Custom VLSI chips improve reliability and increase speed in this high-performance engineering workstation, which combines real-time 3-D color graphics with Unix software and Ethernet network communications.

The IRIS Workstation

Randy Nickel

Silicon Graphics, Inc.

Sophisticated engineering and scientific appplications depend increasingly on high-performance graphics as the primary interface to the user. By allowing the interactive visualization, examination, and analysis of ideas, these tools greatly enhance the creative efforts of professional engineers and scientists. At the same time, engineering workstations have grown in popularity due to their relatively good performance at reasonable cost in applications that are primarily computational in nature. With the combination of a powerful microprocessor, a Unix or Unix-like operating system, bit-mapped graphics, and integrated networking capabilities, these workstations form an attractive alternative to time-sharing on a mini or super-minicomputer.

Most currently available high-performance graphics terminals and workstations, however, use standard components that are designed to be general purpose in nature; often they are incapable of the performance necessary for more demanding applications. Specifically in the area of three-dimensional design, visual simulation or animation, or scientific applications such as molecular modeling, system performance can be inadequate.

A new generation of workstations addresses the performance and functional limitations of more conventional workstations by incorporating custom VLSI circuits in their design. The provision of special-purpose hardware—designed to replace the use of less efficient software running on general-purpose hardware—allows these new workstations to realize order-of-magnitude performance improvements.

Silicon Graphics, Inc., designed its IRIS (Integrated Raster Imaging System) workstation (shown in Figure 1) to combine real-time three-dimensional color graphics with Unix system software and Ethernet network communications. Custom VLSI circuits such as the patented Geometry Engine shown in Figure 2 increase the workstation's performance, reduce cost and power requirements, and increase reliability. The Geometry Engine specifically addresses the problem of creating and manipulating wire-frame and solid geometry in real time. In addition, Geometry Engines can be used for general-purpose computing at rates exceeding six million floating point operations per second.

System description

Conceptually, the IRIS is divided into three pipelined components: the CPU, the Geometry Engine subsystem, and the raster subsystem. The system's CPU, a Motorola 68000 or 68010, manages display lists, runs application programs, and controls the Geometry Engine and raster subsystems. The geometry subsystem provides 2-D and 3-D geometric processing with either 32-bit floating-point or integer formats. All transformations, clipping, and scaling with perspective calculations are performed in the GEs. Geometric primitives for drawing lines, polygons, characters, and parametric and rational cubic curves are supported by the geometry subsystem. The raster subsystem controls up to 24 bit planes of image memory, which can be used in either single- or double-buffered modes.

The IRIS is available in three principal configurations: workstation, terminal, or file server. The workstation can function as a stand-alone unit or communicate to a network of IRIS systems over the Ethernet system. Workstations execute applications software locally and support a complete programming environment. Terminal configurations are intended to be attached to one or multiple hosts over Ethernet. The terminal's host computer can be another IRIS workstation, a Vax, or other computer that exists on the network. RS-232-C communications are also possible. The file server is a nongraphic computer node that supplies data files to other computing nodes on the network.

0272-1716/84/0800-0030$01.00

The IRIS differs fundamentally from conventional bit-mapped graphics workstations because its graphic orientation is toward objects, or geometry, as opposed to the conventional bit-mapped, or image, orientation. As a consequence, manipulations or modifications of the viewed image on the IRIS actually change the geometry or its orientation in space. Conventional displays deal primarily with the pixel or image data, which must be regenerated from the database whenever changes in the viewing orientation are desired.

Geometry Engine. The Geometry Engine comprises four 32-bit floating-point ALUs and a single control store. It can be softly configured to perform one of three basic graphics functions: matrix multiplication, geometric clipping, and mapping to screen coordinates (scaling). Multiple GEs, organized so that each receives its input from the previous engine, can achieve very high processing speeds through parallel computation.

A typical configuration of such a pipeline is illustrated in Figure 3. The first four engines are configured as a 4x4-matrix multiplier that converts 32-bit 3-D user coordinates into a normalized eye space. These coordinates are then clipped to the eye space boundaries by the second set of engines. Finally, the clipped 3-D coordinates are mapped into screen space by the last two GEs. Perspective division is included in this mapping.

Other useful pipeline configurations support high-level operations such as object interference detection, surface shading and visibility computation, and modeling transformations. A single GE pipeline can be quickly reconfigured under control of an application program to perform any of these functions.

System software. IRIS system software includes both the Unix operating system and the IRIS graphics utilities. The networking software is integrated into the IRIS system software. The graphics library is a comprehensive set of user subroutines that gives the applications programmer high-level access to the IRIS hardware. The system allows the programmer to work in world-coordinate space. When configured as a terminal, IRIS applications programs residing on the host computer can be written in C, Fortran, or Pascal languages. When a command is issued to the terminal, the command is either executed immediately or compiled into a display list for later execution. Com-

Figure 1. The IRIS (Integrated Raster Imaging System) workstation.

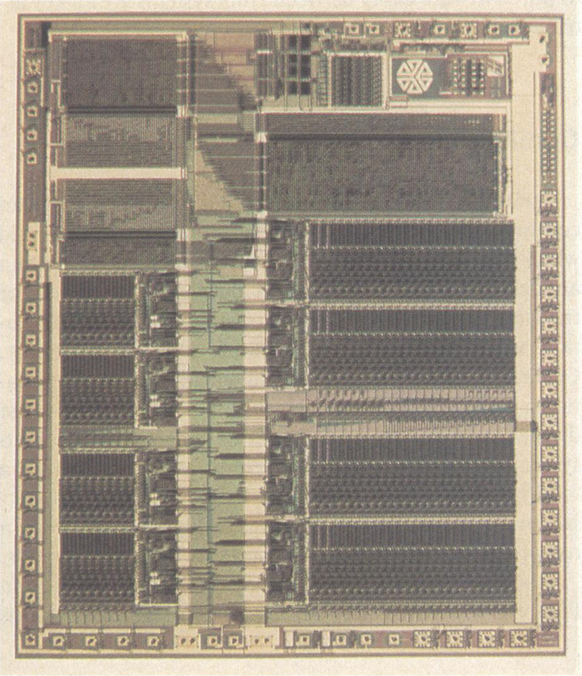

Figure 2. The Geometry Engine, Silicon Graphics, Inc.'s VLSI chip.

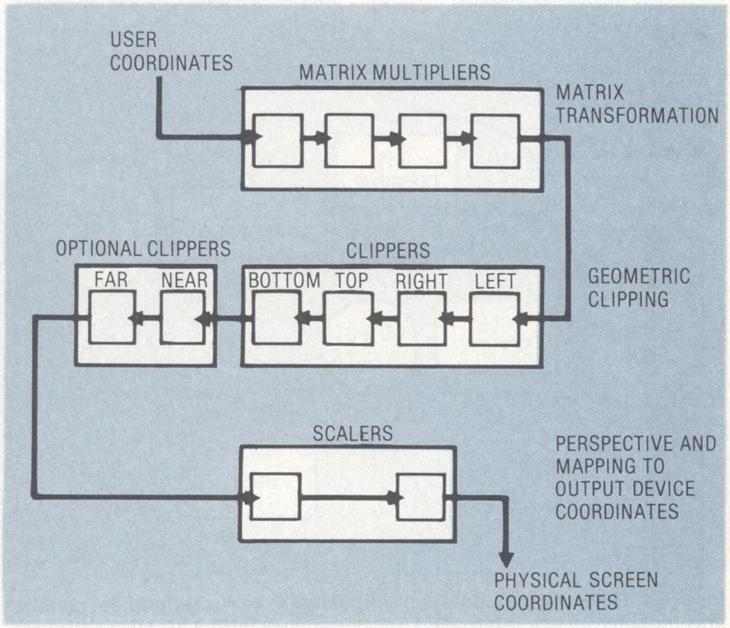

Figure 3. A diagram of the Geometry Engine pipeline.

piled mode offers the best performance as it allows local manipulation of geometry without intervention by the host computer. This capability allows more terminals to be attached to the host than otherwise would be possible. Workstation configurations function the same as terminals except that the application is local to the IRIS as well.

The virtual memory system software provides a powerful set of utilities and software development tools. Demand paging, multitasking, efficient filing, and the ability to port software easily between Unix systems are all key attributes in a Unix environment.

Configurations

Powerful workstations are utilized best when they can communicate efficiently to other computer resources and when they can function in a variety of modes. Certain applications by their nature will always be best suited for the large or mainframe class of computers. These applications will require intelligent terminals attached to the host. Other applications are better suited to smaller workstations, while yet other applications may evolve that are best divided between the two classes of machines. The IRIS system, designed to function either as a terminal or a workstation, can be upgraded from terminal to workstation, if so required. This capability also provides a reasonable transition path for users with changing needs.

Standard Ethernet networking (see the diagram in Figure 4) supports communications between IRIS systems and other systems on the network. This system also facilitates the integration of new technologies and products as they become available.

Workstation configuration. In a workstation configuration, the IRIS provides two 5¼-inch Winchester disks and a streaming-tape drive for local storage. In addition to the standard controller board, an SMD disk controller for high-speed disks can be installed. Both IP/TCP and XNS (Xerox Networking Software) protocols are supported by the Ethernet software and DMA network interface. Three of the four RS-232 ports are available for user peripherals, including an optional dial box and keypad combination. A 10-MHz 68010 running Unix fetches instructions and data with no wait states from 16M bytes of local memory. This dual-ported memory is mapped separately into both the processor and bus virtual spaces.

The 12-chip Geometry Engine pipeline is preceded and followed by 64-word custom FIFO's, allowing steady-state rates of up to 64,000 point transformations per second to be realized. A 16-bit bit-slice processor converts pipeline output into commands to the raster subsystem and also supports special operation modes, including hit testing, cursor tracking, and geometric feedback.

Dedicated MSI circuits update and display a 1024×768 image stored in 24 bits of $1K \times 1K$ bit-plane memory. Lines and rectangles can be drawn at steady-state rates of three and 46 million pixels per second. Characters, whose user-defined masks are stored in a local font memory, are drawn at the rectangle fill rate. Line and rectangle stipple patterns are supported by the hardware with no performance penalty.

A $4K \times 24$-bit color map allows up to 12 planes to be displayed simultaneously and supports both single- and double-buffer display modes. The color map can be bypassed, allowing all 24 planes to display eight bits each of red, green, and blue. Both interlaced and noninterlaced monitors are supported.

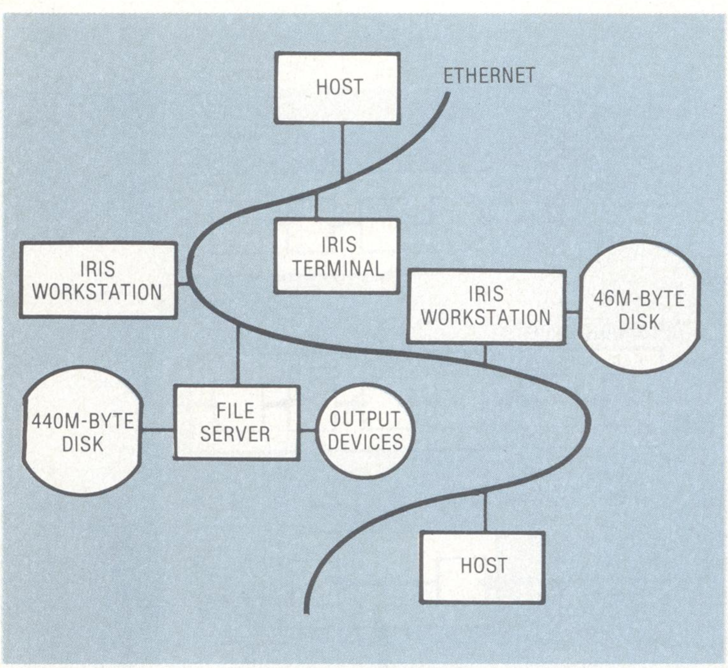

Figure 4. IRIS integrated networking. Hosts, workstations, terminals, and file servers are integrated into a network with a 10-MHz Ethernet system and standard Xerox Networking Software (XNS).

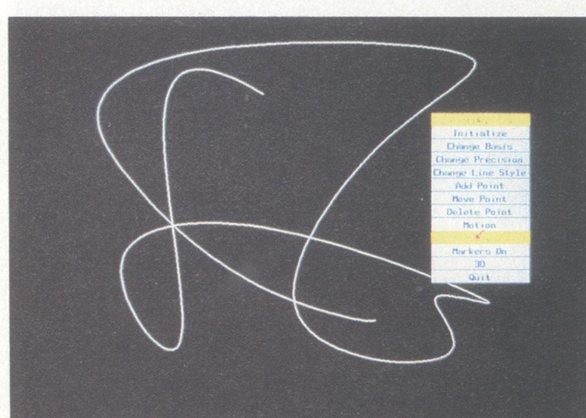

Figure 5. A spline curve generated on the IRIS workstation.

VLSI: Its impact on applications

The extensive use of custom processors in the IRIS workstation improves the performance of applications software through the sheer speed of the chips, and also through the improvement in the quality and efficiency of the code. An example of this improvement can be seen in the spline curve shown in Figure 5; it was generated on the IRIS from one curve command and a set of control points. In compiled mode, the curve can be modified locally through an I/O device such as a mouse. In a terminal configuration, no host intervention is required to regenerate the curve in either 2-D or 3-D space. The specific mathematical function used for the curve is user definable.

Constructing objects such as the turbine blade seen in Figure 6 is greatly simplified. The model can be generated parametrically using the curve and surface generation capabilities of the IRIS graphics library. Since all of the transformation, windowing, and scaling functions are provided by the IRIS hardware, the applications software does not need to be concerned with the complexities of maintaining the orientation and scaling of the geometry.

Two-dimensional applications such as IC design can also make use of the sophisticated capabilities of the IRIS architecture. Hierarchical display lists of arbitrary size are supported. Two-D scaling, rotation, and windowing operations are fast and accurate and require no host intervention. Silicon Graphics developed the schematic editor displayed in Figure 7 to run on the IRIS system for the VLSI circuits that will be used in the next generation of IRIS products.

Interactive solids modeling applications require complex computations that can exceed the capacity of even supermini machines to provide fast response. In these applications, powerful local processing is required. In the robot simulation example shown in Figure 8, the IRIS uses the Geometry Engines, to recalculate the position of the robot arm, to calculate and determine which surfaces

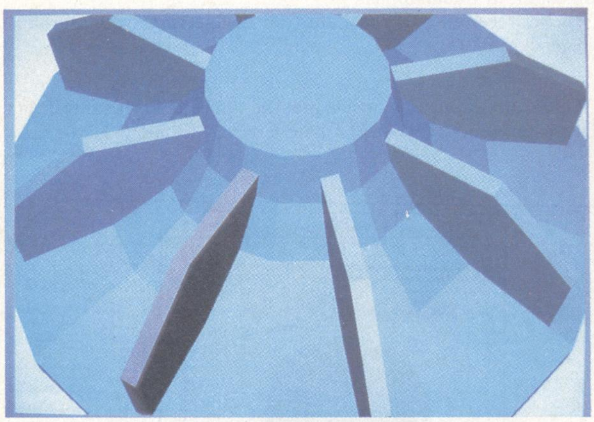

Figure 6. A turbine blade generated with the IRIS graphics library.

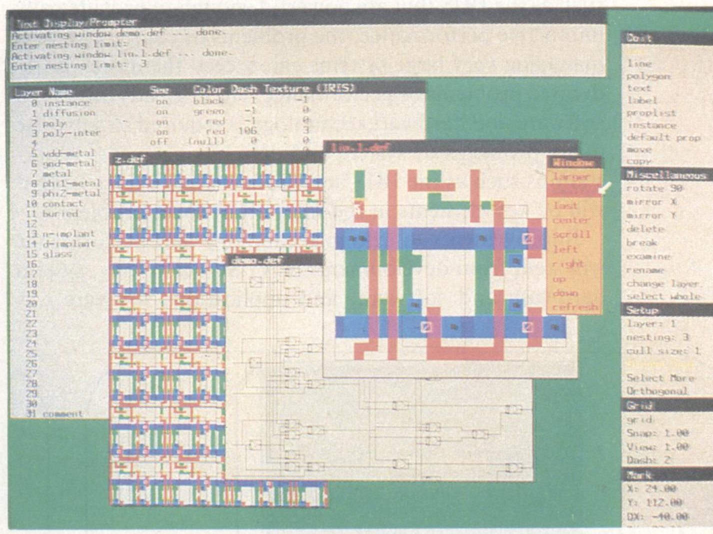

Figure 7. The IRIS schematic editor.

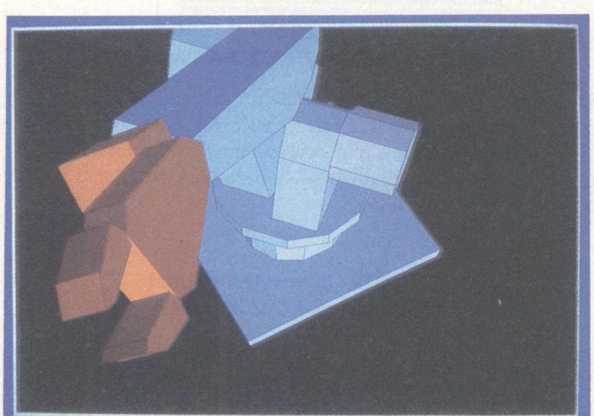

Figure 8. Robot simulated on the IRIS workstation.

Figure 9. Flight simulated on IRIS.

cannot be seen and should be removed, and finally to calculate the shading values of each side or face of the geometry.

An interesting use of the Geometry Engines is interference, or collision, detection, a process that is calculated by clipping the geometry of one object against the geometry of another. In the case of robot simulation, a "collision" between the robot arm and a fixture or perhaps another robot can be determined by the IRIS system. Interactive interference checking in piping applications could be performed in a similar manner.

Many applications can benefit from the ability to simulate motion in real time. Kinematic analysis, visual systems for flight simulators (see Figure 9), or animation of computer-generated images require that geometry be manipulated at real-time rates. In these applications, custom VLSI chips are the only means of providing the necessary performance in a workstation product.

With the advent of workstations and graphics systems such as the IRIS that are powerful enough to provide truly interactive performance, the problems of configuring and managing very large systems can exceed the more clearly defined problems of performance. Here, greatly improved software and hardware technology is required to solve the many problems of system and data management. Workstations such as the IRIS and products that evolve from it will be key elements in the solution of these problems.

Improvements in system technology and performance will stem from developments in VLSI technology, system hardware and software, and applications software. As workstations resulting from these technological advances become more widely available, users will find imaginative new ways to exploit their power. At the same time, chip manufacturing process technology will improve the performance of chip designs. VLSI circuits will be designed to address the specific needs of different applications. Workstations of the future will have basic system architectures augmented by the implementation of specialized processors that not only improve productivity at the user level but also address the complex functional and organizational problems of large distributed engineering and scientific systems. ■

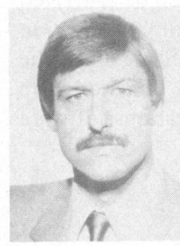

Randy Nickel is marketing manager for Silicon Graphics, Inc., Mountain View, California. Silicon Graphics markets workstations used in CAD/CAM, scientific research, real-time simulation, and animation applications. Previously, Nickel was product manager for small systems at Calma Company, where he was responsible for developing and marketing workstation and terminal products for CAD/CAM systems. Before joining Calma, he worked as a design engineer and consultant in Indianapolis.

Nickel received a BS in industrial engineering from Purdue University in 1973.

Questions about this article may be directed to the author at Silicon Graphics, Inc., 630 Clyde Court, Mountain View, CA 94043.