

# ***O2 Technical Report***

***October 1997  
version 1.0***

© 1997 Silicon Graphics, Inc. All rights reserved. Specifications subject to change without notice.

Silicon Graphics, OpenGL, Image Vision Library, IRIS, and the Silicon Graphics logo are registered trademarks, and O2, O2Cam, IRIX, IrisGL, XFS, and InPerson are trademarks, of Silicon Graphics, Inc.

MIPS is a registered trademark, and R5000 and R10000 are trademarks, of MIPS Technologies, Inc.

UNIX is a registered trademark in the U.S. and other countries, licensed exclusively through X/Open Company Limited.

Ethernet is a trademark of Xerox Corporation. Netscape and Netscape Navigator are trademarks of Netscape Communications Corporation. AppleTalk and QuickTime are trademarks of Apple Computer, Inc. Windows is a registered trademark of Microsoft Corporation. SoftWindows is a trademark of Insignia Solutions, used under license. All other trademarks mentioned herein are the property of their respective owners.

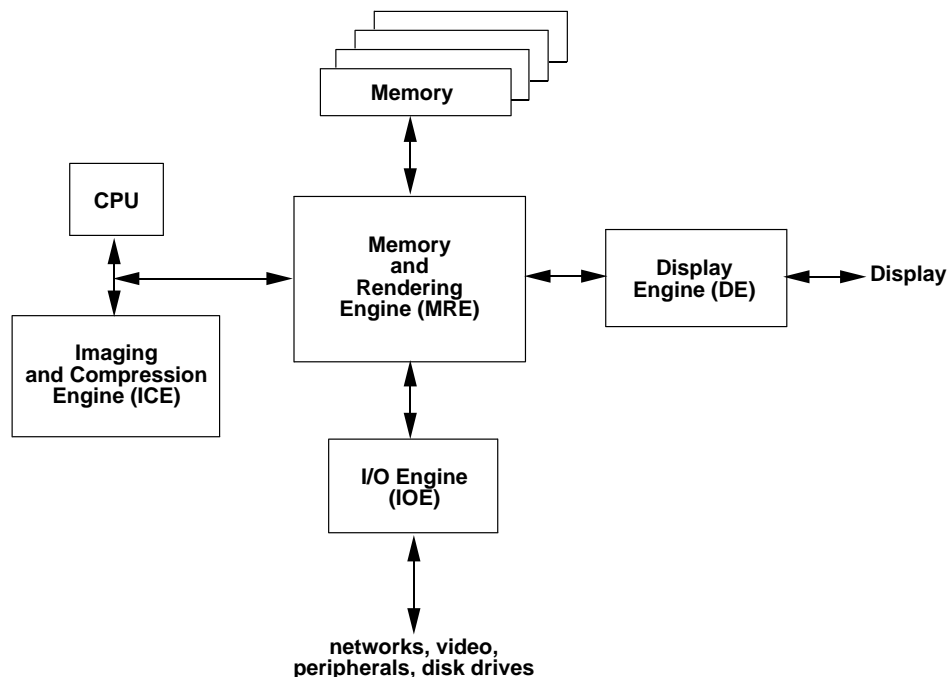
## Chapter 1      Overview

The O2™ system represents a new architectural design philosophy from that of earlier workstations or personal computers. Its Unified Memory Architecture (UMA) design enables integration of advanced technologies such as hardware texture mapping, real time video compression and image processing acceleration.

Traditional workstation architectures use multiple high-speed buses to connect numerous subsystems and their individual local buffers (e.g. frame buffer, texture, z-buffer, or image memories). The UMA design is one in which all buffers in a system are combined into one common memory pool. Each O2 engine, including the CPU, the Imaging and Compression Engine (ICE), the Memory and Rendering Engine (MRE), the Display Engine (DE) and the I/O Engine (IOE) has full access to not only its own data, but also to data from each of the other engines, without requiring a buffer copy.

Benefits of the O2 UMA include:

- Lower overall system cost
- More efficient use of the entire system memory
- More flexible operation because all data is accessible to the CPU
- Increased performance due to reduced buffer copies



**Figure 1-1** Each of the five specialized engines have direct access to all of memory in the O2 system's UMA.

UMA allows each of the processing engines to share all data seamlessly and efficiently. For example, uncompressed video data can be read and stored in main memory. This data can then be used as a texture by the 3D graphics engine merely by passing a pointer; no data copy is necessary. The ability to share data without buffer copies, especially copies across an external bus, removes a major bottleneck in digital media and 3D applications.

The O2 UMA architecture makes features such as the following possible:

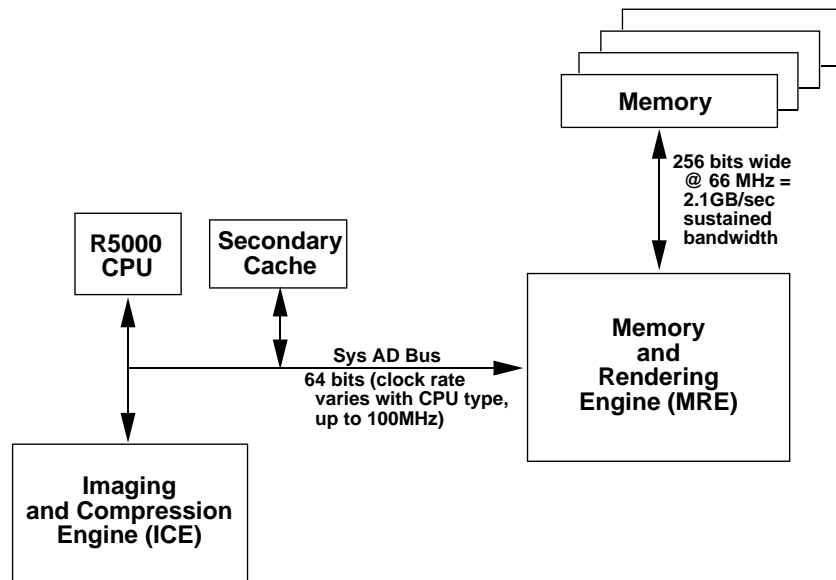
- Removing the need to copy data from one engine to another, thereby eliminating typical system bottlenecks such as texture upload time.
- Eliminating specialized local memory buffers from each engine, significantly reducing overall system cost.
- Efficiently using memory resources - unused memory is returned to the general CPU memory pool for use by other engines or applications.
- Replacing expensive multiple high-speed buses with a single ultra-high-speed connection to memory, significantly reducing total system cost.
- Elimination of the need for dedicated texture memory, instead allocating memory for textures from the general memory pool. Increasing the amount of memory available for texture involves simply adding standard synchronous DRAM.
- Allowing processors to specialize by type of computation, not by algorithm or application.

Applications that combine digital media and 3D data provide capabilities that were previously not possible at the O2 price point. The O2 system's real-time video texture mapping capabilities, for example, have many practical applications ranging from entertainment to military training. Video data of a military exercise can be imported in real time from an infrared sensor, after which it is processed to either reduce noise, or in some training situations, to increase or change the characteristics of the noise. This processed video image can then be used as a texture in a real-time 3D scene. Additionally, the entire training session can be captured in a JPEG compressed movie file, live as it happens, for later review by a collaborator.

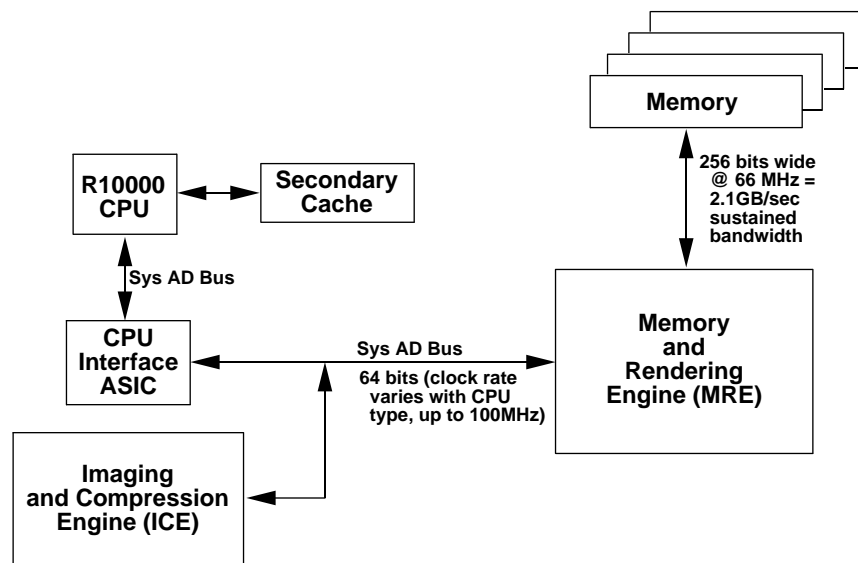
## 1.1 Compute Pathway

The O2 CPU (either a MIPS® R5000™ or R10000™) accesses memory over the memory bus, just like the other engines. The bandwidth adequately serves all of the engines, including the CPU. A 0.5MB or 1MB secondary cache further improves computing efficiency by reducing memory accesses.





**Figure 1-2 R5000 processor** Wide, fast data paths maximize imaging and graphics operations.



**Figure 1-3 R10000 processor** - Higher CPU performance further enhances graphics performance.

Offloading the CPU from some of the more traditional tasks is key to system performance improvements in O2. The other four engines execute the majority of 3D graphics, imaging, video, and compression tasks with minimal CPU involvement. Pixel fill, texture mapping, accumulation buffers, and a whole host of other OpenGL® and video API functions are now performed in other paths; the CPU calculates geometry transforms, clip-check, triangle setup and lighting.

## **1.2 3D Graphics Pathway**

The O2 graphics engine (MRE) accelerates all elements of the OpenGL pixel pipeline. Only the geometry operations mentioned above require CPU involvement.

Hardware-accelerated texture mapping is built into the O2 architecture. With specialized processors and pathways, the system achieves more than 30 megapixels per second trilinear-interpolated, texture-mapped fill rates and 40 megapixels per second using bilinear interpolation. (See Table 4-1 for further information.)

## **1.3 Imaging and Compression Pathway**

A specialized imaging pathway includes a dedicated processor and routines for the accelerated OpenGL Imaging extensions and ImageVision Library® procedure calls. The pathway can execute from 3x3 to 7x7 convolutions at rates of up to 10 megaconvolves per second. A range of color lookup table and window leveling functions are implemented in this pathway. Hardware accelerated roaming, rotating, and zooming are supported via bilinear texture mapping at rates of up to 43 megapixels per second.

Beyond these traditional image processing functions, the imaging path also benefits higher-fidelity 3D graphics. Texture mapping, accumulation buffer operations, subpixel positioning, and other 3D functions are actually image resampling problems for which the O2 imaging path enables higher quality images and better performance.

This same pathway is also used to accelerate compression and decompression of video data, images, and captured screen images. Combined with adjustable frame rates and adjustable resolutions, the compression flexibility maximizes data handling choices.

For example, NTSC- and PAL-sized JPEG compressed files can be read from disk, memory, or over the network and decompressed at 60 NTSC or 50 PAL fields per second at a 4.5:1 or greater compression ratio. They can then be either displayed to a screen or output via the video ports, or both, simultaneously.

In addition, any NTSC- or PAL-sized portion of the screen can be scan converted in real time, compressed, and written to memory or disk in JPEG format, or even sent over the network for decompression and display on another machine. Engineers and professionals can easily capture and share working sessions by taking advantage of this feature.

## 1.4 Video Pathway

In addition to audio functionality, the optional O2 Analog Video Option module provides composite video in and out, S-Video in and out, and a digital video connector. The digital video connector supports either the O2Cam™ or, via a third-party adapter, ITU-R BT.601 (SMPTE259M/ITU-R BT.656) (formerly referred to as CCIR601 (SMPTE 259M/CCIR626)) serial digital video in and out.

The optional O2 Digital Video Option module provides 2 identical ITU-R BT.601 (SMPTE259M/ITU-R BT.656) input channels, 2 identical ITU-R BT.601 (SMPTE259M/ITU-R BT.656) output channels, genlock input, genlock output and general purpose interrupt (GPI) I/O, in addition to audio functionality.

All of these video connections into and out of the unified memory are shared by all five of the engines, as is the data travelling over them.

With the external, dedicated Single Ended Ultra Wide SCSI-3 port connected to striped disks or a RAID array, the O2 video pathway can read or write uncompressed NTSC or PAL video in real time.

## 1.5 Display Pathway

A dedicated display engine manages the color map functions and provides screen capture capabilities. The graphics display engine drives the DAC dot clock at 140MHz for display resolutions up to 1280x1024 at 76Hz. Signals are also provided to drive an optional high-resolution flat panel display board at 5 bits per RGB component at 1280x1024 resolution. This option board also includes a stereo-view connector to support head-mounted display options at 1280x492 at 120Hz.

## 1.6 Key System Features

The O2 architecture, shown in Figure 1-4, can be further described by advances in the following areas:

### *Unified Memory*

- Synchronous DRAM memory system
- Memory-based graphics and digital media rendering

### *OpenGL Rendering Engine (MRE) with Texture Mapping*

- Hardware-implemented OpenGL pipeline from rasterization to frame buffer
- Memory copy engine

### *Image Processing (ICE)*

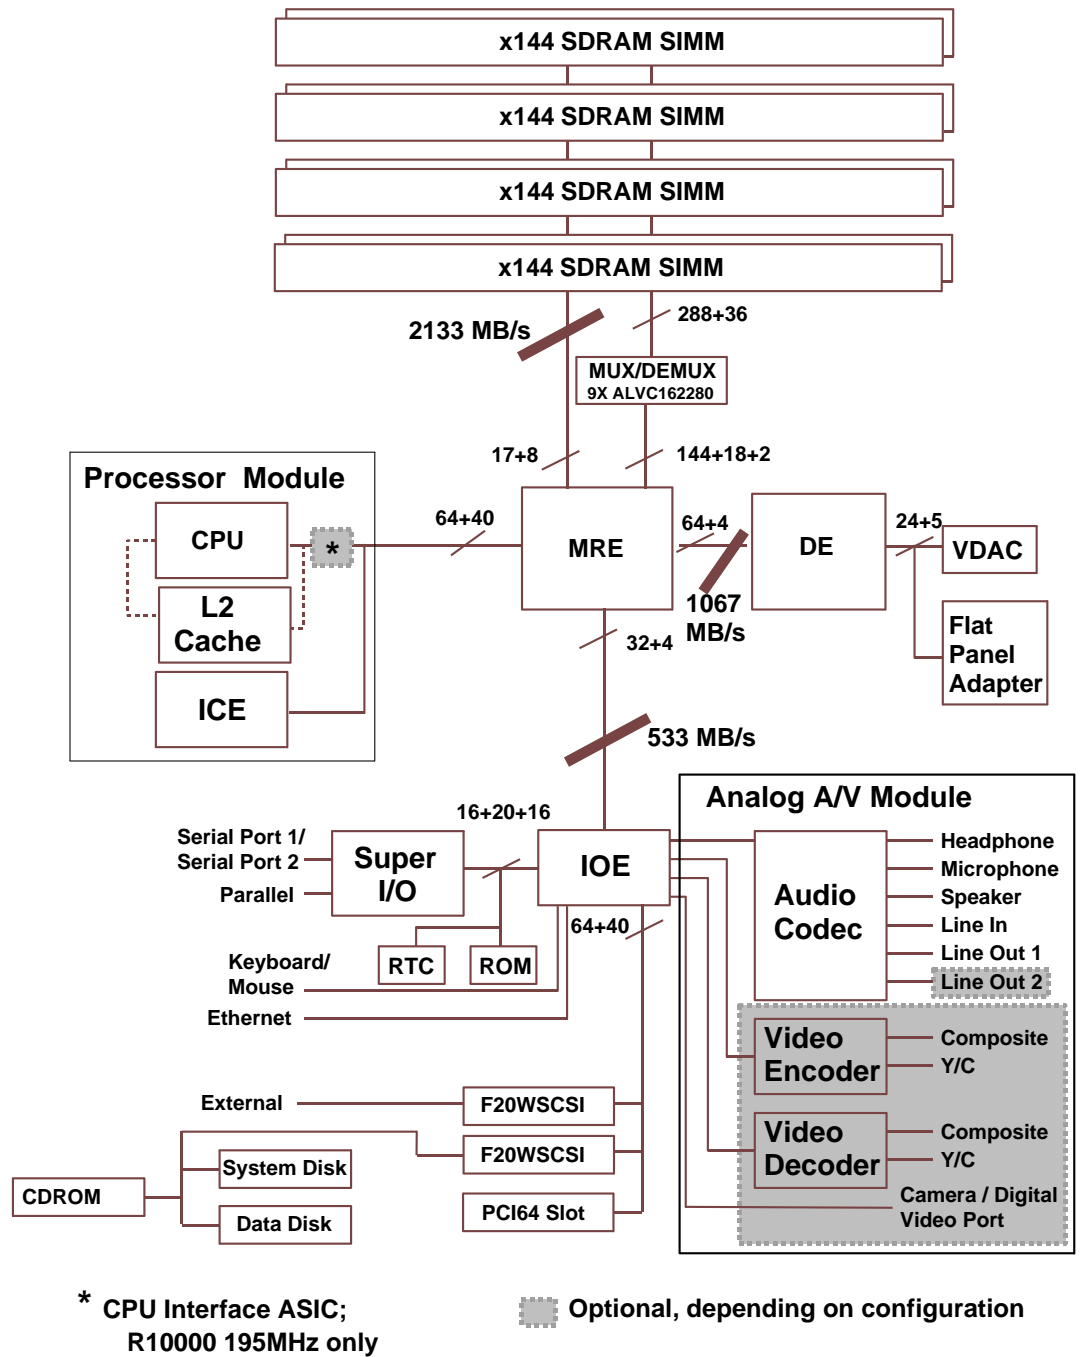
- Accelerated pixel transfer operations: scaling, biasing, component-to-component mapping
- Imaging extensions: convolution, minmax, histogram, color matrix, lookup tables, scale/bias
- Blur and warp
- Pan, rotate, and zoom in real time on arbitrarily large images (limited only by system memory)

### *Integrated Compression/Decompression (ICE)*

- 3 Gops per second Media Signal Processor
- 60MHz Bit Stream Processor
- Single stream, resolution-independent, motion JPEG encode/decode
- Real-time encode/decode (60/50 fields per second) at up to 768x576 resolution

### *Balanced I/O Bandwidth (IOE)*

- Autosensing 10/100baseTX Ethernet™
- Two SCSI buses, one internal, one external, rated at 40MB per second each
- 267MB per second (peak) via PCI64 bus (33MHz, 64 bits)
- Two dedicated video input DMA channels, one output channel



**Figure 1-4** O2 Architecture Block Diagram with optional Analog Video Option module; see Figure 5-5 for a block diagram of the Digital Video Option module.

## Chapter 2 CPU

The MIPS R5000 and R10000 superscalar processors offer the O2 configurations best-of-class processor performance. With high-speed floating-point execution units, the R5000 and R10000 processors provide balanced performance for a broad range of applications.

Primary features of the R5000 processor include:

- 180MHz and 200MHz configurations available
- MIPS IV instruction set
- 2-way superscalar 64-bit architecture †
- 32KB instruction cache
- 32KB primary cache
- Two-way set associative primary cache
- 512KB secondary cache (*180 SC processor only*)
- 1MB secondary cache (*200 SC processor only*)
- One-way direct mapped secondary cache (*SC processors only*)
- Combined multiply/add (MADD) instructions

Primary features of the R10000 processor include:

- 175MHz and 195MHz configurations available
- MIPS IV instruction set
- 4-way superscalar 64-bit architecture †
- 32KB instruction cache
- 32KB primary cache
- 1MB non-blocking secondary cache
- 128-bit dedicated secondary cache bus
- Two-way set associative primary and secondary cache
- Combined multiply/add instructions

*R5000 Microprocessor Product Information* can be found at:  
<http://www.sgi.com/MIPS/products/r5000>

The *R10000 Microprocessor Technical Brief* can be found at:  
<http://www.sgi.com/MIPS/products/r10k>

† IRIX™ 6.3 supports the O32 ABI and N32 high-performance ABI (with 64-bit addressing and arithmetic instructions.) It does not support N64.

### 2.1 Processor Highlights

The MIPS R5000 processor provides O2 systems with:

- High-performance concurrent floating-point and integer operations

- Streamlined 3D geometry calculations with single cycle multiply/add instruction for 3D performance
- Large, fast, on-chip caches (32KB instruction, 32KB data)
- Geometry setup, lighting, and transformation calculations

For applications that require the highest performance possible, the O2 R10000 system is recommended due to its:

- High performance for rendering and interference checking
- Advanced technologies for eliminating bottlenecks: out-of-order execution, branch prediction
- High-speed, non-blocking cache system reduces memory latency for maximized application performance

## 2.2 Processor Specifications

**Table 2-1. Processor Specifications**

	<i>180MHz R5000PC †</i>	<i>180MHz R5000SC</i>	<i>200MHz R5000SC</i>	<i>175MHz R10000SC</i>	<i>195MHz R10000SC</i>
<b>Description</b>	2-way superscalar	2-way superscalar	2-way superscalar	4-way superscalar, non- blocking cache; out-of-order execution	4-way superscalar, non- blocking cache; out-of-order execution
<b>CPU Speed</b>	180 MHz	180 MHz	200 MHz	175 MHz	195 MHz
<b>Primary cache: Instruction Data</b>	32KB 32KB	32KB 32KB	32KB 32KB	32KB 32KB	32KB 32KB
<b>Secondary cache</b>	none	0.5MB	1MB	1MB	1MB
<b>AIM VI</b>	208.3	242.3	232.7	427.2	*
<b>SPECint95 (peak)</b>	3.70	4.82	5.4	7.62	9.30
<b>SPECfp95 (peak)</b>	4.55	5.42	5.7	6.60	8.90

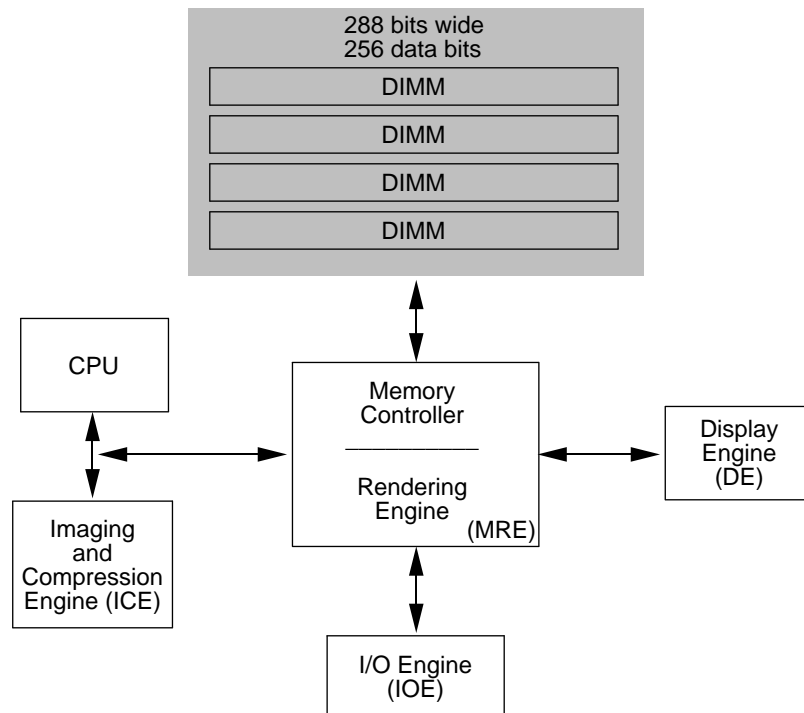
† Available to OEMs only

\* Not currently available

## Chapter 3      *Memory*

The O2 system's high-throughput memory system contributes to the flexible and efficient processing capabilities of all five engines and associated pathways. The CPU, Memory and Rendering Engine (MRE), Imaging and Compression Engine (ICE), Display Engine (DE) and I/O Engine (IOE) all have access to data in memory at an aggregate bandwidth total of 2.1GB per second, ensuring that no subsystem is ever starved for data.

### 3.1 Memory Pathway Structure



**Figure 3-1** All four banks of DIMMs are simultaneously activated by the controller.

Organized as four banks of dual-interleaved DIMMs, the O2 memory subsystem:

- Accesses 256 bits every cycle at 66MHz
- Stores each word in 2 interleaved DIMMs
- Uses SDRAMs with CAS latency of 2, page rate of 1
- Can store from 32MB to 256MB with 16Mbit SDRAMs installed in the banks; from 128MB to 1GB with 64Mbit SDRAMs (available Q4CY97 via a software patch to *IRIX 6.3 Including R10000*)



(Note: 16Mbit and 64Mbit SDRAMs can be installed in the same memory system, but not in the same bank.)

Data coming in and out of memory passes through the memory controller residing on the custom MRE. A critical component to graphics, imaging, and video performance, the memory system's central position in the system streamlines data sharing among the various system pathways. (See Chapter 4 for information about the MRE's graphics capabilities.)

## 3.2 Memory Controller

The memory controller in the MRE simultaneously activates all four banks of DIMMs. Hardwired ECC protection is applied to all system memory, except for frame buffers. ECC protection corrects all single-bit and detects all double-bit errors. To diagnose developing problems early, the operating system logs all single-bit errors during system operation. In addition, memory modules contain built-in self-test (BIST) circuitry that tests the entire memory array and reports errors whenever the system is reset. A fail-safe capability is provided by boot firmware that automatically disables any memory banks that fail BIST.

The O2 memory controller performs sophisticated look-ahead and critical arbitration functions to maximize memory use. By anticipating memory demands and intelligently enabling current requests, periods of inactivity are minimized.

A fixed priority arbitration scheme ensures that time critical functions always have unrestricted access to memory. The priority levels are (from highest to lowest):

- Display Engine (for uninterrupted display updates)
- Audio/video operations (for uninterrupted recording of real-time audio or video)
- Imaging and Compression Engine (for critical real-time video stream compression and decompression)
- Rendering Engine
- CPU

## 3.3 Memory Performance

Capable of 2.1GB per second peak aggregate throughput, the O2 memory system is very flexible in managing demands from all of the critical system pathways. The O2 architecture is designed to enable future headroom as system components such as the CPU increase in performance.

### 3.4 Memory Allocation and Sharing

UMA's free memory allocation scheme allows the O2 memory system to be optimally used by any of the system pathways. Applications that involve complex 3D graphics, large images, or that require a lot of off-screen rendering now have access to all available memory. The *scatter and gather* allocation of memory tiles further enhances performance since memory requests do not involve waits for contiguous blocks of memory.

Once data is stored in memory, it is accessible by any system pathway. Passing control or distributing workload involves only the passing of memory address information; thus actual data movement within the system is minimized by the free access to memory. This unified architecture for handling processor, graphics, and other digital media data translates to improved results in every performance category.

## **Chapter 4      Graphics**

Positioned as the central hub of the system architecture, the O2 Memory and Rendering Engine (MRE) has direct paths to memory and all other engines: CPU, I/O Engine (IOE), Display Engine (DE), and the Imaging and Compression Engine (ICE) (see Figure 3-1). An innovative design, the graphics engine is embedded within the MRE and takes the place of the previous generation raster engine and CPU-intensive implementations. With hardware acceleration of OpenGL and X calls, the O2 graphics path achieves an order of magnitude improvement for the key graphics performance parameters over the previous generation of 3D desktop graphics architectures (see Table 4-1, Graphics Performance Rates).

Even at the entry-level price point, the MRE delivers an unprecedented level of hardware acceleration for graphics operations including triangle and line rasterization, texture mapping, z-buffering, fog, and line antialiasing. The graphics engine can also share memory with and manipulate imaging data from the ICE.

The Unified Memory Architecture complements the MRE's processing capabilities and contributes to the overall throughput achieved by the graphics path. The free memory allocation scheme means that graphics operations are only restricted by the amount of installed memory; the rendering engine has no pre-assigned segments or restrictions, and can efficiently utilize as much or little of main memory as it requires. For example, texture maps, images and off-screen rendering areas are allocated as needed directly from available main memory.

With direct memory access capability, the connection between the graphics engine and the memory subsystem can achieve a peak transfer rate of 2.1GB per second. Shared memory capabilities accelerate overall system throughput since multiple engines can access graphics data for efficiently distributing workloads and executing tasks in parallel.

### **4.1 Rendering Engine Architecture**

MRE's rendering engine:

- Provides hardware acceleration for boosting the performance of most OpenGL and X operations
- Offloads pre-rasterization functions (lighting, transforms, texture vertex assignments, plane equation set-up) to the CPU
- Rasterizes points, lines, triangles, and rectangles atomically from their vertices
- Implements a deep, single-pixel-wide pipeline modeled on the OpenGL rendering pipeline (from the rasterize stage forward)
- Performs copies and clears at up to full memory bandwidth

- Allocates pixel buffers (color, depth, texture) in tiles, and provides address translation buffers for dynamic allocation of pixel buffers.

## 4.2 Rendering Engine Features

The graphics path incorporates most of the OpenGL rasterization pipeline in hardware. System and graphics throughput is dramatically improved compared with previous software-driven, CPU-intensive OpenGL implementations. Specific hardware accelerated OpenGL features are listed in Section 4-7.

The rendering engine features hardware support for the following:

### *Color Pixel Formats*

- 8-bit RGB 3:3:2
- 16-bit RGBA 5:5:5:1
- 32-bit RGBA 8:8:8:8
- 32-bit ABGR 8:8:8:8
- 8-bit color index
- 12-bit color index

### *Internal Frame Buffer Formats*

- 8-bit + 8-bit double buffer format
- 16-bit + 16-bit double buffer format
- 32-bit + 32-bit double buffer format

### *Hardware Texture Mapping Texel Formats*

- 8-bit luminance/intensity (no alpha)
- 16-bit ARGB 1:5:5:5
- 16-bit RGBA 4:4:4:4
- 16-bit luminance/intensity, 8-bit alpha
- 32-bit RGBA 8:8:8:8

### *Pixel Buffers available for Image Display, Storage and Manipulation*

- 2Kx2K 8-, 16-, 32-bit color buffers (on-screen, off-screen, overlay)
- 2Kx2K 32-bit stencil/z-buffer (sz buffer), 8-bit stencil buffer, 24-bit z-buffer
- 1Kx1K down to 1x1 16-, and 32-bit texture maps

### *Point, Line, Triangle, Rectangle Rasterization*

- X and OpenGL point, line, triangle, and rectangle primitives rasterized atomically
- Rasterizer set up from vertex window coordinates
- Bresenham sampling for line and triangle rasterization

- 64Kx64K address space for X coordinates
- 4Kx4K address space for OpenGL coordinates
- 6-bit subpixel positioning for OpenGL coordinates

#### *Window Clipping*

- Five screen masks
- Clip IDs
- OpenGL scissoring

#### *Stippling and Patterning*

- For X and OpenGL lines and spans only
- Opaque stippling for X

#### *Flat and Gouraud Shading*

- RGBA and color index interpolation
- CPU set-up of plane equation parameters for interpolation

#### *Hardware Z-Buffering*

- Z interpolation and OpenGL depth test comparison for 24-bit z-values
- CPU set-up of plane equation parameters for interpolation

#### *Stencil Buffering*

- OpenGL stencil functions and stencil ops for 8-bit stencil values

#### *Texture Mapping*

- 1D and 2D textures (texture-homogeneous coordinates interpolated from plane equation parameters set up by CPU)
- Per pixel floating point perspective divide
- Per pixel mipmap level of detail computation
- OpenGL mipmap and filter modes
- OpenGL texel application modes

#### *Fog*

- Fog values linearly interpolated and applied to RGB color
- Plane equation parameters for interpolation set up by CPU

#### *Antialiasing*

- Partial-pixel coverage generated for lines
- Antialiased lines rasterized as two-pixel-wide lines
- Antialiased line coverage generated from slope- and distance-based coverage tables in ROM
- Filtered antialiased line endpoints

- Coverage applied to alpha
- Host may provide coverage value to be used for any primitive

#### *Alpha Test*

- All OpenGL alpha test functions

#### *Alpha Blending*

- All OpenGL blending functions and ops

#### *Logic Ops*

- All X and OpenGL logical operations

#### *Dithering*

- Supported for 3:3:2 and 5:5:5 RGB formats

#### *Color Plane Masking*

#### *Pixel Transfers (DMA)*

- Pixels transferred from a linear or tiled buffer to the frame buffer
- Source pixels format-converted (including YCrCb to RGB conversion) and sent through rendering pipeline
- Simultaneous blending and copying

#### *Clears*

- Full memory bandwidth byte-aligned flat rectangle fills

### **4.3 Pixel Buffer Tiling**

The MRE speeds the processing of frame buffers by dividing them into smaller sections, called pixel buffers. Dynamically allocated in tiles of 128 rows of 512 bytes which require only 64KB of physically contiguous memory, pixel buffers can be scattered throughout available system memory, thereby minimizing system delays. Pixel buffers are typically used for color buffers (front and back buffer), stencil and depth buffers, offscreen buffers and texture buffers.

The pixel buffer layout is utilized by several APIs, including OpenGL, which accesses it via the *pbuffer* off-screen rendering data structure.

## 4.4 Frame Buffer Requirements

There is no dedicated frame buffer in an O2 system. All color bit planes, as well as overlay planes, stencil-buffer planes, texture and pixel buffer memory is allocated from available main memory.

For example, the maximum amount of memory used by a full-screen (1280x1024) application which uses 32-bit double-buffered pixels would be:

$$1024 \times [ 1280 \times ( 4 + 4 + 4 ) + 1536 \times 1 ] = 17.3 \text{ MB}$$

where:

1024 is the number of scan lines

1280 is the number of pixels per scan line

4 is the number of bytes in the front buffer ( 32 bits )

4 is the number of bytes in the back buffer ( 32 bits )

4 is the number of bytes in the combined stencil/z-buffer  
( 8 bits for stencil buffer plus 24 bits for z-buffer )

1536 is the number of scan line pixels for overlay plane calculations †

1 is the number of bytes in the overlay plane ( 8 bits )

Other than the overlay planes, which are allocated statically, memory for these buffers is allocated on an as-needed basis. Obviously the front buffer is always needed; memory for the back buffer and stencil/z-buffer is allocated as needed for existing windows. Therefore the calculations above are worst case, assuming an application covering the full screen is running.

As can be seen, the majority of main memory remains available for use by the applications and operating system as well as texture and offscreen rendering areas.

† Buffers must always be a multiple of 512 bytes wide. In the case of the value for front+back+stencil/z-buffers, 1280\*16 is a multiple of 512. However, in the case of the overlay buffer, 1 byte \* 1280 pixels/scan line is not a multiple of 512. Therefore, the next higher multiple of 512, 1536, is used in calculating memory requirements for the overlay plane.

Note that the overlay plane is always 8 bits, no matter the bit depth of the front, back, and stencil/z-buffers.

## 4.5 Display and I/O Engine Interfaces

High-speed interfaces to and from other O2 pathways allow the graphics engine to process and manipulate data from virtually anywhere within the system, offering an unprecedented range of options to the user. For example, the O2 graphics engine can texture-map a polygon with image data from sources including the display, algorithmically-generated data from the CPU, results of an OpenGL imaging operation, or a live video feed.

#### 4.5.1 Display Engine

The O2 graphics-to-display interface features include:

- Bus protocol transfer speeds of 66MHz and 133MHz (128 and 64 bits, respectively)
- Buffering and transfer of display data from memory to the display engine in 16x32-byte bursts
- Buffering and transfer of video capture data from the display to memory in 16x32-byte bursts

#### 4.5.2 I/O Engine

The graphics-to-I/O interface features include:

- Bus protocol transfer speeds of 66 MHz and 133 MHz (64 and 32 bits, respectively)
- Buffering and transfer of data from memory to I/O in 8x32-byte bursts
- Buffering and transfer of data from the I/O Engine to memory in 8x32-byte bursts

### 4.6 Rendering Performance Summary

The Unified Memory Architecture and the tight coupling of the Memory and Rendering Engine with the CPU result in the graphics rates listed in Table 4-1.

It is especially important to note that the O2 graphics architecture has the inherent potential to take advantage of future CPU and display technologies.

**Table 4-1 Graphics Performance Rates**

	180 MHz R5000PC †	180 MHz R5000SC	200 MHz R5000SC	175 MHz R10000SC	195 MHz R10000SC
<b>Triangles (50-pixel) flat shade</b>	854K tri/s	859K tri/s	959K tri/s	932K tri/s	1.04M tri/s
<b>Gouraud shaded</b>	560K tri/s	567K tri/s	637K tri/s	649K tri/s	734K tri/s
<b>Gouraud, Lit, Z-buffered</b>	421K tri/s	434K tri/s	485K tri/s	474K tri/s	531K tri/s
<b>Gouraud, Lit, Z-buffered, Nearest-Textured</b>	247K tri/s	250K tri/s	279K tri/s	288K tri/s	326K tri/s
<b>Gouraud, Lit, Z-buffered, Trilinear Textured</b>	248K tri/s	250K tri/s	279K tri/s	288K tri/s	326K tri/s
<b>X lines</b>	1.33M/s	2.2M/s	2.5M/s	2.7M/s	2.8M/s
<b>3D lines Gouraud Z</b>	693K/s	716K/s	799K/s	776K/s	894K/s



<b>Gouraud Fill Rate</b>	66MP/s	66MP/s	66MP/s	66MP/s	66MP/s
<b>Nearest Textured Fill Rate</b>	43MP/s	44MP/s	45MP/s	43MP/s	45MP/s
<b>Bilinear Texture Fill Rate</b>	42MP/s	44MP/s	45MP/s	43MP/s	45MP/s
<b>Trilinear Texture Fill Rate</b>	32MP/s	33MP/s	33MP/s	32MP/s	33MP/s
<b>DMA Pixel Transfers 16-bit</b>	63MP/s	65MP/s	65MP/s	64MP/s	64MP/s
<b>DMA Pixel Transfers 32-bit</b>	48MP/s	65MP/s	66MP/s	50MP/s	55MP/s

† available to OEMs only

## 4.7 OpenGL Extensions

The following sections detail the OpenGL extensions introduced to take advantage of the O2 MRE architecture advances.

### 4.7.1 Pixel Extensions

<i>EXT abgr</i>	For better performance on data generated from IrisGL™-native machines, this extension adds support for pixel data in the order A, B, G, R.
<i>SGIX interlace</i>	Modifies the behavior of most pixel-transferring commands: the source image is a field of an “interlaced” frame; a full frame can be assembled from operations on two independent fields.
<i>SGIX_ycrCb</i>	Allows display of video data in the YCrCb color space directly in graphics, with the graphics pipeline performing color space conversion from the YCrCb to RGBA.
<i>EXT packed_pixels</i>	Support for packed pixels in host memory.

### 4.7.2 Blending Extensions

<i>EXT blend color</i>	Allows a constant to be used as a factor in the blending equation (for blending two RGB images).
<i>EXT blend logic op</i>	Allows a boolean operation to be used instead of the normal sum-of-products in the RGBA blending equation (for reversible drawing).
<i>EXT blend minmax</i>	Modifies the RGBA blending equation to produce the minimum or maximum (rather than the sum) of the source and destination colors.

*EXT blend subtract* Changes the RGBA blending equation to produce the difference (rather than the sum) of the source and destination colors (for comparing images).

### 4.7.3 Imaging Extensions

*SGL color matrix* Adds a 4x4 matrix stack and matrix multiplication to the pixel transfer path (for reordering or duplicating color components, or for color-space conversions).

*SGL\_color\_table* Defines a new, more flexible color lookup table mechanism.

*EXT convolution* Adds 1D and 2D convolution (3x3, 5x5, and 7x7 sizes) to the pixel transfer process (for blurring, sharpening, edge-detection, etc.)

*EXT histogram* Defines pixel transfer operations that count occurrences of specific color component values (histogram) and track the minimum and maximum color component values (minmax).

*SGL pbuffer* Defines GLX pixel buffers for off-screen rendering.

*SGIX fbconfig* Introduces a new way to describe the capabilities of a GLX drawable; removes the “similarity” requirement when making a context current to a drawable; supports RGBA rendering to one- and two-component Windows and GLX Pixmap.

*SGL make current read* Allows OpenGL pixel operations to read pixel data from the buffers of one drawable and draw into the buffers of another.

*EXT import context* Allows multiple X clients to share an indirect OpenGL rendering context; adds convenience routines to get the display for the current context and retrieve the attributes used to create a context.

*EXT visual info* Support for additional X11 Visual types; provides a means to query the X Visual type underlying a GLX Visual and make requests or queries about transparent pixel values.

*EXT visual rating* Particular OpenGL-capable Visuals can be marked as “slow” or “non-conforming.”

#### 4.7.4 General Rendering Extensions

<i>EXT polygon offset</i>	Allows depth values of fragments to be displaced so that lines (or points) and polygons in the same plane can be rendered without interaction.
<i>EXT vertex array</i>	Adds the ability to specify multiple geometric primitives with very few subroutine calls.

#### 4.7.5 Texture Extensions

<i>EXT texture</i>	Defines a variety of “internal formats” for textures (for more control over color resolution or memory requirements); provides a new mechanism for determining supported combinations of texture size and color resolution; defines a new texture environment function.
<i>EXT texture3D</i>	3D texturing (unoptimized software path only for this extension).
<i>SGI texture color table</i>	Adds a color lookup table to the texture mapping process (unoptimized software path only for this extension).
<i>SGIS texture border clamp</i>	Defines a new texture clamping method that ensures that all edge sample values fall completely within the border; the texture image is never referenced.
<i>SGIS texture edge clamp</i>	Defines a new texture clamping method that ensures that all edge sample values fall completely within the texture image; the border is never referenced.
<i>SGIX texture scale bias</i>	Adds scale, bias, and clamp operations to the texture pipeline (unoptimized software path only for this extension).

#### 4.7.6 Texture Buffer Control Extensions

<i>EXT copy texture</i>	Provides the ability to copy pixels directly from the frame buffer (or <i>pbuffers</i> ) to texture memory via a pointer copy.
<i>EXT subtexture</i>	Allows a contiguous portion of an already-existing texture image to be redefined without affecting the remaining portion of the image or any of the other state that describes the texture.
<i>EXT texture object</i>	Supports named texture objects whose contents and parameters may be changed after they are defined.

#### 4.7.7

#### Video and Swap Control Extensions

<i>SGIX dm_pbuffer</i>	Includes a new type of GLXDrawable as a uniform means for OpenGL to access buffers generated by the Video Library (VL), compression and other digital media libraries.
<i>SGI swap control</i>	Allows the user to specify a minimum number of monitor retrace periods to wait between buffer swaps.
<i>SGI video sync</i>	Provides a way to wait until the monitor's vertical retrace occurs.

## **Chapter 5      *Audio and Video***

One of the major design goals of the O2 system was the complete integration of audio, video and compression into the architecture. Examples of the innovative design include native processing of video in non-square pixel format, two full rate video DMA channels direct to memory, real time mipmap texture generation from a video input stream and graphics engine support of YCrCb (4:2:2) pixel type. The Imaging and Compression Engine (ICE) is key in providing much of this functionality. (Refer to Chapters 6 and 7 on Compression and Imaging for further information.)

The O2 architecture treats a video or audio stream in system memory as generic data that can be used by any function in the system. Video manipulations include providing the raster engine access to video data for previewing, using it as a texture map to do accelerated Digital Video Effects (DVE), and keying video with synthetic images with an alpha value. The texture mapping hardware can also expand, shrink, and warp incoming video. Image processing tools or applications can process the video, aided by the hardware acceleration provided by ICE.

Conversely, any graphics, including frame buffer and offscreen rendering areas, can be directed out of the machine as a real-time video stream. The Display Engine (DE) can redirect the images into a DMA channel for subsequent redirection through the video out path in the I/O Engine (IOE). Similarly, any sequence of video-sized images, a decompressed stream off disk or simply a single image from memory, can be transferred via DMA through the IOE.

These advanced capabilities make the O2 platform attractive for:

- Compressed or uncompressed video and audio capture and playback
- Real-time graphics blending and keying
- Multi-layer image and video compositing
- Non-linear video editing
- 3D digital video effects

### **5.1 O2 Audio and Video Options**

All O2 configurations include compression capabilities via ICE. (For a description of the integrated video compression capabilities, see Chapter 6.)

All O2 configurations also provide analog audio capabilities, included in the Audio and Video Option modules. An optional Digital Audio PCI board is also available and can be used in conjunction with any of the Audio and Video Option modules.

Entry level O2 systems come standard with the Analog Audio module. If video I/O capabilities are desired, there are two options available, the Analog Video Option module or the Digital Video Option module.

## **5.2 O2 Audio Functionality**

### **5.2.1 Audio Input and Output**

The audio subsystem on the IOE is contained almost entirely in the highly integrated Analog Devices AD1843 audio codec. Its features include:

- A sigma-delta codec that both digitizes and reconstructs 16-bit stereo pairs
- A 20Kohm impedance microphone input with a programmable +20dB gain block
- Interpolators in the output path, allowing the analog-to-digital (ADC) and the digital-to-analog (DAC) converters to run at different rates
- Sampling rates of 8, 16, 32, 44.1, and 48KHz. Audio streams can be locked to video streams (the audio clock can be locked to a video input clock ensuring that they will view drift relative to each other.)


One stereo input channel and two stereo output channels operate independently using the codec-supplied time base. The interface supports per-sample counting with +/- 4 audio frame accurate time stamping (UST/MS) via the Unadjusted System Time (UST) clock and Media Stream Counter (MSC) for extremely accurate Audio/Video synchronization

### **5.2.2 Analog Audio Module**

All Audio and Video Option modules for O2 systems support identical audio I/O capabilities, with the exception of the Digital Video Option module, which does not include the rear speaker output jack.

The Analog Audio module includes the following features:

- Analog stereo line level audio input (via side RCA jacks)
- Analog stereo line level audio output (via side RCA jacks and a rear 3.5mm stereo phono jack)
- Stereo headphone output (via side 3.5mm phono jack)
- Mono single-ended microphone input (via side 3.5mm phono jack)
- Mono microphone input from O2Cam connector
- Stereo internal system speaker output on the rear of the module



**Figure 5-1** O2 Analog Audio module inputs and outputs

### 5.2.3

#### Digital Audio Board

The Silicon Graphics Digital Audio Board is a half-length PCI board designed to provide users with professional level capabilities.

Features of the Digital Audio board include:

- 8-channel, 24-bit ADAT optical input and output (2 optical connectors can be used for ADAT I/O or SPDIF I/O)
- Stereo AES3 24-bit digital input and output (2x75 ohm BNC connectors)
- Video house sync input (1 BNC connector)
- All inputs and outputs are independent and can be used simultaneously either with independent or synchronized sample rates
- Output sample rates can be slaved to ADAT Optical or AES input clock
- Output sample rates can be slaved to video. Note that this is the video coming into the house sync input, not the native O2 video.

- Nearly arbitrary sample rates from 4kHz to 48kHz
- Jitter attenuation of digital input clocks
- Sample accurate timing information for precise synchronization with other digital media subsystems. Can be slaved to ADAT Optical or AES input clock
- Output sample rates are compliant with:
  - SGI Audio Library (AL) API
  - UST/MSD digital media synchronization support

*Supported Sample Rates include:*

- 32kHz
- 44.056kHz pull-down
- 44.1kHz including VCXO jitter attenuation
- 44.144kHz pull-up
- 47.952kHz pull-down
- 48kHz including VCXO jitter attenuation
- 48.048kHz pull-up

*Master timebases include:*

- internal crystal oscillators
- Video reference (internal or external)
- AES11 input
- ADAT Optical input

*Video Reference Loop Through*

- 75 ohm BNC connectors
- Black burst NTSC or PAL

*Coaxial Digital Audio Input*

- AES-3id 75ohm BNC connector
- AES11-1992 synchronization input (for audio clock rates)
- AES3-1997 professional 2-channel 24 bit digital
- compatible with IEC958, S/PDIF consumer 2-channel digital

*Coaxial Digital Audio Output*

- AES-3id 75ohm BNC connector
- AES11-1992 synchronization output
- AES3-1997 professional 2-channel 24 bit digital
- compatible with IEC958, S/PDIF consumer 2-channel digital

*ADAT Optical Digital Input / Output*

- 12.8Mbps EIAJ RCZ-6901 fiber optic connector, PCI I/O panel

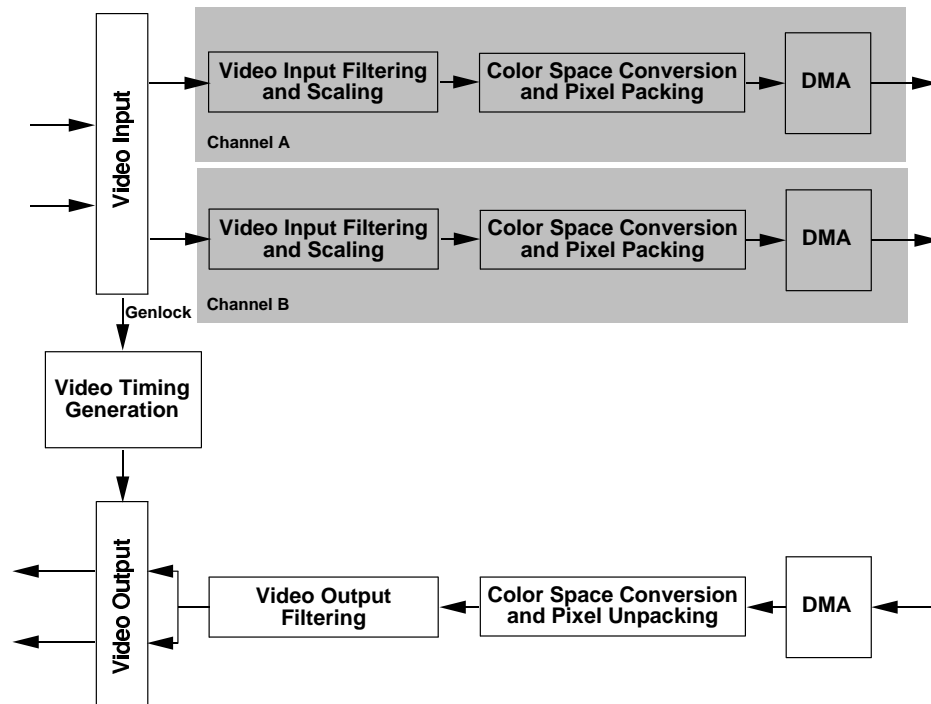


- 8-channel, 24 bit
- Professional interfaces available from third parties include:
  - Analog to Digital Converters
  - AES3-1997 inputs
  - ADAT 8-channel 16-bit tape machine, Roland DAWs, many more...
- compatible with IEC958, SPDIF consumer 2-channel digital

## 5.3 O2 Video Functionality

### 5.3.1 Video Input and Output

Shown in Figure 5-2, the video portion of the IOE equips the O2 workstation with video input and output capabilities equivalent to those found in much more expensive professional video systems. Fast DMA and filtering and scaling features allow real-time capture and display of video data.



**Figure 5-2** The I/O Engine's Video Paths. Integrated video processing functionality delivers performance typical of much more expensive video solutions.

The IOE's video circuitry features include:

#### *Video Input Channels*

- Two simultaneous video input sources: one analog input from composite/S-Video, and one digital input port from the O2Cam or two 4:2:2 serial ITU-R BT.601 (SMPTE259M/ITU-R BT.656) channels
- Clipping
- Can reach into vertical blanking interval to get VITC, closed caption, etc.

#### *Video Input Filtering and Scaling*

- Conversion from non-square pixel format to square PAL or NTSC
- Down-scaling of image to any size
- Mipmap generation
- Color space conversion to RGB with optional dithering

#### *Video Output Channels*

- One video out signal simultaneously sent to both digital and analog outputs
- Separation of RGBA pixels in to two separate 601 output streams, one supporting 4:2:2 YUVA, the other supporting 4:0:0 alpha. The streams are exposed directly on the O2 Digital Video Option module; they can also be accessed on the Analog Video Option module via the O2Cam connector. Software can select video and alpha to appear on the two outputs as well as any combination of the two.

#### *Video Output Filtering*

- Color space conversion
- Conversion from square PAL or NTSC pixel formats to non-square
- Notch filter

#### *Video DMA*

- Linear or tiled (rendering engine-compatible) buffer formats
- Field or interleaved frame modes

#### *Pixel Formats for DMA*

- 32- or 16-bit RGBA on output
- 32-bit YUVA
- 4:2:2 YUV (8- or 10-bit resolution)

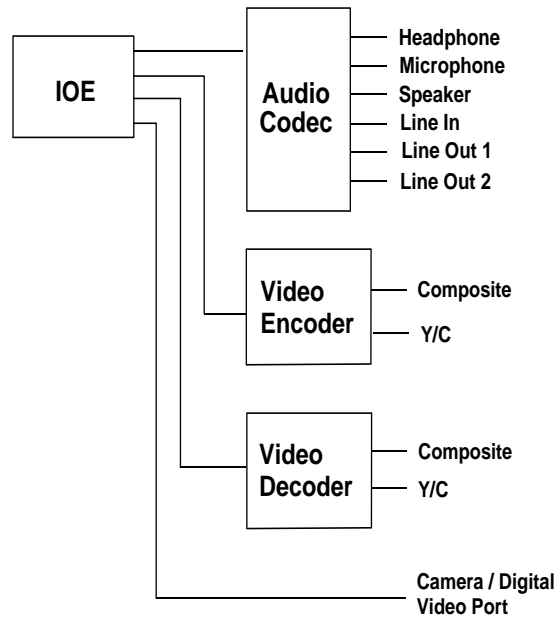
#### *Other Features*

- Timestamp with field/frame counter
- Genlock output to video input port

### 5.3.2

#### Analog Video Option Module


The O2 Analog Video Option module is an 8x4-inch board that performs all of the analog video and audio I/O, and incorporates a video decoder, a video encoder, and an audio codec. The same analog audio circuitry provided on the Analog Audio module is also included with this module.



**Figure 5-3** The O2 Analog Video Option module digitizes all inputs, including multiple video streams, from external devices.

Analog Video Option module features include:

- Composite video in and out (via side RCA jack)
- S-Video in and out (via side mini-DIN jack)
- Optional high-quality digital color O2Cam camera input port, digital video input and output (via rear 68-pin CHAMP connector). Digital video input and output, including SMPTE 259M/ITU-R BT.656, is available via third-party adapters or the Digital Video Option module
- Analog stereo line level audio input (via side RCA jacks)
- Analog stereo line level audio output (via side RCA jacks and a rear 3.5mm stereo phono jack)
- Stereo headphone output (via side 3.5mm phono jack)
- Mono single-ended microphone input (via side 3.5mm phono jack)
- Mono microphone input from O2Cam connector
- Stereo internal system speaker output on the rear of the module



**Figure 5-4** O2 Analog Video Option module inputs and outputs

Connections are made through the side and rear panel of the module, which also includes an extraction mechanism for easy removal.

### 5.3.3


#### Digital Video Option Module

The O2 Digital Video Option module supports 601-sized 8- or 10-bit serial component digital video. This option board supports two identical ITU-R BT.601 (SMTPE259M/ITU-R BT.656) input channels, each capable of passing a full bandwidth 4:2:2, 8- or 10-bit digital video stream to the IOE. It is also possible to receive a single 4:2:2 video stream with a corresponding 4:0:0 alpha stream, however, the alpha image must be DMA'd to a separate buffer in memory.

The Digital Video Option module also supports two identical ITU-R BT.601 (SMTPE259M/ITU-R BT.656) output channels, one capable of passing a full bandwidth 4:2:2, 8- or 10-bit digital video stream from the IOE. The other channel is capable of passing a 4:0:0 alpha stream out; the two outputs can also be configured to produce identical 4:2:2 video streams. These signals result from a single DMA from main memory in which the RGBA pixel is converted in the IOE and provided as YCrCbA to the option board for serialization. This module does not support Link A/B or RGB I/O.

The timing for the two 601 output channels can be either free-running at 270MHz or locked to input one, input two, or the house sync input of the board. Jitter between input horizontal and the recovered 270Mhz clock will be better than 0.74ns. In addition, ancillary data found in the vertical and horizontal blanking interval will be passed through the board undisturbed. Also supported are both input and output GPI triggering, which can be used as general purpose interrupts.

This option board supports digital video I/O only, there are no analog composite, S-Video, or O2 digital camera connections. However, the same analog audio circuitry that is standard on the Analog Audio module, with the exception of the 3.5 mm stereo phono jack on the rear panel, is also included with this module. Digital audio support is through the PCI Digital Audio Option board. (see section 5.2.3 for further information on this board)



**Figure 5-5** O2 Digital Video Option module inputs and outputs


Summary of O2 Digital Video Option module features:

- 2 ITU-R BT.601 (SMPTE 259M/ITU-R BT.656) 8- or 10-bit, 4:2:2 input channels (standard BNC connector)
- 2 ITU-R BT.601 (SMPTE 259M/ITU-R BT.656) 8- or 10-bit, 4:2:2 output channels (standard BNC connector)

Configuration of output streams:

- 4:2:2/4:2:2 (same data on both output #1 and #2)

- 4:2:2/4:0:0 (program on #1, alpha #2)
- 4:0:0/4:2:2 (alpha #1, program #2)
- 4:0:0/4:0:0 (same alpha data on both outputs)
- Genlock input (standard BNC connector)
- Genlock output (loopthrough) (standard BNC connector)
- General purpose interrupt (GPI) I/O (standard BNC connector)
- Analog stereo line level audio I/O (3.5mm each)
- Stereo analog headphone output (3.5mm)
- Mono single ended microphone inout (3.5mm)
- Ancillary data passthrough
- Synchronization of video I/O with analog or digital audio I/O



**Figure 5.6** I/O Digital Video Architecture

Figure 5-6 shows the digital/video architecture of the I/O Engine, which supports two DMA channels available for passing two full bandwidth 4:2:2 video signals into the system memory as well as a single DMA channel for passing 4:2:2:4 out of the system. It will also allow a digital video alpha channel to pass in and out of the O2 along with it's corresponding video stream.

## Technical Specifications of the O2 Digital Video Option module:

### *Pinout assignments for the DB9 connector*

<u>Pin</u>	<u>Assignment</u>
1	Analog Video Genlock In
2	Analog Video Genlock Gnd
3	GPI In
4	GPI In Gnd
5	NC
6	Analog Video Genlock loopthrough
7	Analog Video Genlock Gnd
8	GPI Out
9	GPI Out Gnd

### *General Performance*

- Signal amplitude 800mV +/- 10%
- Rise Time 400 to 1500ps
- Fall time 400 to 1500ps
- Comparison rise and fall time Within 500ps
- Overshoot rise No more than 10% of amplitude, 77mV
- Overshoot fall No more than 10% of amplitude, 77mV
- Jitter Less than 20% of UI, 741ps
- Unit interval 3.704ns
- Cable length on transmit 300 meters for output
- Output configuration Two 8- or 10-bit serialized ITU-R BT.601 (SMTPE259M/ITU-R BT.656) digital video channels. One 4:2:2 program, one 4:0:0 alpha. Channel type is user configurable. The same signal can be mirrored on both outputs. For example, the 4:2:2 source can be assigned to both outputs simultaneously, or the 4:0:0 source can be assigned to both outputs simultaneously.
- Input configuration Two independent 10-bit serial ITU-R BT.601 (SMTPE259M/ITU-R BT.656) digital video inputs


### *Genlock Performance*

- Input signal NTSC/PAL composite video, black-burst, or composite sync signal, or either of the two serial digital video inputs

- Jitter performance      Less than 741ps on the serial digital video output stream when genlocked to the analog video source. The jitter is somewhat less when genlocking to either of the two digital video inputs.

#### *GPI Input / Output*

- Output      Open collector, optically isolated, active low
- Input      Active low, optically isolated



**Figure 5-7** Genlock and GPI cable connectors

## 5.4 A/V Synchronization

Audio and video data can be synchronized via the Unadjusted System Time (UST) clock and Media Stream Counter (MSC). Applications can use the UST/MSC support in all O2 systems to synchronize recording and playback of audio and video signals to within +/- 200 microseconds (a few video lines) of each other, or to external timecode signals such as LTC or VITC.

Applications can also synchronize audio and video signals to incoming or outgoing serial or MIDI signals to within +/- 1 millisecond, which is useful




for applications like field-accurate RS-422 video deck control and audio, video or MIDI timecode-synchronized MIDI sequencing.

UST support support can be used to:

- Record or play audio, video and/or MIDI in sync
- Perform field-accurate laydowns to and captures from a VTR
- Slave audio, video and/or MIDI playback to incoming timecode
- Implement highly accurate punch-in and punch-out capability via a MIDI or serial device
- Measure the total input-jack-to-output-jack delay of a real-time signal processing program

PCI option cards also have access to the O2 system's UST support - the Digital Audio Option board supports even higher-accuracy timestamping than the O2 system's built-in audio. All O2 audio options are able to lock their sampling clock to a reference video signal, providing precise, drift-free synchronization for long-format audio and video material.

## 5.5 O2Cam



**Figure 5-8** O2Cam, available with the Analog Video Option module

The O2Cam digital camera can be used to capture high quality still images and video clips for conferencing, movie and screen captures, as well as other media applications. Many of the bundled media tools can be used to control the O2Cam camera. (See section 10.1.3 for additional information.)

### 5.5.1

#### O2Cam Feature Summary

<b>Sensor:</b>	512 x 480 NTSC color CCD, 1/3-inch format
<b>Standard lens:</b>	F2.0, 6mm with M11 screw mount and built-in IR filter
<b>Output formats:</b>	422 YCrCb pixels with 1 byte luminance, and 1 byte chrominance (Logitech DVCi Specification, Rev 1.3)
<b>Automatic gain and shutter control:</b>	Shutter speed automatically adjusted to achieve average output brightness. Also can be manually adjusted. Shutter speeds are selectable between 1/60, 1/100, 1/125, 1/250, 1/500, 1/1000, 1/2000, 1/4000, and 1/10,000 seconds.
<b>Auto white balance:</b>	Initializes properly at power on for offices with or without windows, incandescent and/or fluorescent lights
<b>Host interface:</b>	DVCi using 68-position 0.8mm Champ connector on attached 8-foot cable. Video is byte wide plus hsync
<b>Microphone:</b>	Built-in directional microphone and preamp produce single-ended line level audio signal
<b>Software control:</b>	The camera can be controlled from applications using the Video Library via tools such as MediaRecorder and VideoPanel. See chapter 10 for more details.

## Chapter 6      *Compression*

The O2 architecture incorporates a versatile pathway for recording and play back of real-time video or still image data, performing compression of that video or image data, making the compressed data available to other engines, and storing compressed data on disk.

The compression pathway supports encode or decode of one stream of NTSC, PAL or 601 resolution video in real-time using industry-standard JPEG compression. Multiple streams can be decoded at lower resolutions and slower frame rates; non-standard aspect ratio streams are supported as well. Images up to 32Kx32K in resolution are supported through this pathway providing enough memory is available.

The O2 architecture includes an integrated processor called the Image and Compression Engine (ICE). ICE currently supports JPEG for video and image compression; microcode for additional formats may become available in the future. This processor also supports OpenGL imaging extensions, which is discussed in Chapter 7. Other compression algorithms, including MPEG-1 and H.261, are supported in software.

### 6.1 Video Compression Pathway

Under software control, the O2 architecture supports the establishment of the following pathway for bringing in video data and compressing that data:

- *Video In*—The I/O Engine (IOE) controls the flow of data from the video inputs to the memory controller.
- *Memory*—The IOE writes uncompressed data to memory. Software can display it on the screen, manipulate it with the Memory and Rendering Engine (MRE) or ICE, use it as a texture, and/or pass it back to Video Out, all without the need for a copy, upload or download.
- *Imaging and Compression Engine (ICE)* — A pointer to the uncompressed video data stored in memory is next passed to the imaging engine.
- *Compression*—The ICE executes the selected compression algorithm on the data.
- *Memory*—The ICE writes compressed data to memory. Once in memory, the compressed video data is accessible by any functional area of the system.
- *Disk*—Compressed data can be written out to disk for storage.

Similarly, compressed data can be taken from disk files and brought back into memory. Once in memory, the data is decompressed on the ICE and can then can be displayed, manipulated, and/or passed to Video Out using the reverse pathway.

The basic compression pathway optimizes video data flow throughout the entire O2 system allowing tasks to be effectively shared. For example, color conversion can be offloaded to the MRE, allowing the compression processing pipeline to be optimally used for its specialized functions while overall processing is accelerated by the parallel workload distribution.

The tightly integrated video pathway, and the resulting real-time throughput, also make it possible to integrate video with other forms of graphics and imaging data. Video can be brought in, wrapped around 3D geometry, and output to the display in real time. Video can also be merged with an image, manipulated, and displayed in real time.

## **6.2 Compression Rates**

With 4.5:1 average compression, O2 workstations can encode or decode one full-resolution (NTSC, PAL or 601) full-rate stream in real time. These throughputs are only possible since the pathways between the video in ports and memory are tightly integrated. By using UMA, video data is handled in the same manner as graphics or imaging data. Since CPU involvement is not required, video processing has minimal impact on other system tasks.

The ICE implements a constant Quantization, or Q-factor (versus a constant bit-rate) algorithm. The actual compression ratio will depend on the nature of the video content.

## **6.3 Comparison With PC Solutions**

The O2 system allows the ICE to be shared like any other system resource. Compression functions can be applied to any image or video stream, and are fully independent of video rates and sizes. Handling multiple smaller “motion image” streams is very easy with the integrated compression solution.

In comparison, PC-based systems treat video as an add-on, with video data being passed from an external video processing module over an I/O channel. The I/O rates, substantially lower than the O2 UMA bandwidth, severely limit throughput and real-time display capabilities. The CPU involvement also carries a higher overhead for video on PC platforms. As a workaround to the bus bandwidth issues, PC add-on board suppliers devised an interconnect cable that bypasses the bus. O2 systems do not need this interconnect, or its associated development and usability issues.

## **6.4 Compression of Images**

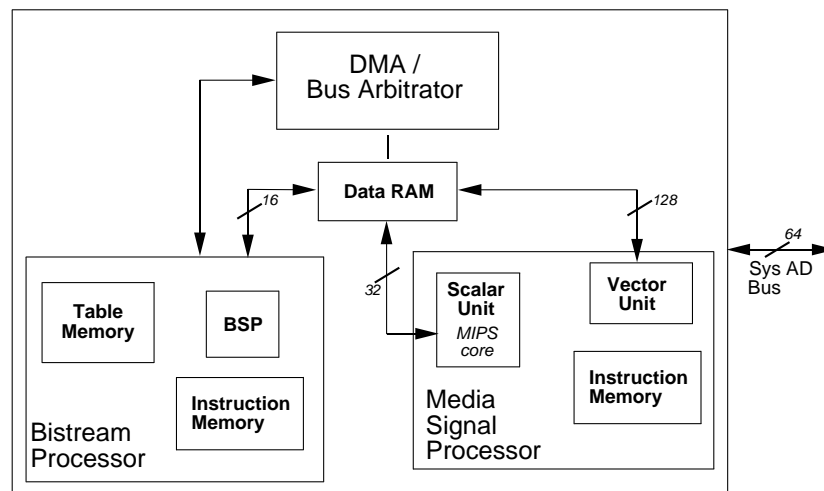
The O2 compression capabilities can be applied to imaging data as well as video data. Images up to the size of 32Kx32K, dependent on the amount of available memory, can be compressed for efficient processing or storage to disk.

## Chapter 7 *Image Processing*

The Imaging and Compression Engine (ICE) processor performs high-speed, high-quality image processing tasks, implementing image and video compression, decompression, and processing functionality in hardware. The processor offers operations such as convolutions, histograms and lookup tables.

### 7.1 Fast and Flexible Image Processing

ICE operates on arbitrary two-dimensional blocks of pixels, including images and fields, frames of video, or arbitrary two-dimensional data. The imaging engine efficiently converts a variety of image protocols, color spaces, and signal domains (e.g. frequency and time) with minimal usage of the CPU. Each image can be operated on independently using different imaging operands. The flow of images through the O2 workstation is efficiently managed by the functional areas illustrated in Figure 7-1.



**Figure 7-1** Imaging and Compression Engine block diagram.

The O2 imaging pathway includes four major processing units:

- A Media Signal Processor (MSP) that performs integer-based signal processing, logical, and mathematical operations
- A bitstream processor (BSP) optimized to perform variable bit-length processing
- DMA engines that move data to and from system RAM
- A host interface that transfers large blocks of data between system memory and internal data RAM

## 7.2 Media Signal Processor

The Media Signal Processor architecture has been tailored to the ICE's specific features and capabilities.

Key features of the MSP include:

- 6KB data RAM
- 4KB instruction RAM
- 32-bit scalar unit
- 128-bit vector unit with one multiply accumulate per clock (performs either 8 16-bit multiplications or 16 8-bit multiplications)
- Dual issue instruction dispatch unit

## 7.3 Bitstream Processor

The bitstream processor (BSP) manages bitstreams of compressed data. The 16-bit RISC-like load-store architecture includes an instruction set with familiar register-to-register operations (e.g. arithmetic operations), instruction stream control (jumps and branches), and memory-to-register data transfers. Additionally, the bitstream processor provides instructions to manipulate arbitrarily aligned tokens in bitstreams of data, and instructions to perform table lookup operations for decoding variable length tokens in a bitstream. The lookup tables currently support JPEG; additional algorithms may be supported with future software releases.

Key BSP features include:

- Accelerated compression standard variable length bit formats
- 16-bit RISC core architecture
- Multi-cycle instruction extensions
- Factory provided Huffman tables for multi-standard support

## 7.4 DMA Unit

Each of the two DMA channels consists of a DMA state machine, control registers, and descriptor memory. The control registers and descriptor memory are programmable by the host, the MSP scalar unit, or the bitstream processor.

Key features of the DMA unit include:

- Flexible address generation that utilizes the system's Unified Memory Architecture, avoiding costly dedicated RAM
- Virtual address support for the MSP access to shared system memory
- Special pixel manipulation features (e.g. Y/C split)

## 7.5 Imaging Operations

Built-in functions in the image processing engine accelerate pixel transfer operations (scaling, biasing, and component-to-component mapping) and OpenGL imaging extensions (convolution, minmax, histogram, color matrix, and lookup tables). After performing these operations, the imaging data is transferred from the imaging engine's buffers to either the frame buffer to a pbuffer (e.g. a pixel buffer, an off-screen rendering area), an off-screen memory area, or into tiled texture buffers. The MRE's graphics engine controls the transfer of buffered imaging data to either the frame buffer or a pbuffer. Because the ICE's Media Signal Processor is a fixed-point processor, all image processing operations are computed in 16-bit fixed-point arithmetic.

All image processing functions are accessed either via the OpenGL imaging extensions (see section 4.7.3) or the digital media image conversion library (dmIC). The Image Vision Library is also available for image processing programming on O2.

### 7.5.1 Convolution

Convolution operators execute linear image filtering in real time. User-definable filters perform a variety of functions, including sharpening and blurring. Applications can optimize image quality and throughput using a wide selection of supported separable and general convolutions filters in 3x3, 5x5, and 7x7 kernel sizes.

### 7.5.2 Histogram and Minmax

Built-in histograms and minmax functions provide statistics on color component values within an image. Histograms count occurrences of specific color component values, and report them in user-defined bins. Using all of table memory, there can be a total of 1024 16-bit bins which is enough for 1 RGBA histogram with 256 bins per component. Minmax keeps track of minimum and maximum color component values. Through OpenGL, applications can analyze the stored statistics and help deliver balanced, high-quality images.

### 7.5.3 Color Matrix and Linear Color Space Conversion

The color matrix operation transforms the colors of each pixel by a user-defined 4x4 matrix. In addition to performing linear color space conversion, the color matrix can also be used to perform color component swapping and component replication.

### 7.5.4 Lookup Tables (LUTs)

Color lookup tables can be inserted in various stages of the pixel path to adjust image contrast and brightness after an image operation. The color LUTs and the histogram share up to a maximum of 2048 bytes of table

memory. This provides enough space for two 8-bit RGBA LUTs, or one 8-bit RGBA LUT, and one 64-bit RGBA histogram. Note that only two tables can be active at any time for hardware acceleration; more than two require CPU assist.

### 7.5.5 Scale/Bias

The scale/bias operation multiplies each pixel component by a scale value, then adds a bias value. This operation is used to perform linear color correction.

## 7.6 Imaging Performance

Current performance specifications for ICE's imaging engine include:

- 42 megapixels per second, bilinear pixel resampling
- 17.4M components per second convolve, 3x3 separable, 8-bit to 8-bit
- 17.6M components per second convolve, 3x3 general, 8-bit to 8-bit
- 28.6M component per second scale/bias, 8-bit to 8-bit
- 28.2M components per second color matrix, 8-bit to 8-bit
- 217.7M components per second blend, 8-bit to 8-bit

Because the image processing engine uses fixed-point arithmetic, convolution, scale/bias, and color matrix operations are accelerated only if all convolution, scale/bias, or color matrix coefficients are between -1.0 and +1.0, inclusive, or if all coefficients are integers.

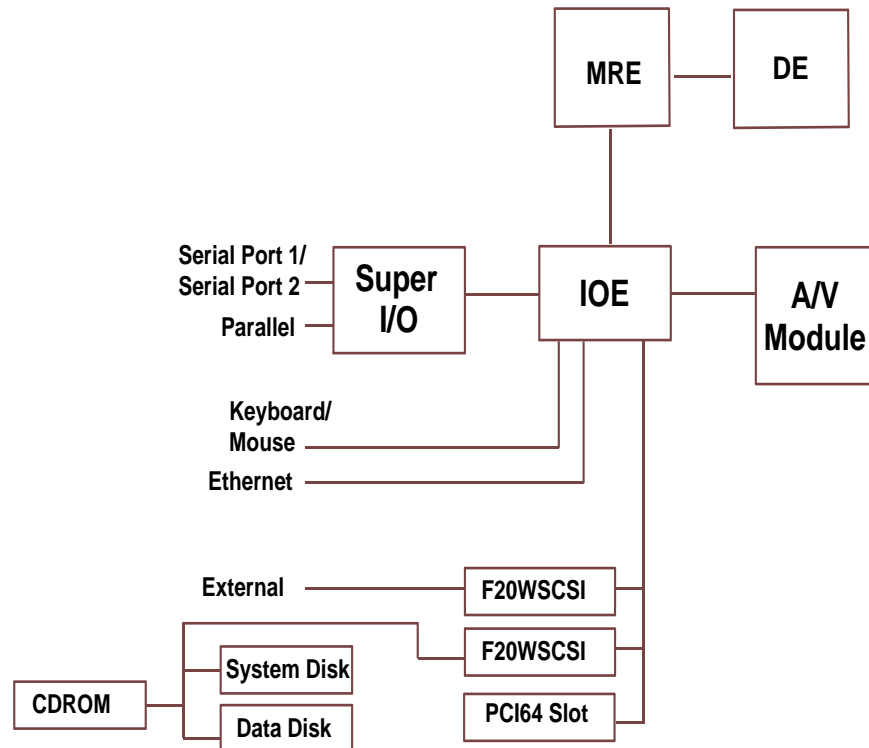
To accelerate an imaging operation with fractional coefficients outside the range of -1.0 to +1.0, prescale all coefficients to the -1.0 to +1.0 range, and then use a post-color-matrix integer scale operation.

Post-operation scaling with floating point can also be accomplished using the alpha blending hardware in the MRE. Please see the technical document entitled *Fast OpenGL-based in place "scale & bias" technique* at <http://reality.sgi.com/mjk> for further information.



## Chapter 8      *The I/O System*

The O2 system I/O Engine (IOE) contains all of the basic I/O interfaces, including keyboard and mouse, parallel and serial connections, audio, video, and Ethernet. The IOE also contains an interface to an external 64-bit Peripheral Component Interconnect (PCI) expansion bus that supports one half-length option card, and to the internal and external Single Ended Ultra Wide SCSI-3 ports. A block diagram of the IOE is shown in Figure 8-1.



**Figure 8-1** O2 I/O interfaces

### 8.1 Fast Ethernet Interface

The Ethernet interface supports both 10baseT and 100baseTX Fast Ethernet connections through a Media Interface Independent (MII) interface. Designed to minimize memory usage and maximize traffic efficiency, the interface features:

- Both half-duplex CSMA-CD and full-duplex circuit switched modes
- Detailed transmit and receive status vectors for each packet
- Internet IP checksum computed for all received packets
- Multicast address filtering

- Automatic transmit padding of short packets to minimum Ethernet length
- Programmable receive interrupt delay time
- Programmable receive gathering of short packets into single buffers
- Flexible transmit and receive buffering
- Functional subset of IOC3 Ethernet interface.

## 8.2 ISA Bus Interface

The IOE features an interface to a PC-style ISA bus used to connect the serial and parallel PC peripherals, an internal Dallas DS1687 Calendar Clock, up to 2MB of Flash-ROM, and a Number-In-a-Can (NIC) serial PROM. This interface is provided for on-system chips only; no cards are supported through the ISA interface (PCI cards are supported as described below). Five DMA channels are available for the serial and parallel interfaces. The interface supports PIO bus cycles and interrupts from the external ports.

The ISA bus connects to a Texas Instruments Super I/O device that offers:

- Two high-speed 16C550C serial ports with hardware RTS/CTS and per port HP IR link support
- Clock prescalers on both serial ports to support MIDI and normal baud rates up to 460.8Kb
- EPP/ECP-1284 high-speed parallel port with all downward-compatible modes

## 8.3 PS/2 Keyboard and Mouse Interface

The O2 system IOE contains two serial PS/2 interfaces that can each support an industry-standard PS/2-style keyboard, mouse, touch-pad, or track-ball type input device. Both serial ports conform to IBM PS/2 Keyboard and Auxiliary (type 1) Serial Port Protocol, with input clock rates as defined by the IBM PS/2 interface specification.

## 8.4 Counters and Timers

A set of counter timers included in the IOE provide a common time base and generate event interrupts. A single 32-bit read-only counter features a 960 nanosecond resolution. Three independently-operating 32-bit event registers generate interrupts, and twelve 32-bit registers are used for audio and video MSC/UST count and timestamp storage.

## 8.5 SCSI Controllers and PCI Expansion Bus

The O2 system has two on-board PCI devices (SCSI controllers) and one 64-bit expansion slot for a half-length option card. For supporting additional devices, the IOE contains a digital PCI host bridge which provides the system interface to a 64-bit 33 MHz Revision 2.1-compliant PCI bus. Posting buffers enable high-performance DMA operation.

## 8.6 Graphics Link

The O2 IOE-MRE link controls all communications between the IOE and the Memory and Rendering Engine (MRE). The queued message-passing interface uses 64-bit wide data and control words, implementing a half-duplex, 32-bit, double-speed, synchronous bus. Bus mastership is negotiated via tokens. The I/O bus arbiter keeps track of outstanding transactions so that the maximum number of memory transactions does not exceed the graphic engine's FIFO.

## 8.7 Standard I/O Connections

The standard I/O configuration includes:

- Two RS-232 serial ports
- One external 68-pin connector Single Ended Ultra Wide SCSI-3 with dedicated controller; another separate controller exclusively for internal SCSI devices)
- One IEEE 1284C parallel port
- One 10/100baseTX Ethernet port, auto-sensing, RJ45 connector
- One HD-15 DDC VGA-style monitor connector
- PS/2 keyboard and mouse connectors
- Audio headphone out and stereo line-out (on the standard Analog Audio and Analog Video Option modules)

## 8.8 I/O Connections on the Analog Video Option Module

I/O interfaces on the Analog Video Option module include:

- RCA style composite video in
- RCA style composite video out
- 6-pin mini-DIN S-Video in
- 6-pin mini-DIN S-Video out
- 68-pin digital video input/output for O2Cam
- Left and right stereo audio RCA in
- Left and right stereo audio RCA out
- Stereo audio RCA out (on rear)
- Microphone in jack (3.5mm, stereo)
- Headphone out jack (3.5 mm, stereo)

## 8.9 I/O Connections on the Digital Video Option Module

I/O interfaces on the Digital Video Option module include:

- two BNC style ITU-R BT.601 (SMPTE 259M/ITU-R BT.656) video input channels
- two BNC style ITU-R BT.601 (SMPTE 259M/ITU-R BT.656) video output channels
- BNC style Genlock/General Purpose Interrupt I/O
- Left and right stereo audio RCA in (3.5mm)
- Left and right stereo audio RCA out (3.5mm)
- Microphone in jack (3.5mm, mono)
- Headphone out jack (3.5mm, stereo)

## 8.10 Optional PCI I/O Cards

The following PCI option boards are available from Silicon Graphics:

- ISDN Basic Rate Interface
- FDDI Single-Attached Copper Adapter
- FDDI Dual-Attach Fiber Adapter
- FDDI UTP/SAS
- 10/100Mb Auto-sensing Ethernet Adapter
- T1/E1 Synchronous Serial card, 2 port RS 232 interface
- T1/E1 Synchronous Serial card, 2 port V.35 interface
- T1/E1 Synchronous Serial card, 2 port X.21 interface
- Single Ended Ultra Wide SCSI-3 Adapter
- Digital Audio Option Board

The following PCI options boards are sold and supported by a third party. SGI does not warranty or service these boards:

- ATM board - FORE Systems, <http://www.fore.com>
- Serial HIPPI - Essential Communications, <http://www.esscom.com>
- Token Ring - Madge Networks, <http://www.madge.com>
- Fibre Channel - Prisa Networks, <http://www.prisa.com>

## Chapter 9      *Display*

The O2 Display Engine (DE) delivers high-resolution color graphics to a non-interlaced display or flat panel display. The primary function of the DE is to fetch several pixel streams from main memory, format and blend them together and output them to the display.

The Display Engine supports:

- 16- or 32-bit normal planes
- 8-bit overlay planes
- Multimode display
- 4608-entry color map
- 256-entry gamma map
- 32x32 three-color cursor
- Programmable video timing
- Stereo glasses support via optional flat panel adapter card
- Presenter 1280x1024 flat panel via optional flat panel adapter card
- 140 MHz dot clock, supporting 1280x1024 resolution at 75Hz refresh rate
- I<sup>2</sup>C interfaces for Display Data Channel (DDC) support and flat panel control
- NTSC and PAL sized screen capture with flicker filter

### 9.1 Display Engine Interfaces

The DE incorporates three ports that provide high-speed data transfer to and from other O2 systems.

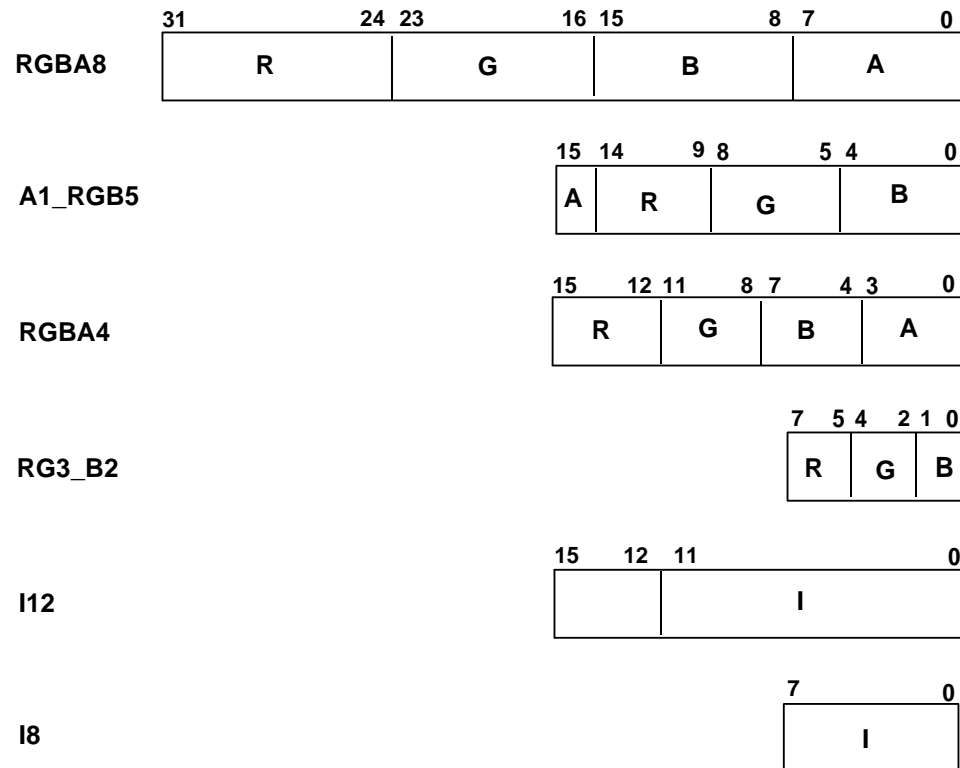
The *Graphics-to-Display Interface* is a point-to-point, burst-oriented protocol with a peak bandwidth of 1GB per second. The data path between the Memory and Rendering Engine (MRE) and DE consists of a 64-bit data bus toggling at 133 MHz.

Using an internal programmable pixel clock output (up to 140MHz), the *DAC/flat panel interface* drives a pixel pipeline and the video timing controller. The interface outputs to the DAC and provides monitor sync signals and a status bit that allows software detection of monitor power. The interface also supplies the horizontal and vertical sync signals required for the optional flat panel display.

The *I<sup>2</sup>C interface* supports the DDC VESA standard that allows two-way communication between the computer and the monitor. This internal interface enables software monitoring of timing information and access to additional monitor functionality.

For more information about the DAC/flat panel interface and the I<sup>2</sup>C interface, refer to the *Digital Monitor Interface Specification*, P/N 001-0179-001.

## 9.2 Pixel and Tile Formats



**Figure 9-1** Pixel formats for DMA (big endian only). Pixels are 8-, 16-, or 32-bits deep.

The O2 stores screen pixels in tiled format to increase 2D spatial locality for screen rendering. Tiles, organized as 128 lines by 512 bytes, can be 8-, 16-, or 32-bits per pixel, are always 64KB in size, and are aligned on 64KB boundaries. The pixel width of a tile is 512 divided by the pixel depth - i.e., 8-bit pixels are tiled 512x128, 16-bit pixels are tiled 256x128, and 32-bit pixels are tiled 128x128.

The frame buffers are built out of an integral number of tiles. For example, a 1280x1024x32 frame buffer exactly occupies 10x8 tiles. The frame buffers, stencil buffer, z-buffer, and texture data are all configured as part of memory. They do not use dedicated areas and are therefore not restricted in size, except by the amount of memory installed in the system. For more information about the typical frame buffer sizes, refer to section 4.4.

### 9.3 Overlay and Normal Planes

The overlay planes are 8-bits per pixel deep, and operate in color index mode. The overlay pixels always index into the last 256 entry window in the 4608 entry color map and always pass through the gamma map; 0x00 values cause the overlay to be transparent.

Normal planes are composed of 64K tiles, addressed by a tile pointer table. The planes can be 8-, 16-, or 32-bits deep.

### 9.4 Monitor

The O2 ships standard with a 17" color monitor; a 20" monitor is optionally available. Features of the monitor include:

- High resolution non-interlaced display with resolution up to 1280x1024
- Multisync
- Sync on green

### 9.5 XMAP, Color, and Gamma Maps

The XMAP block selects between the cursor, overlay, and normal pixels. The color map takes the pixel stream from the XMAP block and performs any necessary color lookup to produce a 24-bit RGB8 pixel stream. It is organized as three 4608x8 maps that are indexed together for I8 and I12 modes, or on a component basis for RGB modes to support X direct color visuals. RG3\_B2, RGB4, and RGB5 pixels are first expanded to RGB8 by bit replication. The gamma map is organized as three 256x8 maps. Cursor pixels always bypass the gamma map, while overlay pixels are always gamma-corrected.

### 9.6 Video Timing and Video Capture

The video timing block generates all of the pixel synchronous timing signals in the display engine. A range of timing signals and functions are controlled, including:

- Horizontal and vertical sync
- 120Hz refresh rates for stereo glasses support (alternating left and right eye images on each frame)
- Simultaneous CRT and flat panel signal support
- Programmable vertical interrupts
- Internal timing for display pipelines

The DE can also send a 140MHz pixel stream back into main memory, where it can be sent to the video output port in the I/O Engine (IOE). The DE video capture (NTSC/PAL mode) consists of a filtering stage and a DMA controller stage.

## 9.7 Software Interface

All of the registers and RAM inside the DE are mapped to a physical address range for software access to precise functionality. All DE registers are read- and write-accessible. Loads and stores are 32-bits, with unused bits software-masked on reads, and ignored on writes.

## 9.8 Video-Style Screen Capture to Memory

The DE can redirect an NTSC- or PAL-sized region of its final output image (after color map lookup, overlay and cursor composition, and gamma ramp) back to main memory in a format useful for video output or compression.

The DE generates one video field per graphics vertical retrace. It generates the video fields either by point-sampling every other graphics line or interpolating every other graphics line with a 1/2, 1/4, or 1/9 flicker filter.

The DE does not resample horizontally. Applications use the Video Library's VL\_SCREEN source node to access the functionality.



## Chapter 10     *Software*

### 10.1 IRIX 6.3 User Environment

The enhanced IRIX 6.3 Desktop User Environment allows users to access iconic views of HTML and FTP sites. Built-in Content Viewers let you view movie, image, Inventor, text, and sound files where you see them in an IconView. Other new features include the following:

- Pasting text on the background to create an icon. e.g. pasting a URL on the background will create a WebJumper icon
- Typing in any IconView window to select icons
- Toolchest access to files or documents on a remote system

The system administration environment is implemented using a task-based user model with a web-based user interface. *Active Guides* step the user through each task, including specifying in the first step all information necessary to complete the task. New tools include support for increasing system security, setting up modems, and managing swap space; versions of system administration tasks available in previous releases have been revised to use the *Active Guide* paradigm.

The objectserver is now obsolete, and no longer ships with IRIX 6.3.

#### 10.1.1     **Intranet/Internet Support**

Silicon Graphics system platforms include complete solutions for today's networks.

New features in IRIX 6.3 User Environment include:

- Cross-platform communication via new Netscape® tools (Netscape Navigator®, Netscape FastTrack Server, Netscape Mail with live Internet hyperlink support; Netscape News Reader)
- OutBox personal Web publisher to easily share documents using the bundled FastTrack server.
- Web-aware Directory Views for viewing icon-based links to audio, video, 3D, images, HTML, or any file on the Intranet or Internet
- Web authoring tools for creating Web-based documents
- Full text indexed (searchable) and Web-viewable documentation, including UNIX® man pages, release notes, help and InSight books. Documentation is available in the traditional formats as well.
- Completely new, Web-based Help and documentation system including browse and search capabilities as well as "How Do I..." help.

### 10.1.2 Standards

The IRIX 6.3 User Environment, built on UNIX SVR4 with Berkeley Extensions, conforms to every major UNIX standard as well as a variety of cross-platform standards. Every IRIX 6.3 system is designed to be compliant with the following standards:

- Display Postscript
- FIPS 151.2
- Motif
- POSIX 1003.1/1003.2
- Tooltalk
- X11R6
- XPG4

IRIX 6.3 systems also achieve compliance with other standards:

- Appletalk® (through Xinet)
- QuickTime™ and AVI file formats
- Triton CDE (through an alternative desktop from Triton)
- XFS™ and EFS; NFS2 and NFS3

### 10.1.3 Media Tools

The Silicon Graphics® desktop ships with a set of powerful digital media tools, giving all users the ability to create high-impact digital media content.

The media tools included with the IRIX 6.3 User Environment include several new features:

- Improved performance with more control and power for creating and editing full-size, full-rate JPEG compressed video
- Unified tools (fewer, more-powerful applications) for simplified media creation
- Common user interface across all tools
- Drag-and-drop interface for every user activity
- Web file formats (Animated GIF, PNG, JPEG, and AVI)

Specific tools that ship with desktop include:

- **MediaRecorder** - provides real-time capture of video from an external source (movies), graphics from the screen, audio from CD, microphone, and video or still images from the O2Cam. Controls help select the source and destination, file format, compression algorithm, frame rate, and size, among other options. Pre-defined settings also make the tool very intuitive for the novice user.
- **MediaConvert** - converts movie, audio, and image files to other file formats or compression algorithms. MediaConvert provides complete control of image sizes, frame and sample rates, and processing of a single

file at a time or in batch mode. A large list of industry standard file formats is available to choose from.

- **MovieMaker** - used for composing and editing multitrack movie files with video, image, and audio. MovieMaker has an intuitive interface, similar to a basic text editor. It also includes filters for adding special audio and video effects, as well as a title palette for adding graphics and text to movies. Output formats include Silicon Graphics, QuickTime, AVI, or MPEG-1 format.
- **FX Builder** - a Silicon Graphics plug-in to Adobe Premiere™ and MovieMaker for editing and building 3D filters and transitions. 3D objects can be animated and textured with movies or images. FX Builder can be accessed from the menu of Adobe Premiere and MovieMaker.
- **Media Player** - a tool with intuitive VCR-like controls for viewing movies or playing audio files on the O2 system. It offers a number of viewing options and transport control functions as well as volume control.
- **SoundTrack** - a program for creating and editing multitrack audio compositions. The user can record (basic nondestructive editing and playback with waveform display), process audio and take advantage of the Adobe Premiere plug-in specification, synchronize audio with digital video files on the system disk, and import and export multiple audio file formats.
- **CD Player / DAT Player** - Both share the same GUI for ease of operation. CD Player enables playing audio compact discs from the SCSI CD-ROM drive through the system. DAT Player enables playing digital audio tapes from the SCSI DAT drive through the system. The tools also support capture from compact disc or DAT for audio production.
- **MIDI Sound Synthesizer** - a General MIDI-compatible sampling synthesizer that is entirely software-based. It has multistage envelopes, resonant filters, and other professional features for rendering audio for music, visual simulation, VRML, and other applications. It has extremely low latency and can be controlled in real time via an external MIDI (Musical Instrument Digital Interface) controller.
- **MIDI Keyboard** - a virtual MIDI controller used for the output of MIDI data (such as note numbers and controller data) on 16 channels to either the internal Silicon Graphics software-based virtual synthesizer (MIDI Synthesizer) or to external MIDI devices attached via a MIDI interface to the O2 system's serial port.
- **MIDI Synthesizer Panel** - an application for controlling the MIDI Synthesizer that allows selection of the sound (preset), volume, and panning setting on each of 16 MIDI channels.
- **Video Control Panel** - provides signal control settings for the external video devices connected to the system.
- **Audio Control Panel** - provides control settings for the external audio devices connected to the system.
- **Video Diagnostics Panel** - includes various test functions for the incoming video signal. A vector and waveform monitor allow precise measurement and monitoring of the video signal.

#### **10.1.4 Collaboration Solutions**

Every Silicon Graphics workstation includes tools to facilitate team collaboration:

- OutBox Intranet solution
- InPerson® video conferencing
- Showcase presentations solution
- Annotator for sharing ideas and information about 3D models

#### **10.1.5 Windows® 95 Support**

O2 workstations shipped with at least 64MB of RAM include a licensed version of SoftWindows™95. For SoftWindows 2.0, contact Insignia Solutions, Inc.

### **10.2 Supported File Formats**

#### **10.2.1 Audio File Formats**

aifc	Audio Interchange File Format w/ compression
aiff	Audio Interchange File Format
avr	Audio Visual Research
bicsf	Berkeley/IRCAM/CARL SoundFile Format
iff	Amiga IFF/8SVX
midi	Musical Instrument Digital Interface
mpeg1	MPEG-1 audio bitstream
next	NeXT/Sun audio
rawaudio	Raw (headerless) audio data
sd2	Sound Designer II
sf2	Creative Labs SoundFont2 Compatible
smp	SampleVision
voc	Creative Labs VOC
wave	WAVE (RIFF) audio

#### **10.2.2 Image File Formats**

bmp	Microsoft Windows Bit Map
cur	Microsoft Windows Cursor
eps	Encapsulated PostScript
fit	File Format for Image Tiling
gif	Graphics Image File
hdf	Hierarchical Data Format
ico	Microsoft Windows Icon
icon	Sun Icon and Cursor
jfif	JPEG File Interchange Format
mint	Apple Macintosh MacPaint
pbm	Portable Bit Map

pcd	Kodak PhotoCD
pgm	Portable Grayscale Map
pic	PIXAR Picture
pict	Apple Macintosh QuickDraw/PICT
pix	Alias Pixel
pfa	Adobe Type I Font
pnm	Portable Any Map
ppm	Portable Pixel Map
ps	PostScript
ras	Sun Raster
rawrgb	Raw pixels packed as rgb
rgb	Silicon Graphics RGB image
rgbA	Silicon Graphics RGB image with generated alpha
rgba	Silicon Graphics RGB image with alpha
rla	Wavefront Raster
rle	Utah Runlength-encoded
rpbm	Raw Portable Bit Map
rpgm	Raw Portable Grayscale Map
rpnm	Raw Portable Any Map
rppm	Raw Portable Pixel Map
soft	SoftImage PIC
synu	SDSC Synu
tga	Truevision Targa
tiff	Tagged Image File Format
viff	Khoros Visualization Image File Format
x	Stardent AVS X
xbm	X11 Bit Map
xwd	X Window Dump

### 10.2.3 Movie File Formats

avi	Microsoft AVI (Audio Video Interleaved) - read only supported compression formats: Cinepak™, Indeo™ 3.2
mpeg1s	MPEG-1 systems bitstream
mpeg1v	MPEG-1 video bitstream
qt	Apple QuickTime supported compression formats: uncompressed, Cinepak, Indeo 3.2, Apple Video, Apple Animation, JPEG
sgimv	Silicon Graphics movie supported compression formats: uncompressed, MVC1, MVC2, JPEG, RLE8, RLE24, RLE32)

### 10.3 Language Support

The following languages are supported via the X Server. Several applications have also been localized for languages including French, German and Japanese.

Arabic, Egyptian  
Arabic, Middle Eastern  
Arabic, North African  
Czech  
Danish  
Dutch (Neth., Belg.)  
English (U.S., Can., U.K., Aust.)  
French (Fr. Belg., Can., Switz.)  
German (Germ., Austria, Switz.)  
Greek  
Islensk (Iceland)  
Italian (Ital., Switz.)  
Hebrew  
Japanese  
Korean  
Norwegian  
Polish  
Portugese (Port., Brazil)  
Russian  
Simplified Chinese  
Slovenski (Slovakia)  
Spanish (Spain, S. Amer., Cen. Amer.)  
Suami (Finland)  
Swedish  
Traditional Chinese  
Turkish

## **Chapter 11      Configurations, Packaging**

### **11.1 System Configurations**

The O2 workstations are available in a variety of configurations, offering users choices combining:

- Processors — R5000 with standard secondary cache or R10000 with standard secondary cache
- *Memory* — 32MB or 64MB base; additional memory can be purchased separately from Silicon Graphics or third-party vendors.
- *Storage* — Internal 2GB or 4GB system disks; the R5000 has one additional disk expansion slot available.
- *Video* — Optional analog or digital video

#### **11.1.1 Entry Level Configurations**

- 32 or 64MB memory, 2GB disk, 12x CD-ROM, 17-inch monitor (32 MB O2 systems do not include SoftWindows95)
- 32-bit double-buffered graphics with hardware z-buffering and texture mapping
- Analog audio I/O and processing; optional video modules are available
- Microphone; optional O2Cam available with Analog Video Option module

#### **11.1.2 Packaging**

- Desktop mini-tower enclosure, 8.5"Wx11.5"Hx10.5"D
- Integrated 12x-speed CD-ROM
- 2 3.5-inch x 1-inch blind-mate SCA disk bays for R5000 systems; 1 bay for R10000 systems
- 1 PCI64 half-length slot

### **11.2 Base System Hardware Features and Options**

All of the individual O2 system components are user replaceable.

<b>CPUs:</b>	R5000, 180MHz, with primary cache †
	R5000, 180MHz, with .5MB secondary cache
	R5000, 200MHz, with 1MB secondary cache
	R10000, 175MHz, with 1MB secondary cache
	R10000, 195MHz, with 1MB secondary cache

† Available to OEMs only

<b>Memory:</b>	32MB to 256MB with 16Mbit SDRAM 256MB to 1GB with 64Mbit SDRAM (Q4CY97)  Synchronous DRAM, 100 MHz, 4 Banks, 288-bit wide
<b>Mass Storage:</b>	Embedded 12x CD-ROM Single Ended Ultra Wide SCSI-3 R5000 systems: 2 1x3.5-inch drive bays (1 standard, 1 expansion) R10000 systems: 1 1x3.5-inch drive bay 2GB, 5400RPM disk drive or 4GB, 7200RPM drive
<b>Networking:</b>	10/100baseTX Ethernet autosensing
<b>Power:</b>	170 Watt power supply
<b>I/O Expansion:</b>	1 half-length PCI slot for optional cards (see section 8.10 for a list of currently available cards)  Presenter/StereoView adapter slot
<b>Audio (standard):</b>	Left and right stereo audio RCA I/O Microphone in jack (5mm, stereo) Headphone out jack (5mm, stereo) Mono internal speaker
<b>Analog Video (optional):</b>	RCA style composite video I/O 6-pin mini-DIN S-Video I/O 68-pin digital video I/O for O2Cam Left and right stereo audio RCA I/O Microphone in jack (3.5mm, mono) Headphone out jack (3.5mm, stereo) Speaker out jack (3.5mm, stereo, on rear) Digital camera with built-in directional microphone
<b>Digital Video (optional):</b>	BNC style ITU-R BT.601 (SMPTE 259M/ITU-R BT.656) video I/O BNC style Genlock/General Purpose Interrupt I/O Left and right stereo audio RCA in (3.5mm) Left and right stereo audio RCA out (3.5mm) Microphone in jack (3.5mm, mono) Headphone out jack (3.5mm, stereo)



**Monitors:** 17-inch, 1280x1024 @75Hz (16-inch viewable)  
20-inch, 1280x1024 @75Hz (19-inch viewable)  
Presenter 1280, 1280x1024 flat panel

**I/O:**

Front Panel: Volume control  
Speaker  
Status LED  
Reset switch  
Power on  
CD-ROM drive

Rear Panel: 2 serial 460kbps ports  
Monitor (HD15)  
Single Ended Ultra Wide SCSI-3 (H-den 68)  
Ethernet 10/100baseTX  
Parallel port (IEEE 1284-C connector)  
Keyboard (PS/2 compatible )  
Mouse (PS/2 compatible)  
System power connector

Side: Right and left audio line I/O (RCA)  
Stereo headphone jack (3.5mm)  
Microphone (3.5mm)  
*also see Audio and Video options above*

### 11.3 System Environmental Specifications

Dimensions: 11.75 x 12.25 x 9.75 inches

Weight: 23 pounds (approximately)

Temperature: 10 to 35 degrees C (operating)  
-40 to + 65 degrees (non-operating)

Humidity: 10% to 80% RH non-condensing (operating)  
5% to 95% RH non-condensing (non-operating)

Altitude: 10,000 feet (operating)  
40,000 feet (non-operating)

Vibration: 0.1 inches displacement with all axes  
Operating: 0.25G, 5-380-5 Hz  
Non-operating: 0.5G, 5-380-5 Hz

Shock:  
Operating: 5G for 11mSec, half-sine pulse in Horizontal plane  
7.5G for 11mSec, half-sine pulse in Vertical plane

Non-operating:	20G for 11mSec, half sine pulse in all planes
Acoustic noise:	35 dBa - single drive powered on, spinning, idle
Heat Dissipation:	< 900 BTU/hr
Power Requirements:	100-132 VAC, 2.5-1.8A 200-264 VAC, 1.3-0.9A 47-63 Hz 3-pin, Earth grounded plug
Regulatory Agency Approvals:	
Electromagnetic Emission:	FCC Part 15, Class A Canada DOC, Class A CISPR 22:1993/EN 55022:1988 Class A VCCI Class 1 EN 50082-1:1992 EN 61000-4-2:1995/IEC 1000-4-2:1995 ESD IEC 1000-4-3:1995 Radiated RF EN 61000-4-4:1995/IEC 1000-4-4:1995 EFT
Product Safety:	CSA 22.2 No. 950-95 IEC 950:1986 + A1. A2 EN 60950: 1988 + A1. A2 UL 1950:1995
Directives:	73/23/EEC Low Voltage 89/336/EEC EMC
Type Approvals:	CSA "Certified" CSA "NRTL/C" TUV "T License"

## ***Chapter 12      Reference WWW Sites***

Following are some relevant WWW links to technical information for O2 and other Silicon Graphics platforms:

- The O2 Unified Memory Architecture White Paper - <http://www.sgi.com/Technology/uma.html>
- O2 Digital Media White Paper - <http://www.sgi.com/Technology/O2dm.html>
- Taste of the Developer's Toolbox - <http://www.sgi.com/TasteOfDT>
- The Developer's Toolbox - <http://www.sgi.com/Technology/toolbox.html>
- The Lurker's Guide to Video - <http://reality.sgi.com/cpirazzi/lg>
- The SGI Audio and MIDI Frequently Asked Questions - <http://www.sgi.com/Technology/Audio>
- Silicon Graphics MIPS Group - <http://www.sgi.com/MIPS>
- Silicon Graphics Silicon Surf - <http://www.sgi.com>
- Current third party PCI adapter card providers:
  - FORE Systems, <http://www.fore.com> (ATM)
  - Essential Communications, <http://www.esscom.com> (Serial HIPPI)
  - Madge Networks, <http://www.madge.com> (Token Ring)
  - Prisa Networks, <http://www.prisa.com> (Fibre Channel)