

AT&T GOES TO 'WARP SPEED' WITH ITS GRAPHICS ENGINE

It can draw and render 3-d images 10 times faster than any competing graphics accelerator, thanks to a parallel architecture that speeds access to the frame buffer

by Stan Runyon



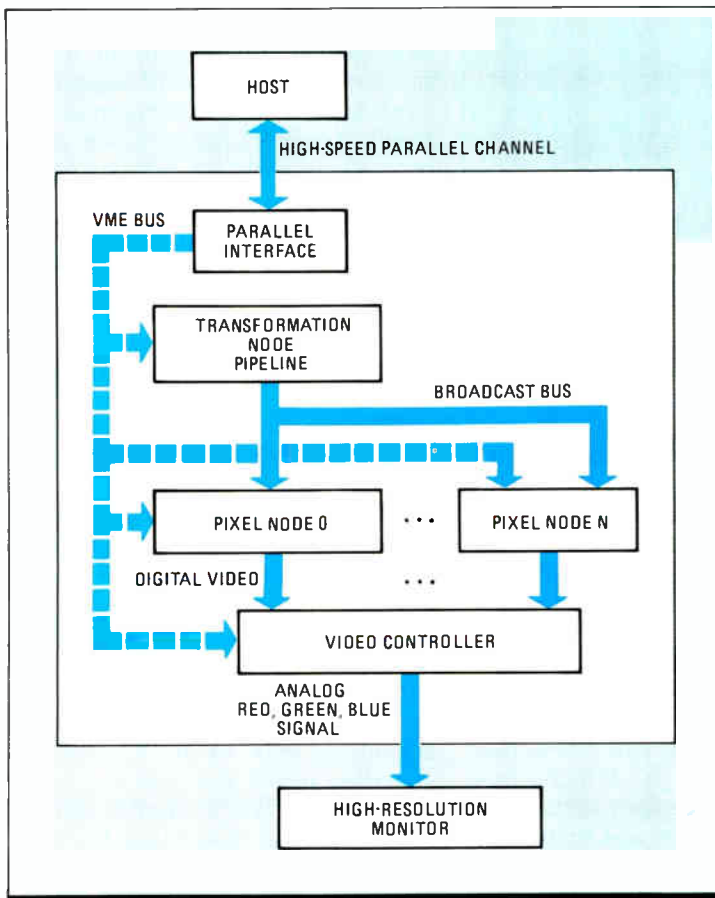
A new high in high-performance, three-dimensional graphics processing is on the way from AT&T Pixel Machines. The Holmdel, N.J. unit of AT&T Co. says its PXM 900 series graphics engine can draw and render 3-d images 10 times faster than any competing graphics accelerator.

The top-of-the-line PXM 964 zips along at 820 million floating-point operations/s. Put another way, it can transform 200,000 vectors/s and display 16 million Gouraud-shaded pixels each second. It is so fast, in fact, that it manipulates and shades color images as fast as any competing graphics work station rendering just wire-frame images.

To achieve the impressive speed, the PXM 900's designers went to a parallel architecture. Their goal was to eliminate one of the major bottlenecks in graphics displays—accessing the frame buffer. The parallel architecture consists of up to 64 distributed digital signal processor chips that take the place of the typical single processor found in other graphics engines. Each of these processors addresses its own portion of the frame buffer, which greatly speeds up access. Distributed addressing also allows for a hefty 48 megabytes of frame buffer—and the larger the buffer, the more realistic the image.

These distributed processors are AT&T's DSP32, a programmable 32-bit floating-point DSP chip. There are from 16 to 64 of them in the different models of the PXM 900 line. In effect, they serve as a multiple-instruction, multiple-data-stream machine—each runs the same algorithm. But the algorithms, which are resident in each processor, run independently, and the data varies from processor to processor.

The PXM 900 is modular: users can add more processors and expand the frame buffer in order



1. PARALLEL ARCHITECTURE. Independent pixel nodes drastically speed up image drawing and rendering.

to boost speed or produce better images. Software compatibility is preserved across the different configurations. Users also can program the machine to match particular applications and so run a variety of algorithms.

Unusual is the 900's ability to handle both image transformations and graphics, or pixel, operations. Typically, the host computer performs the image-processing transformations.

The PXM 900 series is linked by a VMEbus to a host work station, such as a Sun series 3/100 or 3/200. Its speed will suit it to applications involving animation, simulation, modeling, CAD/CAM, and scientific research. AT&T Pixel Machines will introduce the new line at Siggraph '87, July 27 through 30, in Anaheim, Calif. It will go on sale in the fourth quarter at prices ranging from \$45,000 to \$100,000.

Display resolution can go as high as 1,280 by 1,024 pixels, at 32 bits per pixel, in the fully-configured PXM 964 displaying 16 million Gouraud-shaded pixels/s. (Gouraud shading is a widely used algorithm that interpolates light intensities of adjacent pixels on polygon-based images.)

Running the more sophisticated Phong-shading algorithm slows the PXM 964 to 2.25 million pixels/s, but produces more realistic images. The Phong algorithm derives its shading information, not only from the edges as does Gouraud shading, but also from across the surfaces of the polygons themselves for far more realistic images.

Another example of the 964's blindingly fast performance is its ray-tracing speed of 1 million intersections each second. To trace the same number of rays, a VAX minicomputer needs hours; a Cray supercomputer, minutes.

The parallel architecture of the PXM 900 series is a sharp departure from the configuration of present-day graphics engines. Some accelerators use a single processor—but one processor working alone soon runs out of steam because it must handle too much data in too little time: 8 Mbytes for a double-buffered, 1-Kbit-by-1-Kbit, 32-bit/pixel display. Other engines use a series of processors, with each taking care of its own operation, such as drawing lines or moving blocks of data. Performance is usually better than in a single-processor machine; however, processing still proceeds sequentially and that limits maximum speed. Also, the specialized processors of this setup cannot run one another's operations.

The Pixel Machine approach (see fig. 1) attacks the bottleneck with its distributed processing method. Graphics information from the host is fed to a processor pipeline, called the transformation pipeline, which manipulates the data and passes it to an array of processing elements,

or pixel nodes. Each processor in the array is dedicated to managing a portion of a very large frame buffer, 48 Mbytes in the largest configuration, and each connects to its buffer portion via its own bus. Once the nodes have worked over the data, they send it to the video controller through a high-speed backplane, or bus, called the pixel funnel.


The processors in the pipeline and in the nodes are AT&T's DSP32, executing instructions at 10 megaflops. AT&T is just launching an enhanced version of the DSP32 (see p. 00), which should considerably boost the speed of the PXM 900. With individual programmable processors working over individual buses, bandwidth is stretched; there is a better balance between memory and processing; and a variety of algorithms can be implemented directly in the frame buffer—there is no need for a separate floating-point number cruncher. The result is improved efficiency, higher throughput, and greater interactivity.

The transformation pipeline accepts high-level geometric primitives from the host and computes all necessary transformations for manipulating an object in three dimensions—rotation, clipping, scaling, shading, and the like. It maps 3-d object space to 2-d screen space and broadcasts the results to the pixel nodes.

The pipeline consists of either 9 or 18 DSP32 processing elements, connected sequentially. Each processor performs its function and passes the data to the next processor.

The pipeline processors are mounted on a VMEbus-based board, nine to a board. Two pipeline boards can be connected and configured for serial or parallel operation under software control. Since each processor is programmable, the user can reconfigure the pipeline for various functions, such as rotation, scaling, or translation of the image.

Taking the data from the pipeline, the pixel



2. PIXEL NODES. Each pixel node is designed around a signal-processing chip which can be reprogrammed to run various algorithms.

nodes (see fig. 2) compute pixel colors and intensity, with each node controlling its portion of the bank of video random-access memories that make up the frame buffer.

One of the key innovations in the PXM 900 is the pixel interleaving scheme through which the sharing of the frame buffer takes place. This scheme achieves uniform load balancing among the processors by arranging the DSP32 chips in an n by m array, with each processor in the array controlling the n th pixel on the m th scan line. The result: speed increases almost linearly with the number of processors.

The pixel-node array can address up to 32 Mbytes of self-contained frame-buffer memory and 16 Mbytes of off-screen (external) memory. The basic configuration, the PXM 916, contains 16 pixel nodes, with 96 bits per pixel—64 for a double-frame buffer and 32 for a Z buffer, which stores each pixel's screen depth (Z value).

Because the parallel architecture speeds up access to the frame buffer, the 900's designers could implement the deep, 32-Mbyte buffer without sacrificing update speed. With 32 bits per pixel, 24 can be devoted to color, and 8 to overlays, all consuming just 4 megabytes of the buffer.

The Z buffer improves the image even more by contributing a three-dimensional effect to each pixel. Again, because of the buffer's great speed, each pixel can carry 32 bits of Z buffering instead of the usual 8 or 16. The advantage is realized with images in which the edges of objects overlap. Z buffering consumes another 4 megabytes of the frame buffer.

Yet another benefit of the buffer's size is improved animation speed. Instead of the conventional double buffering scheme, the 900 has the capacity to step up to quadruple buffering in the top three models. With double buffering, as the video controller removes data from one frame buffer for display, the display processor can begin to work on the other buffer, getting ready for the next screen update. But once the processor finishes with the backup buffer, it has to wait until the current buffer empties—up to 15 ms at a 60-ns refresh rate. Triple or quadruple buffering eliminates the idle time; the processor can go right to the third and fourth buffers.

Another key element is the PXM 900's pixel funnel. Its job is to accept the information from the pixel nodes' frame buffers, multiplex it, and pass the stream to the video controller. This is no mean feat because of the speeds at which the information must be concentrated. Since the pixel-node array is reconfigurable, so must be the funnel. This is done through software.

The video controller, the last stop for data before the color monitor, receives 32 bits of information from the funnel—8 bits for each of the primary colors and 8 for an alphanumeric overlay. It delivers a 30-bit video signal. With 30 bits, dynamic range is wider than that provided by standard 24-bit video, and the picture quality is better. □

For more information, circle 480 on the reader service card.

AT&T RACED THE CLOCK WITH THE PXM 900

For the designers of the AT&T Pixel Machines' PXM 900 graphics display engine, it was a race against the clock—and a risky race, at that. The goal was a showing at Siggraph '87, the premiere graphics conference. By the time the design team was in place, the July conference was less than a year away.

"The innovative design involved in the parallel architecture—and especially the pixel node and funnel designs—involved considerable risk for the designers, given the relatively short time we had to produce a working product," says Alessandro Piol, marketing director for Pixel Machines. "We were not sure whether this idea would work or not."

AT&T had assembled a diverse design team, drawing from various research and development activities at Bell Labs. James Conant, 33, for instance, Pixel Machines' director of engineering and one of the two leaders of the team responsible for the PXM 900's development, ran a VLSI system design group in Bell Labs' Government Systems Division. The other principal, Michael Potmesil, 34, code-signer of the PXM 900 architecture with Leonard McMillan, 26, came out of the Bell Labs Computer Sys-

tems and Robotics Research Lab. Potmesil has spent the last 12 years working on computer graphics. McMillan and others hailed from Bell Labs' development group for digital signal processors. Before that, McMillan worked on advanced parallel machines at Georgia Tech.

The plan for the PXM 900 was to leverage the speed and power of AT&T's DSP32 signal processing chip into an advanced graphics system. Naturally, the chip's designers were the best source of knowledge for that. As for graphics knowledge, there could be no better fountain than Bell Labs.

Luckily—or cannily—designers on the graphics side had experience in implementing algorithms on parallel machines. Others had been involved with hypercubes, software, image processing, and the like. Marc Howard, 30, the creator of the Pixel Funnel, had 13 years of computer graphics under his belt, and Eric Hoffert knew all about video animation: his work has been displayed in the Museum of Modern Art in New York, and elsewhere.


Did AT&T win its race against the clock? A demonstration at Siggraph, July 27 through 30, in Anaheim, Calif., should tell.



JAMES CONANT



LEONARD McMILLAN



MICHAEL POTMESIL