

# New 3D Engines Redefine the Market

## HP Reaches New Heights While ART Creates Single-Chip Ray Tracer

by Peter N. Glaskowsky

Siggraph always provides the most visual interest of any trade show of the year, along with the most interesting technical information on trends and products in the 3D market. Siggraph 97 was no exception. The most significant introductions at this year's show include HP's PxFI PixelFlow system, with its radical new architecture for high-end 3D rendering, and a single-chip ray-tracing engine from UK startup Advanced Rendering Technology.

While neither of these products will appear in personal computers anytime soon, the rapid rate of evolution in the PC 3D market should bring this technology to the desktop within just three or four years. Processor and system architects need to start preparing now for the radical changes we expect to see in PC graphics.

### PixelFlow Sets New Records for 3D Throughput

Said to be the fastest 3D-rendering engine in the world, HP's Visualize PxFI was the most impressive new product shown at Siggraph 97. Based on the PixelFlow architecture developed at the University of North Carolina in Chapel Hill, the PxFI scales from 4 to 54 rendering units and ranges in price up to more than \$2 million. A fully-configured system offers incredible performance: more than 250 GFLOPS of floating-point power, 12.8 Gbytes/s of internal bandwidth, a peak rendering rate of 100 Mpolygons/s, and a fill rate of 675 Mpixels/s for textured anti-aliased pixels or an astounding 53 Gpixels/s for flat-shaded 500-pixel triangles.

It's worth pointing out that these performance measurements are not directly comparable to those of mainstream 3D chips, which may appear to be faster in some ways than each of the PxFI's rendering units. The simpler rendering engines used in PCs are designed for peak performance on scenes consisting of about 10,000 polygons. The PxFI is designed for interactive display of CAD models with millions of polygons, a task no PC-based solution can handle.

The PixelFlow architecture is the latest in a series of advanced designs from UNC Chapel Hill. PixelFlow assigns each pixel on the screen to its own pixel processor, and all processors render the scene in parallel using a SIMD instruction-dispatch technique. Even greater parallelism is achieved by providing multiple sets of pixel processors, called flow units, that each operate on a subset of the scene database. Figure 1 shows how flow units are combined to form the complete PxFI rendering subsystem.

After each flow unit has completed its share of the rendering task, its output, in the form of per-pixel color values in a separate frame buffer, is composited with the frame buffers

from all the other flow units, and Z values are compared to select the proper values for each pixel in the final output. After hidden-surface removal is completed, a final rendering pass is performed, lighting and texturing the previously computed per-pixel color and orientation values using the Phong shading algorithm. Because the PxFI performs this step after removing hidden surfaces, its rendering performance depends only on screen resolution—not polygon count or shading technique, as with conventional renderers.

### Multiple Flow Units Share Rendering Tasks

To provide such high performance, each flow unit fills a large PC board with 32 Enhanced Memory Chips (EMCs) plus two 180-MHz PA-8000 RISC processors for geometry processing. Each EMC has 256 pixel processors with 384 bytes of local RAM, for a total of 8,192 processors per flow unit. Each pixel processor includes a 0.5-MFLOPS floating-point engine for setup, shading, and texturing calculations, yielding 4 GFLOPS per flow unit in addition to the floating-point power of the PA-8000 processors. Each flow unit also has 64M or 128M of synchronous DRAM for the display list and PA-8000 program storage plus another 64M of DRAM for texture storage.

Up to nine flow units can be installed in a single chassis, and up to six chassis can be connected together in a single rendering system. Data is moved among the units over a 1.6-Gbyte/s geometry network (GN) and a 12.8-Gbyte/s image-composition network (ICN).

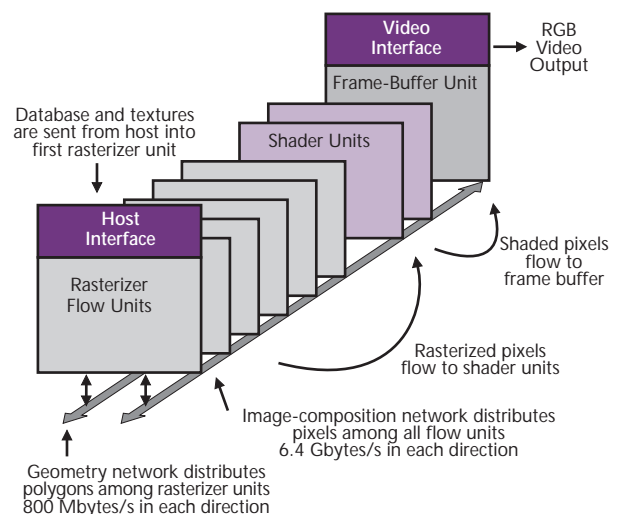


Figure 1. The PixelFlow architecture uses from 4 to 54 flow units, each with 8,192 SIMD pixel processors, to render, shade, or display 3D graphics. A fully configured system can render over 100 Mpolygons/s with a fill rate of more than 53 Gpixels/s for large triangles.

### For More Information

Hewlett-Packard's Visualize PxFI Web site at [www.hp.com/wsg/visualize](http://www.hp.com/wsg/visualize) contains additional information about the product's design, features, and performance.

The University of North Carolina at Chapel Hill maintains its own site for the PixelFlow architecture, which may be found at [www.cs.unc.edu/~pxfl](http://www.cs.unc.edu/~pxfl).

Advanced Rendering Technology's Web site, found at [www.art.co.uk](http://www.art.co.uk), provides an overview of the AR250 and RenderDrive as well as a gallery of sample images rendered on a register-level simulation of the chip.

Flow units can be assigned under software control to handle rasterizing, shading, or frame buffering. The initial rendering task is distributed among the rasterizer units, then pixel data is sent to the shaders, where shading and texturing is performed, and finally to the frame buffer(s).

The PxFI is compatible with OpenGL but offers only a fraction of its available performance with conventional OpenGL software. To achieve the PxFI's full potential, applications must use a set of extensions that support display-list processing. This is essentially the same problem that faces developers of software for Microsoft's Talisman or any other block- or pixel-oriented rendering system.

Talisman itself was also in the news at Siggraph. Fujitsu became the second company, after Trident (see MPR 6/2/97, p. 16), to announce plans for a single-chip Talisman renderer. Like Trident, Fujitsu is betting on Talisman's ultimate success despite the failure of the original Talisman reference design.

As mainstream PCs adopt these advanced architectures, software developers will need to switch from immediate-mode to retained-mode (display list) rendering models. ISVs are likely to resist this change, since it reduces their ability to add unique value to their products with custom rendering algorithms, but the performance argument will ultimately be compelling enough to force the change. We expect this to happen within the next three years.

### ART Debuts a Single-Chip Ray Tracer

A completely different approach to rendering has been taken by Advanced Rendering Technology (ART). The British startup firm has announced a chip containing a complete ray-tracing engine, the first time ray-tracing has been reduced to a single-chip design. The new AR250 includes a fixed-function floating-point geometry unit to perform the ray tracing itself, plus a simple programmable engine that supports various shading algorithms to produce the final bitmap output.

The AR250's basic specs are as impressive in their own way as HP's PxFI's. Each chip operates at a peak rate of 4 GFLOPS, performing 80 million intersection tests per second; this is fast enough to render a 10-million-polygon image in HDTV

resolution in just 30 seconds, roughly 15 times the speed of a 266-MHz Pentium II processor.

Because it is not meant for real-time rendering, the AR250 has no display controller and cannot be used as a PC's graphics adapter. Instead, the chip is intended for use in standalone rendering systems. ART's own RenderDrive combines 4, 16, or 64 AR250s with a local processor and a network interface. RenderDrive is shared by PCs and workstations in a LAN environment, receiving 3D scene models in the standard Pixar RenderMan .rib format and returning rendered images over the LAN.

The company has not announced pricing or availability for the AR250, being manufactured by LSI Logic, but says that RenderDrive will ship in October for \$19,950 (with four AR250s) to \$170,000 (with 64 chips). This is certainly much more expensive than the commodity PCs with which RenderDrive competes, but it still represents an excellent price/performance value if the AR250 performs as advertised.

Even the high-end configuration would take several seconds to render an HDTV-quality frame, but we expect the power of a 64-chip RenderDrive system to be reduced to a single device within three years. A few more years of progress will permit real-time ray tracing in consumer-level hardware, giving end users yet another choice of basic architectures for PC 3D.

### Conventional Architectures Also Make Progress

Chromatic Research revealed more details of its forthcoming Mpact 2 media processor (see MPR 11/18/96, p. 1) at the Hot Chips conference. The processor core gains floating-point support, and an integrated 3D-rendering engine promises performance on a par with any existing mainstream 3D chip when the chip ships later this year. Clock rates are expected to exceed 100 MHz, further boosting throughput.

Initially, Mpact 2 will use its floating-point capability just for setup processing, relying on the host processor for geometry calculations. Eventually, some of the geometry operations may be performed on Mpact 2, but we expect that the wider use of Pentium II processors will make this unnecessary.

Mpact 2's rendering engine can output one shaded or textured pixel per clock. Texture sampling requires one clock per texel, so bilinear texturing—requiring four texel samples per pixel—should proceed at 25 Mpixels/s or better, confirming Chromatic's predictions of competitive performance.

Interestingly, Chromatic did not include dedicated logic for the final depth-sorting operation, finding instead that the programmable core was capable of performing this task more than fast enough to keep up with the rest of the 3D pipeline.

### The End of the Conventional Pipeline

With most new 3D products using unconventional approaches to 3D rendering, we conclude that the conventional 3D pipeline is doomed. It will be a few years, at least, before conventional solutions are completely displaced by new architectures, but this trend is now inescapable. The effects on OEMs, software developers, and end users will be substantial; it is past time to start preparing. □