IRIS 4D/70 Superworkstation Technical Report

SiliconGraphics
Computer Systems

Abstract

Superworkstations combine the compute power of a fast minicomputer with real-time three-dimensional graphics. By giving users personal access to this combination of compute and graphics power, superworkstations allow individual users to visualize complex data in real time. This major advance is made possible by merging several key technologies:

- RISC computing architecture
- Custom VLSI graphics processors
- Efficient software development environment

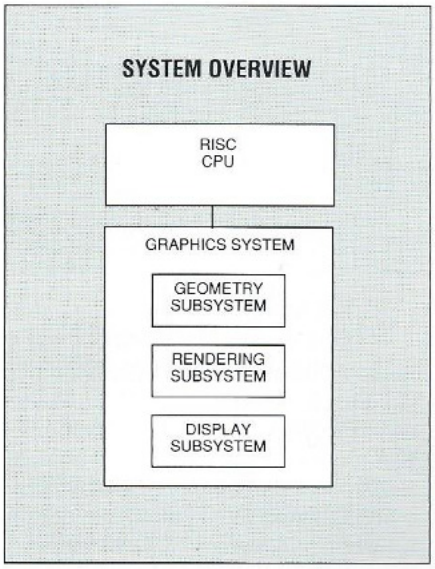

The initial application of RISC (Reduced Instruction Set Computer) technology increases the computing power available in superworkstations by two to five times over the previous CISC (Complex Instruction Set Computer) technology. Although workstations have been able to display wireframes and flat-shaded polygons in real time, they have not offered the performance required to display smooth-shaded images in real-time. Aided by a new generation of VLSI graphics processors, the IRIS 4D/70 offers the performance increases that allow users to display complex, smooth-shaded images in real-time. The IRIS 4D/70 combines these leading-edge technologies with an efficient software development environment to deliver dramatic improvements in object simulation and viewing freedom.

Introduction

New technologies have come together to create a new class of computing platforms called superworkstations. These superworkstations give users personal access to computing power that is an order of magnitude greater than traditional workstations. At the same time, superworkstations give users greater freedom and imaging fidelity to visualize their work. The superworkstation fills a gap that has existed between the traditional workstations and the most recent generation of large scale supercomputers.

Over the past several years, graphics workstations have gained widespread acceptance in the scientific and engineering communities. This is principally due to the workstations' ability to provide the individual scientist or engineer with enough compute and graphics power to analyze and visualize a large number of real world problems. The ability to deliver such power was made possible by several innovative technologies: 32-bit microprocessors, advances in dedicated graphics hardware, multi-windowing environments, and networking.

By blending these technologies into one coherent system, the IRIS 4D/70 Superworkstation from Silicon Graphics expands the horizons of compute environments. The focus of this report is on the integration of the powerful technologies that allow the IRIS 4D/70 to extend its lead in 3D graphics superworkstations into a new dimension: visual computing.

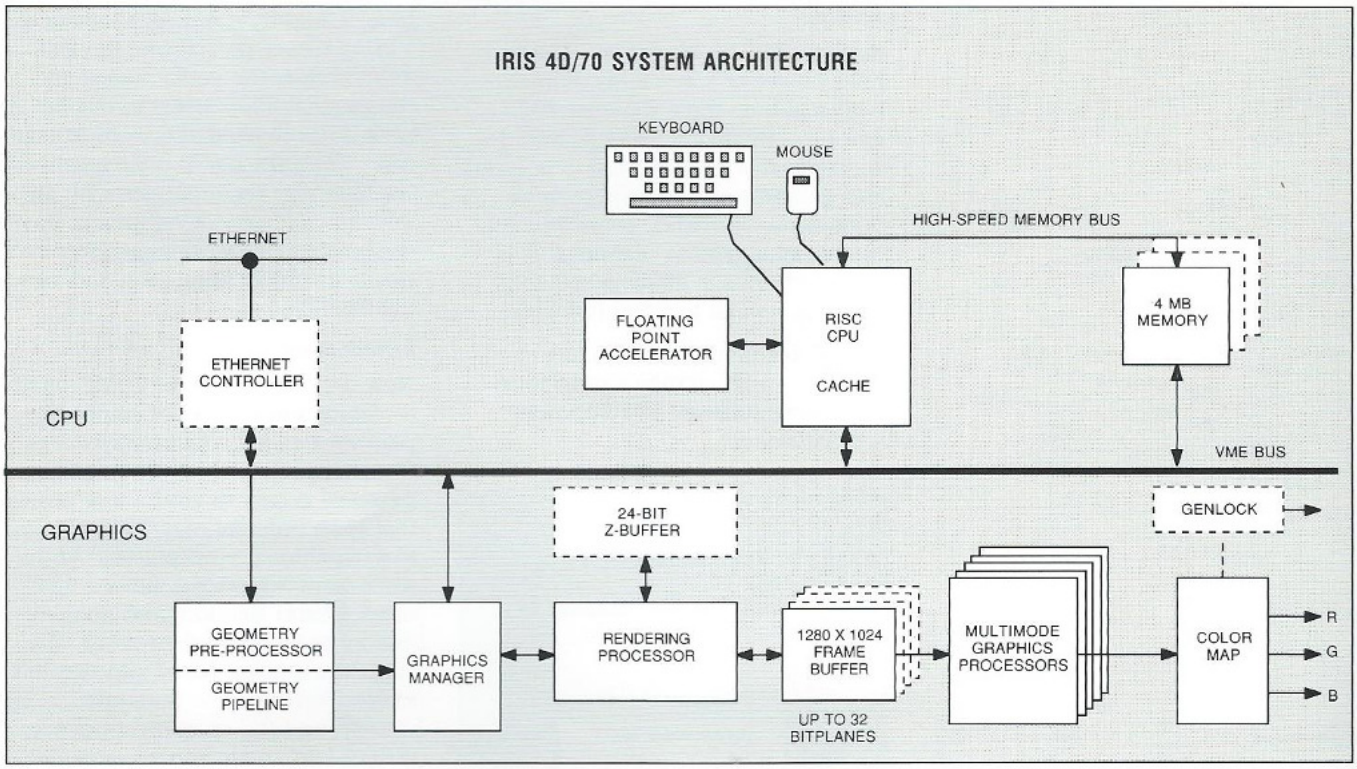

1. A Balanced Fusion: RISC and IRIS 4D/70 Supergraphics

The IRIS 4D/70 fuses enhanced graphics performance with significantly increased computing power enabled by RISC computing technology. The resulting balance between the RISC CPU and proprietary graphics engines ensures constant performance at peak rates.

The choice of RISC technology as the computing platform for the IRIS 4D/70 guarantees a rapid growth path in future computing performance. The latest member of this new generation of superworkstations, the IRIS 4D/70, surpasses the computing performance of the fastest Motorola 680x0 CPUs available today. In future generations, faster clock speeds, larger cache memories, and larger and faster main memory will yield even higher levels of performance.

Reduced Instruction Set Computer (RISC) technology radically reduces the instruction set of the CPU to a well-defined minimum. The resulting instruction set contains the absolute essentials necessary to invoke continuously recurring instructions such as load, add, and store. For these processing primitives, a streamlined architecture is specifically designed to achieve high peak performance at a relatively low hardware cost. Advanced compilers intelligently convert instructions from high-level programming languages into RISC machine code for efficient execution.

Massive parallelism and efficient pipelining allow custom-designed VLSI graphics engines to operate in an optimum fashion on problems that have been efficiently converted from the complex user domain into fundamental primitives at the machine level.

The graphics technology developed by Silicon Graphics allows graphics problems to be intelligently converted into their fundamental structural components. Any arbitrarily defined solid object is converted through a sequence of processing steps from 3D polygons to single pixels. First, 3D convex or concave polygons are transformed, projected, clipped, and scaled by a pipeline of custom VLSI Geometry Engines. The resulting polygons are further decomposed by a dedicated processor and standardized into simple trapezoids with ordered vertices. The initial complex problem has now been subdivided into a number of smaller pieces. Parallel proprietary rendering engines perform the computationally intensive task of shading each piece by efficiently determining the location, color and depth of each individual pixel.

1.1. RISC Technology Computing

RISC technology computing, the use of Reduced Instruction Set Computers as core elements in a computing system, represents the latest major advance in microprocessor computing technology. RISC technology will be widely employed in future generations of scientific and engineering workstations. Leading the world in this technology is the R2000 series of microprocessors from MIPS Computer Systems. Silicon Graphics' latest superworkstation, the IRIS 4D/70, uses the MIPS microprocessor and takes full advantage of its computing power.

RISC technology is the discipline of designing computers with reduced, i.e. simplified, instruction sets such that the majority of instructions can be executed in one machine cycle. The result is an instruction set where most instructions are relatively unencoded and fixed in length– usually a single 32-bit computer word. These types of instructions are easily decoded and executed. The resulting simplicity of the decode, execute, and other CPU functions enables pipelined instruction streams. The silicon real estate that is freed up because of the simplicity of the instruction fetching and decoding circuitry is devoted to logic-intensive, fast arithmetic, and RAM-intensive translation look aside buffers to provide fast virtual memory support and fast instruction execution.

1.2. Speed Advantage versus Memory Requirements

RISC technology computing results in a significant speed advantage over a Complex Instruction Set Computer (CISC) such as a VAX 11/780. Compared to an "average instruction set computer" such as a Motorola 68020 or 68030, the RISC CPU processes "real world" applications at approximately twice the speed.

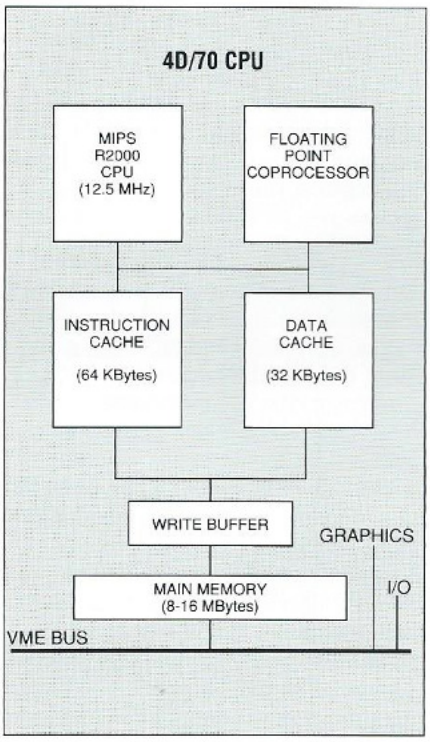

Due to the simplicity of the instruction set, the object code generated for a RISC computer will almost always be larger than equivalent (and slower) object code generated for conventional computers. However, the increased speed of program execution more than compensates for the increased program size.

Realizing the potential of any CPU requires a high level of balanced integration with other equally high-powered system components. The IRIS 4D/70 delivers the full performance potential of the MIPS processor to the user by packing all of the CPU circuitry and memory on a single VME board. The IRIS 4D/70 continues Silicon Graphics' tradition of optimizing system throughput and performance. The architecture of the RISC computing system is illustrated in the diagram below.

To provide additional performance, small high-speed memories (caches) are placed between the processor and main memory. These separate instruction and data caches provide for the increased instruction bandwidth needed to realize the maximum potential of the RISC technology. The caches are direct mapped for added speed. The dual cache scheme ensures that instructions and data will not compete for the same cache locations in the direct mapped implementation. The data cache is write-through, so that main memory always has the latest copy of any data. A write buffer keeps the write-through cache from having to wait for the slower main memory to complete its write cycles. This architecture organizes various types of memory in a balanced hierarchical structure resulting in major savings in access time.

Running the industry standard integer benchmarks, the 12.5 MHz RISC CPU is ten times faster than a VAX 11/780. This fast integer performance of the IRIS 4D/70 is augmented by the MIPS R2010/12.5 floating point chip. This chip performs over 1100 double-precision floating-point operations per second on the industry standard Linpack benchmark.

The IRIS 4D/70 is the second in Silicon Graphics' family of RISC-based superworkstations that will offer continually greater levels of performance in the months to come. As MIPS Computer Systems introduces new and faster processors, these processors will be incorporated into the IRIS-4D product family.

The R2010/12.5 floating-point coprocessor provides an unprecedented level of IEEE Standard floating-point performance in a single, full-custom CMOS VLSI chip. The R2010/12.5 floating-point coprocessor is tightly coupled to the RISC R2000/12.5 processor to form a seamless integration of floating-point and fixed-point instruction sets. Because each processor receives and executes instructions in parallel, floating-point instructions execute at the same single-cycle-per-instruction peak rate as fixed-point instructions. Like the R2000/12.5 RISC processor, the R2010/12.5 uses a load/store-oriented instruction set, with single-cycle loads and stores. Floating-point operations are started in a single cycle, and their execution is overlapped with other fixed-point or floating-point operations.

1.3. Optimizing Compiler

Ordinary compilers translate each program statement written in a high level language such as FORTRAN or C into its equivalent in machine level instructions. The translated or "compiled" program can then be executed on the particular machine for which the compiler was written.

Since ordinary compilers typically don't examine the compilation process in detail, they end up creating far more machine instructions than necessary to do the job the programmer intended. This is because each high level instruction is translated atomically, without regard to the previously generated machine code and without consideration for the instructions to follow. Each high-level language instruction is treated like a small independent unit that has to allocate its own resources and set up and maintain its own registers without any notion of the potential repetition within the surrounding high-level language statements. The resulting redundancy of the machine code greatly reduces run-time efficiency.

An optimizing compiler, on the other hand, intelligently analyzes the context of the machine code it is producing. Performance gains can be obtained by removing redundancies, creating shared arguments and utilizing other sophisticated optimization techniques to generate tight, streamlined code during compilation.

On reduced instruction set hardware, the compiler provides another performance-critical function. By generating the machine code for the RISC processor, the compiler reduces the functional complexity of high level instructions to the set of machine primitives. For instance, the compilation decomposes complex processes such as procedure calling into a coherent set of simple machine codes tailored to the optimized RISC architecture. A RISC technology compiler intelligently utilizes the processor-cache design in order to take maximum advantage of the machine architecture. Even complex instructions, when compiled, execute faster on the RISC architecture than on an equivalent complex instruction set processor.

1.3.2. The IRIS 4D/70 Compiler Suite

In addition to optimizing the program at the code generation level, the IRIS 4D/70 compiler uses a collection of high-level optimization techniques that are unusually advanced for a computing platform in its price range. Developed in conjunction with both the RISC processor and the operating system, the resultant compiler is both run-time and compile-time efficient. Such an optimizing compiler is an essential element of a superworkstation.

The IRIS 4D/70 compiler suite currently contains compilers for C and FORTRAN 77, with other compilers planned for release in the near future. The FORTRAN compiler adheres to the ANSI standard, with common extensions. The C compiler conforms to the de facto standard.

Compilers for both languages make use of a common code generation and optimization back end. The front end language translators generate an intermediate language, called U-code, developed at Stanford University, specifically designed to provide uniform compilation and optimization support. In addition to providing a single, special purpose optimization vehicle, it also allows new language translators to be added to the compiler suite with the guarantee that they will take full advantage of both the optimizer and the underlying hardware architecture. The compiler suite is thus modular, flexible and intrinsically easier to enhance and support than other approaches.

The U-Code compiler-optimizer consists of several layered components. At compile-time, the user sets a parameter that indicates how many cascading layers of optimization to invoke. Invoking all three layers typically results in run-time savings of 30% to 50%. The highest layers use a very complete set of sophisticated machine-independent techniques for optimization. Intermediate layers optimize further, exploiting the advantages gained through the use of the RISC architecture. The lowest levels take advantage of the MIPS processor/coprocessor's specific cache and register design to generate streamlined machine code.

Sophisticated optimization algorithms, such as those found in the IRIS 4D/70 compiler suite, can lengthen compilation time unless the compiler itself is designed to be efficient. Binary forms increase the compile-time efficiency for all layers of the compilation process. The compiler is self-compiling and therefore self-optimizing. The IRIS 4D/70 compiler makes extensive use of run-time statistics to ensure that the resulting run-time performance gain justifies the compile-time effort.

1.3.3. Run-Time Efficiency

The IRIS 4D/70 compiler improves run-time efficiency on several levels corresponding to the different functional layers in the optimizing compiler. The result of any optimization is transparent to the user.

At the highest level, run-time efficiency is gained through a variety of advanced techniques similar to those found in optimizing compiler research and supercomputing platforms, but not previously found in workstation environments. Techniques provided by a machine-independent global optimizer include copy propagation, common subexpression elimination, invariant code motion, redundant arguments removal and strength reduction.

At the lowest level, the optimizing compiler takes full advantage of the load/store nature of RISC machines by making register allocation very effective. By globally analyzing usage patterns, the optimizer decides which variables should reside in registers, and schedules instructions based upon the characteristics of the MIPS processor. Since the optimization process is layered, intermediate computational results generated by higher level optimizations are also analyzed for register allocation. The thirty-one general purpose registers in the MIPS CPU make the register allocation portion of the optimization particularly versatile.

At the highest level, cross module register allocation is performed based on an analysis of interprocedural usage. While a difficult task to perform, it can yield significant performance enhancements, and is unique on the IRIS 4D/70 Superworkstation. Together with the in-line procedure merging available at this level, it allows high level programmers to use as many procedure calls as needed, knowing that they are computationally inexpensive. All three layers of optimization eventually make use of a machine code generator. This code generator further optimizes the code by folding in expressions, deleting dead code, and performing branch and label optimization which are best suited to the MIPS instruction set.

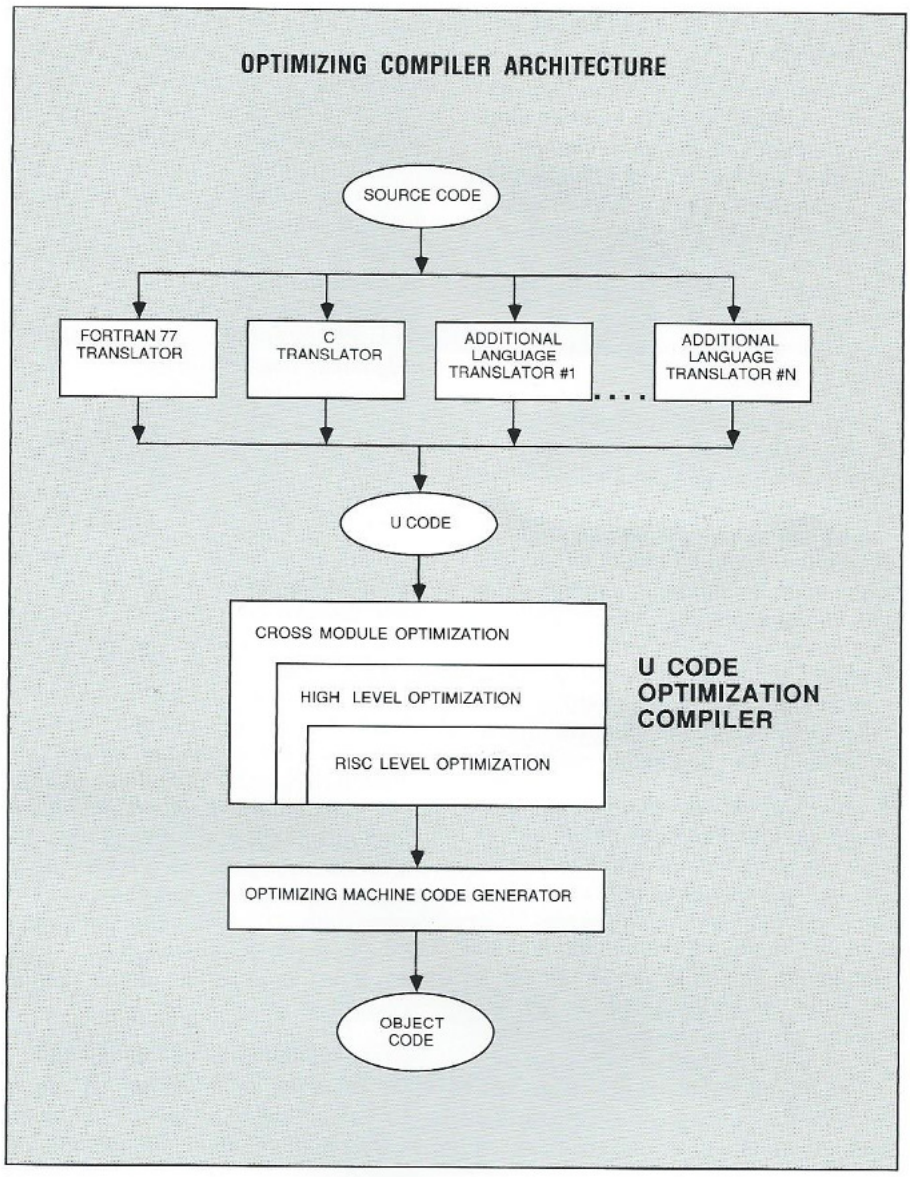

The IRIS 4D/70 optimizing compiler represents a significant advance in workstation technology, using a very complete set of higher level optimization techniques while at the same time exploiting the benefits obtained from the MIPS chip set. In the supercomputing environment, the importance of having an excellent optimizing compiler tuned to the hardware has long been recognized. For the first time Silicon Graphics brings this technology to a superworkstation.

2. Supergraphics Enabled By VLSI Technology

With the introduction of the IRIS-4D family, Silicon Graphics enters a new era: the four-dimensional universe is modeled accurately and dynamically on a superworkstation. This dramatic technological leap inaugurates new dimensions in solids modeling, visual simulation, molecular modeling, animation and other domains with demanding visual computing requirements.

The IRIS 4D/70 incorporates significant technological advancements, based on increased parallelism, several new proprietary and semi-custom VLSI processors, and faster memory. Integrated into a Graphics Subsystem, these silicon solutions enable the system to meet constant, demanding deadlines to draw full frames in a fraction of a second. By being able to render solid 3D objects in this extremely short time span, the quality of visualization is greatly improved.

Equally important is the introduction of a hardware-based multiple windowing environment that augments the user's freedom to concurrently view information. Simultaneous viewing and interaction with data, text and geometric objects in various forms of visual representation such as 3D solids, wireframes and diagrams is supported. Dynamic imaging fidelity of such high quality can only be delivered by the most advanced VLSI technology integrated into a geometry based computing system.

2.1. Graphics Subsystem

As an application program executes, the RISC processor identifies high-level graphics commands contained in Silicon Graphics' comprehensive Graphics Library and sends them with their associated data to the Graphics Subsystem for display. Thus, the Graphics Subsystem is responsible for offloading the CPU from most graphics related activity. All the CPU has to do is to fire and forget the graphics command; the CPU immediately resumes its own activities.

The Graphics Subsystem of the IRIS 4D/70 is a complete, self-contained module designed from a systems perspective. It includes a number of proprietary VLSI processors, and resides on three triple-high by quad-wide VME boards. The Graphics Subsystem handles the transforming, rendering and displaying of graphical information. These operations are performed locally through the extensive use of pipelined custom VLSI processors, parallel finite-state processors, fast memory and a powerful 32-bit microprocessor. Conceptually, graphics information is processed by three distinct sections of the Graphics Subsystem. These sections are called the Geometry Subsystem, the Rendering Subsystem and the Display Subsystem.

2.1.1. Technological Advancements

The following list highlights the innovations in graphics technology represented by the IRIS 4D/70 Superworkstation:

- A proprietary Geometry Pre-Processor controls context switching for concurrent multiple windows.
- Two 4 x 4 matrix sections of the Geometry Pipeline operate in parallel to boost the throughput to 140,000 clipped 3D vector transformations per second.
- A 68020 microprocessor with 1 Mbyte private memory offloads the RISC CPU and acts as the manager of the Graphics Subsystem. It also converts concave or convex polygons into trapezoidal primitives.
- The Rendering Processor generates pixel addresses that lie between two points on a scan line, or uses Bresenham's algorithm to compute the pixels lying between two vertices.
- Four proprietary processors serve as iterating interpolators that iteratively interpolate the color intensities and depth values for each pixel in parallel.
- The frame buffer is partitioned into three functionally separate domains:

  The Image Planes contain 1280 by 1024 pixels with a minimum of 8 bits per pixel. They can be expanded to 24 bits per pixel.

  The Window Planes are 8 bits deep; they are dedicated to storing information pertaining to the multi-mode windowing environment and additional over and underlay planes for pop-up menus.

  Optional Depth Planes feature an increased accuracy of 24 bits per pixel. In conjunction with a depth interpolator, they rapidly perform hidden surface removal at hardware speeds.
- Proprietary Multi-Mode Graphics Processors read the content of the frame buffer in five parallel streams. They can display any combination of color-indexed or full color RGB images in either single- or double-buffer mode in up to 64 windows simultaneously.

2.2.1. Geometry Subsystem

The Geometry Subsystem interfaces with the RISC host, receives world coordinate data, performs the specified transformations and sends the screen space vertices off to the Rendering Processor. The major components of the Geometry Subsystem are the Geometry Pre-Processor, the Graphics Manager and a pipeline of 17 Geometry Engines, buffered on both sides by Geometry Accelerators.

Geometry Pre-Processor

The Geometry Pre-Processor (GP) intelligently interfaces the Graphics Subsystem with the RISC processor. It feeds coordinate data into the Geometry Pipeline, analyzes the graphics commands, and delineates the windowing contexts. A context is defined by all parameters necessary to describe the state of a single window process. Any necessary context-switching is accomplished by the GP. It communicates in a feedback loop with the Graphics Manager (GM), notifies the GM of context-switches, and prompts the GM to save the previous window's state. A new windowing process starts immediately. When the new context commands have been completed, the GP reactivates the pipeline with the stored context's parameters.

Geometry Pipeline

The data stream supplied to the pipeline is a high-level graphics instruction set mixed with coordinate data. The Geometry Pipeline's efficiency results from the ability to perform matrix transformations, clipping, and scaling functions that can be processed by fast pipelined hardware. The Geometry Engines (GE) are Silicon Graphics' custom 2.0 micron NMOS VLSI technology processors, operating at 10 MHz. Each GE contains four parallel 32-bit floating-point ALUs and a microcoded control store. Geometry Engines efficiently perform the crucial task of mapping the users' three-dimensional object space into

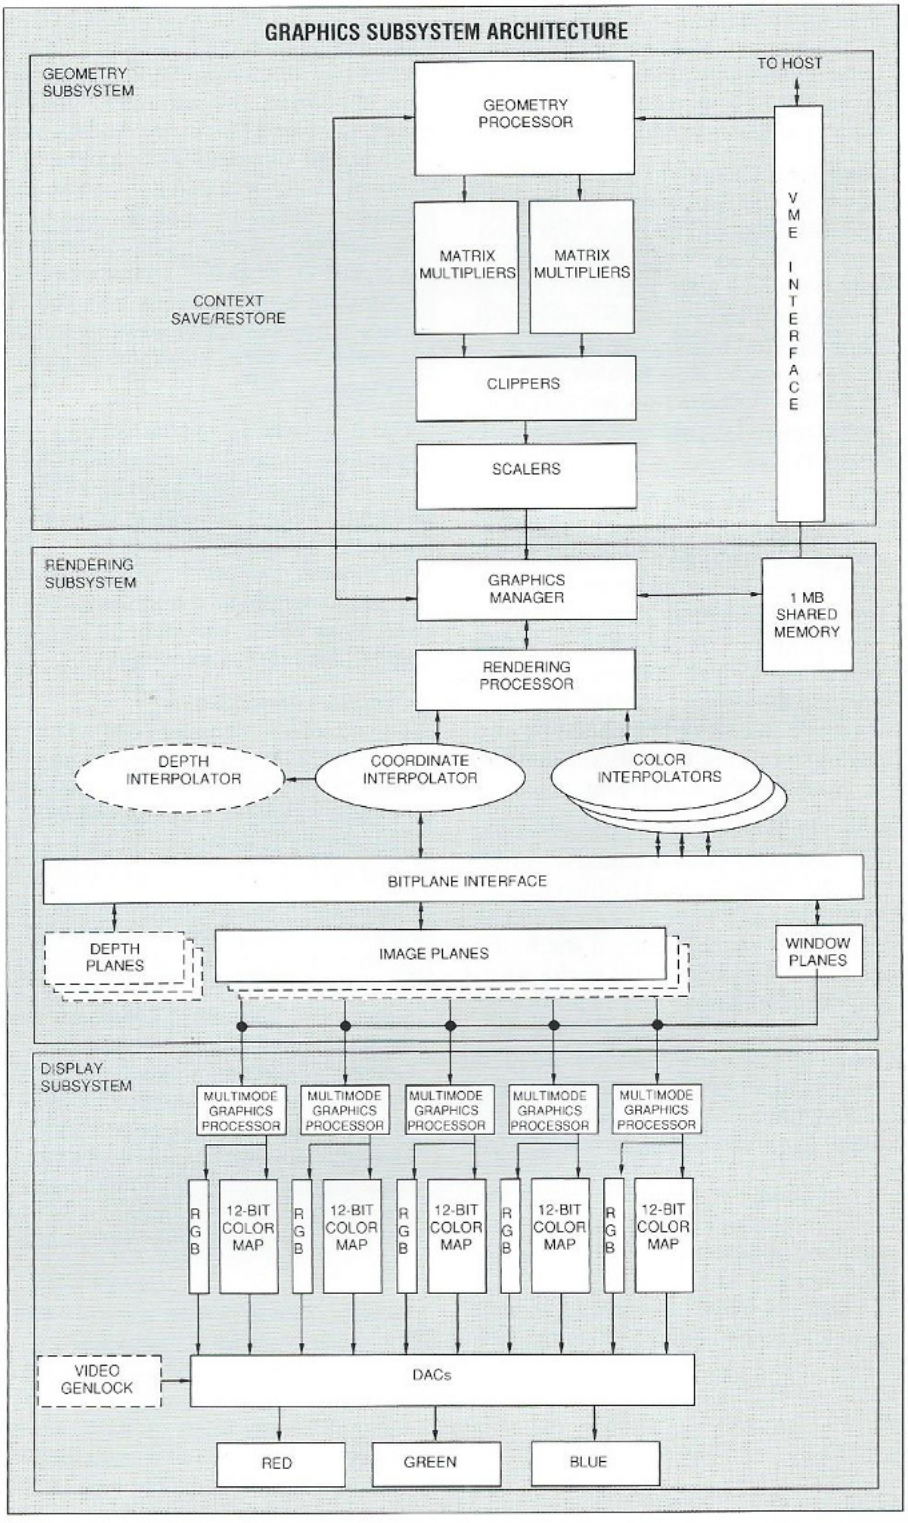

the two-dimensional screen space. They can also generate parametric cubic curves and rational bi-cubic surfaces in addition to transforming points, lines and polygons.

Seventeen GEs are arranged in a pipeline; the matrix section, originally consisting of four engines, has been doubled for additional performance gains of up to sixty percent. Since all instructions for matrix multiplication, clipping, and scaling are stored in ROM on each GE, a pointer into the ROM-based instruction set selects the functions to be performed by a particular GE.

The Geometry Accelerator (GA), another custom VLSI processor running at 10 MHz, provides First-In-First-Out (FIFO) buffering of data and high-speed floating-point conversion from IEEE or DEC floating-point format to GE format. Located at the head and the tail of the pipeline, the GAs increase system performance by supplying/receiving data at constant rates to/from the Geometry Engines.

The Matrix Section is represented by the first eight GEs, organized as two parallel pipes in the Geometry Pipeline. Each system of four GEs is capable of manipulating a stack of eight 4 x 4 matrices and performs the 3D transformations within a homogeneous coordinate system. The existence of this stack supports the hierarchical organization of graphic objects. It may be controlled directly with load, store, and multiply instructions.

The Clipping Section is composed of six GEs, where each GE clips against a specific bounding plane. This allows for clipping to the left, right, top, bottom, near and far planes. In hit-testing mode, the Clipping Section detects intersections between the cursor and graphic objects through the use of a special "hit-testing" matrix. This matrix typically corresponds to the cursors bounding box. If a hit occurs, the associated coordinate data passes through the pipe, otherwise it is suppressed.

The final two Geometry Engines constitute the Scaling Section that maps the coordinates received from the clipping subsystem into coordinates understood by the Rendering Subsystem. Perspective and/or intensity mapping of the Z-coordinates results in values that may be interpreted by the Display Processor.

Graphics Manager

The Graphics Manager (GM) consists of a 68020 processor, running at 16 MHz, with 1 Mbyte of associated local memory. This DRAM is shared with the RISC CPU via the VME bus, and is used to store local graphics data. The GM also stores the Graphics Library's command interpreter, window-contexts, communication protocols and other status data of the Graphics Subsystem. The 68020 organizes the logistics for all resources of the Graphics Subsystem and supervises the activities of the Geometry Pipeline. It communicates with the host processor and handles interrupts between the Geometry Subsystem and the CPU.

The Graphics Manager synchronizes the switching of contexts in cooperation with the Geometry Pre-Processor. When a window is only partially visible, display information must usually be clipped against non-rectangular boundaries. The GM breaks partially obscured windows into a series of rectangles, and stores this information in a table. This table reflects the current status of the window manager and is updated whenever a window is created or modified. Up to 64 separate windows can be maintained simultaneously. The GM also loads color maps, specifies cursor and icon positions to the Multi-Mode Graphics Processors, and passes requested pixel information from the image planes back to the application process on the host.

In addition, the Graphics Manager handles the important job of fracturing convex or concave polygons received from the Geometry Engines into simple trapezoids with ordered vertices. The trapezoids constitute the basic geometric units which the Rendering Processor uses to determine beginning and end points on each scan line.

2.2.2. Rendering Subsystem

The IRIS 4D/70 features custom-designed rendering engines to speed up the compute-intensive tasks of shading polygons and depth cuing lines. The optional Z-buffer is supported by an additional processor to rapidly solve the hidden surface problem.

Rendering Processor

The Rendering Processor (RP) consists of a finite state machine with 32 Kbytes (4K x 64 bits) of soft microcode, a 16-bit ALU, a 16-bit multiplier, tables for Modula 20 conversion, and reciprocal values for numbers up to 1279. The latter are used to quickly generate the addresses for the interleaved structured frame buffer. The RP receives the projected screen coordinates of the vertices of the trapezoids. Bresenham's algorithm is applied to calculate location and color of the pixels that best approximate the boundary lines of the trapezoidal polygon. This algorithm is particularly fast. Simple multiplications and additions determine the slope of the line and the delta value from the previous coordinate. The resulting scan-line endpoints are passed to the interpolators.

Iterating Interpolators

Each of these five processors is a proprietary VLSI chip. Three of these chips perform linear interpolations in parallel to generate the color of each pixel that lies between two points on a single scan-line. A second, similar semi-custom processor computes the frame buffer address of these same pixels in parallel. If the Z-buffer option is installed, an additional interpolation processor is used to compute the depth values of individual pixels within a single span. By employing the interpolators as rendering engines, the IRIS 4D/70 draws over 5500 Gouraud-shaded, Z-buffered polygons per second.

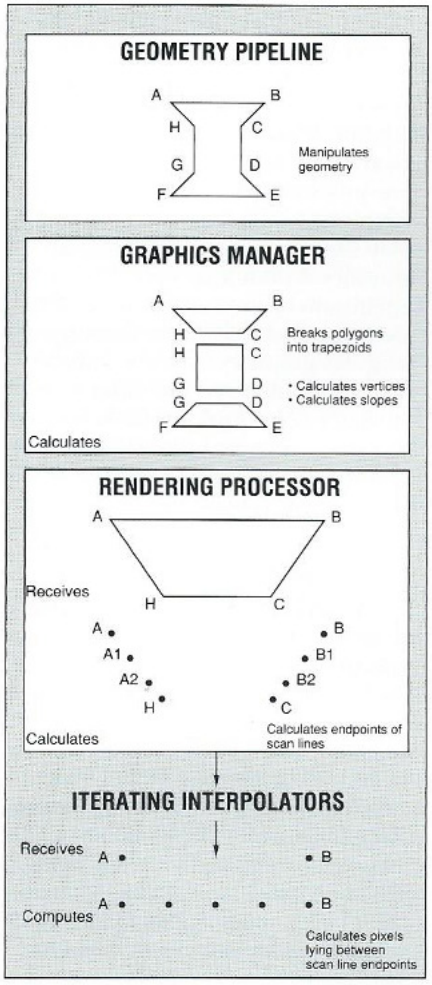

The data received from the interpolators is multiplexed by the Bitplane Interface, and used to update the appropriate sections of the frame buffer. The frame buffer consists of three functionally separate types of memory: Image Planes, Window Planes and Depth Planes.

Image Planes

This memory stores up to 24-bits of color information for each pixel of the 1280 x 1024 high-resolution display. The standard configuration is equipped with eight planes of color information and can be easily expanded by inserting extra Single Inline Memory Module (SIMM) plug-on memory boards. When the IRIS 4D/70 is fully configured, 24 Image Planes reside together with the Window Planes in five memory modules, which are each made of eight 64K x 4 bit Video RAM (VRAM) components. Thus, in the maximum configuration, 160 VRAMs are used to support 24 Image Planes and 8-bit Window Planes at a 1280 x 1024 resolution. All memory is dual-ported and configured in an interleaved fashion to allow for fast parallel read out.

Window Planes

These eight bitplanes hold the data to define the state of windows visible on the display. The data consists of window identification tags (4 bits) and overlay and underlay colors (4 bits). The window planes support the display of multiple 3D images in a concurrent display environment. Unlike the image planes that receive data from the interpolators, color values for the over and underlays are not depth-cued.

Depth Planes

Optional Depth Planes facilitate the Z-buffer method of hardware hidden surface removal. The Z-coordinate associated with each pixel is stored in 24-bit Depth Planes. When hidden surface removal is enabled, the Z-coordinate of each incoming pixel is compared to the current depth value already stored. If the new Z-coordinate is smaller than the old depth value, it is closer to the user's viewpoint, and therefore visible. In this instance, the new pixel data is used to update both the image planes and the Depth Planes at this location. Conversely, if the incoming Z-coordinate is larger, it must be farther away from the viewport. Therefore, the pixel must be hidden from view and is discarded.

2.2.3. Display Subsystem

The Display Subsystem receives pixel information from the frame buffer, routes it through the appropriate display mode, and sends it to the DACs for display.

Multi-Mode Graphics Processors
Five multiplexed Multi-Mode Graphics Processors (MGP) concurrently read the contents of the Image and Window Planes. Information from the Window Planes indicates in which color mode the pixel data currently being converted to video is displayed. The IRIS 4D/70 displays images simultaneously in single-buffered RGB, double-buffered RGB, single-buffered color-index, and double-buffered color-index color modes.

Made possible by massively parallel access to the frame buffer memory, Multi-mode Graphics Processors allow simultaneous display of up to 64 different windows in various color modes. The user gains the flexibility to display both static and dynamic imagery at will, without having to consider their color representation scheme. Combined with the context-switching capacity of the Geometry pre-processor, the Multi-Mode Graphics Processors enable the display of real-time smooth shaded graphics within a complex windowing environment for the first time. The MGPs also position and draw the rectangular graphics cursor or a full-screen crosshair cursor. A 64 x 28 bit auxiliary color look-up table for overlays and underlays is also implemented directly in these chips.

Digital-to-Analog Converters
High-speed Digital-to-Analog Converters (DACs) drive the red, green and blue guns of the display. When the Graphics Subsystem operates in RGB color mode, the DACs receive up to twenty-four bits of color information for each pixel. Eight of these bits are directly assigned to each of the red, green and blue DACs yielding more than 16 million colors. The DACs are multiplexed to handle the input from the five pixel pipes. When a particular window operates in color-index mode, pixel data packets are used as indices into a 12-bit in, 24-bit out color map before being sent to the DACs. This map defines 4096 simultaneously visible colors from a palette of 16.7 million. The pixel-mapping feature of the color-index mode allows screen colors to be quickly modified by simply changing the values stored in the color maps.

Color Interface
Besides controlling the interleaved timing of the parallel Multi-Mode Graphics Processors, the Color Interface controls the timing of the SYNC and BLANK signals that structure the scanning pattern of the display device. Three oscillators provide the appropriate timing pulses to the output devices. The first oscillator supports both, 60 Hz non-interlaced and 30 Hz interlaced, 1280 by 1024 display resolutions with a 107.352 MHz clock. The second oscillator provides a 12.272 MHz clock to support RS-170, which is used to produce NTSC video signals. The third oscillator supports the EURO standard with a 15 MHz clock.

An optional genlock board is available to synchronize the display of the IRIS 4D/70 with external video sources.

3. Computing Environment

3.1.1. Graphics Library

The IRIS 4D/70 contains an extensive Graphics Library (GL) that allows programmers to take full advantage of the technology provided by the graphics hardware. All commands are accessible through both high-level C and FORTRAN subroutine calls. The GL is designed to simplify development of highly interactive, real-time graphics applications. The GL commands can be categorized as commands for interaction, manipulation, or drawing and rendering.

Interaction
The Graphics Library gives the programmer all the necessary tools to create applications with a high degree of user interaction. Supported input devices include the mouse, keyboard, digitizing tablet, and dial & button box. Combined with the window manager and other sophisticated user interface techniques, developers can create a highly responsive computing environment that provides a customized interface to the application program. Programs can be written to provide immediate response to user input and to offer high levels of productivity.

Manipulation
The advantages of the Graphics System are fully attained through the use of a number of manipulation routines in the GL. Programmers can create their own object, world, and viewer coordinate systems and map them with orthographic or perspective projection to any sized viewport on the graphics screen. Objects can also be translated, rotated, scaled and displayed in real-time with the use of double-buffering commands. These commands allow direct control of the graphics hardware.

Drawing and Rendering
The Graphics Library provides primitives that allow the programmer to create graphic objects such as points, lines, arcs, circles, polygons, parametric curves, rational bi-cubic surface patches, and Gouraud shaded, Z-buffered solids. Complex objects can be easily created by invoking combinations of these commands. The user does not need to be concerned about the generation itself but can concentrate on building larger graphic scenes and interactively moving the objects. All of these commands can be called by a program and executed in immediate mode, or they can be compiled in a hierarchical display list to provide instancing capability.

3.1.2. Software Compatibility

The Graphics Library (GL) of the IRIS 4D/70 provides full source-level compatibility with the GL on previous Silicon Graphics superworkstations. The GL of the IRIS 4D/70 is a superset of the previous GL. Code that exists on IRIS 3000 workstations can be ported to the IRIS 4D/70 with minimal effort. New routines let the user take advantage of the gains in hardware technology that have been applied to the IRIS 4D/70's design. These new routines activate faster Gouraud shading, concave polygon handling, overlay and underlay bitplanes and a hardware cursor.

To provide optimal use of the new graphics technology, extensions have been made to some of the existing GL routines. For instance, the higher resolution of the image planes and the increased accuracy of the Z-buffer have been taken into account.

3.2. Operating System, Programming Tools, and Networks

An enhanced version of UNIX System V.3™ is available on the IRIS 4D/70, incorporating many features of the Berkeley 4.3 UNIX release. Additionally, the operating system includes many local system enhancements to support real-time graphics. Silicon Graphics has extended this software environment to permit easier connectivity and improved system throughput. Enhancements include the Extent File System (EFS), a unique UNIX file structure handler that provides a two-fold improvement in file handling performance over the standard System V architecture. EFS has been modified so that it writes files sequentially on the disk. This strategy lowers seek times and leads to a dramatic improvement in system performance.

Easier connectivity is achieved through a port of UNIX 4.3 BSD TCP/IP. This industry standard communications protocol is implemented on an Ethernet™ controller that is available as an option for the IRIS 4D/70. Silicon Graphics also supports the Network File System™ that supports transparent access to files across a network. Adherence to these industry standards permits easy integration of the IRIS 4D/70 Superworkstations into existing computing environments.

A superior program development environment is available on the IRIS 4D/70 Superworkstation through the System V.3 operating system and the IRIS *Edge*, a window-based graphical interface to DBX (a standard UNIX 4.3 BSD debugger). *Edge* provides an easy to use, menu oriented, graphical interface with automatic scrolling, color highlighting of source text, and the ability to recompile within the debugging environment. With *Edge*, the programmer can watch the source code as it executes in one window and view the results in a separate window.

3.3. Physical Device

The IRIS 4D/70 Superworkstation was designed with special attention to its physical functionality and appearance. The design team chose an innovative package that features two towers connected on a common base. The twin-tower arrangement is the result of board reduction, modularization of components and aesthetic considerations. As a result of this design effort, the IRIS 4D/70 resides comfortably in an office environment.

Thanks to VLSI circuit technology, Silicon Graphics engineers were able to squeeze more computing power into a relatively smaller space. One of the two towers contains a 12-slot card cage for the RISC-based CPU, the graphics subsystem and peripheral controllers. The second tower houses the power supply and up to four stacking storage peripheral modules, such as 170 Mbyte hard disks or streaming tape drives.

The chassis accommodates 12 slots of triple-high, full-depth VME cards. Spanning an area of 366.7 mm by 400 mm, these large cards provide over 3.5 times the usable board space of standard dual-height VME cards. This minimizes the interconnect wiring and related interface circuitry found in system designs partitioned across multiple smaller boards. With the more extensive use of high-speed on-board busses, potential performance losses due to lower speeds and higher traffic on the backplane bus are further diminished.

Silicon Graphics also provides an assortment of adaptor frames to configure third party VME cards into these slots.

The twin-tower packaging of the IRIS 4D/70 is designed for low maintenance cost by providing convenient access to the card cage. Disk and tape drives are housed in modules so that they may be latched externally onto the system unit during field installation or, at a later time, by the user. Different modules either directly plug together for SCSI peripherals, or connect through a cable for ESDI or QIC-02 interfaces. Expansion and service is easily facilitated; if required, the user can disengage latches and store the drive module in a secure location. Except for a small power indicator lamp, all controls and access panels are concealed behind easily accessible covers.

The system is powered with a 1000 W supply, the largest available for the standard 20A 115 VAC wall receptacle. This gives the user maximum flexibility in configuring this system with specific drive and card combinations and allows the system to be installed without special wiring.

Cooling in the card cage is accomplished via two 160 cfm (free airflow) brushless DC crossflow blowers selected for quietness and efficiency. The drive modules and the power supply are separately cooled. The twin-towers feature orthogonal air circulation to reduce thermal cross-loading.

The system unit is a full Faraday cage and is designed to meet UL 478, CSA 154, VDE 0871, and FCC Class A specifications.

On the desktop, the user faces a 19" Hitachi monitor mounted on a tilt and swivel base for ease and flexibility. The monitor features high-resolution, with 1280 x 1024 picture elements, and scans non-interlaced at 60 Hz for flicker free viewing. It is color balanced at 6550 degrees Kelvin to preserve a superb color fidelity.

The keyboard is a Keytronics model arranged in the standard IBM RT format. It complies with DIN 66234, one of the most demanding European human factor standards. The mouse is manufactured to Silicon Graphics' specifications by Mouse Systems; it operates optically at 200 cpi resolution and plugs into the side of the keyboard to reduce desktop cable clutter.

The system unit fits comfortably beneath a standard desk. The volume of the package is visually reduced by the twin-tower packaging. To accentuate the special identity of the IRIS 4D/70, Silicon Graphics' designers selected a new color palette. The machine's coating blends dark grey, raspberry and beige colors into a pleasing harmony.

The unique combination of form, function, color and performance adds up to a new look for a new technology: The IRIS 4D/70 embodies the principles of visual computing.

By fusing the most advanced RISC technology with proprietary VLSI graphics processors, a finely tuned superworkstation has been crafted that sets the standards for a new type of visual computing. The IRIS 4D/70 Superworkstation will assist professionals in a multitude of disciplines to comfortably solve their tasks by channeling the power of human imagination into visual representations that can be dynamically observed, analyzed and controlled. The speed of visualization teams up with the speed of mind in a dramatic race to engender unprecedented powers of productivity.

2011 Stierlin Road
Mountain View,
California 94043

SiliconGraphics
Computer Systems