 MHz	7.375 MHz

Table 1: Indy Video and Indigo^2 Video Analog Sampling Rates

3.3 Analog Output

Indy Video and Indigo^2 Video have one analog output channel in addition to the SGI digital output which supports the Cosmo Compress option.

Composite and S-VHS analog output is converted from 4:2:2 YUV by a Philips 7199 encoder and reconstructed with a 5-MHz low-pass filter and sin x/x correction.

3.4 SGI Digital Video Expansion Port

The digital expansion port supports two channels of real time, SGI digital video data. One 60-pin connector supports both channels. One channel is for input only; the other channel is programmable for input or output.

Data in the port is sampled as square-pixel 4:2:2 YUV with 8-bits per channel:

- Square-pixel video adheres to the CCIR 601 and CCIR 656 standards with three exceptions: the number of clocks during blanking is nonstandard, the active pixel area is modified to accommodate the square-pixel sampling formats, and the voltage level of the signal is TTL rather than ECL. For true CCIR 601 video I/O, the Galileo with 601 option is required.
- For square-pixel sampling of the video, the clock is twice the pixel rate: 24.54545 MHz for NTSC and 29.5 MHz for PAL. For CCIR 601 sampling, there is a slight (~10%) scaling of the image in the horizontal dimension only when viewed on the graphics monitor. Throughput to the video outputs is not affected. The scaling is proportional to the ratio of the square-pixel clock rate and 13.5 MHz. Applications can use Graphics Library routines to correct for the scaling.

3.5 Scan Conversion

Indy Video and Indigo^2 Video maintain all video as 4:2:2 YUV. The only time this YUV data is converted to RGB is for display on the graphics screen or for the analog RGB output. All graphics data brought from the screen into Indy/Indigo^2 is directly converted from RGB to 4:2:2 YUV. See Section 3.8 Color Space Conversion.

3.5.1 Digitized Video to Graphics

After the analog video input has been digitized and/or processed, it may be converted to RGB and displayed on a workstation monitor.

Digitized video is converted to RGB in a dedicated frame buffer. Note this frame buffer is not shown on the block diagram. It resides in the scan converter block and is not shown because it is not accessible by applications or end users. The buffer may be displayed at full 24-bit resolution (8 bits per component) or split into two half-resolution, 12-bit planes (4 bits per component). All three windows (one 24-bit and two 12-bit) may be displayed at the same time, or two 24-bit windows may be displayed.

Although the video may be displayed on the workstation monitor at 12-bit resolution all 24 bits of data are available to Indy Video and Indigo^2 Video and the system. Note the following restrictions apply to all video windows:

- Video windows cannot overlap and a 24-bit window cannot share the same horizontal scan line on the graphics screen as another 12-bit window or 24-bit window.
- The graphics screen is 1280x1024 pixels, so only one full size PAL window can fit on the graphics screen at a time. Windows can be reduced in size to fit on the screen.

Video in graphics windows may be reduced, zoomed and panned:

- Each window may be reduced by six integer factors, from 1/2 to 1/7. To minimize aliasing, the data is filtered with a box filter before it is decimated. Note both 12-bit resolution windows must be decimated by the same factor.
- Each window may be zoomed by any integer factor from 2 to 7. Indy Video and Indigo^2 Video create the zooms by building a square array of replica pixels around each pixel in the source video.
- Each window may be panned to any location in the source image.

For de-interlacing, Indy Video and Indigo^2 Video use a selectable de-interlacing filter that uses the average value of the two immediately-adjacent lines. Another option for custom applications is to use the filter to generate black lines for the missing field at every 60th (or 50th) of a second. Inserting black lines reduces the intensity of the image by one-half.

3.5.2 Graphics to Digitized Video

The graphics subsystem output may be converted to video in 1x or full screen modes:

- In 1x mode, a video resolution window is converted to video (640x486 or 768x576).
- In full screen mode, an image twice the resolution of an NTSC image, or 5/3 the resolution of a PAL image is converted to video.
- Everything that appears on the graphics screen is converted, including cursors, pop-up menus, dialog boxes, and so on, but a graphics application can turn off these events as needed.

An application may convert the graphics pixel-by-pixel or line-by-line:

- Pixel-by-pixel conversion, in which each graphics pixel is converted to a video pixel, should be used for images up to the size of the video format. The data passes through a flicker-reduction filter during the conversion. Indy Video and Indigo^2 Video store the filtered data in the scan converter's dedicated frame buffer, from which the data may be read at a rate determined by the video timing.
- Line-by-line conversion should be used for images that are larger than the video format. Indy Video and Indigo^2 Video reduce aliasing by filtering each line of the graphics data horizontally, first, then vertically. The filtering also minimizes artifacts produced during decimation.

  Flicker reduction, which is selectable, is performed during the vertical filtering.
- Two separate areas of the graphics screen may be converted to video simultaneously if they do not share horizontal lines.

Indy Video and Indigo^2 Video use four field buffers for scan rate conversion of two channels of video. The phase relationship between graphics and video vertical intervals is monitored to determine when fields should be skipped or replaced. A digital phase detector determines when graphics frames must be dropped. Indy Video and Indigo^2 Video pass status information about the dropped frames to the workstation CPU, allowing applications to render motion in the video image more smoothly.

Framelocking the graphics frame rates to input can be performed which allows for smooth graphics output to video without repeat or skip frames.

3.6 Frame Synchronization

Indy Video and Indigo^2 Video use a dedicated buffer for frame synchronization; the buffer may also be used for storing still images. The buffers store video fields in the internal 4:2:2 YUV format, 8 bits per component. They can input and output synchronized video to and from the functional units and to the workstation CPU.

3.7 Access to the CPU

4:2:2 YUV data must be used for reads and writes between the Indy Video and Indigo^2 Video frame buffer and the CPU. The video from the Indy Video and Indigo^2 Video frame buffer may be full resolution frames or decimated by factors of 2 and 4. Vertical blanking information may be obtained without disturbing the active video. When writing data from the CPU to the Indy Video and Indigo^2 Video frame buffer data may be full resolution frames or 1/2 or 1/4 size (previously decimated) frames which can be zoomed to full resolution on Indy/Indigo^2.

3.8 Color Space Conversion

Data is color space-converted during graphics-to-video and video-to-graphics conversions, but not during video input and output or reads and writes to the field buffers. Because of the 4:2:2 format of the video, the U and V components are sub-sampled during interpolation and decimation (every other pixel is converted, instead of every pixel). The sub-sampling causes aliasing artifacts which Indy Video and Indigo^2 Video remove with anti-aliasing filtering.

3.9 Video Crosspoint

The crosspoint/switch matrix allows the following sources to connect to any of the following destinations:

Source	*Destination*
Analog video in	Analog Video out
Digital video in 1	Digital Video out
Digital video in 2	Host memory (frame buffer input)
Host memory (frame buffer output)	Alpha blender foreground
Alpha blender pixels	Alpha blender background
Alpha blender alpha	Graphics 1 (24-bit resolution)
Graphics A (24-bit resolution)	Graphics 2 (12-bit or 24-bit resolution)
Graphics B (24-bit resolution)	Graphics 3 (12-bit resolution)

3.10 Keys

Indy Video and Indigo^2 Video generate linear, 8-bit keys and passes the 8-bit values to the alpha blender for mixing (see Section 3.11). Changes and updates to the keys are double-buffered during vertical blanking.

Keys are created on a pixel-by-pixel basis. Indy Video and Indigo^2 Video apply mathematical formulas to the Y and U-V values of each pixel in a selected portion of an image, producing an alpha value for that portion. It blends this value with the corresponding alpha of the remainder of the image:

Pixels may be chosen in three ways:

- Whole image
- X-Y location in the image

- Random (pixels chosen by a random numbers generator)

Note: For more information please refer to the Video Library Programmer's Guide.

3.11 Alpha Blending

In a mix, the foreground and background each have an alpha value. The most common sources of the values are the key generator and the graphics display bus. The alpha blender uses the two values to composite the mix; it can generate the composite in 16 ways.

Dedicated registers containing the YUV values of a flat field background may be used for fades to and from any color.

Note: For more information please refer to the Video Library Programmer's Guide.

Section 4 Software Features

Software for Indy Video and Indigo^2 Video follows Silicon Graphics successful strategy of providing common functionalities in a library of device-dependent and device-independent routines addressed through an API. This software is included in the Digital Media Development Kit.

Libraries insure that upgrades are compatible with current releases, that applications developed for one product can be ported easily to other products, and that functions provided by one library interface reliably and consistently with the functions of other libraries.

The software consists of the following:

- Video Library™
- Video tools
- Video daemon
- Kernel (device driver)
- X window system server

4.1 Video Library

The library provides an API that is common to all Silicon Graphics video products. Through the API, video applications developed for one Silicon Graphics product can be designed to run on all other Silicon Graphics products. Applications designed for a simpler product can be made to run on more complex products with little or no modifications; applications designed for more complex products can return error codes when they run on simpler products and calls are made to an unavailable parameter.

The basic model for the library is the video stream. Library routines manage the video stream by defining a source device, a destination device and the path between the two. Parameters called controls modify the path; all devices should recognize a subset of common controls.

Applications written to the square-pixel Video Library API will be able to access the other Silicon Graphics libraries, such as Compression Library™ and Graphics Library™.

4.2 Video Tools

The tools, built from calls to Video Library, provide a range of basic functionalities, such as displaying video in a window, capturing video frames and outputting video. GUI control panels and a command-line interface set the device parameters.

To illustrate how to use the library, the source code for each tool and control panel is also provided.

The basic tools are outlined as follows:

vlinfo — Use the video info tool to display information about video devices available through the VL, such as the name of the server, number of devices on the server, and the types and ID numbers of nodes, sources, and drains on each device.

videoout — Use the video output tool to output video from a video-sized window or a nearly full-screen area of the graphics screen.

videoin — Use the video input window tool to view video-in-a-window.

vidtomem — Use this tool to capture a single frame (the current video input) or a specified number of frames, depending on the hardware limits for burst capture, and write the data to disk on hardware that supports the video-to-memory path. Capture size can also be specified. The data can be color space converted or left as raw data, which can be used by the *memtovid* tool.

memtovid — Use this tool to output single frames (images) to video out on hardware that supports the memory-to-video path.

videopanel (vcp)

Use this graphical user interface to set controls, such as hue or contrast, on devices. The panel resizes itself dynamically to reflect available video devices.

vintovout — Use this tool to send video input to video output. There are four optional parameters:

- *-n devicenum* — The number of the video device to use.
- *-v inputmodenum* — The number of the video input mode to use.

-o outputnodenum	The number of the video output node to use.
-I	Print the node and path numbers for use with the command line interface.

The tools *vlinfo*, *vidtomem*, and *memtovid* are command line tools. In addition to their man pages, these tools are explained in the *IRIS Utilities Guide*.

4.3 Video Daemon

All data passes through the video daemon except the digitized video. The daemon handles both device-dependent and device-independent tasks, such as multiple device management, client access to other devices, event dispatch, and system configuration and device state information.

4.4 Kernel

The daemon communicates with the Indy Video and Indigo^2 Video boards through a device driver located in the IRIX kernel. The device driver is responsible for translating requests from the daemon into requests that can be understood by the hardware.

4.5 X Window System Server

The window system server is responsible for window positioning, loading color map tables and clipping the graphics screen around the video.

Section 5 Specifications

Table 4 lists Indy Video and Indigo^2 Video specifications.

Board Function	Specification	Value
Compatibility	Television standard	NTSC PAL
Sampling	Square-pixel NTSC	12.27 MHz, 640 pixels/line
	Square-pixel PAL	14.75 MHz, 768 pixels/line
	4:2:2 YUV	8-bits/component, U and V sub-sampled by 2
Input connectors	Composite	2 RCA
	Y/C (S-VHS)	1 S-VHS (4-pin DIN)
Output connectors	Composite	1 RCA
	Y/C (S-VHS)	1 S-VHS
Composite Input	Color separation	Chrominance trap only
	Aperture filters	Selectable
	Coring	Selectable up to 3 LSB
	Hue control	180
	Chroma gain	Automatic
	Luminance gain	Automatic
	Differential phase	4
	Chrominance-luminance gain	4%
Composite output	Operating modes	Genlocked or stand-alone
	Horizontal phase	3 µsec advance to 1 µsec delay
	Horizontal blanking	Fixed at the following rates: • 11.4 µsec NTSC square pixel • 11.9 µsec PAL square pixel
	Vertical blanking	Adjustable through all vertical interval lines
	SC-H phase	180
	Differential phase	1
	Differential gain	1%
	Chrominance-luminance gain	2%
	Chrominance-luminance delay	20 nsec
	Frequency response	4.2 MHz - 2 dB
	Blanking level output voltage	0.25 V
	S/N ratio	48 dB

Board Function	Specification	Value
Y/C (S-VHS) input	Aperture filters Coring Hue control Chroma gain Luminance gain Differential phase Differential gain Chrominance-luminance gain Chrominance-luminance delay Frequency response	Selectable Selectable up to 3 LSB 180 Automatic Automatic 4 4% 4% 40 nsec 4.2 MHz -3 dB (Y only)
Output jitter		<20 nsec in genlock
Frame buffers	Number Bandwidth to CPU Readback to CPU decimation Data format	4 1.5 Mpixels/sec maximum 2x or 4x, with or without anti-alias filtering 4:2:2 YUV, 8 bits/component
Graphics-to-video conversion	Reduction ratios, with anti-alias filtering Flicker reduction Color space conversion accuracy	2:1 NTSC square-pixel; 5:3 PAL square-pixel Selectable 1 LSB

Table 2: Indy Video and Indigo2 Video Specifications