3D Chips Break Megatriangle Barrier

Better Designs and Processes Help Crank Out 1M Polygons/s in Mainstream PCs

by Peter N. Glaskowsky

'97 PC TECH FO

For most of us, there are still 180 shopping days until Christmas. For the makers of 3D chips, however, the season has already begun. By the end of this month, most PC OEMs will have chosen graphics cards for systems to be sold in 4Q97. Motherboard designs have already been settled, but the arrival of AGP has disturbed the normal pattern of motherboard design: more systems will use AGP graphics cards instead of built-in graphics controllers.

Recent 3D chip announcements from 3Dlabs, ATI, Number Nine, Nvidia, S3, and Trident come just in time for OEM consideration. These new parts all promise dramatic increases in 3D speed and quality over last year's offerings, and all of these vendors claim to have unique advantages over their competitors.

It's a close call, but on the basis of our analysis and early results from the new Ziff-Davis 3D WinBench benchmarking program (see sidebar, p. 6), we select ATI's 3D Rage Pro as the best of the new bunch. The Rage Pro design emphasizes fast floating-point setup calculations, efficient caches, and a fast, flexible local-memory controller. These features provide the performance and differentiation OEMs need.

Market Forces Demand Rapid Technical Progress

The PC 3D market is now about two years old, but all the necessary hardware technology to get it started was in place more than five years ago. The result has been unusually rapid progress in 3D-chip performance, about three times faster than the norm for CPU design. In just two years, mainstream 3D chips have evolved from 3D "decelerators"—chips slower than software-based 3D rendering—into devices with very respectable performance.

The 3D chips announced this spring achieve their performance improvements by addressing the two primary contributors to 3D performance: computational resources and memory bandwidth.

The calculations necessary for 3D can be divided into three general categories, as Figure 1 shows. The 3D pipeline consists of scene management tasks, geometry and setup calculations, and rendering. Scene management, the highest-level function, prepares a database of 3D objects along with information on lights and the position of the virtual "camera." The geometry category may be further divided into geometry transforms and setup operations. Rendering consists of applying shading and texturing to pixels.

It's clear that scene definition must remain on the host processor and rendering must be handled by dedicated logic. The controversy lies in where to break the pipe between the host CPU and the 3D chip. With 3D chips evolving faster than host processors, that breakpoint has been moving steadily up the line. Each step increases the amount of work the 3D chip performs and frees the host to manage more sophisticated and dynamic 3D scenes.

Today, the best mainstream 3D chips, including all but one of the products in this article, split the pipeline halfway through the geometry section. The host handles scene management and geometry transformations, while the 3D chip handles setup and rendering.

The result: a twofold increase in polygon throughput over previous solutions that used host-based setup processing. Rendering performance has increased by an even larger factor, reflecting both the increased polygon counts as well as the demand for higher visual quality. Quality in 3D rendering is mostly a matter of texture filtering. The simplest form of texture mapping samples one texture pixel (texel) for each screen pixel, but this method yields an objectionable shimmering effect. Bilinear filtering produces a much better effect but results in four times the memory traffic.

The increased demand for bandwidth can be met in part by on-chip texture caches, but for now the best solution is to provide wider, faster local memory arrays in the graphics subsystem. This feature, along with on-chip setup processing and faster rendering engines, distinguishes the best from the rest of the 3D-chip pack.

3Dlabs Continues to Move Into Mainstream

Best known for its high-performance Glint chips for high-end PC CAD workstations (see MPR 3/5/96, p.16), 3Dlabs has been steadily building market share with Permedia, one of the fastest mainstream parts available. Permedia's most popular configuration pairs it with 3Dlabs' Delta setup accelerator, which offloads the straightforward but time-consuming setup calculations from the host. The combination, known as Permedia NT, is aimed at the low end of the CAD market.

As a two-chip solution, Permedia NT offers good performance but a relatively high cost. The conventional solution to this sort of problem is higher integration, and this is exactly what 3Dlabs has done. The new Permedia 2 (P2) chip, announced at the recent PC Tech Forum, combines an enhanced Permedia rendering engine with the same setup processor used in Delta, achieving high performance while staying within the typical $30–$35 budget OEMs have allocated for graphics chips.

The P2 is similar to the original Permedia design, as Figure 2 shows, but offers twice the performance of the earlier chip. Added to the new device is a 230-MHz RAMDAC and a flexible video subsystem that connects to external video I/O devices such as MPEG decoders and NTSC/PAL video encoders. Three independent DMA controllers keep commands and data moving over the chip's new bus interface, which supports 33-MHz and 66-MHz PCI clock rates as well as 66-MHz (1×) AGP.

The video-output channel offers independent timing control; it can display a subset of the main screen or a completely separate screen. This allows the P2 to be used in a "living-room PC," displaying conventional computer data on a monitor while simultaneously displaying video on a TV. The P2 goes a step beyond the usual method of implementing this feature, providing a fixed-function video window in local memory, and instead treats any video source as just another bitmap in texture memory. As a result, the P2 can apply 3D display techniques to video, enabling special effects like mirroring and blending multiple video streams.

This philosophy of applying 3D technology to 2D graphics extends to relatively "old" technology like font rendering. Fonts may be cached in local texture memory, with the 3D engine handling font rasterization. The result is unusually fast text display.

The P2 is typical of the new crop of 3D chips in its memory interface. Two to eight megabytes of 100-MHz SDRAM or SGRAM can be connected to the P2's 64-bit local memory port. The P2 uses the block-write and write-per-bit features of SGRAM, if available, so the task of clearing out Z and frame buffers during 3D rendering proceeds at an effective rate of 3.2 Gbytes/s rather than the basic 800-Mbyte/s rate of the memory bus.

Local-memory bandwidth is a good first approximation for rendering performance, and the P2 doesn't disappoint on that score. 3Dlabs claims the part is capable of 40M-pixel/s rendering rates for Z-buffered, bilinear-filtered polygons. The effective polygon rate will be limited by host CPU performance for geometry calculations, but the P2's setup and rendering engines support a peak throughput of 1M triangles/s. The 3D WinBench score for the P2 is quoted at 131 for a 233-MHz Pentium II system with AGP.

ATI Rage Pro Spans Wide Performance Range

Currently number two in the PC 3D market after S3, ATI has taken a decisive technical lead over S3's best with the new Rage Pro. The new part improves performance over previous Rage chips by increasing the speed of its local 64-bit memory interface, upgrading the standard PCI interface to support 2× AGP, and adding an on-chip setup engine and motion-compensation logic to accelerate MPEG-2 decoding for DVD playback.

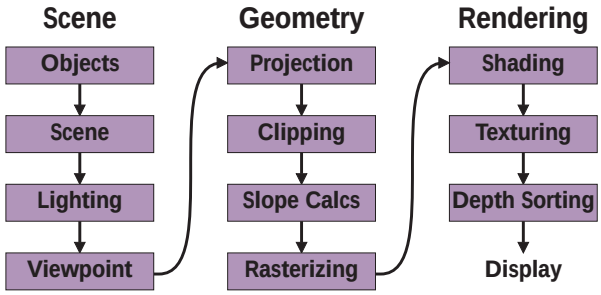

Figure 1. The basic 3D pipeline consists of three distinct tasks. Scene management prepares a database of 3D objects. Geometry processing translates object coordinates into pixel locations. Rendering applies color and texture and removes invisible pixels.

Performance is right on target for the high end of the mainstream market, with a peak throughput of 1M triangles/s and 50M pixels/s for bilinear-filtered, textured triangles. Rage Pro scores 180 on 3D WinBench on a 266-MHz Pentium II system, the highest result reported for any chip described here. This result is an estimate based on the performance of preproduction hardware, however, and should be considered roughly equivalent to the second-best score of 171 reported by Nvidia for its Riva 128.

ATI was the first vendor to announce support for the 2× mode of AGP, achieving a theoretical peak transfer rate between the chip and main memory of 533 Mbytes/s. However, this feature is probably of more real benefit from a marketing standpoint. The peak bandwidth for PC main

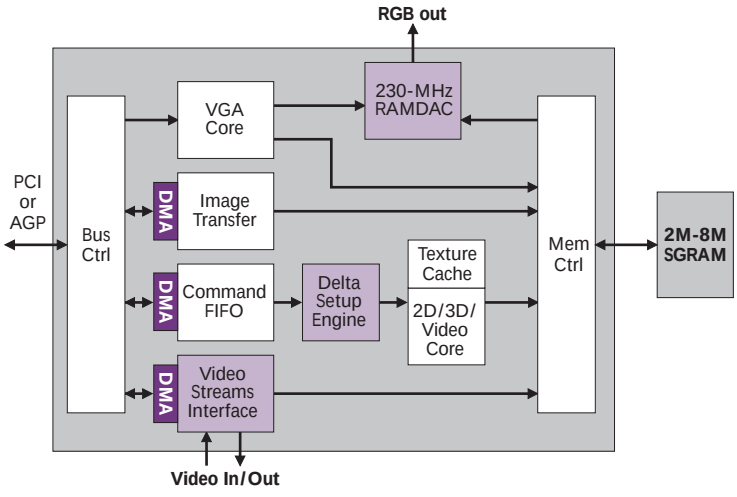

Figure 2. The 3Dlabs Permedia 2 is based on the original Permedia architecture (white blocks) but adds a setup engine, video interface, and RAMDAC (light purple blocks). The memory interface is 64 bits wide and operates at 100 MHz with SDRAM or SGRAM devices.

3D WinBench Arrives

The biggest challenge facing 3D chip vendors and OEMs is not designing a fast 3D chip but figuring out what "fast" means. Until recently, there have been no useful 3D benchmarks for PCs. Even workstation-oriented benchmarks like Viewperf are being made increasingly irrelevant by progress in 3D hardware.

The Ziff-Davis Benchmark Operation (ZDBOp) supplies a wide variety of PC benchmarks, offering complete suites of both synthetic and application-based performance-analysis tools. 3D benchmarking proved to be unusually challenging due to the rapid evolution of PC 3D hardware—performance increased 16× in the past two years—and the wide variety of quality features such as antialiasing, texture filtering, and translucency that often mean different things on different chips.

As it turns out, Intel was also working on a 3D benchmark. Last year, Intel transferred its "Mona Lisa" benchmark to ZDBOp, and it became the basis of what is now 3D WinBench. Many significant changes were made, but the basic principles remain the same: at a fixed resolution (640 × 480 with 16-bit color), several 3D scenes are rendered. The scenes have various quality features enabled and range in complexity from 900 to 6,300 triangles per frame, typical of today's game software.

The rendering speed for each test is measured in frames per second. The overall result is referenced to a 200-MHz Pentium system without 3D acceleration, which is assigned a score of 10. The benchmark also calculates the peak triangle and pixel rates for the 3D subsystem under test, and it allows the user to determine whether each quality feature the 3D chip claims to support is actually working properly.

3D WinBench currently runs only under Windows 95 and Direct3D, pending a fully compatible version of Direct3D for Windows NT. The emphasis on frames-per-second is a more subtle limitation. Real applications will never attempt to render frames more rapidly than the screen can display them—typically 60 or 75 Hz. If the 3D hardware is faster than necessary, it simply waits for the screen to "flip" before starting to render the next screen. 3D WinBench disables this delay and the related double buffering of the screen.

We hope to see a slightly different approach used for future versions of 3D WinBench. Instead of measuring frames per second for scenes of a fixed size, it would be better to measure maximum scene size for a fixed frame rate, probably 60 Hz. This would produce the critical number that software vendors need to know: how complex scenes can be and still be rendered in real time.

More information on 3D WinBench is available from ZDBOp's Web site: *www.zdbop.com*.

memories today, using 66-MHz SDRAM, is the same 533-Mbyte/s figure. Main memory is shared with the host CPU, of course, and the CPU has a higher priority for access to memory. As a result, an AGP graphics chip gets only whatever bandwidth is left over after the CPU takes what it needs for system and application software. In most systems, this available bandwidth is unlikely to exceed 200 Mbytes/s, well within the capability of the 1× AGP interfaces of other 3D chips.

Support for 2× AGP will become more important in 1998. PCs with 100-MHz main-memory arrays will boost main-memory bandwidth to a peak of 800 Mbytes/s, leaving up to half of that available for 2× AGP graphics chips.

Local memory in the graphics subsystem will remain the primary location for textures being rendered to the screen, and it carries the highest demand for bandwidth. Rage Pro supports up to 8M of 100-MHz SGRAM through a 64-bit port, matching the bandwidth of the Permedia 2. The Rage Pro can drive true-color 1600 × 1200-pixel, 85-Hz displays using the internal 230-MHz RAMDAC.

Optionally, the part may be configured with up to 16M of Samsung's dual-ported WRAM. In this configuration, the serial output ports on the WRAM chips are connected to an external RAMDAC that operates at up to 250 MHz. This option provides higher performance in high-resolution display modes, since video-refresh traffic is removed from the local-memory bus, yielding as much as 50% more local-memory bandwidth than SGRAM.

The WRAM option also allows Rage Pro to store more texture data in local memory, reducing bus traffic and boosting display quality. Given the cost of an external RAMDAC, the cost premium for WRAM chips, and the extra 8M of local memory, this high-end configuration adds about $90 to the basic bill of materials. Even if few such cards are sold, ATI benefits from giving its customers the option of offering entry-level, midrange, and high-end performance from a single chip and one set of software drivers.

Rage Pro is unique among current 3D chips in supporting a form of texture compression based on the vector-quantization algorithm for still images. This VQ compression yields an effective compression ratio of 8:1. Typical 4M configurations will have less than 1M left for texture storage after reserving space for a double-buffered frame buffer plus Z buffer; texture compression expands this to an effective 7.5M. Texture compression also reduces memory-bandwidth demands by the same factor, effectively improving the performance of the local memory array.

Eventually, texture compression will be directly supported by mainstream APIs like Microsoft's Direct3D. We expect the technique to be widely used, resulting in substantially enhanced visual quality for 3D scenes. For the time being, however, applications must be written with specific support for VQ-compressed textures. Microsoft is considering support for VQ, as well as its own TREC texture-compression algorithm, in DirectX 6.0, due early next year. In the meantime, ATI will be pursuing ISVs directly.

ATI has chosen a quadratic approximation technique for perspective correction during texturing operations. This method is somewhat less precise than the full floating-point divides used by most other 3D chips; the difference in quality is noticeable but not crippling.

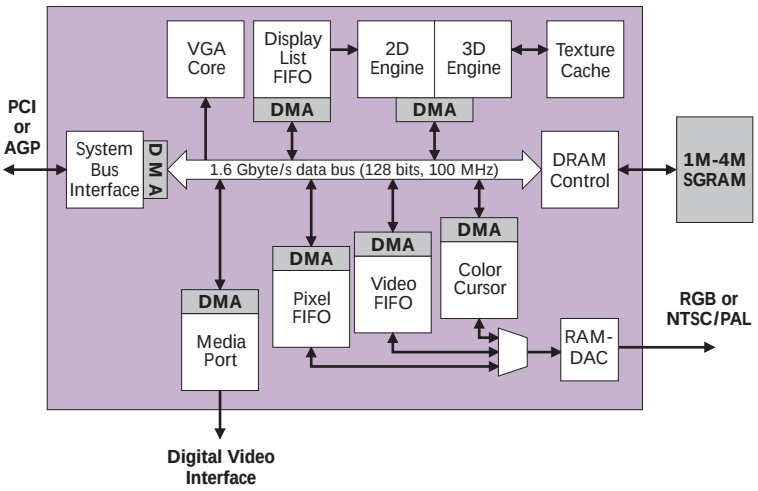

Figure 3. Nvidia's RIVA 128 is designed around a high-speed backbone bus that connects the AGP system interface to local memory. Most internal subsystems include a DMA controller and can operate in parallel within the limits of the available bandwidth.

Number Nine's Number Three Should Be a Hit

It may surprise some to learn that veteran graphics-card company Number Nine has been selling 3D chips since 1994. The first two parts, the Imagine 128 (see MPR 8/1/94, p. 5) and Imagine 128 Series 2, offered only limited 3D support and found few design wins outside of Number Nine itself.

The company's new Ticket To Ride (T2R) chip, however, may allow the company to get some sales outside its own design group. The T2R design emphasizes high-end performance at a high-end price. Like most other recent 3D chips, the T2R includes a floating-point setup engine; Number Nine rates the engine's performance at 700 MFLOPS and 1.2M triangles/s.

The T2R has a flexible 128-bit memory interface compatible with conventional DRAM, VRAM, SGRAM, and WRAM at up to 100 MHz. The part allows much larger memory arrays than the other parts described here, up to 48M for a combined DRAM/VRAM configuration. This may seem like a lot of RAM for a graphics card, but 3D can require plenty of local storage for color, alpha, Z, and texture buffers, especially in high-resolution display modes.

Number Nine also raises the bar for 3D-chip pricing, asking $125 for the T2R in 10,000-unit quantities. The chip does not include an integrated RAMDAC, raising the subsystem price to closer to $135. This is a high price to pay for good 3D performance, but fast 3D chips represent one of the best ways for OEMs to improve system performance. Spending an extra $100 on a better 3D subsystem provides better 3D performance than spending it on a faster speed grade of Intel processor.

With a 3D WinBench score of 120, the T2R is slower than competing parts on Direct3D, but the company is putting more emphasis on the professional OpenGL market (for which benchmark results are not yet available). The company recognizes that the chip's high price puts it well out of reach of high-volume applications, but unlike some vendors, Number Nine is likely to make money on whatever T2R sales it is able to achieve.

Nvidia Strikes Back With Riva 128

One of the first PC 3D chips, the NV-1 from Nvidia (see MPR 7/10/95, p. 13) failed to achieve widespread success. The chip included 3D and audio support without Sound Blaster compatibility, a combination that few OEMs needed. Nvidia also emphasized the NV-1's native ability to display curved surfaces directly without requiring the typical intermediate step of tessellation into discrete triangles. Software vendors needed to make significant changes to their code to accommodate this feature, and few did. While the NV-1 was quite good at conventional triangle rendering, this fact was often overlooked.

Late last year, Nvidia decided to scrap plans for an NV-3 derivative that would remain software- and hardware-compatible with the NV-1 and started over with a new architecture optimized from the ground up to support Microsoft's Direct3D API.

The result is the Riva 128, a very fast 3D-rendering accelerator with a fast integrated setup engine, as Figure 3 shows. Three things set the Riva 128 apart from its competitors. Of all

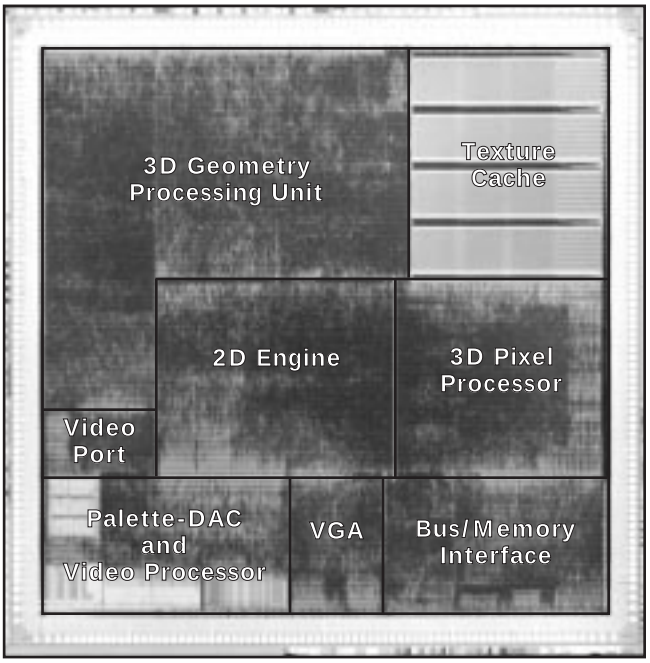

Figure 4. Nvidia's RIVA 128 includes 3.5 million transistors and is built on a 0.35-micron gate-array process with 12K of embedded SRAM for the texture cache and various internal FIFOs.

mainstream parts, the Riva 128 offers the widest and fastest local-memory bus (128-bit, 100-MHz SGRAM), integrates the largest amount of on-chip cache (12K), and has the highest transistor count (3.5 million).

To ensure access to the fab technology necessary to make such a complex part, Nvidia renewed its relationship with NV-1 partner SGS-Thomson. The Riva 128, shown in Figure 4, is made in a 0.35-micron process with five metal layers and measures approximately 75 mm^2.

The chip illustrates the complex and often contradictory demands made of 3D-chip designers. The setup engine in the Riva 128 is an impressive accomplishment. It performs 50 floating-point operations per cycle at 100 MHz, achieving a peak throughput of 5 GFLOPS and five million triangles per second, and it can sustain a rate of 1.5M triangles/s. Unfortunately, this is faster than necessary. Even a 266-MHz Pentium II cannot perform geometry calculations for more than about 750,000 triangles per second under normal application conditions. While the extra performance gives the Riva 128 some margin for future processors, the extra transistors give it a somewhat higher manufacturing cost than necessary for use with today's CPUs. Also, it is worth noting that 3Dlabs achieves similar polygon rates with its 100-MFLOPS Delta setup engine, showing the often loose relationship between peak performance ratings and effective throughput.

THE PHOTO GROUP

Trident VP Yong Yao announces the company's new 985DVD 3D chip at the PC Tech Forum.

Nvidia's decision to provide a 128-bit path to local memory yields an extremely high pixel-fill rate of 100M pixels/s for bilinear-filtered textures, two to three times faster than other mainstream parts. As a result, the Riva 128 can provide equivalent performance at higher resolutions—but the chip is hampered by a 4M limit on its local memory array. The same logic applies to the integrated 206-MHz RAMDAC. While the RAMDAC supports up to 1600 × 1200-pixel displays, true double-buffered 3D with Z sorting is not possible at resolutions above 800 × 600 pixels. Nvidia acknowledges these contradictions and expects to greatly increase the amount of local memory supported by the next version of the part, planned for this fall.

S3 Defends Turf With Virge/GX2

S3 has sold more 3D chips than any other company, but S3's offerings have not fared well in performance comparisons with the competition. Unfortunately, the Virge/GX2 falls even further behind in both performance and features, giving ATI a chance to overtake S3 in sales volume.

Among all recently introduced mainstream 3D chips, the Virge/GX2 is unique in not having a setup engine. This forces the host processor to perform all setup calculations in addition to geometry processing, reducing the peak potential throughput to less than 300K triangles/s despite the chip's rendering capacity of about 1.6M triangles/s. Like Nvidia, S3 also chose to support only 4M of local memory, and Virge/GX2 includes only a 170-MHz internal RAMDAC, rendering the Virge/GX2 unsuitable for more capable and profitable graphics-card configurations.

These decisions are likely to cost S3 some market share, but the company should still retain its sales-volume leadership for 1997, due to its brand-name recognition and ATI's unwillingness to serve third-party add-in card vendors.

The Virge/GX2 offers one feature found on few other mainstream chips, however. An integrated NTSC/PAL encoder with three-tap flicker filter allows the Virge/GX2 to drive an RGB monitor and a TV at the same time, making the chip a good choice for PCs used as part of a home theater system. This is a relatively small market, and we believe the die space used for this feature should have been used instead to provide better 3D performance. Most other chips covered here support video outputs using external encoders.

Trident Reveals Three New Chips

The newest member of Trident's 3DImàge family, the 985 is also a third-generation part, building on the success of the earlier 975 and 975DVD. At the recent PC Tech Forum, Trident's Yong Yao unveiled plans to advance the family further in the second half of the year with the 985DVD, which will double 3D-rendering performance and add motion-compensation logic to accelerate DVD playback.

Unlike most 3D chips, the 985DVD is designed to meet the needs of both desktop and mobile systems. Trident designed the chip with analog RGB outputs for desktop monitors as well as an LCD controller for notebooks and an NTSC/PAL output for an external video encoder. The 985DVD will provide a 2× AGP system interface with full support for sideband signals and read pipelining, plus a 64-bit, 100-MHz controller for up to 8M of local SGRAM.

The 985DVD's polygon throughput is rated at 1M triangles/s, while the chip's rendering speed is rated at 70M pixels/s for bilinear-filtered texture fill. The triangle rate is limited by the speed of available host processors, while the pixel rate probably represents a peak figure that few real-world applications can achieve. Trident estimates the new chip's 3D WinBench score at over 150, based on the 985's score of 120.

Like the other chips discussed here, the 985DVD will be made on a 0.35-micron process. The chip has four metal layers and 2.4 million transistors on a 140-mm^2 die. Trident will sell the part for $45, yielding an estimated 40% margin over the $32 manufacturing cost estimated by the MDR Cost

Model. Because the part achieves a relatively low transistor density—only one-third that of Nvidia's Riva 128, for example—we believe Trident will be able to shrink this part significantly and increase its margins.

Yao also announced plans to build a single-chip implementation of Microsoft's Talisman 3D-rendering architecture. The new part, due out early next year, folds together the tiler and compositor functions from the original Talisman reference design (see MPR 8/26/96, p. 5) into one logic block and uses embedded DRAM for the compositing buffer. Trident's implementation will not include the media processor or digital audio/video features of the Talisman reference design, but it will be possible to use a separate media processor, such as Philips' TriMedia, to provide a complete multimedia subsystem in two chips.

Because of plans at Trident and other companies to bring single-chip Talisman solutions to market by mid-1998 and Cirrus Logic's cancellation of the original tiler and compositor chip designs, Microsoft has canceled its own plans for the reference design. This pushes back the first phases of the Talisman schedule by six months or so, but the gating factor for Talisman's mainstream success has always been silicon technology. Talisman will still have an opportunity for mainstream success in 1999 as Talisman-aware games begin to appear and Talisman graphics cards drop down into the sub-$300 range.

3D Vendors Share Common Goal

If it seems that most PC 3D chip vendors are reading from the same book, it's because they are. A year ago, Microsoft told the industry to prepare for chips capable of rendering a million polygons per second—and now, sure enough, that's what every chip company claims to have achieved.

Only a few parts can sustain this rate—and some fall well behind—but at least it's a rough estimate of where the 3D hardware industry is today. Most vendors support 100-MHz SGRAM and include on-chip setup engines. Chip companies that fall behind, like S3, are at risk of being overtaken by more agile competitors.

It appears that 3D chip vendors have now caught up with the normal Moore's Law curve for transistor count and clock rate. These factors will improve at only the usual 2× every 18 months from now on. Overall 3D performance will continue to improve at four times this rate, and the difference will come from enhanced 3D architectures.

In the next year, 3D geometry acceleration will become a standard feature on all 3D chips, solving the computation problem. Solving the bandwidth problem will require larger internal SRAM arrays for texture caches and embedded DRAM for on-chip Z and smaller frame buffers.

Price and Availability

Table 2 shows price and availability information. To contact 3Dlabs (San Jose, Calif.), call 408.436.3455 or access the Web at *www.3Dlabs.com.* To contact ATI (Thornhill, Ontario, Canada), call 905.882.2600 or access the Web at *www.atitech.com.* Contact Number Nine (Lexington, Mass.) at 617.674.0009 or *www.nine.com.* Contact Nvidia (Sunnyvale, Calif.) at 408.720.6100 or *www.nvidia.com.* Contact S3 (Santa Clara, Calif.) at 408.588.8000 or *www.s3.com.* Contact Trident (Mountain View, Calif.) at 415.691.9211 or *www.trid.com.*

Off-chip memory will continue to be necessary for larger frame buffers. High-resolution true-color displays require 4M–8M frame buffers, more memory than vendors can afford to put on a single chip within the next few years. Texture storage will also exceed the capacity of on-chip DRAM; most chips in 1998 will need another 4M–8M of local memory just for this purpose.

ATI's Rage Pro performs the best on 3D WinBench and is a strong choice for mainstream PCs. Nvidia's chip is right behind Rage Pro in performance, and its more powerful setup engine will scale better in future systems. Number Nine's expensive T2R is best for workstation applications. Trident's 985DVD is a solid desktop offering that is also a strong notebook product. Permedia 2 is unique in its growth path to even faster 3Dlabs products. S3's GX2 lacks a setup engine and thus will not perform as well. All of these chips, however, offer a significant boost in 3D performance over current 3D products, providing a happy 1997 holiday season for PC consumers. M

In the next issue, we will describe recently announced 3D chips with more innovative and advanced architectures, including parts from 3Dlabs, Rendition, S-MOS, and VideoLogic.

	3Dlabs Permedia 2	ATI Rage Pro	Number Nine T2R	Nvidia Riva 128	S3 Virge/GX2	Trident 3DImàge 985DVD
Bus interface	AGP 1×	AGP 2×	AGP 1×	AGP 1×	AGP 1×	AGP 2×
Fastest memory type	SGRAM	WRAM	WRAM	SGRAM	SGRAM	SGRAM
Memory width	64	64	128	128	64	64
Memory clock rate	100 MHz	100 MHz	100 MHz	100 MHz	100 MHz	100 MHz
Texture cache	n/a	4K	8K	8K	None	4K
Setup engine	Yes	Yes	Yes	Yes	No	Yes
Peak triangle rate	1M	1M	1.2M	5M	1.6M	1M
Peak pixel rate	40M	45M	35M	100M	35M	70M
3D WinBench score	131	180*	120	171	n/a	150*
RAMDAC bandwidth	230 MHz	230 MHz	None	206 MHz	170 MHz	230 MHz
Video in/out ports	In/Out	In/Out	No	In/Out	In/Out	In/Out
Dual display	TV	TV	No	TV	TV	TV/LCD
Retail price	$35	$30	$125	$30	$30	$45

Table 1. These six chips share several features, including AGP support, fast local memory, and 0.35-micron process technology. Wider memory buses and larger texture caches distinguish the best from the rest. n/a: not available *estimates (Source: vendors)