AMD Eighth-Generation Processor Architecture

ADVANCED MICRO DEVICES, INC.
One AMD Place
Sunnyvale, CA 94088

Combining Eighth-Generation 32-bit x86 Microarchitectural Performance with 64-bit x86-64 Architectural Capability

AMD's eighth-generation processor architecture, code named "Hammer," ushers in a complete set of microarchitectural advances while fundamentally altering the next-generation of x86-based system architecture. The result is not only a generational leap in processor performance, but also a tremendous level of scalability in a system's delivered performance. The microarchitecture offers native support for 32-bit x86 software and is the first to feature support for the 64-bit x86-64 architecture. This combination of scalable performance for today's x86 software while enabling the use of the x86-64 architecture will lead to a new level of delivered performance, making the Hammer architecture a true eighth-generation x86 architecture.

x86-64 Architecture

AMD's 64-bit strategy allows the latest in processor innovation to be brought to the existing installed base of 32-bit applications and operating systems, while establishing an installed base of systems that are 64-bit capable. The Hammer microarchitecture provides a flexible upgrade path by implementing support for x86-64 architecture while still offering native support for 32-bit x86 software. This is accomplished in a straightforward manner because the Hammer microarchitecture already utilizes key data and address paths that are 64 bits wide and will incorporate a 48-bit virtual address space and a 40-bit physical address space. The support of x86-64 architecture is designed to have minimal impact on the Hammer die size and no impact on processor frequency scaling. Future improvements to the processor core will accelerate both 32-bit and 64-bit applications. High-performance systems using 64-bit capable processors from AMD are planned to be amongst the highest performing 32-bit systems ever built.

A complete discussion of the x86-64 architecture is beyond the scope of this document. For a complete discussion please refer to the *x86-64 Technology White Paper* at http://www.x86-64.org/documentation_folder/white_paper.pdf. Detailed technical information on the x86-64 instruction set architecture is also available at www.x86-64.org.

AMD's Eighth-Generation Processor Microarchitecture Overview

The Hammer processor architecture design is optimized toward the principal goal of delivering next-generation performance to the customer. Achieving this goal requires the ability to strike a balance between next-generation microarchitectural per clock cycle performance and the ability of the microarchitecture to further scale in frequency in a given process technology.
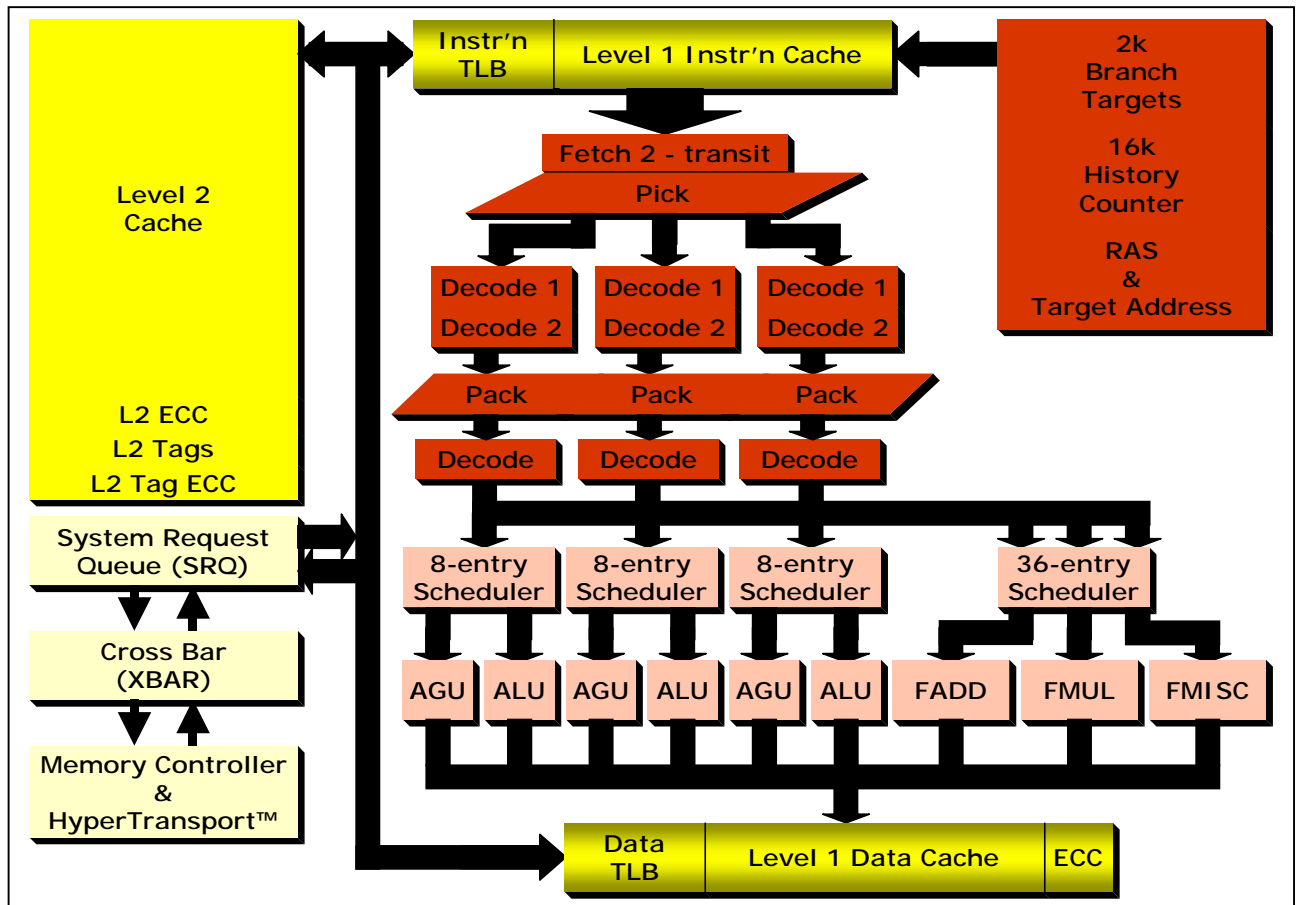

Figure 1: Hammer Microarchitectural Block Diagram

The clearest example of this design philosophy is seen in the changes to the base pipeline in the Hammer processor from the previous generation. The pipeline's front-end instruction fetch and decode logic id refined to deliver a greater degree of instruction packing from the decoders to the execution pipe schedulers. Accommodating this change requires a redefinition of the pipe stages in order to maintain a high degree of frequency

scalability, resulting in two additional pipe stages when compared to the seventh-generation microarchitecture. The end product is a 12-stage integer operation pipeline and a 17-stage floating point operation pipeline. The overall microarchitecture is shown in Figure 1.

Along with frequency improvements from the lengthened pipeline to allow for microarchitectural enhancements, the Hammer processor will be initially produced using a .13micron SOI (Silicon on Insulator) process technology. The microarchitecture is implemented to scale well in frequency with process scaling below the .10micron level. On the other end of the performance equation are the key features that improve the ability of the Hammer microarchitecture to recognize a higher IPC (Instructions executed Per Clock) than previous generation microarchitectures.

Integrated DDR DRAM Memory Controller

As processor microarchitectural capabilities have advanced, one of the greatest limitations on performance gain has become system architecture ability to deliver sufficient memory bandwidth to the processor core while reducing access latency. The Hammer microarchitecture directly addresses this bottleneck by integrating a memory controller into the processor, revolutionizing the way an x86-based processor accesses main memory. The result is greatly increased bandwidth available directly to the processor at a reduced latency.

The Hammer microarchitecture incorporates a dual-channel DDR DRAM controller with a 128-bit interface capable of supporting up to eight DDR DIMMs (four per channel) as seen in the Hammer microarchitectural diagram in Figure 2. The controller will be initially designed to support PC1600, PC2100, and PC2700 DDR memory using unbuffered or registered DIMMs. This translates into available bandwidth to the processor of potentially up to 5.3GB/s with PC2700 memory. This direct interface can significantly reduce the memory latency seen by the processor and latency will continue to drop as the processor frequency scales. Furthermore, this can allow for even more aggressive bandwidth utilization by hardware and software prefetching to further reduce the effective memory latency seen by the processor.
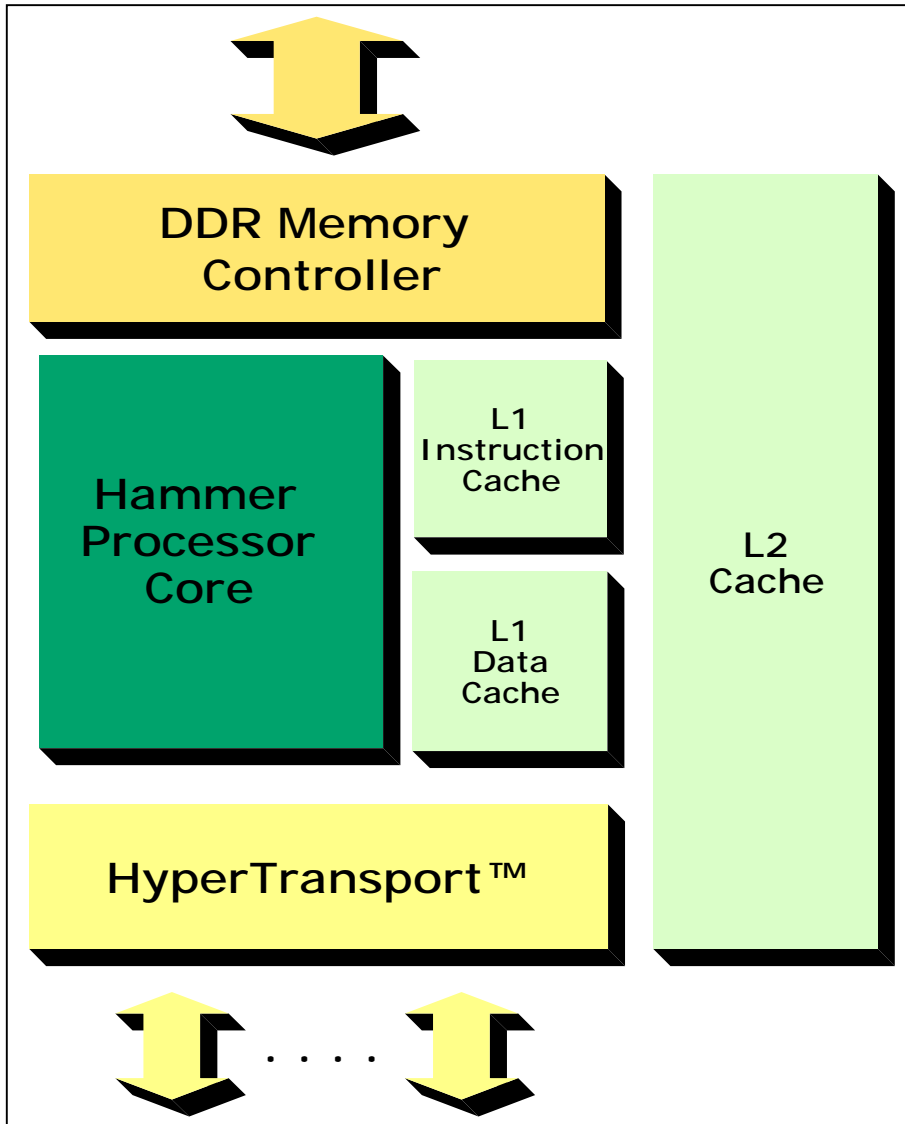
Figure 2: Hammer Microarchitecture Functional Unit Diagram

The integrated Hammer memory controller has an even more dramatic effect when powering multiprocessing systems. The controller results in an outstanding advance in x86 system architecture scalability by enabling "glueless" multiprocessing where the available memory bandwidth to the system scales with the number of processors. In Figure 3, an example of a four-processor multiprocessing system is shown. In this configuration, the system is able to support up to 32 DIMMs capable of delivering an extraordinary 21.3GB/s of available memory bandwidth to the system with PC2700 memory.
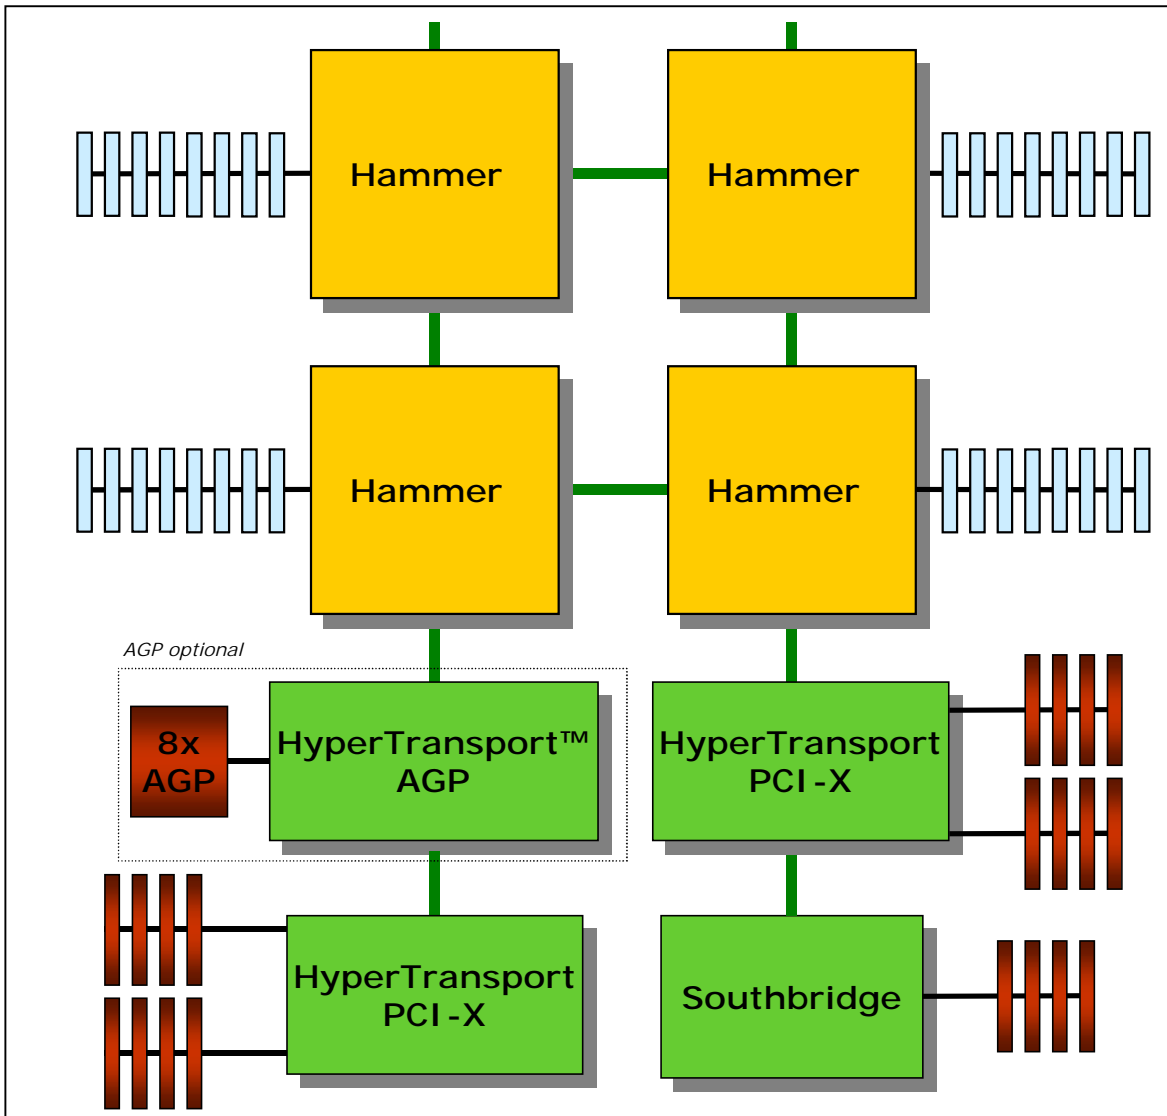
Figure 3: Four-Processor Hammer Processor System Architecture

The integrated memory controller results in much higher bandwidth when compared to AMD's previous processor generations and significantly lowers latencies from processor to memory, and yields a considerable increase in delivered performance as advances are made in memory technology.
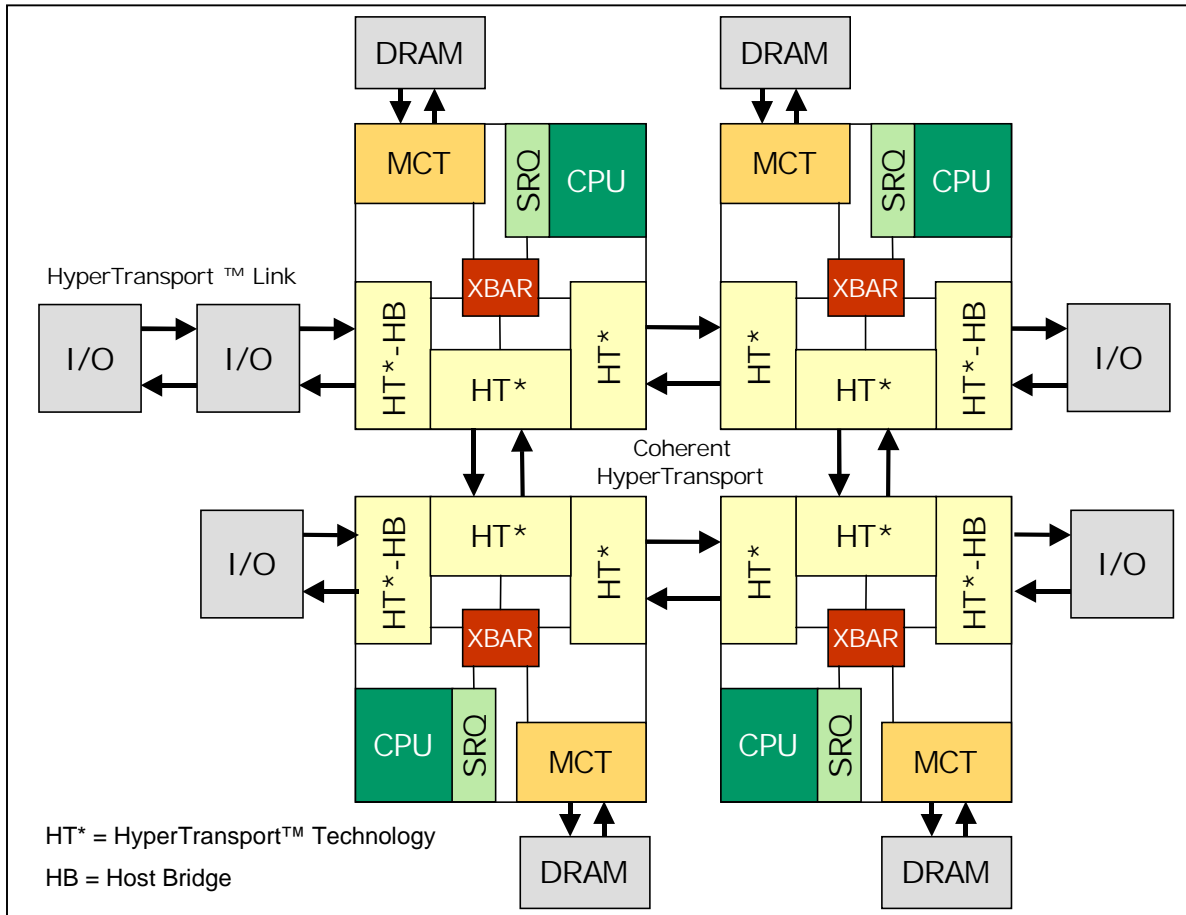
Figure 4: Four-processor Hammer Processor System with Northbridge Blocks

HyperTransport™ Technology and Processor Northbridge Architecture

Similar to the improved scalability in memory bandwidth resulting from the memory controller's integration into processor, HyperTransport™ technology links are built into Hammer processors to provide scalable bandwidth interconnect links between processors and I/O subsystems. The links within the Hammer processor are configured for 16-bit communication in each direction at operating frequencies of 1600MT/s (megatransfers per second), for a bandwidth of 3.2GB/s in each direction.

Figure 4 is an example of a four-processor configuration illustrating how AMD Hammer processors interface between each other and I/O subsystems via HyperTransport technology. The HyperTransport technology links between processors utilize a coherent

protocol, whereas the I/O links follow a non-coherent protocol. A full discussion of HyperTransport technology is beyond the scope of this document but specifications are available from the HyperTransport Consortium.

With three HyperTransport technology links and a memory controller integrated into the processor, it becomes necessary to route command and data information to the responsible interface as efficiently as possible. To address this need, a crossbar communications architecture was developed. The "XBar" shown in the Northbridge details of Figure 4, routes command and data information between the memory controller, the three HyperTransport technology links, and the processor's own SRI (System Request Interface.) Figure 5 provides a more detailed look at the command and data paths, including internal bus widths between the Hammer processor and its Northbridge components. It also shows how a single processor with two Hammer processor cores would be interfaced into the Northbridge by both interfacing a single SRI.
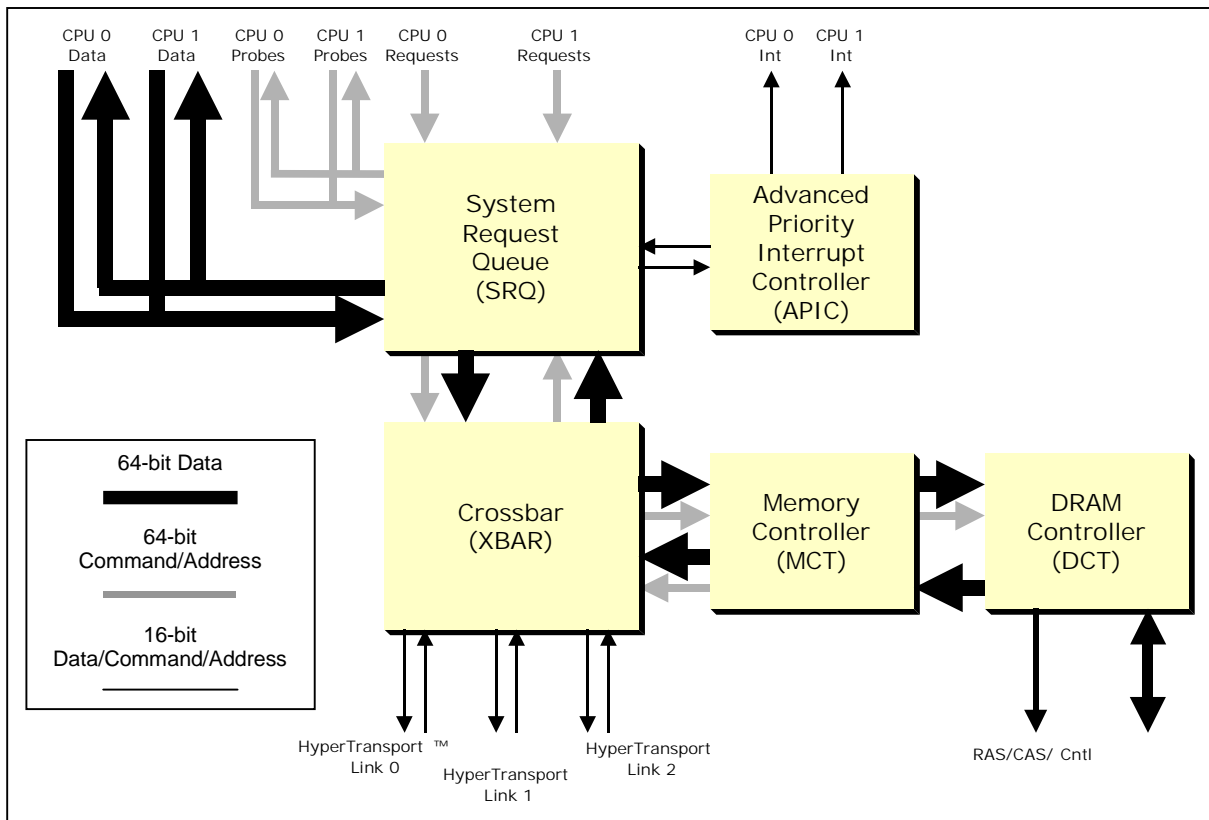

Figure 5: Hammer Processor Northbridge Functional Block Connectivity

TLB Subsystem and Branch Prediction Enhancements for Large Workloads

The Hammer microarchitecture's TLBs (Translational Lookaside Buffers) are configured as shown in Table 1.

Table 1: Hammer Processor Microarchitecture TLB Sizes and Associativities

	Size	Set Associativity	Size Differential from Seventh-Generation
L1 Instruction TLB	40-entry[1]	Fully associative	+16 entries
L1 Data TLB	40-entry[1]	Fully associative	None
L2 Instruction TLB	512-entry	4-way	+ 256 entries (2X)
L2 Data TLB	512-entry	4-way	+ 256 entries (2X)

[1]1.32 4K page entries and 8 2M/4M page entries

In addition to having larger TLB entry sizes with reduced latencies compared to previous generations of AMD processor microarchitectures, the Hammer microarchitecture features a flush filter allowing multiple processes to share TLB without SW intervention.

Hammer processor branch prediction is similarly augmented to offer higher performance, particularly on larger workloads, through an increase in the number of bimodal counters in the global history counter to 16K, which is four times the number in AMD's seventh-generation microarchitecture.

Cache and TLB Reliability Features

ECC (Error Correcting Code) protection is featured for the L1 data cache, L2 cache data and tags, and DRAM with hardware scrubbing of all ECC protected arrays. The DRAM ECC protection also includes support for Chipkill ECC, which provides a means for recovering from errors involving an entire DRAM chip. All non-ECC protected arrays are generally read-only data, which can be recovered, so they are parity protected. Those arrays include the L1 instruction cache and all of the TLBs.

32-Bit x86 Instruction Set Architecture Support

The AMD Hammer processor microarchitecture features support for all 32-bit industry-standard architectural extensions supported by previous AMD processor generations, including Intel's MMX™ and AMD's 3DNow!™ Professional technology (combining Enhanced 3DNow! technology and SSE). In addition, it introduces support for all instructions necessary to be fully compatible with SSE2 technology.

Meeting AMD Eighth-Generation Goals

The AMD Hammer processor architecture has accomplished three overall design goals:

- Establish the x86-64 instruction set architecture.
- Set the precedent for eighth-generation 32-bit x86 performance.
- Build a scalable system architecture that meets the needs of multiple processor generations by integrating the memory controller into the processor microarchitecture and enabling a highly scalable system bus via HyperTransport technology.

The result is a single, highly scalable architecture that provides next-generation performance across industry segments by having a flexible upgrade path from 32- to 64-bit x86 architecture. It is designed to deliver superior performance on today's and tomorrow's applications.

AMD Overview

AMD is a global supplier of integrated circuits for the personal and networked computer and communications markets with manufacturing facilities in the United States, Europe, and Asia. AMD produces microprocessors, flash memory devices, and support circuitry for communications and networking applications. Founded in 1969 and based in Sunnyvale, California, AMD had revenues of $4.6 billion in 2000. (NYSE: AMD).