NEC Sets OpenGL 3D Speed Record

New TE4 Chip Set Outruns Windows NT CAD Competition

by Peter N. Glaskowsky

NEC's new TE4E is the fastest Windows NT OpenGL 3D accelerator in the world—but it's available only in Japan. NEC's Express5800 workstation, equipped with the TE4E graphics card based on NEC's TE4 chip set, achieves a score of 32.4 on the ProCDRS-02 benchmark and is expected to hit 35 with the next release of the silicon. The latter score is about 40% higher than the next-fastest Windows NT 3D-CAD accelerator, as Figure 1 shows.

The Express5800 workstation is not available in the U.S. market, and NEC has no plans to bring it here. NEC attended the recent Workstation Leadership Forum to drum up OEM interest in the TE4 series. If NEC can sell the TE4 into workstations from major U.S. OEMs, it will certainly provide serious competition for other graphics vendors in this already overcrowded market segment.

New Design Based on Proprietary Architecture

The TE4 chip set is the fourth generation of NEC's TE series, which dates back to the early 1990s. Originally used in a Unix-based workstation family, the TE series was retargeted to the growing Windows NT workstation market in 1995.

As Figure 2 shows, the TE4E has a fairly conventional organization. The GA400 geometry processor acts as a bridge between AGP and a local 64-bit, 66-MHz PCI bus. This PCI bus is shared between an ATI Rage chip for VGA compatibility (eliminating the need for a separate graphics card) and NEC's own rendering engine.

This rendering engine, dubbed PEC (for pixel engine core), derives much of its performance advantage from its wide, fast memory interfaces. Four different memory arrays are connected to the PEC, including a 256-bit-wide texture memory and three local-memory arrays totaling 288 bits in width. The aggregate bandwidth from these arrays is 6.8 GBytes/s, fast enough to sustain a 200-Mpixel/s drawing rate for trilinear-textured polygons, or 400 Mpixels/s without textures.

These pixel-drawing rates may seem low compared with those of today's best PC-gaming accelerators, some of which can draw textured pixels at 366 Mpixels/s. These gaming chips, however, rely heavily on on-chip caches to support texture mapping. Most 3D games are designed to provide good temporal locality for texture accesses. Caching is less effective in CAD applications, since 3D models are rarely optimized to improve rendering performance.

Even the best PC gaming chips have no more than 3.2 GBytes/s of peak bandwidth to their single-ported local memories. This single port carries all texture, frame-buffer, Z-buffer, and display-refresh traffic. As a result, mainstream PC 3D chips can't match the sustained throughput of professional 3D products such as the TE4.

Because of its many wide memory interfaces, the PEC comes in an immense 1,296-pin BGA package. The PEC is currently made in a 0.25-micron process and operates internally at just 100 MHz, leaving some headroom for future performance improvements as better ASIC fab technology becomes available.

The use of Mitsubishi's 3D-RAM memory chips (see MPR 8/2/99, p. 12) for the frame, Z, and window-ID buffers also contributes to the effective performance of the TE4E. Each 10-Mbit 3D-RAM chip has an integrated 32-bit ALU that handles the final steps in the 3D-rendering pipeline as data is written into the chips. Operations such as depth testing, alpha blending, and stenciling that would require a read-

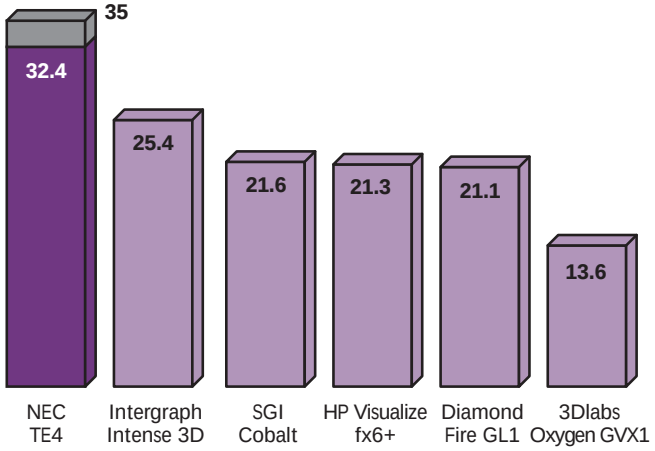

Figure 1. The TE4E outruns the top five Windows NT 3D-CAD accelerators on SPEC's ProCDRS-02 benchmark. The gray area represents NEC's promised future performance improvement.

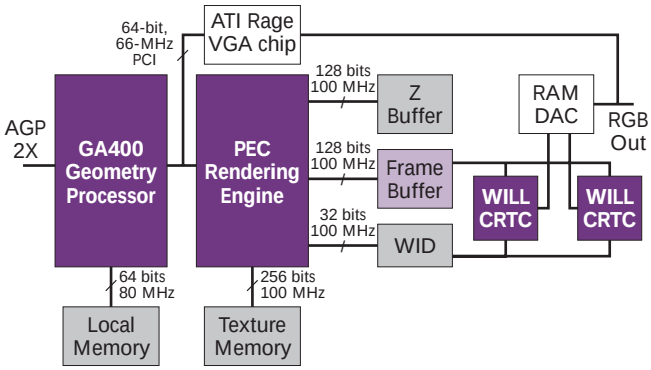

Figure 2. The TE4E uses four proprietary ASICs, five memory arrays, and commodity VGA and RAMDAC chips.

Price & Availability

The TE4E graphics card is currently available only in NEC's Express5800 workstations, which are sold only in Japan. Prices for these workstations begin at ¥1.6 million (about $13,750).

More information is available from NEC on the Web: *www.nec.co.jp/english/today/newsrel/9906/2101.html*.

modify-write cycle on conventional SDRAMs can be accomplished in a single write cycle on a 3D-RAM.

Since 3D-RAMs are dual-ported, display-refresh traffic does not interfere with rendering operations. Instead, the 3D-RAM refresh ports are attached to two WILL (window ID controller) chips, additional ASICs that buffer and reformat pixel data going to the board's RAMDAC.

Geometry Accelerator Boasts High Throughput

Most 3D CAD work is done in wireframe mode, and CAD benchmarks are weighted heavily toward this mode. The limiting factor for wireframe drawing is the rate at which vertex data can be transformed. Today's best professional-3D cards use hardware geometry engines to perform these transformations, removing the host CPU as the bottleneck in this process, and the TE4E is no exception to this rule.

NEC's GA400 geometry engine, which was developed with RealVision, provides 9 GFLOPS of performance, allowing it to process about 12.5 Mpolygons/s with up to two light sources. Figure 3 shows the chip's AGP and PCI interfaces, each of which provides a peak bandwidth of 533 MBytes/s. NEC says this is enough bandwidth for 16 Mpolygons/s. NEC is working on designs using 4× AGP, which will provide enough bus bandwidth to handle about 32 Mpolygons/s.

Another 2 GFLOPS are performed in the PEC's on-chip setup engine, providing the same effective polygon throughput on the simpler setup operations. Setup processing is also crucial to good benchmark scores; it is the setup engine that generates screen coordinates for wireframe drawing and textured polygons alike.

The GA400 has its own local memory, used to store vertex arrays and transform matrices. OpenGL provides a mechanism to allow sets of vertex data to be saved and reused from one frame to the next. This reuse reduces the work needed for subsequent frames, thus improving overall performance.

Dual-Board Configuration Doubles Throughput

For applications where the performance of a single TE4E board isn't enough, NEC is working on a dual-board configuration known as the TE4M. Because 3D rendering is readily parallelizable, the TE4M should achieve nearly twice the TE4E's sustained polygon and pixel throughput.

The two cards will be linked by ribbon cables, allowing them to share polygon and pixel data. Both GA400 chips work in parallel to perform geometry processing, then vertices are distributed to one or both cards as needed for rendering. NEC has not yet announced its plans regarding the TE4M, but we expect it to ship by the end of this year.

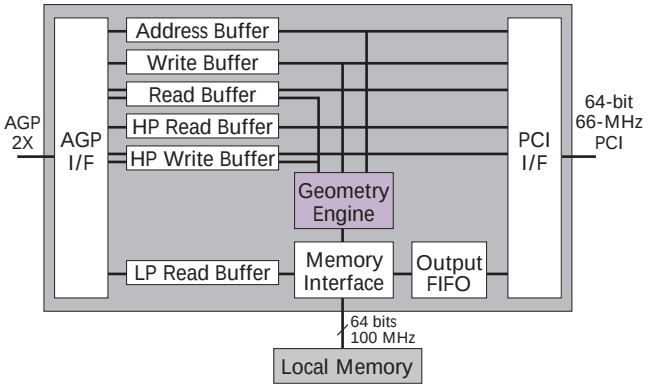

Figure 3. The GA400 geometry engine incorporates a buffered AGP-to-PCI bridge and supports 64M of local SDRAM memory.

Business Opportunity Is Real, but Fleeting

Though NEC is clearly interested in the U.S. OEM market, it recognizes that this market is already overcrowded with products from several independent professional-3D vendors, including 3Dlabs, Evans & Sutherland, IBM/Diamond, Intergraph, and Real3D. These products are sold to system OEMs without 3D accelerators of their own.

With the bulk of the demand for PC workstations being met by OEMs with their own proprietary solutions (HP, IBM, Intergraph, SGI), the remaining market for independent 3D products is small, just a few hundred thousand units a year. Though this represents several hundred million dollars in revenue, the R&D and manufacturing costs associated with high-performance 3D products are high, and only the fastest products command high prices.

Despite the small size of this opportunity, NEC isn't the only vendor hoping to break into this field in the next year. PixelFusion, Raycer, and SP3D have each announced they are working on professional 3D accelerators.

NEC's performance may lead the industry right now, but its competitors are not standing still. Even the dual-board TE4M could be left far behind by future members of 3Dlabs' Jetstream or Intergraph's Wildcat product families. Both of these companies have announced plans for 3D accelerators based on their current chip sets, but with four times the performance of their current boards.

Despite the stiff competition, we believe NEC should continue looking for U.S. customers. The company's performance lead is already substantial, and it is working on even faster products. Because the TE4 series is fully OpenGL compatible, creating a U.S. product is more a matter of translating the documentation than creating new software. These costs should be small compared with the potential revenue NEC would gain from even a moderately successful product. M