

Intel Innovates With Integrated Graphics

New 810 Chip Set Combines 3D Acceleration, New Core-Logic Architecture

by Peter N. Glaskowsky

Intel's new 810 chip set combines core logic, graphics, flash memory, and a random-number generator. Many of these elements are new—but as for the 810's architecture, it's déjà vu all over again.

It's been four years since Cirrus, Opti, VLSI, and Weitek tried—and failed—to change the way mainstream PCs are designed. In 1995 these companies introduced chip sets using a unified-memory architecture (UMA), in which a single bank of DRAM provided both main and graphics memory (see MPR 6/19/95, p. 1). Intel later brought out its own UMA-capable chip set, the 430VX, (see MPR 2/12/96, p. 4). The UMA revolution had fizzled by this time, however, and the VX was rarely used in its UMA configuration.

UMA didn't disappear entirely. Cyrix's UMA-based MediaGX (see MPR 3/10/97, p. 1) led to the emergence of the sub-\$1,000 PC market. Although the MediaGX has faded away, SiS, Trident, and VIA have taken up the cause with UMA chip sets for low-end PCs and notebooks.

The 810 (formerly code-named Whitney) allows Intel to strike back at its Asian competition with a low-cost solution that offers the power of the Intel brand name. With graphics technology derived from Intel's low-end 740 graphics chip, a new chip-partitioning scheme, and no ISA bus, the 810 will enable OEMs to offer Celeron PCs at prices below anything seen to date. We expect these systems to sell well, but a few of the old UMA problems are still lurking in the 810.

Design Debuts New Hub-Based Hierarchy

The 810 uses a new hub-based hierarchy that we expect to see on future Intel chip sets. As Figure 1 shows, the 810 still uses two primary chips, but the connection between the two is now a proprietary interface instead of a standard PCI bus.

The 82810 Graphics and Memory-Controller Hub (GMCH) connects the CPU bus to up to 512M of 100-MHz SDRAM main memory and the new hub interface. The GMCH supports 66- and 100-MHz processor buses. The 810's integrated graphics controller is also on the GMCH, along with analog RGB and digital flat-panel monitor ports.

The hub interface provides 266 MBytes/s of peak bandwidth, twice that of the PCI bus used in earlier chip sets to connect the north and south bridge chips. The GMCH data sheet shows 11 signal pins plus a differential strobe signal, but it provides no other clues to the operation of this interface. We assume it transfers 8 bits on both edges of a 133-MHz clock.

The hub interface connects to the 82801 I/O Controller Hub (ICH), which connects to audio, disk, USB, super-I/O, and PCI peripherals. The ICH has another new

Intel-proprietary interface for the 82802 Firmware Hub (FWH), which contains flash memory for the system BIOS and a unique new hardware random-number generator.

The hub architecture allows Intel to mix and match components among its chip sets. We expect the 820 (Camino) chip set, due in September, to have a similar architecture, and Intel may reuse the ICH and FWH chips from the 810 with the 820's new GMCH, which should have a 4× AGP interface in lieu of the 810's on-chip graphics engine.

Traffic between the 810's integrated peripheral interfaces and the CPU or main memory should move faster in the 810 than in previous chip sets. The PCI bus in earlier designs could cause substantial latency for such transfers. PCI devices, however, are now effectively further away from main memory, and this may slightly reduce the performance of any PCI-bus peripherals added by OEMs or end users.

UMA Architecture Limits Overall Performance

UMA, the most fundamental feature in the 810 design, is also its most serious limitation. Intel calls its UMA implementation Dynamic Video-Memory Technology (DVMT), but in most respects it is a classic UMA design. With a single DRAM array for both system and graphics memory, the 810 inherits some problems from earlier UMA chip sets.

The primary problem with UMA is the competition for main-memory bandwidth between the CPU and the graphics subsystem. Conventional PCs use a separate memory array for graphics, offloading traffic for drawing operations and

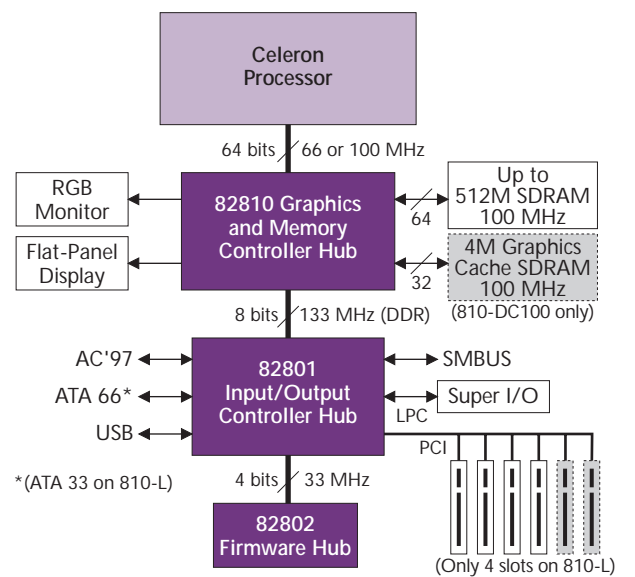


Figure 1. Intel's 810 chip set consists of two primary ASICs plus a flash memory chip that also includes a random-number generator.

752 Offers Discrete 3D Option

Customers who desire Intel 3D-graphics technology with greater performance and more flexibility than provided by the 810 chip set can select Intel's 752 graphics chip, a discrete graphics accelerator that works with any AGP-equipped chip set. At 100 Mpixels/s, the 752 renders about 25% faster than the 810. The discrete chip supports up to 16M of 133-MHz SDRAM local memory, and it has all the key features of the 810's graphics subsystem, including anisotropic texture filtering and a flat-panel display interface. Like the 810, the 752 handles only 16-bit-color 3D rendering. To these features, the 752 (née Portola) adds a video I/O interface port.

The Intel 752 is priced at \$19.50 in 10,000-unit quantities, about the same as Nvidia's faster and more capable Vanta (see MPR 4/19/99, p. 17). The 752 offers one thing Nvidia can't provide, however—the Intel brand name. Without that, the 752 would be unlikely to attract many buyers. As an Intel product, the 752 is likely to match the modest popularity of its predecessor the 740.

display refresh. UMA systems use one array for both. Applications that make heavy use of graphics take a triple hit on main-memory bandwidth—to run the program, draw the graphics, and refresh the display.

When the 810's graphics subsystem is configured to drive a high-resolution true-color monitor, display-refresh traffic (169 MB/s for a 1,024 × 768-pixel screen with 24-bit color refreshed at a 75-Hz rate) uses up about a third of the available main-memory bandwidth and substantially reduces the system's overall performance. The MediaGX uses frame-buffer compression to mitigate this impairment, but Intel does not use this technique in the 810.

The bandwidth lost to display refresh creates a hidden cost for UMA systems. To achieve performance comparable to that of a non-UMA machine, the 810 requires the buyer to move up about one CPU speed grade. A non-UMA Celeron-433 system should be about as fast as a comparably equipped machine that combines a 466-MHz Celeron with the 810 chip set. The price difference between Celeron speed grades, typically \$10–\$30, could wipe out the cost savings offered by the 810. Like those early UMA systems, the 810 is likely to fare poorly when compared against systems with the same CPU, as is common practice in computer magazines.

The 810 has a very good implementation of UMA at least. The chip set uses a four-level arbitration scheme to manage main-memory accesses. The CPU normally has the highest priority, followed by the graphics engine. Intel says the 810's best-case average latency to main memory than its mainstream 440BX chip set.

Isochronous tasks such as display refresh are usually satisfied by the remaining available bandwidth. If CPU or

graphics operations hold off refresh accesses long enough to deplete an on-chip FIFO below a trigger point, however, the refresh controller can issue high-priority requests to refill the FIFO. This override function ensures that the display never blanks out, even with the most demanding applications.

An optional display cache consisting of 4M of 100-MHz SDRAM—supported only by the DC100 version of the chip set—stores command and Z-buffer data during 3D rendering, boosting 3D performance by up to 30%. This reduces bandwidth demands on main memory when 3D applications are running but has no effect on 2D operations. Other versions of the 810 store these data in main memory.

Intel has done a good job of minimizing a common problem in early UMA products. These early chip sets permanently allocated part of main memory to the graphics frame buffer at boot time, making it unavailable to the operating system. With the 810, just 1M of memory is permanently allocated to graphics at boot time for VGA emulation. A 1,024 × 768 Windows desktop in 8-bit color can fit into the default allocation. The memory needed for larger resolutions (up to 1,600 × 1,200), greater color depth (16 or 24 bits), and 3D support is allocated dynamically once the operating system loads. The allocation can amount to 10M, or 6M if the display cache is present. When the extra memory is no longer needed for graphics, it is returned to the operating system.

Graphics Performance and Features Are Average

With a rendering rate of just 80 Mpixels/s, the 3D engine in the 810 is slower than many of today's least expensive discrete graphics chips and just 20% faster than Intel's stand-alone 740, announced nearly 15 months ago (see MPR 2/16/98, p. 1). (The new 752, which shares its 3D core with the 810, is slightly faster at 100 Mpixels/s; see sidebar.)

Because of the close coupling between the CPU, graphics engine, and memory controller, the 810's performance on benchmarks and real applications is more encouraging. Intel reports 384 3D WinMarks on the Ziff-Davis 3D WinBench test for the display-cache configuration of the 810 with a 466-MHz Celeron processor and the fastest available PC100 SDRAM main memory. This score is well below what can be achieved by midrange discrete graphics chips but is better than the performance of competing integrated solutions.

The new chip set definitely lags behind the competition in rendering quality. Like the 740, the 810 is limited to 16-bit color for 3D graphics. Most other graphics-chip families have moved on to 24-bit true color, which is better for nongame 3D applications such as Web-based shopping and searching, computer-aided design, and data visualization.

To its credit, the 810 (and the 752) supports a 16-tap anisotropic texture filter like that found in Nvidia's newest high-end chip, the TNT2. When this filter is used, the 810's rendering speed drops to just 20 Mpixels/s, too low to be useful for real-time rendering except at low screen resolutions.

The 2D engine in the 810 should meet the needs of most users, if only because 2D acceleration technology has long

since passed the point of diminishing returns. Being so closely coupled to main memory and the CPU also helps the 810's performance in image-editing and digital-video applications. Pixels generated by image-processing or video-decompression algorithms running on the host processor can be moved to the frame buffer just as quickly as to main memory, since the frame buffer is part of main memory. This eliminates the delay for accessing discrete AGP or PCI graphics chips.

Compatibility with the latest flat-panel displays is provided by a digital I/O port that complies with all the major flat-panel interface specifications. This port can drive flat panels of up to 1,280 × 1,024 pixels, somewhat less resolution than offered by competing chips from ATI and others. Like most such chips, the 810 requires external line drivers; ATI now offers graphics chips with integrated drivers that reduce the overall cost of implementation as well as board area.

Flash Chip Features Random-Number Generator

A unique feature of the 810 is its inclusion of a true random-number generator (RNG). Intel announced it was working on RNG technology last fall (see MPR 10/5/98, p. 16); the 810 is the first implementation of this research.

The circuitry for the RNG function is found in the FWH chip, and is located there because Intel plans to add other security features (such as secure data storage) to future versions of this chip. The RNG consists of a ring oscillator with an operating frequency modulated by the thermal noise from a resistor. The output from this circuit receives additional processing in hardware. An Intel-supplied software driver performs a final hash operation to yield a bitstream in which each successive bit is unrelated to those that precede it—the mathematical definition of randomness.

The 810's RNG operates at a relatively low 75 Kbits/s, too slow to be used directly by some programs. Instead, Intel recommends using the output of the 810's RNG to seed a conventional pseudo-random-number generator (PRNG). By reseeding the PRNG frequently, users can foil any effort to deduce the pattern of the PRNG's output. This ensures the cryptographic security of the resulting sequence of numbers, aiding electronic commerce, statistical analysis, and other applications that rely on randomness.

Three Configurations Are Offered

The top-of-the-line 810-DC100 has the largest set of features, including the display cache, an ATA-66 disk controller, manageability features such as remote wakeup, and support for six PCI slots. Its better 3D performance makes it the preferred selection for users who might occasionally play 3D games.

The middle configuration, known simply as the 810, drops the display cache. This makes it a better choice for most corporate desktops, where 3D is rarely if ever used.

The least expensive version, known as the 810-L for its low-end orientation, provides only an ATA-33 disk controller. It also lacks the 810's manageability features and can manage just four PCI slots.

Pricing and Availability

The 810 chip set is available in three versions: the entry-level 810-L (\$26.50), the 810 (\$30.50), and the 810-DC100 (\$34). The optional 4M of display cache SDRAM would add about \$6 to the cost of a DC100-based system.

The 82810 chip comes in a 421-contact BGA, while the 82801 uses a 241-contact BGA, and the 82802 is available in a 40-pin TSOP or a 32-pin PLCC package.

All versions are shipping now, and Intel expects systems using the 810 chip set to ship in early June.

All versions of the 810 provide an integrated audio controller with an AC-97 audio codec interface. This controller further enhances the cost advantages of the 810 by cutting another few dollars off the parts bill.

The 810-L is just \$1 more expensive than Intel's 440ZX Celeron chip set, which does not include graphics or audio (see MPR 1/25/99, p. 4). Even the cheapest PCI graphics chip would make the ZX solution more expensive than the 810. OEMs could still save money by going with Intel's \$15 440EX (or similar Asian chip sets) and a cheap graphics chip, but such a system would compare poorly with the 810 on both performance and features and would command a lower price.

New Design Leads to New Tradeoffs

The 810's mediocre performance is a consequence of its implementation, not its architecture. Integration can be used to achieve high performance, as proved by SGI's Visual Workstation systems (see MPR 2/15/99, p. 12)—but only at a high cost. Low-cost integration leads to compromises. In the 810, the compromise is the competition for main-memory bandwidth. Ordinary office-productivity applications will work fine on the 810, but applications that stress memory bandwidth, graphics performance, or both may not.

Though it can't match the performance of discrete solutions, the 810 offers clear advantages in overall cost and system complexity. These factors are more important to some users than performance, a point proved by the success of similar products from SiS and VIA. There will be more competition in this space over the next year. ATI and S3 plan to sample core-logic products with integrated graphics by the end of 1999. Since these vendors have better 3D cores than Intel, SiS, or VIA, we expect them to achieve better 3D performance.

Major PC OEMs will begin shipping 810-based PCs this quarter at prices that should prove attractive to end users. Users associate the Intel brand name with quality and reliability. These factors, combined with a low price, will help compensate for the performance shortcomings of 810-based systems. As long as Intel and its OEM partners are careful in their product positioning, the 810 is likely to be successful despite its performance. 