Napalm Ignites Graphics Market

New 3dfx VSA-100 Chip Offers Improved Quality, Unprecedented Scalability

by Peter N. Glaskowsky

At Comdex last month, 3dfx dropped a bomb on the PC graphics industry with its introduction of not one but three 3D-accelerator product lines based on a single new chip code-named Napalm. The new VSA-100 graphics chip forgoes some of the special features of its competition, such as geometry acceleration and support for double-data-rate (DDR) memory, to deliver as much raw 3D-rendering power as the company can squeeze in. Indeed, the chip's new architecture can deliver more performance than any single-chip rendering engine, because 3dfx has added scalability support that allows up to 32 of the devices to work in parallel.

The VSA-100, which is not due until 2Q00, does include a few quality-oriented features that have been conspicuously absent from previous 3dfx products and adds other new features enabled by the massive pixel-drawing rates of the parallel configurations. The VSA-100 is 3dfx's first graphics chip to support 32-bit color 3D rendering, a 24-bit depth buffer, and an 8-bit stencil buffer. The chip also supports a 4× AGP interface and 64M of local memory.

The company will use the VSA-100 in at least three board configurations—the single-chip Voodoo4 and two- and four-chip Voodoo5 products. The Voodoo5 boards support 3dfx's T-buffer, a simplified version of the accumulation buffer found on some high-end visualization systems, which can be configured to provide full-scene antialiasing, motion blur, depth-of-field, and other special effects.

Figure 1 shows a four-chip Voodoo5 configuration, which relies on a separate AGP–AGP bridge ASIC. The two-chip board has both devices connected in parallel to its AGP connector, a configuration not originally considered in the AGP specification but one that other vendors, including ATI, have implemented successfully.

Quantum3D will deliver products with 4 to 32 VSA-100 chips operating in parallel to support extremely fast high-resolution graphics subsystems for professional visual-simulation applications such as flight simulators. The maximum configuration, which will be implemented across multiple boards, will offer a peak rendering speed in excess of three billion antialiased pixels per second. This is the sort of scalability sought by SGI and Nvidia in their codevelopment plans announced earlier this year (see MPR 9/13/99, p. 4).

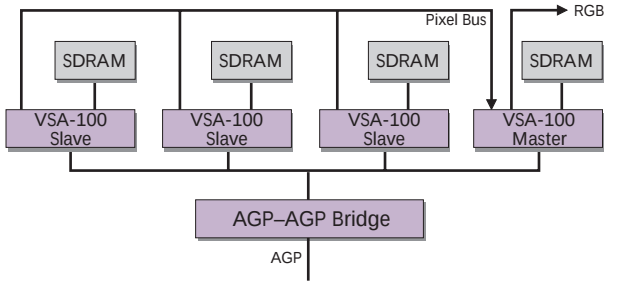

Figure 1. The 3dfx Voodoo5 graphics card will include four VSA-100 graphics chips, a bridge ASIC, and four banks of local memory.

Focus Expands Beyond Fill Rate

Because 3dfx derives much of its revenue from sales to gamers through retail and OEM channels, the company has always been tightly focused on delivering what gamers want.

In the early days of PC 3D, 3dfx was instrumental in delivering superior visual quality to its customers, since hardware-accelerated games looked much better than those driven by software alone. The company's current Voodoo3 products offer few quality features not found in the original Voodoo, however. It has left out quality-oriented features such as 32-bit color and high-resolution texture support in order to devote more transistors to faster 16-bit graphics.

This focus put 3dfx at a competitive disadvantage against chips with longer feature lists. In response, 3dfx has equipped the VSA-100 with support for 32-bit color, textures up to 2,048 × 2,048 pixels in size, full-screen antialiasing (in multichip configurations), and other quality features that have previously been available in competing 3D chips.

The company has gone one step beyond its competitors in providing hardware support for certain special effects. These effects are enabled by what 3dfx calls the T-buffer. The T-buffer is similar to the accumulation buffer used by SGI and other high-end 3D subsystems, but it works at the triangle level rather than on the complete frame buffer, making it more flexible. The technique works by rendering slightly different images into two or more T-buffers, then combining these multiple buffers into a single image.

Each VSA-100 can manage two T-buffers. Two chips can implement a four-sample antialiasing filter by offsetting each buffer by half a pixel in each of four directions from the nominal pixel center. This form of antialiasing cuts the effective fill rate by 75% but greatly improves visual quality. Demonstrations of this technique by 3dfx show that an antialiased 640 × 480-pixel display can look about as good as a non-antialiased display with 1,024 × 768-pixel resolution.

By rendering a moving object to a slightly different position along the direction of motion in each buffer, 3dfx's T-buffer can simulate motion blur. Depth-of-field blur can be simulated by rendering objects at the virtual focal plane to the same location in each buffer, while objects that are nearer or more distant are rendered with successively greater X-axis and Y-axis displacements.

Texture Compression Boosts Texture Size

The VSA-100 includes texture-compression support for two sets of algorithms. One set, developed by S3 and included in Microsoft's Direct3D, is already supported by some games. The other set, designed by 3dfx, is said to offer superior visual quality and be easier to implement than the S3 methods. These new algorithms have been released into the public domain by 3dfx, which hopes they will be adopted by Microsoft and other 3D-chip companies. Both sets of algorithms make practical the use of larger, more detailed textures, improving visual quality.

The company chose to omit one feature found in competing chips that will eventually become more important than raw pixel-drawing speed. Nvidia's latest 3D chip, the GeForce 256 (see MPR 9/13/99, p. 4), and S3's Savage2000 include integrated geometry-processing engines that can greatly improve polygon throughput. But 3dfx believes that geometry acceleration will not become sufficiently important to end users before its own geometry-accelerated products become available, in about a year. Without a geometry engine, the VSA-100 is smaller (14M transistors vs. the GeForce's 23M transistors), making multi-chip configurations more affordable.

Though the VSA-100 runs at the same 166-MHz clock rate of the Voodoo3, the new chip offers twice the V3's performance on single-textured triangles. Pipeline improvements and texture compression will yield additional improvements on real-world applications, estimated by 3dfx to be up to 60% over the older design. Even so, 3dfx may have fallen short in single-chip rendering performance. The GeForce's peak pixel-drawing rate is 44% faster than that of the VSA-100 on both 16- and 32-bit drawing operations.

By the time the VSA-100 ships in volume, Nvidia may have released its planned successor to the GeForce 256, a 0.18-micron part code-named NV-15. Though this new part will be substantially faster than Nvidia's current offering, we do not expect it to match the scalability features of the VSA-100, and this is where the true power of the new Voodoo lies. Because of its support for multichip configurations, the VSA-100's potential for rendering throughput outstrips that of all competitors.

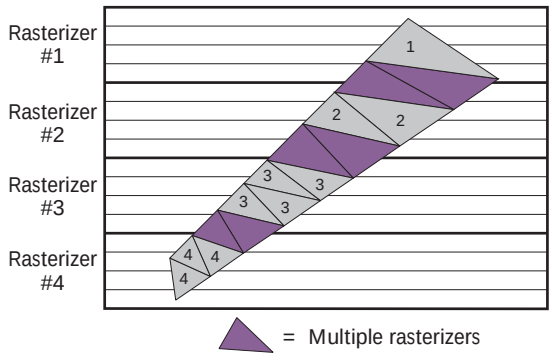

Figure 2. A four-way Voodoo5 configuration spreads rasterizing duties across all four VSA-100 chips, each of which handles up to 128 contiguous scan lines.

Price & Availability

The 3dfx (*www.3dfx.com*) shipments of VSA-100-based products should commence by 2Q00. Board pricing will range from $179 for a single-chip Voodoo4 PCI or AGP card with 32M of memory to a $599 four-chip Voodoo5 AGP card with 128M of memory.

The VSA-100's scalability is achieved through an updated form of 3dfx's scan-line interleaving (SLI) technology, used in the earlier Voodoo2. Two-board Voodoo2 configurations rendered even-numbered scan lines on one board and odd-numbered scan lines on the other. In the VSA-100's new SLI approach, each chip can render from 1 to 128 contiguous scan lines while the other chips are rendering the same number of lines in parallel. One VSA-100 chip receives video data from the other VSA-100s in the subsystem and generates the final display. As Figure 2 shows, this division of labor—similar to that found in the Glint R4 chip from 3Dlabs (see MPR 11/16/98, p. 20)—means that some polygons need not be processed by all of the VSA-100 chips, which increases overall efficiency.

Gamers Just Can't Get Enough

Though the VSA-100's 2D and video engines are only slightly improved over those in the Voodoo3, they are still quite adequate for most buyers. We expect the single-chip Voodoo4 cards to maintain 3dfx's dominance of retail graphics-card sales in 2000. The Voodoo4 should also expand 3dfx's role in the OEM market, mostly in high-end enthusiast PCs.

Unfortunately for 3dfx, the Voodoo4 is unlikely to extend the company's reach into mainstream systems. OEM buyers will see no benefit from the VSA-100's scalability, and the new quality features will likely be perceived as merely matching the feature sets of less-expensive 3D chips from ATI, Nvidia, and S3.

The four-chip Voodoo5 board is likely to be the most interesting to 3D-gaming enthusiasts. No other PC graphics card is likely to surpass this configuration's 1.3 Gpixel/s of peak rendering performance in 2000. For gamers, the lack of geometry acceleration is the only meaningful omission from the Voodoo5, but it is very unlikely that geometry acceleration will become a required feature for any PC games during the Voodoo5's lifespan.

Even with an estimated street price of $599 for four-chip Voodoo5 boards, 3dfx is likely to sell several hundred thousand of them. The company sold a comparable number of Voodoo2 SLI board sets to gamers in 1998 for $400 or more, and with no 2D support on Voodoo2, these customers also had to buy a separate 2D-graphics card. Clearly, there is a strong niche market for high-performance (and high-priced) 3D accelerators, and 3dfx remains the only company willing to design products to serve it. M