New Chips, New Ideas at Siggraph '99

Professional 3D Cards Debut While Experts Debate Tomorrow's Technology

by Peter N. Glaskowsky

Siggraph '99, held last month in Los Angeles, marked the debut of several new high-end PC 3D products and offered valuable insight into future trends in the 3D market. Intense3D, Evans & Sutherland, and 3Dlabs announced new professional 3D cards, continuing the historical trend of rapid improvement in this dynamic but overcrowded field.

Siggraph (*www.siggraph.org*) also continued its tradition of showcasing the latest ideas in graphics. Industry professionals and academic researchers described new rendering concepts that could appear in 3D workstations as soon as next year, and in desktop PCs shortly thereafter.

New 3D Cards Break Records

NEC's record-breaking TE4E professional 3D-graphics accelerator card (see MPR 8/2/99, p. 20) has a new challenger. Intense3D's new WildCat 4110, announced at Siggraph, achieved a ProCDRS-02 score of 36.3, surpassing NEC's claimed score of 35. Neither result has yet been certified by SPEC, the benchmark's issuer, but NEC's product has been shipping since July.

The new Intense3D card is due to ship this month, and it's sure to substantially undercut the price of NEC's offering. Intense3D (*www.intense3d.com*) expects systems equipped with the new card to retail for as little as $4,000, 70% below the price (in Japan) of NEC's Express5800 workstation—currently the only way to get a TE4E.

The basic specifications of the 4110 are an excellent fit for its target market. Intense3D's proprietary geometry engine can process up to 6 Mpolys/s, while the partner 3D-rendering chip can send up to 143 Mpixels/s to the screen.

Intense3D also demonstrated a dual-pipeline member of its Wildcat family, which is due to market in early 2000. This configuration achieved a ProCDRS-02 score of 60.5. NEC says its TE4M—essentially two TE4E cards that cooperate on geometry and rendering tasks—should achieve a ProCDRS-02 score of 65. Stay tuned; this contest isn't over.

Evans & Sutherland Looks Ahead to 4000

Thirty-year 3D veteran Evans & Sutherland (E&S) has introduced its new Realimage 4000 family, based on a new architecture the company calls Hyperpixel that is designed to provide a high degree of scalability. E&S has not yet announced any board products in the new family, but it expects to announce products in 4Q99 that will have up to four geometry processors teamed with one, two, or four rendering engines.

E&S (*www.es.com*) says the ultimate high-end Realimage 4000 configuration will achieve a fill rate of 550 Mpixels/s and a polygon rate of 18 Mpolys/s. The excellent rendering performance of this product is due in part to E&S's use of Mitsubishi's 3D-RAM5 memory chips (see MPR 8/2/99, p. 4), which include on-die ALUs.

As with E&S's previous Realimage products, multiple rendering chips in the 4000 series are interleaved on a pixel-by-pixel basis. While other vendors prefer wider vertical or horizontal stripes or more complex tiling schemes, E&S believes its method provides the best overall results. Pixel interleaving requires all rendering chips to process every polygon. The alternatives hope to benefit from culling polygons that don't cover an individual stripe or tile, but E&S says this benefit is not worth the cost of the needed culling and clipping calculations.

A unique feature of the Hyperpixel architecture is its support for a 128-bit pixel format, which includes 32 bits of Z-coordinate data and double-buffered 8-bit overlay and 8-bit alpha planes. Graphics supercomputers such as SGI's Infinite Reality use long pixel formats to enable rendering techniques more advanced than those typically found on PC-class 3D accelerators. We expect this concept to see wider use in more affordable rendering architectures.

3Dlabs Changes Direction of Jetstream

When 3Dlabs (*www.3dlabs.com*) announced its Jetstream professional 3D accelerator family at the 1998 Microprocessor Forum (see MPR 11/16/98, p. 20), it planned to create a high-end configuration using a binary-tree hierarchy of its Glint G3 geometry accelerators with Glint R4 rendering chips at the leaf nodes. In this design, the first G3 was to handle all vertex transforms, while the final layer of four G3s took care of lighting calculations, supporting a sustained throughput of 44 Mpolys/s with eight light sources.

This plan fell victim to the unanticipated difficulty of routing a board set with the equivalent of 14 separate on-board AGP interfaces in this arrangement, and 3Dlabs was forced to change its architecture to a simpler approach. The new structure uses eight G3/R4 pairs interconnected by a total of fifteen 4×-AGP interfaces, but the chip pairs are arranged in a simpler, easier-to-route sequence. With eight G3s to share the transform and lighting calculations instead of five, it expects the new design to outperform the old. It expects to ship boards using the new design in mid-2000, a slip of some six months from the original schedule.

The company has also been making some progress on other products. It introduced the Oxygen GVX210 at Siggraph, marking the new high end of the company's product line. The introduction also marks the arrival of 3Dlabs' Glint G2 geometry-accelerator chip, a straightforward evolution of

the company's now-elderly Gamma 1. The G2 is rated at 6.3 Mpolys/s, some 40–60% faster than its predecessor. The GVX210 features one G2 and two Glint R3 rendering chips that use stripe interleaving to achieve a claimed pixel-drawing speed of 460 Mpixels/s, nearly as fast as the high-end Realimage 4000 configuration due out next year.

Though 3Dlabs pioneered the transfer of geometry processing from the host CPU to dedicated chips, it remains aware of the significant calculating capacity of Intel's current Pentium III processors. It has released a beta version of the multithreaded PowerThreads SSE driver software for its Oxygen VX1 card, which lacks a geometry chip. The new drivers use Intel's SSE instructions to accelerate transform and lighting calculations and can take advantage of multiple processors in systems so equipped. According to 3Dlabs, this MP capability is unique in the industry and gives an average performance boost of 40% on 3D processing in dual-CPU systems.

Even with multithreading support, 3D accelerators that rely on the host for geometry processing are substantially slower than those with hardware geometry engines. With geometry acceleration now appearing on mainstream chips such as S3's Savage2000 and Nvidia's GeForce 256 (see MPR 9/13/99, p. 5), we expect end users to demand this feature on professional products as well. Intelligent load balancing between the CPU and the graphics card will be a key goal of these future products, ensuring that all available processing power is used to improve 3D performance.

X3D Looks to the Future of the Web

The Web3D Consortium described its plans for the next generation of its virtual-reality modeling language (VRML). The consortium received the blessings of the International Standards Organization (ISO) for the VRML 97 standard and hopes to repeat this achievement with the next major version in 2002.

Part of the effort is a Web integration strategy based on the extensible markup language XML. The extensible 3D (X3D) specification, due out in draft form by the end of this year, will allow VRML content to be expressed within Web pages through the use of XML tags, much as text formatting and hyperlinks are today expressed in Web pages through HTML tags.

Dozens of companies are now participating in the Web3D Consortium (*www.web3d.org*). Some are developing code destined for public release with Open Source (free) or Community Source (free for noncommercial use) licenses. Blaxxun, Shout, and Sun have already shown demonstration applets that can display X3D content using small (~50K) Java applets, as long as the host system is equipped with a standard low-level 3D API such as Direct3D or OpenGL.

Evolution Today, Revolution Tomorrow

Some of the most intriguing presentations were to be found at the Siggraph/Eurographics Workshop on Graphics Hardware, held in conjunction with the main Siggraph conference. This event has a long history as a venue for the presentation of academic and commercial research, especially the work of European researchers. This year, for the first time, the workshop included a number of presentations on the technology behind commercial products.

Stanford professor Pat Hanrahan provided the workshop's first-day keynote speech. He offered his vision of the next major generation of 3D hardware. This generation, he said, will be distinguished by its inclusion of programmable rendering engines. Current software rendering systems such as Pixar's popular RenderMan, used to produce high-quality 3D animation for the film and broadcast industries, rely on libraries of custom code to produce a wide variety of lighting and surface-texture effects. Next-generation 3D-rendering hardware, Hanrahan says, must provide similar flexibility.

This trend was visible in Hot3D presentations by ART (*www.art-render.com*) and PixelFusion (*www.pixelfusion.com*). Both companies provided more details on their respective products, which include programmable cores that support custom rendering algorithms. ART's focus on high-quality ray tracing for film and television applications (see MPR 9/15/97, p. 9) mandated the inclusion of a programmable rendering engine, and much of the company's business comes from its ability to custom-tailor its products to the unique needs of its customers.

Some of the technical presentations in the Siggraph conference also focused on this area. Though most 3D applications are moving in the direction of photorealistic rendering, photorealism isn't a universal goal. Artists and engineers alike can benefit from nonphotorealistic rendering techniques that enhance aesthetic appreciation or technical understanding. Procedural rendering is essential to achieve these effects.

Two Siggraph presentations are of particular interest in this connection. A team of engineers from Brown University, Adobe, and Pixar developed an algorithm that renders conventional 3D models using the artistic style of Theodore "Dr. Seuss" Geisel. Geisel's style was chosen because it is "such an extreme departure from the domain of conventional computer graphics." Distinctive rendering algorithms are crucial in the 3D-graphics industry because they offer software developers a way to differentiate their 3D applications from those of the competition.

Another team developed a renderer in the distinctive style of Jean "Moebius" Giraud, using it to produce animations inspired by the artist's "City of Fire" portfolio. While far from photorealistic, such imagery can permit more effective visual communication.

The interest in unconventional approaches to 3D graphics points out the future potential of this market. Today's no-nonsense CAD accelerators may be tomorrow's most sophisticated artistic tools. M

Peter N. Glaskowsky chaired the Hot3D track of the Siggraph/Eurographics Workshop. The Workshop presentations are available online at www.merl.com/hwws99.