Cahners
MICROPROCESSOR REPORT
www.mpronline.com
THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

PC 3D Prepares for PlayStation

Vendors Brace for Meteoric Impact of Computing Appliances

By Peter N. Glaskowsky {01/17/00-06}

After several years of turmoil, the PC 3D industry settled down somewhat in 1999. Sales grew steadily, but with 3D chips already accounting for almost all sales of PC graphics controllers at the start of the year, there was no dramatic market growth to accommodate.

Only one major vendor (Number Nine) went out of business, and no new companies were able to push their way into the top ranks of the business. Figure 1, based on data provided by 4th Wave *(www.fourthwave.com)*, shows that 1999's top five vendors (ATI, NVIDIA, S3, Matrox, and 3dfx) controlled about two-thirds of the market. Integrated-graphics chip sets (a mix of products from Intel, SiS, and VIA) accounted for about 22% of PC unit volume.

The year, however, was just the **Calm Before the Storm**. Sony, though not a player in the PC graphics space, unveiled its next-generation PlayStation 2 system and made it clear that it intends to greatly expand its influence in 3D gaming, the technology driver for the PC 3D market. At Comdex, Sony CEO Nobuyuki Idei likened Internet-enabled video-game appliances such as the PlayStation 2 to meteorites that could render the PC extinct as a mass-market computing and entertainment platform.

Compared with this apocalyptic future, 1999 may one day be seen as a golden age for the PC graphics industry. Users have had several great 3D chips to choose from, and many of the companies making these chips have found business models that allow them to make money. The PC 3D market may get tougher next year, but 3D vendors are getting stronger and smarter.

Brute Force Offers Serious Speedups

Most 3D-chip makers used brute force to stay on the 3D treadmill, achieving a 4× performance boost over the course of the year. Moving from 0.35-micron process technology to 0.25-micron allowed vendors to double transistor counts and increase clock rates by about 50%, while architectural enhancements contributed the balance of the speedup.

Though most 3D chips provided a single rendering pipeline in 1998, dual-pipeline chips became the norm in 1999. Dual-pipe designs came in various forms, some—

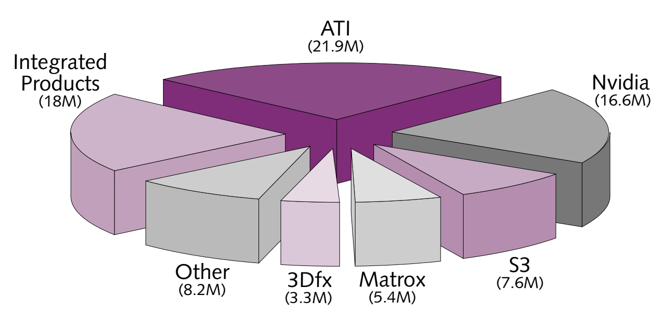

Figure 1. According to 4th Wave *(www.fourthwave.com)*, ATI sold about 27% of the 3D-graphics accelerators in the PC market in 1999.

such as 3dfx's Voodoo3—capable of drawing one pixel per clock with two different textures applied to each. Others, notably NVIDIA's Riva TNT2, can draw two single-textured pixels or one dual-textured pixel in each clock period. The fastest of these chips offer clock rates of 183 MHz. NVIDIA's new GeForce 256 has four parallel pipelines but slightly lower clock speeds. This trend toward higher parallelism is likely to continue for the next several years.

Both 3dfx and ATI achieved further multiples of previous performance levels through support for multiple graphics chips per card. For 3dfx, this strategy represented a return to a feature that made its earlier Voodoo2 the most successful 3D-gaming card of its day. The Voodoo2's scan-line interleaving feature allowed gamers to buy two identical cards to get twice the effective performance. The VSA-100 includes an enhanced version of this feature that allows up to 32 chips to work together, giving it both the **Highest Price** and **Highest Performance** titles in the **Unlimited PC Class**. Claiming the **First Gigapixel Card** award, 3dfx introduced a four-chip Voodoo5 board that should be the first product on the market to draw more than one billion pixels per second when it ships early this year. Would-be competitor Gigapixel, meanwhile, announced its first chip—rated at just 0.1 gigapixel/s—earning it the **Behind the Curve** award.

ATI put two of its Rage 128 Pro chips on a single AGP card to create the Rage Fury Maxx. The Maxx board overcomes the gamer's standard objection to ATI graphics chips—relatively weak rendering speed—giving ATI its best shot at the PC gaming market in years, especially if the company can follow it up with a competitive single-chip implementation.

Geometry Acceleration Offloads the CPU

Today's PC microprocessors are a significant obstacle to improved 3D acceleration, due to their limited capacity for 3D-geometry processing. Indeed, one of the distinguishing characteristics of PC microprocessors as a class is their weak floating-point performance compared with that of workstation CPUs. Even with the addition in 1999 of SIMD floating-point accelerators, such as SSE from Intel and AltiVec on PowerPC, the continuing emphasis on integer processing, combined with price pressure from PC makers and buyers, has prevented CPU processing power from keeping up with the rapid improvements in 3D-chip rendering speed.

The answer, which has been obvious for the past few years, is to offload repetitive geometry-processing operations to dedicated silicon. These operations require little flexibility and almost no memory, making them much easier to implement than CPU cores with comparable throughput.

Several companies in the professional-3D market had previously introduced standalone geometry chips, but none of these reached the consumer market, despite attempts by RealVision and Rendition. It fell to S3 and NVIDIA to introduce the first mainstream graphics chips with integrated geometry acceleration. S3's Savage2000 was the first announced, but NVIDIA's GeForce 256 shipped first, breaking all previous records for single-chip 3D acceleration. Because GeForce leads all shipping products in both geometry and rendering speed, we give the GeForce 256 our **Analyst's Choice Award** for **Best 3D Accelerator of 1999**.

Geometry acceleration can be used to increase the frame rate for scenes of a given complexity, or to increase the complexity of a scene without reducing the frame rate. The former option requires little specific software support within the application, since most applications update the display as fast as possible. Today's 3D games already run at or near the maximum frame rate of PC displays, however. Increasing the update rate from about 60 Hz to 120 Hz or more offers only slight benefits at best; the difference is usually unnoticeable. Geometry-enabled chips will have the greatest impact in systems with relatively slow CPUs. Buyers who can afford Intel's or AMD's best processors will benefit least from geometry acceleration.

To take full advantage of geometry acceleration, applications must be designed to work well over a very wide range of geometry performance. Multiple sets of models for the objects in the scene must be provided, or one set must be designed that can be rendered at various levels of detail. Only a few of the newest PC 3D games include such support, generally made visible to the user through a manual "complexity" adjustment with two or three settings—enough to allow geometry acceleration to make a difference, but not enough to take full advantage of it.

Researchers in the 3D market are working on more effective solutions. Multiresolution meshing technology can adapt the complexity of a model according to the available performance and the model's visibility on the screen. As an object occupies more of the screen, it can be drawn with greater detail, requiring more geometry-processing resources. This adaptation can be very gradual, with not just two or three steps but hundreds.

It is likely to take another year or two for adaptive mechanisms to be developed and become prevalent in PC games and other software. In the meantime, chips such as GeForce will be fully utilized by only a fraction of the available titles. We expect to see more 3D chips with geometry acceleration announced in 2000 as the technology becomes more affordable and more useful.

Integration Reduces Cost, Performance

Low-end PCs have never offered good 3D performance. In the early days of the PC 3D market, low-end systems simply didn't include 3D accelerators, or they used chips that claimed to accelerate 3D while actually decelerating real-world applications. As the market matured, low-end systems received the high-end chips of one to two years past. With the rapid pace of 3D-chip development, this meant that low-end PCs offered, at best, 25% the 3D performance of high-end systems.

The flurry of integrated-3D chip sets in 1999 has led to even lower price points—and an even wider performance gap Intel introduced four versions of its integrated chip set, Acer and SiS released two each, and VIA and Trident announced desktop and mobile versions, respectively, of a mutually-developed product. ATI and S3 are expected to release similar products in the coming year. These chip sets represent a very good deal for OEMs and end users, but their low prices are primarily attributable to reduced manufacturing profit margins rather than architectural advantages. Integrated-graphics chip sets are much more expensive to manufacture than similar products without graphics, yet they command a small premium at best over the nonintegrated alternatives.

Some companies adopted very aggressive pricing to gain market share. Intel's 810 and 810E chip sets inherited the large, slow graphics core from the discrete 752 graphics chip, which won our **Mercy Killing** award for being canceled before it could die on its own in the open market. Even with this die-size penalty, the 810 sells for prices only slightly higher than Intel's own nonintegrated low-end chip sets. This pricing tactic means the 810's integrated graphics are nearly free—earning the **Fair Market Value** award. To OEMs, the 810 is **An Offer They Can't Refuse**, despite its anemic performance in 3D games and 2D-only productivity applications alike.

Some integrated-3D products, notably those from Acer and SiS, offer much better performance while preserving competitive prices. Unfortunately, end users may find it difficult to distinguish the good from the bad. Indeed, the success of these products may depend on finding a way to communicate their performance advantages to the end-user community. Some end users will never care about graphics performance, of course, so the key will be to identify and reach those who do.

Vertical Integration Changes Shape of Market

The rapid strides in 3D technology in 1999 were echoed by interesting changes in business models. In 1998, 3Dlabs bought competitor Dynamic Pictures, chiefly to acquire DP's add-in-card business. This purchase proved to be the first in a series of mergers and acquisitions in 1999.

As its chief rival, NVIDIA, began to encroach on 3dfx's sales to independent board vendors, 3dfx bought STB to guarantee its access to retail sales channels. STB's brand gave way on store shelves to eye-catching 3dfx-branded packaging. This unified branding strategy allowed 3dfx to use national TV advertising to drive demand for its products.

At midyear, S3 purchased Diamond Multimedia, gaining access to Diamond's 3D-card business as well as to Diamond's Rio MP3 player and other consumer-multimedia products. The deal seems to have had little effect on S3's ability to sell graphics chips, but it greatly diversified the company's product line and reduced S3's dependence on its ability to produce graphics chips. We give S3 the **Houdini** award for escaping the 3D-chip box. (S3 also purchased the assets of Number Nine in December. Number Nine, one of the oldest companies in the PC-graphics business, discontinued development of its own graphics chips earlier in the year and couldn't survive on graphics-board manufacturing alone.)

PC-3D Market Stabilizes

There are fewer names on our PC 3D vendor list this year—36 vs. 52—but we have raised the bar for inclusion. The companies still on the list have developed, or are believed to be developing, 3D accelerators for the PC market. Most of the companies dropped since last year are still around but are now focused on embedded applications. (The full list is available online at *www.MDRonline.com/3d*).

Exactly half of these companies (including almost all the major discrete 3D-chip makers except NVIDIA) now have integrated board-manufacturing operations, as indicated in boldface. Product names are indicated within parentheses; currently shipping products are also in boldface.

PC 3D vendors and products: **3Dfx** (**Voodoo**); **3Dlabs** (**Glint, Permedia**); Acer Labs (**Aladdin TNT2**); **Advanced Rendering Technology** (**AR250**); ARK Logic (**ARK 8100**); ArtX (Nintendo Flipper, **Acer Aladdin 7**); **ATI** (**Rage**); BitBoys (**Glaze3D core**); **Creative Technology** (Silicon Engineering Mojo); **Evans & Sutherland** (**Lightning, Tornado**); GigaPixel (**GP-1** and GP-2 cores); **Hewlett-Packard** (**Visualize**); **IBM** (**GXT3000P**); IGS Technologies (**CyberPro 5300**); **Intel** (**810 chip set**); **Intense3D** (**Wildcat**); IXMicro (**TwinTurbo 128-3D**); **Matrox** (**G400**); Micron (Rendition Vérité); **Mitsubishi Electric** (**VolumePro**); **NEC** (**TE4**); Neomagic (**MagicMedia 256XL+**); NVIDIA (**RIVA, Vanta, GeForce**); **PixelFusion** (Fuzion); PowerVR (**PowerVR**); **Primary Image** (**P10**); RealVision (**GA400**); **S3** (**Savage**); **Silicon Graphics** (**Cobalt**); Silicon Magic (**SM3110**); Silicon Motion (**Cougar**); SiS (**530, 540, 620, 630 chip sets**); **Sony** (**PS2 Graphics Synthesizer**); Stellar (**PixelSquirt core**); Trident (**Blade3D**); VIA (**Apollo MVP4 chip set**)

Among the PC 3D chip vendors that have dropped out of the market since last year, most remain in business with other product lines. Number Nine (purchased by S3) was the only major chip maker to cease operations in 1999. The competitive landscape is likely to get even tougher in 2000, as 3D-chip makers are squeezed between integrated products at the low end and consumer-electronics products such as Sony's PlayStation 2, limiting growth in the PC gaming market. (Sony also plans to use its Graphics Synthesizer in PC workstation products, so it is included here.)

3D Events of 1999

SGI rolls out Windows NT workstations based on its proprietary Cobalt chip set (*MPR 2/15/99-02*, p. 12).

S3, once given up for dead, roars back with three major product introductions in 1999: Savage4 (*MPR 2/15/99-03*, p. 14), mobile Savage/MX and Savage/IX (*MPR 8/2/99-msb*, p. 4), and Savage2000 (*MPR 9/13/99-msb*, p. 4).

Intel's Streaming SIMD Extensions (SSE) greatly improve 3D-geometry throughput on Pentium III processors (*MPR 3/8/99-01*, p. 1), but mainstream 3D chips such as the Savage2000 and NVIDIA's GeForce 256 (*MPR 9/13/99-msb*, p. 4) offer even higher performance for real-world applications.

Ratification of the MPEG-4 standard (*MPR 3/29/99-05*, p. 18) paves the way to object-oriented broadcast programming with embedded 3D objects, but widespread use of the new standard is not expected for three to five years.

Sony unveils details of its Emotion Engine and Graphics Synthesizer chips for the PlayStation 2 (*MPR 4/19/99-01*, p. 1), which will rival high-end PCs for performance and quality in 3D games and may cut into the business of many PC 3D-chip makers.

3dfx, ATI, Matrox, and NVIDIA announce new mainstream PC graphics chips with 3D performance greatly improved over that of their predecessors (*MPR 4/19/99-05*, p. 17).

Intel rolls out its 810 chip set with integrated 3D graphics plus the discrete 752 graphics chip (*MPR 5/10/99-04*, p. 17). Despite poor 3D performance, the chip set is attractively priced and is widely used. The 752 is later canceled (*MPR 8/23/99-msb*, p. 4).

SiS claims higher performance for its 540 and 630 integrated-graphics chip sets (*MPR 5/31/99-06*, p. 22).

VIA purchases the Cyrix and IDT x86 processor-design operations, positioning itself to develop CPUs with integrated 3D to compete with Intel's anticipated Timna processor (*MPR 7/12/99-02*, p. 5). SiS later inks a deal with Rise to develop similar products (*MPR 11/15/99-msb*, p. 4).

Microsoft releases version 7 of its DirectX multimedia API set (*MPR 7/12/99-06*, p. 14), the first version of DirectX to support hardware geometry acceleration, features later to be included in S3's Savage2000 and NVIDIA's GeForce.

NEC announces the TE4E professional graphics accelerator (*MPR 8/2/99-06*, p. 20), which sets performance records on SPEC's ProCDRS-02 benchmark.

Bitboys announces the Glaze3D 1200, the most aggressive embedded-DRAM graphics-chip design to date (*MPR 8/23/99-msb*, p. 4). Each chip includes 9M of DRAM and four rendering pipelines. Performance is estimated at 600 Mpixels/s, significantly faster than any previously announced PC 3D chip.

The annual Siggraph conference (*MPR 9/13/99-04*, p. 18) sees new 3D products from Intense3D and Evans & Sutherland, plus a redesign of 3Dlabs' Jetstream family.

Nvidia and Acer Labs introduce the Aladdin TNT2, a P6-bus chip set with integrated 3D based on Nvidia's TNT2 chip (*MPR 10/6/99-msb*, p. 4). The Aladdin TNT2 is later joined at Acer by the Aladdin 7, a Socket 7 chip set based on a new design from secretive startup ArtX (*MPR 12/6/99-msb*, p. 4).

The announcement of Intel's 820 chip set (*MPR 10/6/99-06*, p. 30) is delayed to November. The new chip set marks the debut of Direct RDRAM and the arrival of AGP 4×, but the high cost of 820-based systems limits end-user acceptance. The 820's arrival is preceded by that of the 840, a workstation chip set with dual RDRAM channels (*MPR 10/25/99-07*, p. 28).

Mitsubishi announces the VolumePro vg1000, the second generation of its volume-rendering chip family (*MPR 10/25/99-08*, p. 30).

Apple buys Raycer Graphics, a pro-3D startup that never reached market (*MPR 11/15/99-msb*, p. 4). In separate but contemporary announcements, Gigapixel released the GP-1, its first 3D core, and Intel purchased the remaining assets of former partner Real3D.

After years of strong sales of 2D-only parts, NeoMagic bows to the inevitable and releases a low-power notebook-graphics chip with 3D acceleration (*MPR 12/6/99-msb*, p. 4). The polar opposite of the NeoMagic chip, 3dfx's VSA-100, is also announced (*MPR 12/6/99-02*, p. 12). The VSA-100, code-named Napalm, will run almost as hot as its namesake but offers unmatched scalability.

Late in the year, Intel got into the merger and acquisition fever by buying Real3D, the Lockheed Martin spinoff that provided the key 3D technology for Intel's 740 and 810. Real3D had hoped to establish itself as an independent supplier of 3D technology and chips for the PC and arcade-game market, but its only real success came from arcade-game efforts (notably as a supplier to Sega) that simply weren't enough to support it. Real3D was never able to crack the PC market, in part because it gave its best technology to Intel and in part because potential customers knew it gave its best technology to Intel. We give Real3D the **Damned If You Do** award for 1997–1999.

ATI, which began the year with the greatest share of the market, already had its own board business. Despite the best efforts of its competitors, ATI has retained its market leadership and receives the **Staying Power** award.

Sales by Intel, NVIDIA, and S3 are swelling, however, making life harder for ATI, which may need **Viagra** to keep it on top in 2000.

Only NVIDIA remained above the fray, despite losing its Diamond and STB sales. Whether by accident or by design (rumor had it the company was seeking to buy a board vendor, but no deal ever materialized), NVIDIA has become the largest independent graphics-chip supplier. NVIDIA is effectively the sole supplier to many independent graphics-card makers, a role that has greatly increased NVIDIA's visibility and revenue. Other 3D-chip companies are trying to gain a share of NVIDIA's business, but for now NVIDIA is doing better than ever.

Video Games Threaten PC Superiority

The biggest threat to the continued dominance of PC graphics chips comes from video games and related information appliances, which are certain to be widely adopted by the same home users who buy today's gaming PCs. Despite the potential advantages of the PC platform—including high-resolution displays, high-capacity main memory and hard disks, and all the other benefits of their higher materials-cost budgets—PCs are unreliable and difficult to use. These two shortcomings render PCs vulnerable to predation by smaller, simpler, and more reliable appliance-like computers.

The threat is being taken seriously by PC 3D vendors. New products are being developed to help keep consumers and developers interested in the PC platform, and new marketing messages are being written to address the threat directly. Chips with geometry acceleration from NVIDIA and S3 are the PC industry's answer to Sony's Emotion Engine. The Voodoo 5, for its part, proves that PC chipmakers can still afford to outspend designers of console games.

If these chips are the PC industry's sole defense against dedicated video-game consoles, they may be no more than an electronic **Maginot Line**. Indeed, the battle for gaming superiority may be over. Console games are already the preferred gaming platform for the vast majority of avid gamers, and games for Sega, Nintendo, and the original PlayStation already outsell PC games by a wide margin.

Video games, in turn, may be the electronic equivalent of the **Trojan Horse**. As next-generation platforms arrive with vastly improved performance, multiplayer-gaming features, and other attractive attributes such as Internet access and DVD-movie playback support, these new machines will be snapped up by current console-game owners. They will also be adopted rapidly by even traditional PC owners and consumers who have little interest in games, and should account for the vast majority of home computing devices by 2010.

PC 3D will nevertheless continue to account for much of the progress in the overall 3D market and the vast majority of product introductions. While new video-game consoles may come out once a year on average, each year sees dozens of new PC 3D cards. This variety keeps the PC 3D market interesting, but it may not be sufficient to ensure growth. The pressure on 3D-chip designers is already intense, and it will only get stronger in the coming years. ◇