HP DELIVERS PHOTO REALISM ON AN INTERACTIVE SYSTEM

Realistic 3-d images can be viewed from any angle without recomputing the TurboSRX's algorithms

Leapfrogging the competition in graphics and computational performance is the aim of every graphics-work-station company, and the latest to jump ahead in the race is Hewlett-Packard Co. It made the leap by implementing its precision architecture processor on a fast new chip and introducing new graphics architecture and hardware. The Palo Alto, Calif., company has produced a work station performing 14 million instructions/s, with three-dimensional graphics that look as real as photographs. The system, which offers interactive capabilities as well, draws images up to 10 times faster than previous top-of-the-line HP work stations.

HP's technical computer group in Ft. Collins, Colo., is introducing the model 835 TurboSRX as the new high end of its HP 9000 work-station line. The 835 refers to the new 14-mips precision-architecture processor; the new high-performance graphics component is called TurboSRX. A second model, the HP 9000 350 TurboSRX, mates the new graphics subsystem with HP's MC68020-based work station. Prices start at $91,500 and $70,000, respectively—comparable to other high-end graphics work stations.

The TurboSRX systems are the first to offer a new hardware-assisted rendering technique called radiosity, which models the way light reflects between all surfaces of a displayed image. The result is a rendered object, with the photorealism of ray tracing, which can be viewed from any angle. Ray tracing, which the TurboSRX offers as a software utility, requires a complete rerendering of the object for each new view—radiosity does not.

These work stations are also the first to offer sixth-order nonuniform rational B-splines. Affectionately called Nurbs in the graphics world, these B-splines deliver faster graphics interactivity on the HP work stations than any competitive system can provide, the HP designers say.

A B-spline is the mathematical equivalent of a set of draftsman's French curves, which allow a designer to create complex curves—the wing of an airplane, for example—by fitting simpler curves together. The sixth-order Nurbs that HP is offering on the TurboSRX systems produce greater continuity in joining simpler curves together and therefore more realistic-looking surfaces than the fourth-order Nurbs offered on previous HP graphics work stations.

The model 350 TurboSRX and model 835 TurboSRX are especially good for such scientific and design applications as molecular modeling, solids modelling, animation, and imaging. The model 835 system, which is scheduled to ship in October, will compete directly with high-end graphics work stations and desk-top supercomputers offered by Apollo, Ardent, Raster Technologies, Silicon Graphics, and Stellar (see p. 97). In the midrange of performance, the model 350 TurboSRX, with the same graphics performance but lower computational performance than the 835, will compete with products from Apollo, Silicon Graphics, Sun Microsystems, and Tektronix. The 350 TurboSRX will ship in late June.

One distinction of the TurboSRX systems from competitive offerings is their ability to implement and accelerate radiosity algorithms in hardware. Radiosity is a new alternative to ray tracing. Although it requires a lot of computing power, ray tracing has been used since the late 1970s to provide photorealistic images of solid-object models. Equally realistic images produced by the radiosity technique can be manipulated interactively, something impossible with ray-tracing.

The radiosity algorithms were developed by students at Cornell University under the guidance of Donald Greenberg, director of the computer graphics program and made available to HP as part of a cooperative

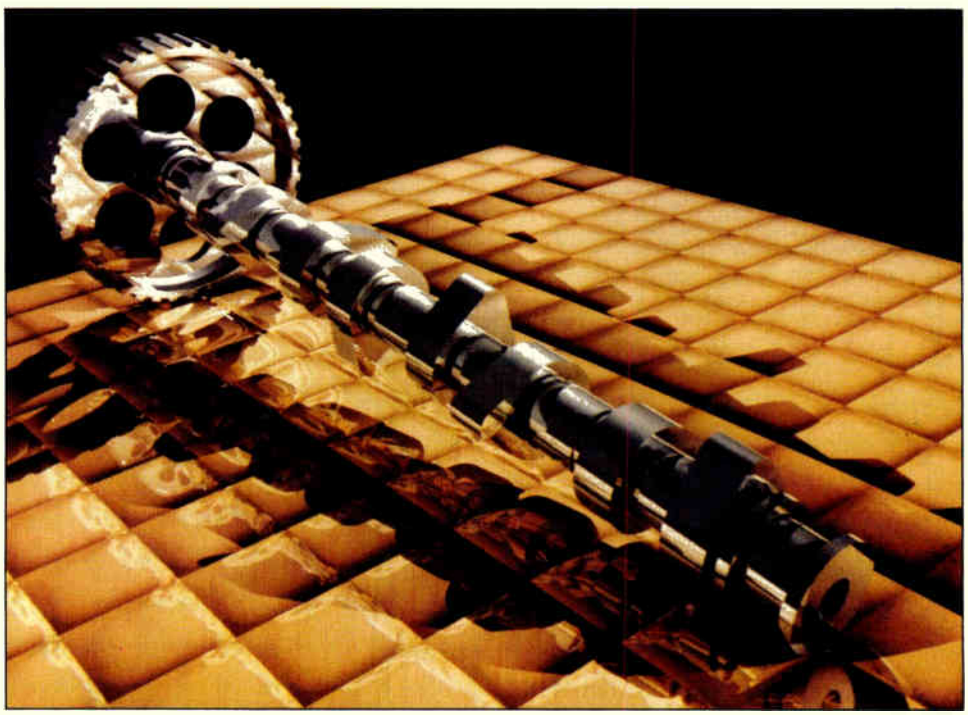

1. This image of a camshaft from an HP 835 TurboSRX shows off the system's photograph-like realism and provides an example of its ability to render curved surfaces with complex lighting and shading.

technology exchange agreement. The algorithms are based on fundamental energy equilibrium methods used in thermal engineering. Light leaving a surface—the radiosity of the surface—consists of light emitted, reflected, or transmitted through the surface. The radiosity method models the interaction of light travelling between these reflecting surfaces. The radiosity algorithms are accelerated with proprietary hardware in the TurboSRX graphics subsystem. To calculate the light relationships between all polygonal surfaces, the graphics accelerator "creates a set of simultaneous equations describing all possible combinations of light between all the polygons in a given scene," says Harry Baeverstad, R&D section manager at HP.

Ray tracing and radiosity, though producing excellent rendered images, are both very computational-intensive operations, especially if the scene contains large numbers of reflective surfaces. An automobile surrounded on three sides by mirrors, for example, could take over a week to render on a graphics work station. However, the great advantage of an image rendered with the radiosity technique is that once the radiosity of a scene is computed, the viewer can rotate, zoom, or view the object from any angle, in other words, interact with it in real time. That is impossible with ray tracing because each movement of the object requires a complete recalculation of the ray trace.

By comparison, in the radiosity technique, light reflections will not change based on the viewing angle. "While you can view a scene from different angles in a radiosity environment, you cannot change the geometric relationship of objects in the scene because that changes how the light interacts between the objects," Baeverstad says.

"The main advantages of our implementation [which includes hardware-assisted radiosity and software-implemeted ray-tracing] is that it is totally compatible with Starbase, our graphics library interface," he explains. "Any display list used to render an object on the TurboSRX can also use either the ray tracer or the radiosity functions without having to change its data base at all."

Instead of having to perform ray tracing in software or write a special applications program to provide radiosity, an applications software package such as Geomod from Structural Dynamics Research Corp., a General Electric Co. subsidiary in Cincinnati, Ohio, a user only needs to make a call to the Starbase library for either of the two rendering functions, ray tracing or radiosity. The system automatically does the rest.

Just as radiosity permits the user to interact with a fully rendered image in real time, HP's B-spline capability allows the user to interact in real time while creating an image. A B-spline is the mathematical equivalent of a draftsman's French curve. To create a complex curve such as the profile of an airplane wing using French curves, a draftsman selects different curves from a template and traces the curves together to form the desired shape. To create the complex curve mathematically in a computer graphics system, applications software uses different polynomials, each of which defines a unique curve, or B-spline.

The software places these B-splines together to form a final curve and displays the result on a CRT screen just as the draftsman selects French curves from his template and places them together. This process is termed a piecewise approximation of a curve. To create a surface—such as that of an airplane wing—the B-spline is replicated along the *z* axis. The higher the order of the polynomial used in each B-spline, the higher is the degree of continuity between the joints of different B-splines.

A nonuniform rational B-spline enables an application program to specify the joints, where one curve meets another, to be nonlinearly spaced. The capability allows the application program to break a B-spline into even smaller pieces. For example, in the design of an automobile trunk, one B-spline defines the curve sloping downward toward the bumper. To add a wind spoiler to the trunk, a nonuniform B-spline jutting upward where the trunk would normally curve downward to the bumper is inserted in the bumper's B-spline curve.

Finally, a rational B-spline surface is one expressed as a ratio of two polynomials. "Without this capability, the system cannot produce true conic sections—cylinders, spheres, cones, etc.," says Baeverstad. Using the HP system, a designer could not create a model of an automobile cam shaft, for example, without using a rational B-spline (see fig. 1).

HP's original SRX was the first system with B-spline functionality in its graphics engine. It represents surfaces as polynomial equations rather than as a collection of polygons which approximate the surface. Before the SRX, an applications program, such as a solids-modeling package like Geomod from SRDC, had to translate the B-spline representation in its data base into a representation which can be displayed with polygons on the CRT screen.

Not only did this require much computing power and a large memory to contain the polygon representation, but it also meant that the user could not

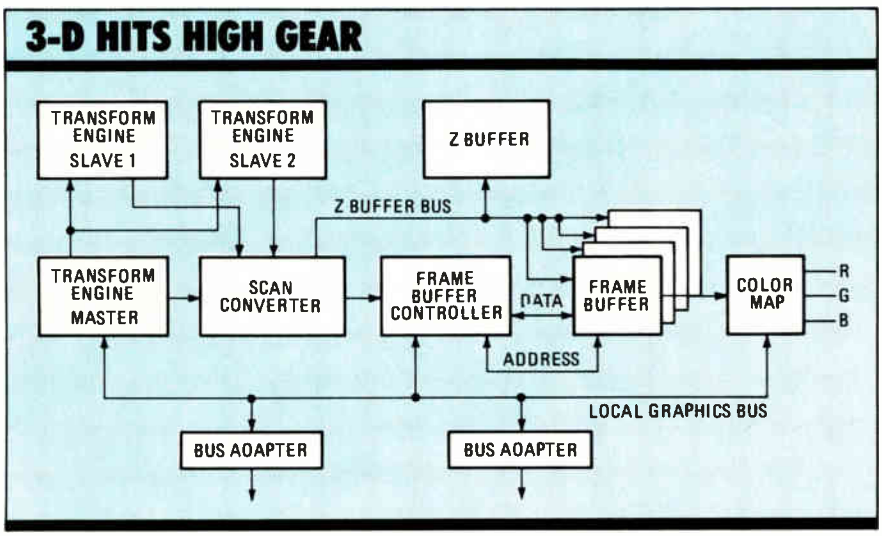

2. HP's new TurboSRX graphics processor contains the graphics transform engine, scan converter, Z buffer, frame buffer controller, frame buffer, and color map for fast solids modeling.

interact with the system in real time. Now applications programs using B-spline representation can command the TurboSRX to draw B-splines.

For example, a program to create a sphere with a polygon every 30 seconds of arc requires 9.3 Mbytes of memory and 35.5 seconds of computing time on a work station without B-splines. By contrast, the TurboSRX with B-splines takes 792 bytes of memory and 4.62 seconds of computing time.

A graphics processor containing the graphics transform engine, scan converter, frame buffer, Z buffer, frame-buffer controller, and color map is the main element of a Turbo SRX (see fig. 2). The processor is a separate assembly that fits into the desk-high cabinet of a HP9000 model 350 or model 835 work station.

During operation, the central processor of the associated HP9000 work station sends segments of a graphics image display list—a polygon, a line, or other primitive—through the bus adapter, over the local graphics bus, into the master transform engine of the graphics processor. If the master is busy with a previous segment, it passes the incoming segment to one of the two slave transform engines, both of which operate in parallel with the master.

After the CPU processes the image using world coordinates, the transform engines convert the image to the 1,280-by-1,024-pixel display-coordinate system. These engines also perform all the computation for B-spline, light source, and color. Each of the three transform engines is a new HP-designed 177,000-transistor custom chip called Treis (transform engine in silicon). The Treis chips are implemented in HP's own n-MOS III process.

As a graphic element comes into the transform engine, it is processed and sent to the scan converter, which performs six-axis interpolation of the graphic primitive or polygon. Where the transform engine specifies a graphic primitive in screen coordinate space, the scan converter specifies the pixel needed to represent the primitive on screen in x, y, and z coordinates—three of the six axes of interpolation. The other three are red, green, and blue colors, which the scan converter also interpolates; for example, the different colors on intermediate pixels resulting from a red pixel on one edge of a polygon and a blue pixel on another.

The frame buffer controller routes the pixels out of the scan converter into the correct frame buffer bank. At the same time, the scan converter routes z-axis data to the Z buffer. As the frame buffer controller draws pixel data into the frame buffer, it uses the Z-buffer data to determine the relative position of one surface with another to remove automatically one surface hidden by another. The frame buffer also controls data transfer's timing from the color map to the CRT display. The system comes with a new color-map chip developed by HP and implemented in its own 108-MHz, 1-μm CMOS gate array. Four of these color-map chips are used in the system. Each chip accepts 8 bits in and provides a 24-bit digital-to-analog conversion, which emits an analog red, green, or blue signal. *–Jonah McLeod*

For more information, circle 486 on the reader service card.

TECHNOLOGY TO WATCH **COMPUTER GRAPHICS**

STELLAR'S GRAPHICS MACHINE STRESSES INTERACTIVITY

The single-user supercomputer generates real-time dynamic graphics at 500,000 3-d vectors/s

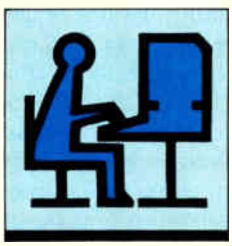

One of the worst-kept secrets in recent years, the GS1000 graphics supercomputer from Stellar Computer, Inc., made its official debut this week. The long-rumored system [*Electronics*, Feb. 18, 1988, p. 22] combines on a desktop the power of a minisupercomputer with high-performance real-time 3-d graphics—and all for a price of about $100,000.

The system executes more than 20 million instructions/s and can perform as many as 40 million double-precision floating-point operations/s. In the graphics realm, the important numbers are its ability to draw 500,000 3-d vectors/s and display 150,000 Gouraud-shaded triangles/s. That parlay gives the GS1000 computational performance equal to or better than that of early minisupercomputers combined with graphics performance that tops today's best work stations.

The GS1000's multiple-processor architecture hinges on a superwide 512-bit data path. More than half of 61 high-density CMOS application-specific integrated circuits implemented in the system are found in the extrawide data paths. Another key feature is the multistream processor (see fig. 1), which has four full internal streams or instruction paths and consists of a three-chip set. On top of its graphics performance, the system executes more than 20 million instructions/s

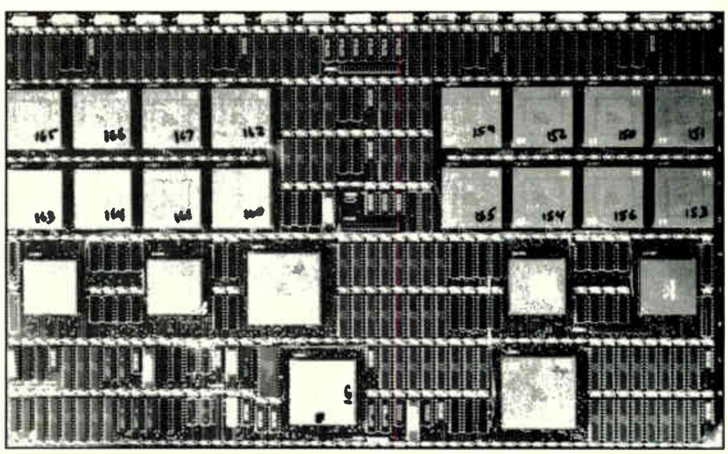

1. This multistream processor board from the Stellar GS1000 graphics supercomputer uses many of the big ASICs developed for this system.