S3, Matrox, 3Dfx Target 3D Mainstream

Improved Rendering Engines Offer Impressive Performance

by Peter N. Glaskowsky

It's been an unusual year for the 3D-chip industry. In 1997, most of the important new 3D engines were announced in April or May—though it would be a few more months before some of them shipped. This year the announcements began in February with the Intel740 (see MPR 2/16/98, p. 1) and will continue through August at least.

Why the change? In part, it's because vendors have existing products that are worth protecting. Introducing new parts before they're in production would reduce demand for the parts that are still available. It may be that this year's pattern of chip introductions will be repeated in the future; the many simultaneous releases in 1997 may turn out to be the unusual thing.

The extended series of chip announcements this year will be covered in a series of articles. We begin with three fast new mainstream chips and will continue in the next two issues with chips that are even faster, have interesting new features, or simply haven't been announced yet.

In this issue, we cover S3's Savage3D, a chip that more than makes up for all the disappointing 3D chips S3 has released over the past two years; Matrox's MGA-G200, the first serious 3D chip from a company that has long offered the best 2D chips on the market; and 3Dfx's Voodoo Banshee, the first integrated 2D/3D accelerator from the company that virtually defined 3D acceleration for PC games.

S3 Comes Through

One company that almost didn't show up in the 3D market last year was industry giant S3, which sold more graphics chips than any other vendor in 1996 and nearly duplicated that performance in 1997.

Unfortunately for S3, it failed to develop a competitive 3D chip in 1997, offering only the relatively anemic ViRGE/GX2, which was no match for arch-rival ATI's Rage Pro and was far slower than Nvidia's impressive RIVA 128.

The ViRGE/GX2 wasn't even as successful as the earlier ViRGE/DX, which was available for prices well below $10 and could be used with low-cost EDO DRAM, enabling $40 add-in cards with basic 2D, 3D, and video functions. Although these low-end chips kept S3's unit volumes high, its revenues sank.

During the past year, 3D acceleration shifted from a desirable option to a mandatory feature for any mainstream graphics chip. It became clear that S3's continued survival depended on coming up with a competitive 3D accelerator.

Fortunately, it looks like S3 has done just that. The new Savage3D, announced at the Computer Game Developer Conference to emphasize the chip's target market, has all the features any game could need and enough performance to satisfy all but the most demanding users.

In designing the Savage3D, S3 traded in the old ViRGE architecture for a new tile-based renderer similar to that in the Intel740. Figure 1 shows a block diagram for this chip, illustrating what S3 calls "dual rendering pipelines," because of parallel shading and texturing functions, though the Savage3D can render only one pixel per clock. Like the 740, the Savage3D renders the full screen in one pass. The Savage3D does not require polygon sorting, which caused software problems for tile-order rendering architectures like PowerVR and Microsoft's Talisman.

The Savage3D still benefits from tiling's increased locality of reference for texture-memory fetches, allowing it to achieve a sustained throughput of 125 million pixels per second with bilinear- or trilinear-filtered texturing. This is over 10 times the pixel throughput of the ViRGE/GX2.

Key Performance Features Are Provided

Polygon processing is also significantly enhanced, since the Savage3D includes S3's first 3D setup engine (a feature that appeared on all competing chips last year). It remains to be seen how the Savage3D's polygon throughput will stack up against competing products, however. Like other vendors, S3 offers only an unrealistic peak rating—five million polygons per second—that will never be achieved in a real system, due to CPU-processing and bus-bandwidth bottlenecks.

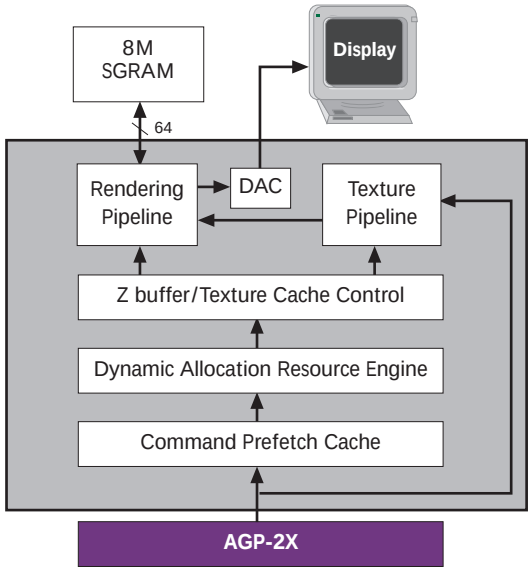

Figure 1. S3 says its Savage3D has "dual rendering pipelines," but its design is comparable to that of other single-pipeline chips.

The Savage3D also includes S3's first full 2× AGP implementation with sideband addressing and pipelined AGP execute-mode texture fetching. This should provide the chip with enough bandwidth to textures stored in system memory to keep up with its high-performance rasterizer.

The other key to good rendering performance is fast local memory, one area where the Savage3D may fall behind its competition. While many other vendors have gone to 128-bit local-memory controllers, the Savage3D offers only a 64-bit memory bus for 125-MHz SGRAM. This provides only 68% of the peak bandwidth of the 128-bit, 100-MHz SGRAM array on Nvidia's RIVA 128, a 1997 product. S3's offering is limited to 8M of local memory; while this should be enough for most customers, it will block S3 from OEM accounts that require a 16M option for high-end systems.

Savage3D Offers High-Quality Rendering

The new chip offers quality features to match those of most other chips announced this year. The Savage3D uses a full 32-bit rendering pipeline for true color plus alpha, with a dithering mode to allow 16-bit 3D games to achieve some of the benefits of 32-bit rendering. Both 16-bit and 24-bit Z buffering is supported, providing better depth resolution for applications that need it. The chip also assists with more advanced features, such as bump mapping and anisotropic texture filtering, though these techniques will require some software support as well.

S3 developed its own texture-compression technology for the Savage3D, and Microsoft licensed the 3D solution for inclusion in version 6 of DirectX. Other 3D vendors, however, have been reluctant to license S3's technology; most have developed their own compression algorithms instead. This has created some confusion for software developers, who must author texture maps for each texture compression engine they wish to support. Microsoft's support for S3's approach will give the Savage3D a leg up on the competition.

S3 has long been a leader in digital-video support on mainstream graphics chips, and the Savage3D offers one of the best feature sets for this application. It includes MPEG-2 motion-compensation logic, hardware support for multiple video windows, and a 60-MHz VIP video port (the fastest announced to date) that provides an interface for hardware MPEG-2 decoders at HDTV data rates. An integrated video encoder allows the Savage3D to drive NTSC and PAL displays directly, saving the cost of the discrete encoder required by many other graphics chips.

Matrox Joins the Fight

If any mainstream graphics-chip vendor had an even weaker 3D story than S3 in 1997, it was Matrox. S3's ViRGE family didn't have the features or performance to win the most important OEM contracts, but Matrox's Millennium barely offered any 3D support at all.

Matrox's excellent 2D performance and its enviable reputation for quality allowed it to hold on to many of its OEM customers for business PCs. Eager for a share of the growing consumer 3D gaming market but without 3D technology of its own, Matrox was forced to produce a 3D-only graphics card based on NEC's PowerVR PCX2. The PowerVR strategy was a necessary step to make the Matrox brand visible to 3D gamers, but it couldn't match the profit potential of home-grown technology. This year, Matrox has developed the MGA-G200, a new graphics chip with good 2D and 3D performance, and the cards based on it represent a much better deal for both Matrox and its customers.

The company's new MGA-G200 graphics chip is based on a completely new architecture. Though the part has only a 64-bit local memory bus, Matrox has chosen to call its design a "128-bit DualBus" architecture. While this may sound like the chip has a pair of 128-bit buses, the phrase actually refers to the independent 64-bit input and output buses inside the chip. Matrox says this approach allows more efficient use of the chip's external memory than competing approaches that use a 128-bit bidirectional bus—but the 64-bit memory bus prevents the G200 from matching the performance of true 128-bit implementations.

Matrox has kept many details of the G200's design to itself. For example, no information has been released about

Pricing and Availability

S3's Savage 3D is sampling now, with production planned for 3Q98. S3 (*www.s3.com*) has priced the chip at $35 in 10,000-unit quantities.

Matrox's MGA-G200 is in production now. Matrox (*www.matrox.com*) sells G200-based expansion cards and offers the chip on the merchant market only to OEMs and motherboard makers.

3Dfx's Voodoo Banshee is also sampling now, with production scheduled for 3Q98. The chip will be sold for $38 in large volumes. See *www.3dfx.com*

	S3 Savage3	Matrox MGA-G200	3Dfx Banshee
Bus Interface	AGP 2×	AGP 2×	AGP 1×/2×
Local Memory Type	SD/SGRAM	SD/SGRAM	SD/SGRAM
Memory Width	64 bits	64 bits	128 bits
Memory Clock Rate	125 MHz	143 MHz	100/125 MHz
Maximum Memory	8M	16M	16M
Core Clock Rate	125 MHz	84 MHz*	100/125 MHz
Peak Pixel Rate	125M	84M*	100M/125M
Peak Triangle Rate	5M	n/a	4M
3D WinMarks	861*	812*	n/a
RAMDAC Speed	250 MHz	250 MHz	250 MHz
Video In/Out Ports	In/Out	In/Out	In/Out
Availability	3Q98	Now	3Q98
Volume Price	$35	n/a	$38

Table 1. These three chips will all be popular in consumer systems, though 3Dfx's Banshee is likely to have the best overall performance. n/a: not available. (Source: vendors except *MDR)

the G200's internal clock speed. Our testing shows the G200 can render large polygons at almost 84 Mpixels/s; if the part can draw one pixel per clock, as most current 3D chips do, its clock rate should be about 84 MHz.

With up to 16M of local memory and a 230-MHz or 250-MHz RAMDAC on the die (depending on speed grade), the G200 can drive an external display at up to 1,920 × 1,200-pixel resolution with 24-bit color, making the G200 suitable for the 2D needs of the most advanced users. 3D graphics, with double-buffered 32-bit color plus Z, are supported at up to 1,280 × 1,024-pixel resolution in the 16M configuration, allowing the chip to be used in some professional applications.

Average Speed Combined With High Quality

Matrox has not released detailed 3D benchmark results for the G200. Our own testing on a 400-MHz Pentium II system running under Windows 95 and DirectX 5.0 shows the part is about as fast in 3D WinBench as Nvidia's RIVA 128. Visual quality is substantially better, however, due to more precise internal calculations. The G200 also supports 32-bit color and single-cycle trilinear filtering for the growing number of applications that can take advantage of these features. While Matrox also claims antialiasing support for the G200, the drivers shipped with the product do not implement this feature correctly, cutting its 3D WinBench score by some 34 points (4%).

The G200's lower performance, compared with that of other 1998 3D chips, is likely to keep the chip out of the top ranks of 3D-gaming accelerators, but it has more than enough 3D performance for the typical home and business buyers who represent Matrox's most important market.

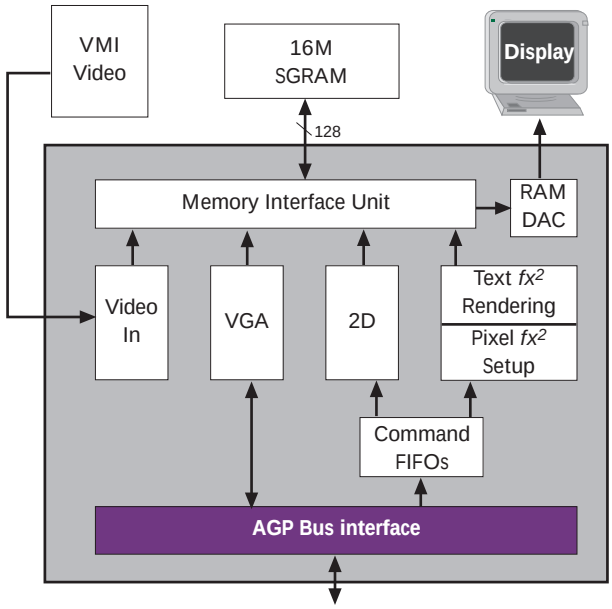

Figure 2. The 3Dfx Voodoo Banshee includes a 3D core similar to a single-texturing Voodoo2, plus strong 2D and video features.

Much of Matrox's business has come from the company's emphasis on 2D performance for productivity applications, and the G200 should not disappoint on this score. Matrox promises better 2D performance from the G200 than any previous graphics chip—though it may expect stiff competition from other chips to be released this year.

Voodoo Banshee Wails on 2D, 3D

Hoping to capitalize on the success of its Voodoo2 chip set and erase the memory of the much less successful Voodoo Rush, 3Dfx has developed a single-chip 2D/3D/video accelerator it calls Voodoo Banshee.

Banshee's 3D design was taken directly from Voodoo2, though Banshee includes only one texturing engine instead of the Voodoo2's pair, as Figure 2 shows. With a slightly higher core clock frequency (100 MHz vs. 90 MHz), Banshee will offer better performance on applications that need only one texture applied to each polygon (including the critical 3D WinBench benchmark), though peak multitexturing performance is cut nearly in half. Eliminating the single-clock dual-texturing capability of Voodoo2 ensures that Banshee will not compete with the Voodoo2 for high-end systems. Banshee also does not support Voodoo2's scan-line interleaving (SLI) feature, which permits a further doubling of its fill rate when two Voodoo2 cards are used together. Although a pair of Voodoo2 cards costs $500–$600 at retail, 3Dfx reports that 30% of these cards are sold in pairs, illustrating the demand for supreme 3D performance.

Banshee's 100-Mpixels/s peak rendering rate is slower than the 125-Mpixels/s of S3's Savage3D, another single-texturing chip with which Banshee will compete directly for the AGP slots of 3D gaming systems. Like the Savage3D and Matrox's G200, Banshee can perform trilinear texture filtering at full speed, eliminating the performance penalty found on most 1997 chip designs for this quality feature. Banshee's AGP implementation may give it another slight disadvantage. The chip supports only 1× AGP mode, though sideband addressing is supported.

More important, however, Banshee has a 128-bit local memory interface that runs at up to 100 MHz, providing more peak bandwidth than the Savage3D's local memory and allowing Banshee to sustain its performance even for memory-intensive operations. All on-chip graphics subsystems are 128 bits wide, including the VGA core—an industry first, 3Dfx says. While other graphics designers have provided 128-bit engines for Windows GUI acceleration, a wider VGA core provides better performance on old DOS-based games. These applications aren't as important as new Windows 3D titles, but 3Dfx wanted Banshee to have outstanding gaming performance for all games, not just the newest.

As part of this strategy, 3Dfx says Banshee's 2D engine is effectively as fast as it can be, coming within a few percentage points of the null-driver theoretical maximum speed, and is about 8% faster than Matrox's G200 on the 2D WinBench 98 Business Graphics test.

Banshee also improves on Voodoo2 with support for up to 16M of local SGRAM and a 230-MHz or 250-MHz RAMDAC; like Matrox, 3Dfx offers two speed grades, depending on the maximum resolution needed. The faster RAMDAC is designed for displays with up to 1,920 × 1,440-pixel resolutions. 3Dfx appears to have done a good job of leveraging its 3D technology to produce a top-notch 2D engine. Banshee's 2D core has its own setup engine—similar in function to a 3D setup engine but simplified for 2D graphics—to accelerate 2D polygon drawing, and it uses the SGRAM block-write commands to accelerate area fills.

By combining 2D and 3D in one part, 3Dfx has given Banshee some important capabilities not found in the 3D-only Voodoo2. In particular, Banshee can display 3D graphics in a window on the Windows desktop, allowing Banshee to support programs that use both 2D and 3D windows.

Like Voodoo2, however, Banshee lacks some other useful 3D features. Neither product supports AGP execute-mode texturing, restricting 3D texture maps to the space available in the local memory and requiring some additional attention from 3D programs. Banshee also lacks a true 32-bit rendering mode; though it renders internally in 32 bits, the frame buffer must use a 16-bit color format.

Rounding out the part, 3Dfx included digital-video features, such as a bidirectional VMI port and digital RGB outputs, to support external video encoders.

The company is planning to release an improved version of Banshee shortly after the first production shipments in 3Q98. The newer part will have a 2× AGP interface, operating at 125 MHz internally and on its local-memory bus, matching the peak 3D fill rate of S3's Savage3D and yielding further speedups on 2D operations as well. This faster Banshee will also offer 230-MHz and 250-MHz RAMDAC options.

New Chips Should Satisfy Most Users

These three chips are enough to satisfy the vast majority of graphics buyers. S3's Savage3D and 3Dfx's Voodoo Banshee are evenly matched overall. The Savage3D is likely to cost slightly less at the chip and board levels and will have the full force of S3's marketing muscle behind it. Matrox's MGA-G200 should be popular in business systems and with consumers who know and respect the Matrox brand name. The G200 is also shipping today, giving it a head start on the S3 and 3Dfx offerings.

Experienced 3D gamers, who represent the strongest market for high-performance 3D accelerators, are more likely to choose Banshee. These users will appreciate the chip's Voodoo heritage, and its 128-bit memory and superior 2D performance will eliminate any strong reason to choose a competing part. Many games are written to get the most out of the Voodoo architecture, and for these games, nothing less than Voodoo will do. M