

M I C R O P R O C E S S O R

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

LOOKING BEYOND GRAPHICS

Nvidia's New GPU Architecture Energizes High-Performance Computing

By Tom R. Halfhill {10/05/09-01}

A typical PC has at least three cooling fans...and one case heater. That “heater” would be the graphics-processing unit (GPU)—usually a separate, highly specialized coprocessor dedicated to graphics. It does a brilliant job when the PC is running a graphics-intensive game or playing

a video. At other times, it's largely underutilized, radiating unused power as heat.

In fact, a discrete GPU is the most underutilized component in a PC. Although it's capable of amazing things, it spends much of its time performing routine chores like scrolling the screen. Yet it has the potential to be the swiftest processing engine in the system—and it's already there, just waiting for something to do.

For four years, Nvidia has waged a campaign to redefine the role of GPUs. Not that graphics aren't important. Pushing pixels has been good business for Nvidia, the world's leading graphics-processor company. Since the early 1990s, Nvidia's GPUs have offloaded most of the graphics processing from CPUs in millions of personal computers and video-game machines. Nvidia continues to design its GPUs for superb graphics performance.

Since 2005, however, Nvidia has cultivated another fast-growing market—GPUs for computing applications beyond graphics. At first, this was called “general-purpose GPU” (GPGPU) computing. Today, Nvidia calls it GPU computing. In the professional market, it's called high-performance computing (HPC).

By harnessing the massively parallel-processing resources originally designed for 3D graphics, clever programmers can apply GPUs to a much broader range of computing tasks. Some of those tasks, such as video transcoding, still involve graphics to some degree. Whereas a CPU might spend an hour or more converting a recorded video for uploading to

YouTube or burning a DVD, a GPU can tear through the job in minutes. The GPU can also enhance the video frame-by-frame, using consumer versions of sophisticated software once available only to NASA, law-enforcement agencies, and surveillance experts.

Other GPU-computing applications are data-intensive tasks having little or nothing to do with graphics. Examples are seismic-reflection analysis for oil exploration and financial analysis for market trading. (Using a cluster of Nvidia GPUs, Bloomberg recently speeded up its custom bond-pricing software by 800%.) Additional examples, such as medical imaging, combine high-resolution graphics with heavy-duty number crunching.

On these kinds of workloads, an ordinary GPU can blow away the latest multicore CPUs. Link several GPUs together in a workstation or a cluster of systems, and you've got a “desktop supercomputer”—a brute-force number-crunching machine at a fraction of the price and operating costs of a room-size supercomputer. AMAX Information Technologies, a U.S. company, recently built an Nvidia-based cluster for the College of Ocean and Earth Science at Tongji University in China. The cluster will be used for geophysics and seismology.

Originally, Nvidia used the term Compute Unified Device Architecture (CUDA) to describe its GPU-computing platform. (See [MPR 1/28/08-01](#), “Parallel Processing With CUDA.”) Later, CUDA became a blanket term for Nvidia's GPU architecture, run-time platform, and software-development tools. Now Nvidia is raising the bar.

Editor's Note

Microprocessor Report adapted and expanded this article from an In-Stat white paper written for Nvidia. The white paper was researched, written, and reviewed to the same standards as *MPR* articles. The white paper is available on Nvidia's website at www.nvidia.com/fermi.

Nvidia's next-generation GPU architecture, code-named Fermi, adds powerful new features for CUDA general-purpose computing. Fermi processors will continue to shoulder the graphics workloads in PCs and videogame consoles, but they are taking the largest step yet toward becoming equal-partner coprocessors with CPUs.

New Features for GPU Computing

On September 30, at its GPU Technology Conference in San Jose, California, Nvidia demonstrated the first silicon of a chip based on the Fermi architecture. If the project proceeds on schedule, Fermi GPUs could hit the market this year. Here's a summary of the improvements:

- Error-correction codes (ECC) protect the external DRAM, L1 data caches, L2 cache, and register files. ECC is unimportant for graphics but vital for financial customers and others running mission-critical software. No other GPU has ECC.
- Fermi-class GPUs can perform double-precision floating-point math up to eight times faster than existing Nvidia GPUs. Double precision is now exactly half the speed of single precision. A new fused multiply-add instruction improves precision and conforms to the latest IEEE 754-2008 floating-point specification.
- The memory-address space (both virtual and physical) has been unified and expanded from 4GB to one terabyte (1TB), vastly increasing the capacity for large data sets.
- Thanks to unified memory addressing and other improvements, Fermi can run object-oriented C++ code, not just procedural C. This will make programming the GPU easier.
- Fermi GPUs will have as many as six I/O interfaces to external DRAM, each 64 bits wide. Although Nvidia has not announced which types of memory Fermi will support, these interfaces will provide greater bandwidth when coupled with the latest GDDR5 DRAM.
- The memory hierarchy is fully cached, including new L1 caches shared by tightly coupled groups of 32 processor cores. For each group, programmers can partition 64KB of SRAM to serve as a hardware-managed L1 cache or a software-managed shared memory.
- Fermi can run concurrent "CUDA kernels" or global functions, whereas existing Nvidia GPUs are limited to

running multiple threads within one kernel. This feature adds another level of parallelism.

- Nvidia has revamped the entire instruction-set architecture (ISA) to make it an easier target for high-level compilers. Among other enhancements, all instructions are predicated, eliminating the need for separate branch instructions in some code sequences.
- An improved instruction scheduler can issue two operations per clock cycle, versus one operation per clock before.
- A new load/store unit can execute 16 operations per cycle, greatly improving I/O performance. Atomic read-modify-write instructions are 5 to 20 times faster.

Some of these enhancements are unimportant for 3D graphics but were requested by Nvidia's existing and prospective GPU-computing customers. Indeed, some features (such as ECC) would actually *reduce* graphics performance. To avoid compromising the GPU's central role as a graphics coprocessor, Fermi has provisions for disabling or bypassing features that are irrelevant for graphics. For instance, GeForce-branded GPUs for consumer PCs will omit ECC; Tesla- and Quadro-branded GPUs for workstations and supercomputers will include it.

GPU Computing Takes Off

Fermi surpasses the latest GPU announced by Nvidia's leading competitor, AMD. Conceptually, Fermi's closest competitor is Intel's Larrabee. Expected to debut next year, Larrabee chips will attack the market from the opposite direction. They will use general-purpose x86-architecture cores for both 3D graphics and GPU computing—a radically different approach from adapting a specialized 3D-graphics architecture for general-purpose computing. Until the first Larrabee chips appear, it's not clear whether the x86 can make such a leap. (See [MPR 9/29/08-01](#), "Intel's Larrabee Redefines GPUs.")

No wonder this market is attracting big-gun competition. By Nvidia's estimates, the total available market for GPU computing is worth about \$1.1 billion a year. Figure 1 breaks down the market into broad categories. By comparison, the total available market for GPUs in desktop PCs is worth \$1.5 billion to \$2.0 billion a year. However, GPU computing is still in its infancy. It's sure to grow as GPUs become more powerful, software-development tools improve, programmers become more experienced, customers find new applications, and more companies realize that GPU computing gives them a competitive edge.

Fundamentally, Fermi processors are still graphics processors, not general-purpose processors. The system still needs a host CPU to run the operating system, supervise the GPU, provide access to main memory, present a user interface, and perform everyday tasks that have little or no data-level parallelism. However, Nvidia's new features and improvements will transform Fermi GPUs into hybrid coprocessors, capable of tackling a wider range of applications—for consumers and professionals.

Understanding the Fermi Architecture

Nvidia's terminology can be confusing for those accustomed to general-purpose CPU architectures like the Intel x86. The confusion isn't deliberate. It's partly a consequence of the graphics heritage of Nvidia's architectures, but another reason is Nvidia's massively parallel approach to multithreading. Whereas an Intel x86 processor core with Hyper-Threading typically executes two threads per core, an Nvidia GPU frequently works on thousands of threads, switching among them on every clock cycle.

CUDA architectures have more in common with the highly specialized, massively parallel CPU architectures that have occasionally appeared since the mid-1990s. Like CUDA, these architectures typically have large numbers of very simple processor cores and massively threaded programming models. However, almost all the specialized architectures soon failed in the marketplace. Some met resistance because programming was too difficult. Others were doomed because their parent companies were underfunded or poorly managed startups. Sometimes they disappeared only because no one gave them a chance.

Although Nvidia is overcoming many of the technical hurdles of those architectures, a bigger factor contributes to CUDA's success. CUDA comes from a large, mature company, and it piggybacks on the high-volume consumer graphics market. In effect, games are subsidizing CUDA's development while the GPU-computing market becomes self-sufficient. Furthermore, the multicore CPUs from Intel, AMD, and other companies are driving new demand for better parallel-programming languages and tools. CUDA is riding all these waves.

Fermi is a significant advance over previous GPU architectures. Nvidia's first architecture expressly modified for GPU computing was the G80, introduced in November 2006 with the GeForce 8800. The G80 departed from the traditional approach of using dedicated shaders, separate vertex/pixel pipelines, manually managed vector registers, and other features that were efficient for graphics but unwieldy for general-purpose computing. The G80 was also Nvidia's first GPU architecture programmable in C—an essential feature for compute applications.

G80-class GPUs had 128 programmable shaders, later called streaming processor cores and now known simply as CUDA cores. Eight CUDA cores were grouped into a "streaming multiprocessor" whose cores shared common resources, such as local memory, register files, load/store units, and thread schedulers. G80-class chips had 16 of these streaming multiprocessors (16 SMs x 8 CUDA cores per SM = 128 cores per chip). With time slicing and fast thread switching, a streaming multiprocessor can run thousands of parallel threads on these cores.

In June 2008, Nvidia introduced the GT200 architecture. The GT200 retained eight CUDA cores per streaming

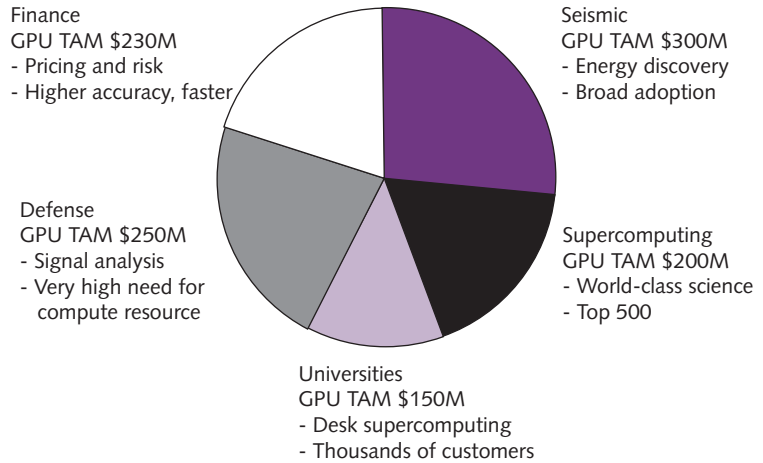


Figure 1. Nvidia estimates that the total available market (TAM) for GPU computing is at least half as large as the desktop-PC market for GPUs. The upside potential looks greater. Whereas the PC market is maturing, GPU computing barely existed four years ago and is growing fast.

multiprocessor but bumped the number of those multiprocessors to 30, for a total of 240 CUDA cores per chip. GT200 chips are sold as the GeForce GTX280 (for consumer PCs), Quadro FX5800 (for workstations), and Tesla T10 (for high-performance computing). These were also the first Nvidia GPUs to support double-precision floating-point operations—unnecessary for 3D graphics but vital for many scientific and engineering programs.

Fermi supersedes the GT200 architecture. It has 32 CUDA cores per streaming multiprocessor—four times as many as the GT200 and G80. Initially, Fermi GPUs will have 16 streaming multiprocessors, for a total of 512 CUDA cores per chip. This expansion alone would significantly boost throughput, but additional enhancements deliver even more performance.

Fermi, From the Ground Up

The block diagrams that follow illustrate the Fermi architecture. Figure 2 starts at ground level, showing an exploded view of a single CUDA core and its relationship to the streaming multiprocessor to which it belongs. Although a CUDA core resembles a general-purpose processor core, like those found in a multicore x86 CPU, it's much simpler, reflecting its heritage as a pixel shader.

Each CUDA core has a pipelined floating-point unit (FPU), a pipelined integer unit, some logic for dispatching instructions and operands to these units, and a small queue for holding results. And that's all.

Unlike the processor cores typically found in general-purpose microprocessors, CUDA cores lack their own register files or L1 caches. Nor do they have multiple function units for each data type (floating-point and integer). In fact, they don't even have a load/store unit for accessing memory. CUDA cores are very simple processing engines designed to execute only one instruction at a time per thread and then

quickly switch to another thread. They are optimized not for single-threaded performance as individuals but for multithreaded performance when operating en masse.

Figure 3 shows a Fermi streaming multiprocessor with 32 CUDA cores and additional elements. This diagram explains why CUDA cores can get by without their own register files, caches, or load/store units—those resources are shared among all 32 CUDA cores in a streaming multiprocessor. Those 32 cores are designed to work in parallel on 32 instructions at a time from a bundle of 32 threads, which Nvidia calls a “warp.” (This organization has implications for the CUDA programming model, as we’ll explain below.)

At the top of Figure 3, notice the L1 instruction cache—a new feature in Fermi. (Nvidia hasn’t disclosed the size of this cache.) Each streaming multiprocessor also has 64KB of configurable local memory that programmers can partition as L1 data cache and general-purpose shared memory. Each partition can be either 16KB or 48KB. In contrast, the G80 and GT200 architectures had only 16KB of shared memory per streaming multiprocessor, and the memory couldn’t operate as hardware-managed L1 caches.

Massively Parallel Multithreading

Another shared resource in a streaming multiprocessor is the instruction-scheduling and dispatch logic. Fermi improves

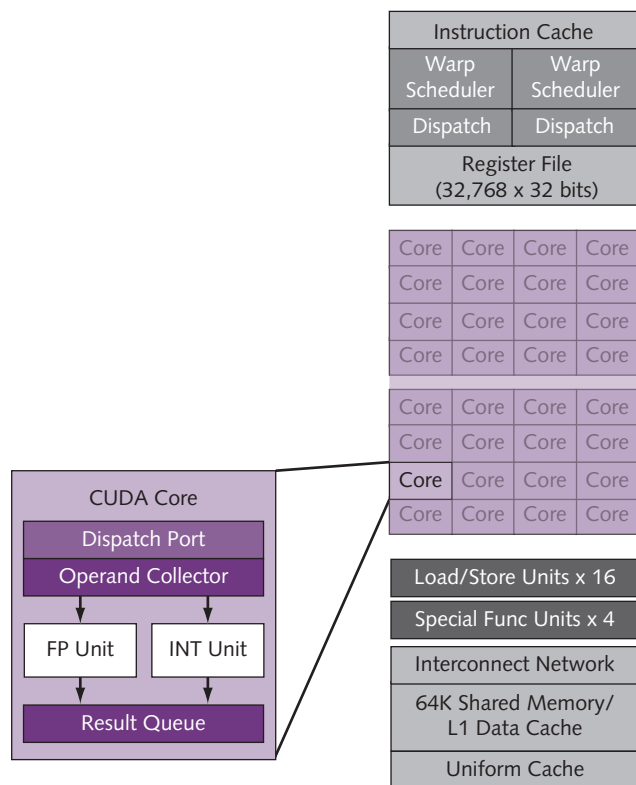


Figure 2. CUDA-core block diagram. CUDA cores evolved from pixel shaders, so they are very simple processor cores. The cores lack their own general-purpose register file, L1 caches, multiple function units for each data type, and load/store units for retrieving and saving data.

on the scheduler in previous CUDA architectures by issuing two instructions per clock cycle instead of one. However, this capability isn’t quite the same as the dual-issue scheduling in conventional superscalar microprocessors.

In a Fermi GPU, the dual-issue pipelines are decoupled and independent. They cannot issue two instructions per cycle from the same thread. Instead, they issue two instructions per cycle from different warps. Each streaming multiprocessor can manage 48 warps. Because each warp has 32 threads, a streaming multiprocessor can manage 1,536 threads. With 16 streaming multiprocessors, a Fermi-class GPU can handle 24,576 parallel threads.

Multithreading on this scale looks intimidating, but programmers needn’t manage the threads explicitly—CUDA does that. Also, the threads in a warp should have no interdependencies. In theory, CUDA makes deadlocks between threads impossible, overcoming a major source of bugs in multithreaded code. (See *MPR 4/30/07-02*, “The Dread of Threads.”)

Of course, there aren’t enough CUDA cores to execute instructions from every thread on every clock cycle. Each

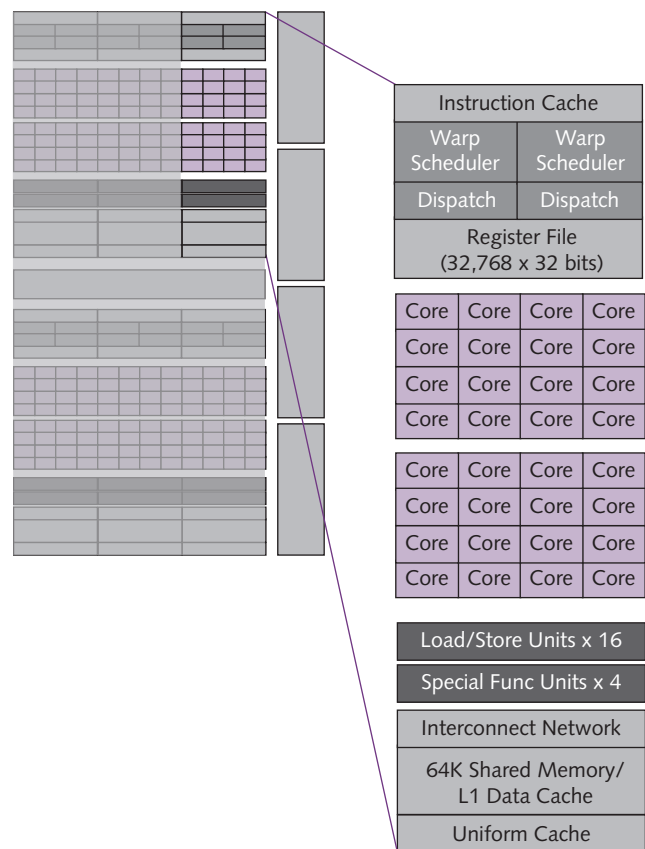


Figure 3. Streaming-multiprocessor block diagram. In the Fermi architecture, each streaming multiprocessor has 32 CUDA cores—four times as many as the previous GT200 and G80 architectures. All 32 cores share the resources of their streaming multiprocessor, such as registers, caches, local memory, and load/store units. The “special function units” (SFUs) handle complex math operations, such as square roots, reciprocals, sines, and cosines.

core executes only one instruction per cycle, so instructions from “only” 512 threads can execute at a given moment. Switching among threads is instantaneous, so instructions from 512 different threads can execute on the next clock cycle, and so on. This massively parallel threading is the key to CUDA’s high throughput. Figure 4 illustrates how the warp schedulers and dispatch units interleave the warps and the instructions from different threads.

Another shared resource in a streaming multiprocessor is a new load/store unit, which can execute 16 load or store operations per clock cycle. It does even better when using a special “uniform cache,” seen at the bottom of Figure 3. Matrix-math operations often load scalar values from sequential addresses belonging to a particular thread, and they also load a common value shared among all threads in a warp. In those cases, a streaming multiprocessor can load two operands per cycle.

Figure 5 is the highest-level view of the Fermi architecture. All 16 streaming multiprocessors—each with 32 CUDA cores—share a 768KB unified L2 cache. By the standards of modern general-purpose CPUs, this cache is relatively small, but previous CUDA architectures had no L2 cache at all. Fermi maintains cache coherency for all the streaming multiprocessors sharing the L2 cache.

Thread scheduling is managed in hardware, not in software. Fermi has an improved “GigaThread” scheduler that can manage thousands of simultaneous threads and switch contexts between graphics and compute applications in as little as 25 microseconds—ten times faster than Nvidia’s previous schedulers. Switching among threads within a graphics or compute instruction stream requires only one clock cycle, but alternating between graphics and compute workloads takes longer, because the caches must be flushed and refilled.

Fermi’s Memory Hierarchy

The memory hierarchy of a Fermi GPU is somewhat different from the better-known hierarchy of a general-purpose CPU. For one thing, a GPU has a large frame buffer—as much as a gigabyte of DRAM in today’s systems. Originally, the frame buffer was dedicated to graphics data. But programs can stash any data in this buffer, a valuable feature for general-purpose computing. For nongraphics applications, Nvidia sometimes refers to the frame buffer as “local DRAM.”

Another important difference between GPUs and CPUs is their proximity to system memory. CPUs have a direct path to main memory through an integrated memory controller or the north bridge of the system chipset. GPUs must reach out

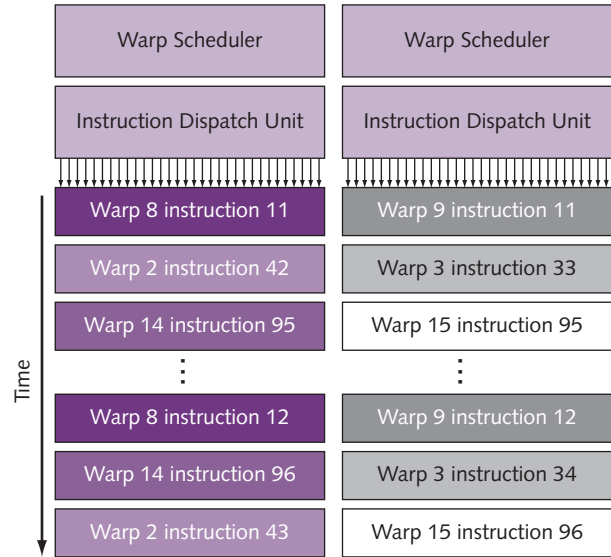


Figure 4. CUDA multithreading in the Fermi architecture. To achieve massively parallel multithreading, CUDA bundles 32 threads into a “warp” and executes one instruction from each thread per clock cycle. Each streaming multiprocessor can manage 48 warps. The warp schedulers and dispatch units interleave the threads and warps to keep the CUDA processor cores busy. The GPU can switch from one thread to another on each clock cycle.

over the PCI Express bus by sending a request to the host CPU. The long latency inherent in this transaction requires GPU programmers to use local DRAM, caches, and shared local memory more wisely. Figure 6 illustrates the memory hierarchy.




Figure 5. Fermi architecture block diagram. This top-level view of the architecture shows the 16 streaming multiprocessors, the six 64-bit DRAM interfaces, the host interface (PCI Express), and the GigaThread hardware thread scheduler.

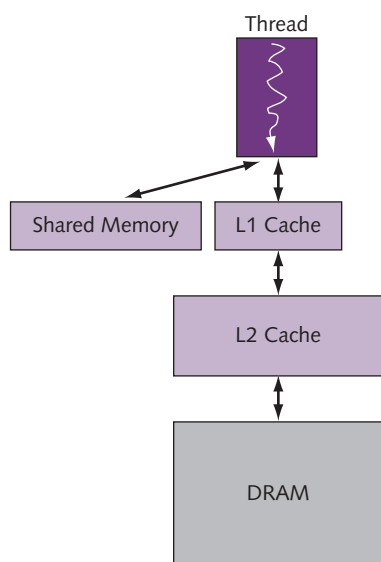


Figure 6. Fermi's memory hierarchy. The shared memory and L1 cache are local to each 32-core streaming multiprocessor and shared by all those cores. The L2 cache is coherently shared among all 16 multiprocessors in the GPU. In this diagram, "DRAM" refers primarily to the GPU's local DRAM (frame buffer), attached directly to the GPU. Additional DRAM is available in the system's main memory, but it requires a lengthy transaction over the PCI Express bus.

High-level languages like C and C++ usually have no concept of memory hierarchy. All memory looks the same to them (flat). Left alone, compilers would allocate memory for variables inefficiently. To solve this problem, Nvidia added new keywords to its CUDA implementation of C and C++. These keywords precede variable declarations, allowing programmers to specify where the program will store the data referenced by the variable. This solution avoids unwanted references to distant memory. Variables declared without a special keyword will be local to a thread.

For example, the `_device_` keyword specifies that a variable stores data in the GPU's local DRAM. The `_shared_` keyword specifies the shared memory of a streaming multiprocessor, so it limits access to a block of threads. The `_global_` keyword makes a variable available to all blocks of threads, so the data probably ends up in system memory. However, all types of memory (except shared memory, which is already local) is cacheable, which minimizes slow references to system memory. Figure 7 illustrates the relationships between local variables, shared variables, global variables, and the memory regions referenced by the variables.

The six DRAM interfaces seen in Figure 5 provide fast access to local DRAM, essential for graphics processing as well as for general-purpose computing. Each 64-bit interface can work independently. Together, they can access up to 6GB of DRAM in the first implementation of the Fermi architecture. Nvidia hasn't announced which types of DRAM will be supported or the bus speeds of the chips. Currently, Nvidia's GTX280 uses GDDR3 DRAM to deliver 141.7GB/s of peak bandwidth. With GDDR5 DRAM, Fermi

will certainly exceed that performance. Nevertheless, as with all modern microprocessors, memory bandwidth remains the most likely performance bottleneck.

A Closer Look at Fermi's Features

Fermi's enhancements can be grouped into four broad categories: new features intended exclusively for general-purpose computing; improvements to memory and I/O; features that boost instruction throughput; and features that aid software development. They run the gamut from hardware to software. Some enhancements are useful for accelerating 3D graphics as well as for GPU computing. Although Fermi has additional new features exclusively for graphics, Nvidia hasn't disclosed them at this time.

The most-wanted new feature for GPU computing is ECC. It's also a key feature that will set Nvidia's GPUs apart from the existing competition. AMD's latest GPU has error *detection*, but not error correction. On September 23, AMD introduced the first chip (code-named Cypress) in its new Evergreen family of GPUs. Cypress is shipping in ATI Radeon HD5870 and HD5850 graphics cards. It implements error-detection codes (EDC) instead of ECC. EDC detects one- and two-bit errors during transmissions over the GDDR5 bus to the frame buffer and then automatically triggers a retransmission.

In contrast, ECC adds error-checking and correction hardware, so it can correct most errors without retransmission. Nvidia's implementation of ECC can detect one-bit and two-bit errors, and it can correct one-bit errors. Fermi adds ECC to the I/O pins of the chip's external DRAM interfaces, but it doesn't stop there. Nvidia has implemented ECC throughout the Fermi architecture—to the register files and shared memories of the streaming multiprocessors, and to the coherent L2 cache shared by all the multiprocessors. Few chips implement ECC as thoroughly as Fermi does.

ECC protects data transfers against transient soft errors, such as those induced by electromagnetic interference and the radiation constantly bombarding the Earth from outer space. Soft errors are growing worse as transistors keep shrinking in step with Moore's law. (Nvidia is manufacturing the first Fermi chips in TSMC's 40nm CMOS fabrication process—the same process AMD is using for Cypress.) If an error is too severe for ECC to correct, the program can retransmit the data or display an error message to the user. ECC is particularly important for data-center applications, where the failure rate of an individual component is amplified by the large scale of the system.

With existing GPUs, developers resort to various workarounds for the lack of ECC. Some programmers implement ECC in software, severely impairing performance. Others tell users to run their programs two or more times on the same data, then compare the results for consistency. By adding ECC hardware to its next-generation Fermi processors, Nvidia is giving notice that it's serious about mission-critical computing.

There's always a price to pay for ECC, even when implemented in hardware—and especially when ECC is implemented as extensively as it is throughout Fermi. Nvidia estimates that performance will suffer by 5% to 20%, depending on the application. For critical programs requiring absolute accuracy, that overhead is of no concern. It's certainly much better than the overhead imposed by existing work-arounds. For less-critical applications, the performance hit could be a killer. That's why Nvidia allows programmers to disable ECC in Tesla-branded GPUs for the professional market and won't enable ECC at all in GeForce GPUs for the consumer market.

Double Precision Gets a Boost

Another Fermi enhancement that focuses primarily on high-performance computing is faster double-precision floating-point math. Traditionally, GPUs have been optimized for single-precision (32-bit) floating-point, because it's sufficient for 3D graphics. However, many scientific and engineering applications need the greater accuracy and range of double-precision (64-bit) numbers.

For Nvidia, these competing interests pose a difficult trade-off. Implementing 64-bit FPUs and datapaths throughout the processor would approximately double the amount of floating-point computational logic and wiring. The chip would be larger (e.g., costlier) and use more power, yet it would be no better for its primary market—consumer 3D graphics.

As a result, GPUs didn't even support double-precision floating point until very recently, and the initial support has been lackluster. The first Nvidia GPU to have double precision was the GT200-based GeForce GTX280, introduced in June 2008. With Fermi, Nvidia is taking a big step forward. The new architecture is up to eight times faster with double-precision floating point and enhances the accuracy of some single-precision operations, too.

Fermi's floating-point datapaths are still 32 bits wide, minimizing the amount of logic and wiring throughout the chip. To perform 64-bit operations, these 32-bit paths are paired, and Nvidia has significantly improved their efficiency. A 64-bit floating-point operation is now exactly half as fast as a 32-bit floating-point operation. Although that relationship seems only logical when pairing two 32-bit datapaths, today's GPUs are much slower, because the synchronization is less efficient.

Specifically, a Fermi-class streaming multiprocessor can execute 16 double-precision fused multiply-add (FMA) instructions per clock cycle. That's eight times faster than the FPUs in today's GT200 chips. An FMA performs two operations with a single instruction ($D = A \times B + C$). Therefore,

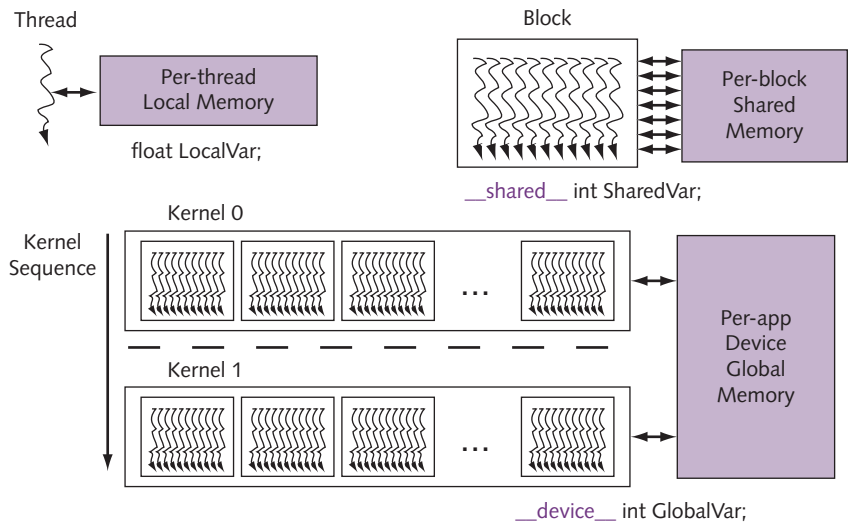


Figure 7. Memory allocation with CUDA. By using special CUDA keywords added to C, programmers can specify the memory regions in which variables store their data. This is important not only for scoping the variables, but also for keeping local data close to the CUDA processor cores. The wider the scope, the more distant the memory.

a Fermi-class streaming multiprocessor can execute 32 double-precision floating-point operations per cycle. With 16 streaming multiprocessors per chip, a Fermi GPU can execute 512 billion double-precision floating-point operations per gigahertz.

Striving for Teraflops

Currently, the streaming multiprocessor in a GeForce GTX280 chip runs at 1.29GHz, limited by an older 65nm fabrication process. Manufactured at 40nm, a Fermi-based GPU should easily exceed that clock frequency. (Nvidia has not yet disclosed the clock rates for Fermi GPUs.)

At 1.5GHz (our very conservative estimate), the GPU would deliver 768 double-precision gigaflops—and twice as many single-precision gigaflops. (Fermi's enhanced FMA instruction supports 32-bit operations, unlike previous generations.) At 2.0GHz, double-precision performance would reach the magic teraflops threshold, once the province of multimillion-dollar room-size supercomputers.

At its core clock frequency of 850MHz, AMD's Cypress (in the Radeon HD5870) reaches 544 gigaflops (double precision) and 2.7 teraflops (single precision). Clock for clock, Cypress looks faster than Fermi, but it appears to be more clock-limited. Also, the difference between single- and double-precision performance is much greater—about 5x vs. 2x. Both architectures will soon break the double-precision teraflops barrier.

Intel's Larrabee aspires to similar levels of performance. Using FMA instructions in its 16-lane vector-processing units, a Larrabee x86 core can perform 32 single-precision floating-point operations per clock cycle. Therefore, at 1.0GHz, the maximum theoretical throughput is 32 single-precision gigaflops per core. Larrabee's SIMD lanes are 32 bits wide and are

paired for 64-bit operations, so double precision is half as fast as single precision (just like Fermi). Intel hasn't announced specific implementations of Larrabee, so the number of cores per chip and clock speeds are unknown. At 1.0GHz, a Larrabee chip would need 32 cores to reach one single-precision teraflops. Alternatively, 16 cores running at 2.0GHz would reach the same threshold.

In public presentations, Intel has estimated the relative (not absolute) performance of Larrabee chips with as few as eight cores and as many as 64 cores. *Microprocessor Report* expects to see different designs for different markets, with a 32-core chip initially anchoring the high end of the Larrabee line. The largest designs will probably challenge the floating-point performance of Nvidia GPUs based on the Fermi architecture. Whether Larrabee can deliver that performance at a similar cost and with similar power consumption remains to be seen. And its 3D-graphics performance (apart from raw floating-point performance) is another subject altogether. (See [MPR 9/29/08-01](#), "Intel's Larrabee Redefines GPUs.")

Better Multiply-Adds

In addition to boosting floating-point throughput, the Fermi architecture improves precision. In fact, GPUs can now perform floating-point operations more accurately than x86 processors do. Fermi maintains full precision during the intermediate calculations of an FMA—even for single-precision operations—instead of truncating or rounding the intermediates. Previously, only double-precision FMAs were handled this way.

(Precise single-precision arithmetic is important for 3D-graphics as well as for high-performance computing. It reduces the "sparkling" effects marring the edges of objects when imprecise calculations change the positions of pixels from one frame to another.)

```
Computing y = ax + y with a serial loop:
void saxpy_serial(int n, float alpha, float *x, float *y)
{
    for(int i = 0; i<n; ++i)
        y[i] = alpha*x[i] + y[i];
}
// Invoke serial SAXPY kernel
saxpy_serial(n, 2.0, x, y);

Computing y = ax + y in parallel using CUDA:
__global__
void saxpy_parallel(int n, float alpha, float *x, float *y)
{
    int i = blockIdx.x*blockDim.x + threadIdx.x;
    if( i<n ) y[i] = alpha*x[i] + y[i];
}
// Invoke parallel SAXPY kernel (256 threads per block)
int nblocks = (n + 255) / 256;
saxpy_parallel<<<nblocks, 256>>>(n, 2.0, x, y);
.....
```

Figure 8. CUDA kernels are modified versions of conventional functions written in C. The `__global__` keyword is a special CUDA extension—it tells the compiler this function is a CUDA kernel. The other changes tell the compiler to divide the workload into blocks, each with 256 threads.

In accordance with the latest IEEE 754-2008 specification, Fermi's floating-point hardware supports all four rounding modes (nearest, zero, positive infinity, and negative infinity). There's also new hardware support for handling "subnormal" numbers—those that fall between the smallest-possible positive number and largest-possible negative number of the floating-point number system.

Today's GPUs simply flush subnormal numbers to zero, losing accuracy. An x86 CPU throws an exception in these cases and spends thousands of clock cycles handling the exception in software. Fermi handles subnormals in hardware, allowing them to gravitate toward zero without sacrificing accuracy or performance.

Nvidia's enhanced FMA instruction and full support for IEEE 754-2008 put the Fermi architecture on par with other high-performance microprocessor architectures and exceed the capabilities of some of them. Intel's latest Streaming SIMD Extensions (SSE4.2) for the x86 lack an FMA, and the x86 incurs rounding errors when performing its version of a multiply-add (MADD). However, Intel is introducing new FMA instructions for the 16-lane vector-processing units in Larrabee. Intel's Itanium architecture also has FMA, as does IBM's Cell Broadband Engine (but only for double precision, not single precision).

Along with much faster double-precision throughput, the greater precision of these operations makes the Fermi architecture a serious contender in the realm of high-performance computing for science and engineering. On September 30, Oak Ridge National Laboratory announced plans to build a huge new supercomputer based on Fermi. Applications will include energy research and climate modeling.

New Concurrency for Global Kernels

Another significant performance improvement in the Fermi architecture is the ability to run concurrent global functions, or "CUDA kernels," as Nvidia calls them. Existing CUDA architectures can execute multiple threads within a kernel, but they can't run more than one kernel at a time. Figure 8 compares a function written in standard C code with a CUDA kernel rewritten for parallel execution.

Concurrent global kernels add another level of parallelism to CUDA. Previous CUDA architectures are limited to exploiting data-level parallelism when a single algorithm operates on a single data set. A global kernel may call a sequence of algorithms in a step-by-step fashion; and the data may reside in two or more arrays; and the function may spawn thousands of threads. However, at any given moment, the streaming multiprocessors can execute only one global kernel. Executing another global kernel requires waiting for the previous global kernel to finish.

The Fermi architecture overcomes that limitation and allows up to 16 independent global kernels to execute concurrently. This feature adds a degree

of instruction-level parallelism to CUDA's forte, data-level parallelism. Programmers who can wrap their brains around these concepts will see their code run a lot faster. (See [MPR 3/31/08-01](#), "Think Parallel.")

One remaining limitation is that concurrent global kernels must belong to the same program. Fermi still cannot manage application-level parallelism. For example, as is the case today, the GPU must switch contexts to handle the separate instruction streams of graphics and compute tasks. Nevertheless, the additional parallelism of concurrent kernels within a program is a welcome enhancement to CUDA's prodigious capacity for data-level parallelism.

Table 1 summarizes the salient differences between the Fermi architecture and Nvidia's previous GT200 and G80 architectures.

Instruction-Set Improvements

Nvidia has overhauled the Fermi instruction-set architecture (ISA) to make it friendlier for high-level language compilers. For instance, all instructions are now fully predicated. An instruction can decide whether or not to execute by checking a status bit, which may be set or cleared by a previous instruction. In some code sequences, predication eliminates the need for separate compare and branch instructions, allowing a compiler to generate more-streamlined code.

Additionally, the Fermi ISA improves the efficiency of "atomic" integer instructions. These are read-modify-write instructions that read a value in memory, modify the value by performing an arithmetic or logical operation, and then quickly write the result back to the same memory location. Atomic instructions must be uninterruptible. Otherwise, an incomplete sequence of operations would leave the data in an unpredictable state, with possibly serious consequences for the program.

Nvidia GPUs have included atomic instructions since 2007, but they were relatively slow. Fermi makes them 5 to 20 times faster. The method is unorthodox—atomic instructions are handled by special integer units attached to the L2 cache controller, not by the integer units in the

Feature	Nvidia Fermi	Nvidia GT200	Nvidia G80
Introduction	2009	2008	2006
Single-Precision FP	512 fused multiply-add ops per cycle	240 multiply-add ops per cycle	128 multiply-add ops per cycle
Double-Precision FP	256 fused multiply-add ops per cycle	30 fused multiply-add ops per cycle	—
IEEE 754-2008 Compliance	Full, SP and DP	DP only	—
Streaming Multiprocessors	16	30	16
CUDA Cores Per SM	32	8	8
Warp Schedulers Per SM	2	1	1
Special Function Units Per SM	4	2	2
Shared Memory Per SM*	48KB or 16KB (configurable)	16KB	16KB
L1 Instruction Cache Per SM	Yes	—	—
L1 Data Cache Per SM*	16KB or 48KB (configurable)	—	—
Total CUDA Processor Cores	512	240	128
L2 Cache	768KB, shared	—	—
Error-Correction Codes	DRAM, shared memories, L2 cache, registers	—	—
Concurrent Kernels	Up to 16	—	—
Fully Predicated ISA	Yes	—	—
Memory Addressing	40 bits (1TB), unified	32 bits (4GB), divided	32 bits (4GB), divided
Memory I/O Interfaces	6 x 64 bits, 6GB DRAM	8 x 64 bits, 1GB GDDR3 DRAM	6 x 64 bits, 768MB GDDR3
C++ Programmable	Yes	—	—
Transistors	3.0 billion	1.4 billion	681 million

Table 1. Comparison of Nvidia's three CUDA-capable GPU architectures. The G80, introduced in 2006, opened the door to CUDA and was Nvidia's first architecture modified for GPU computing. The GT200 architecture followed in 2008. Although it improved performance, it retained several limitations of the G80. The new Fermi architecture overcomes most of those limitations and adds features not seen before in GPUs, such as error-correction codes (ECC). *The shared memory and L1 data cache in each streaming multiprocessor is 64KB of SRAM. Programmers can partition this memory into 16KB and 48KB segments, using either segment as cache or as closely coupled shared memory.

CUDA cores. Executing atomic instructions in the cores would introduce too much latency as the data is shuttled across the chip.

Efficient handling of atomic instructions is important for software developers. Two emerging standards for parallel programming have embraced the concept of atomic operations: OpenCL and DirectCompute.

OpenCL is an open industry standard managed by the OpenCL Compute Working Group, which is part of an industry consortium called the Khronos Group. DirectCompute is an application programming interface (API) from Microsoft and was formerly known as Compute Shader or DirectX 11 Compute. DirectCompute is part of the larger DirectX 10 and DirectX 11 APIs in Windows Vista and Windows 7. Developers can use OpenCL or DirectCompute instead of (or in addition to) Nvidia's CUDA tools.

For More Information

Nvidia publicly disclosed the Fermi GPU architecture on September 30 at the GPU Technology Conference in San Jose, California. Nvidia CEO Jen-Hsun Huang demonstrated the first working silicon of a Fermi chip. Nvidia hasn't announced specific implementations of the architecture but hopes to ship devices this year. For more information, visit www.nvidia.com/fermi.

Although some observers view OpenCL and DirectCompute as competition for CUDA, Nvidia insists the technologies are complementary. Nvidia has voiced support for all GPU programming models and works closely with Microsoft and the OpenCL group. On September 28, Nvidia released the industry's first public OpenCL GPU drivers for Windows and Linux.

The Fermi architecture allows compilers to optimize the memory layout for arrays when using OpenCL or DirectCompute. Efficient memory layout is a crucial feature, because arrays are central to the programming model. All data sets on which the parallel threads operate are stored in arrays.

Programmers Get a Break

Just a few years ago, GPUs weren't even programmable in C. GPGPU pioneers had to struggle with an unconventional assembly language and try to fit the square pegs of their general-purpose code into the round holes of a graphics-oriented architecture. Nvidia's CUDA development tools allowed programmers to write code in fairly conventional C, albeit with a few missing features and some extensions specific to CUDA.

The Fermi architecture makes the GPU programmable in object-oriented C++, not just in procedural C. Moreover, Fermi supports the entire C++ language, not a subset. Programmers can use virtual functions, try/catch statements (which handle exceptions), and make system calls, such as `stdio.h` (the standard input/output channel). Nvidia provides a CUDA C/C++ compiler, but any third-party tool vendor can target the CUDA architecture.

A key factor in bringing C++ to CUDA is Fermi's new unified memory addressing. Previous CUDA architectures allocated different memory regions for local memory, shared memory, and global memory. The fragmented memory map required different load/store instructions for each region and prevented Nvidia from fully implementing pointers in C and C++. In particular, C++ relies heavily on memory pointers, because every "object" in this object-oriented language requires at least one pointer. A typical program creates hundreds of objects.

The Fermi architecture fixes those problems by unifying everything in a single memory space. Load/store instructions can now access any type of memory within this unified

space. At run time, Fermi uses a new address-translation unit to convert pointers and other memory references to physical memory addresses. In addition, Fermi expands the memory-addressing range from 32 bits (maximum 4GB of memory) to 40 bits (maximum 1TB of memory). Together, these changes bring higher-level languages to CUDA, simplify the programming model, and vastly expand the memory available for large data sets—vital features for high-performance computing.

Unlike most other compilers, CUDA compilers don't translate source code directly into native machine code. Instead, they target a low-level virtual machine and Parallel Thread eXecution (PTX) instruction set. The PTX virtual machine is invisible to users and delivered as part of the GPU's graphics driver. Few people even know it's there, but it's very important to CUDA software developers. PTX provides a machine-independent layer that insulates high-level software from the details of the GPU architecture.

When users install a CUDA-enabled program, the driver automatically translates PTX instructions into machine instructions for the specific Nvidia GPU in the system. It doesn't matter whether the GPU is based on the G80, GT200, or Fermi architecture—the installed program will be optimized for that GPU. In this way, CUDA programs run natively (full speed) but are delivered in a machine-independent format that's ready for installation on any CUDA-capable system. Nvidia says more than 100 million PCs have CUDA-capable GPUs.

CUDA Works With Visual Studio

Besides bringing C++ to CUDA, the Fermi architecture with PTX 2.0 makes it easier to use other high-level programming languages and compilers. Fermi supports Fortran—as well as Java (via native methods) and Python. On September 29, The Portland Group released a beta version of its CUDA Fortran compiler. Java is widely used in the financial industry, and Python is popular for rapid application development. PTX 2.0 also adds features for OpenCL and DirectCompute, such as new bit-reverse and append instructions.

Fermi enables better tool integration, too. For the first time, programmers will be able to use Microsoft's Visual Studio to write CUDA software. Nvidia has announced Nexus, a new tool suite that integrates with Visual Studio, so programmers can write CPU and GPU code in a single development environment. Nexus includes a debugger, a performance profiler, and a graphics inspector that works with APIs like Direct3D. Because Fermi supports exception handling at the hardware level, programmers can set breakpoints in their CUDA source code and single-step through a program while debugging.

Visual Studio is a popular development environment, so CUDA programmers will no doubt welcome Nexus. (Nvidia plans to release a beta version on October 15.) This integration is a significant achievement, and it adds yet another option



for writing high-performance computing software on GPUs. (See [MPR 7/28/08-02](#), “Tools for Multicore Processors.”)

Everyday users will be oblivious to these behind-the-scenes improvements for software development. However, the Fermi architecture is certain to attract new developers and broaden the base of CUDA software. In anticipation of the Larrabee launch, Intel is strengthening its software-development resources, too. Intel recently acquired two companies specializing in parallel-programming tools. (See [MPR 9/14/09-01](#), “Summer Shopping Spree.”)

AMD is pursuing the market for high-performance computing as well. ATI Stream is similar in concept to CUDA but isn't as mature and hasn't attracted as many software developers. (See [MPR 12/22/08-01](#), “AMD's Stream Becomes a River.”)

Redefining the GPU Market

With the Fermi architecture, Nvidia is making a big push to move its GPUs beyond consumer graphics and games. Not that games aren't important—they remain the driving force behind the evolution of discrete GPUs. Serving the consumer market will continue to be Fermi's day job. But GPUs are becoming more than mere coprocessors for offloading graphics chores from the system's host CPU.

For consumers, GPU computing can speed up media-intensive programs that run sluggishly on even the latest multicore CPUs. To take just one example, digital video is wasting untold hours of time as PCs grind through a lengthy transcoding task. CUDA-enabled transcoders like Elemental Technologies' Badaboom can reduce the waiting period to minutes. And cleaning up the video frame-by-frame was unthinkable until MotionDSP's vReveal came along.

The wide appeal of these programs may persuade consumers to pay closer attention to the specs of a new PC. Until now, users who disdained games could be satisfied with a low-end graphics processor integrated with the system chipset. GPU computing opens up entirely new possibilities, potentially making a discrete GPU a must-have feature.

For professionals, GPUs are even more compelling. Whereas the performance of multicore CPUs is rising incrementally, high-performance computing on GPUs is growing by leaps and bounds. As tools improve and developers get the hang of programming these beasts, the gains can exceed two orders of magnitude.

And time is money—sometimes figuratively, sometimes literally. If an energy company can expand its oil and gas reserves by analyzing its seismic probes more quickly, it can pump more crude and capture more market share. If an investment-management company like Bloomberg can calculate the optimum pricing for bonds overnight, while the

markets are closed, it can gain a crucial edge for the next day's trading.

Fermi Advances the Art

One obstacle—the much-discussed difficulty of writing parallel software—is gradually being overcome by better GPU architectures and development tools. Nvidia's Fermi architecture is a leap forward in this regard. It's not just faster. By supporting C++ (and other high-level languages), OpenCL, DirectCompute, and Visual Studio, it will be easier to program, too. As developers look around and see their competitors porting applications to GPUs, resistance to new programming paradigms will fall, and we'll all get faster software.

Perhaps the most exciting prospects are the applications yet to be discovered. Every leap in computing power leads to new applications that were impractical before. Unlocking the power of parallel processing has been a dream of computer scientists for decades. The advent of affordable chip-level multiprocessing may be the catalyst. If one lesson can be drawn from the history of computing, it's that the world has an insatiable appetite for computing power. Someone will always find a way to use it.

Nvidia already has the momentum in this field, ahead of AMD's graphics processors and Intel's unseen Larrabee. The Fermi architecture has the potential to increase that lead. AMD is pursuing a somewhat different strategy with its future “Fusion” chip architecture. Fusion will integrate a graphics processor with the x86 CPU on the same chip—an option unavailable to Nvidia, which doesn't make x86 processors. CPU/GPU integration has various trade-offs, but it will probably boost graphics performance in lower-cost systems. Intel's Larrabee remains an unknown factor, too, and it will offer different trade-offs than conventional GPU architectures.

Nvidia's vulnerability may be its attempt to combine world-class graphics performance with general-purpose compute performance in one chip. With three billion transistors, a Fermi GPU will be more than twice as complex as Nvidia's existing GT200 chips. At some point, if not now, features intended to boost compute performance may compromise the chip's competitive position as an affordable graphics processor. At that juncture, the architectures may have to diverge—especially if the professional market grows as large as the consumer market.

For now, however, Fermi significantly advances the state of the art in this field. It's up to Nvidia to implement the architecture in a good chip design, manufacture the parts without hiccups, and ship the finished products—before competitors can steal back the momentum. ♦

To subscribe to Microprocessor Report, phone 480.483.4441 or visit www.MPRonline.com