AMD Athlon™ Processor Floating Point Capability

The Most Powerful, Architecturally Advanced Floating Point Engine Ever Delivered in an x86 Microprocessor

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

Introduction: The Most Powerful Floating Point Engine for x86-Based Platforms

The AMD Athlon™ processor is designed to power the next generation in x86 computing platforms, delivering the ultimate performance for cutting-edge applications and an unprecedented visual computing experience. To deliver this ultimate visual computing experience, AMD has designed the AMD Athlon processor to include the most powerful floating point engine for x86 platforms. This white paper details the next-generation features and technology that set the AMD Athlon processor's floating point engine apart from the floating point units (FPUs) of previous-generation x86 processors.

The Need for a High-Performance Floating Point Engine

High-performance floating point units—once considered to be necessary only for scientific, mathematical, and engineering applications running on workstations—are now used in a variety of other applications, such as modeling physical objects in motion, applying special lighting effects to complex digital images, and enabling more realistic and smoother movements of objects in 3D games.

With the proliferation of 3D-intensive software, many computing platforms require powerful floating point capabilities. At home, the popularity of 3D entertainment and gaming titles has skyrocketed as software developers have tapped the power of floating point to deliver ever increasing levels of realism and detail in their latest titles. Robust floating point performance has become essential in the workplace where the latest CAD, desktop publishing, digital content creation, spreadsheet, business graphics, mathematical, and statistical applications rely on floating point units to perform complex calculations. Exceptional floating point performance continues to be a critical capability for scientific and engineering workstation applications.

To meet the floating point (FP) performance demands of cutting-edge, FP-intensive software applications, AMD equipped the AMD Athlon processor with the most powerful and architecturally advanced floating point execution engine ever implemented in an x86 microprocessor[1]. A centerpiece of the seventh-generation AMD Athlon processor architecture, AMD's "no-compromise," superscalar floating point engine complements the capabilities of the AMD Athlon processor's enhanced 3DNow!™ technology (see "Enhanced 3DNow!™ Technology" white paper) and enables FP performance that is unparalleled in the x86 computing industry.

[1] As measured by the industry-standard SPECfp_base95 benchmark; see Appendix for details.

A Few Floating Point Basics

Floating point arithmetic is simply a way to represent a real number with a "floating" decimal point rather than a "fixed" decimal point. It enables more precise mathematical calculations than possible with *integers*, or whole numbers that lack a decimal point. Floating point techniques are especially useful for geometric calculations common in CAD/CAE processing, high-precision mathematical calculations, and 3D graphics applications involving physics, geometry, and triangle setup.

Processor-based floating point arithmetic has been thoroughly standardized. The IEEE-754 and 854 standards specify the numeric formats, value sets, and how the basic floating point arithmetic works. The basic floating point types include:

- Single-precision format – 32-bit data width, 6 to 7-digit precision
- Double-precision format – 64-bit data width, 15 to16-digit precision
- Extended-precision format – 80-bit data width, 18 to 20-digit precision.

Early x86 processors interfaced to a separate, dedicated x87 floating point unit, also called the "math co-processor," to perform numerically intensive calculations. Beginning with the 486 processor, the x87 floating point unit was integrated onto the processor die, delivering on-chip floating point functionality and performance to application developers.

The next innovation in floating point technology for x86 processors was AMD's 3DNow! technology, introduced in May 1998 as a key feature of the AMD-K6®-2 processor. AMD's 3DNow! was the first 3D-enhancement instruction set for the x86 architecture, as well as the first SIMD (Single Instruction Multiple Data) floating point technology for x86 processors. In essence, 3DNow! technology is a set of 21 instructions using superscalar SIMD floating point capabilities to significantly improve the processor's ability to handle floating point calculations, dramatically improving 3D graphics, audio, and other multimedia performance. SIMD enables AMD processors to execute multiple single-precision floating point calculations per clock cycle.

3DNow! technology was designed to complement the role of the x87 floating point unit by enhancing the processor's overall ability to accelerate certain algorithms that lend themselves to SIMD single-precision floating point computation, such as physics and geometry operations. In fact, a primary rationale for developing 3DNow! was to provide specific acceleration for those applications that are well-suited to vector processing of single-precision floating point code. Applications requiring greater than single-precision processing are better handled by the x87 FP engine.

Today's software developers require exceptional floating performance across all computing platforms to power new generations of FP-intensive applications, from 3D imaging and modeling to CAD/CAE to speech recognition to sophisticated scientific and technical applications. The seventh-generation AMD Athlon processor now delivers the x86 industry's most powerful floating point engine to meet these processing needs and to enable superior FP performance for the next generation of x86 platforms.

Anatomy of the Powerful AMD Athlon Processor™ Floating Point Engine

The AMD Athlon processor features a three-issue, superscalar floating point design based on three pipelined, out-of-order floating point units, each with a one-cycle throughput. These three execution pipelines make the AMD Athlon processor's floating point engine both *superscalar* and *superpipelined*. The term superscalar means that the processor has multiple instruction pipelines that can operate in parallel, enabling the processor to execute multiple instructions in a single clock cycle. Superpipelining means that there are many small pipeline stages—each performing a simple, dedicated task—to optimize processor clock frequency.

The AMD Athlon processor is the world's first x86 processor with these capabilities, delivering unprecedented performance to satisfy the most demanding applications on x86 platforms. Figure 1 (AMD Athlon architecture block diagram) details the three FP execution pipes used in the AMD Athlon processor.

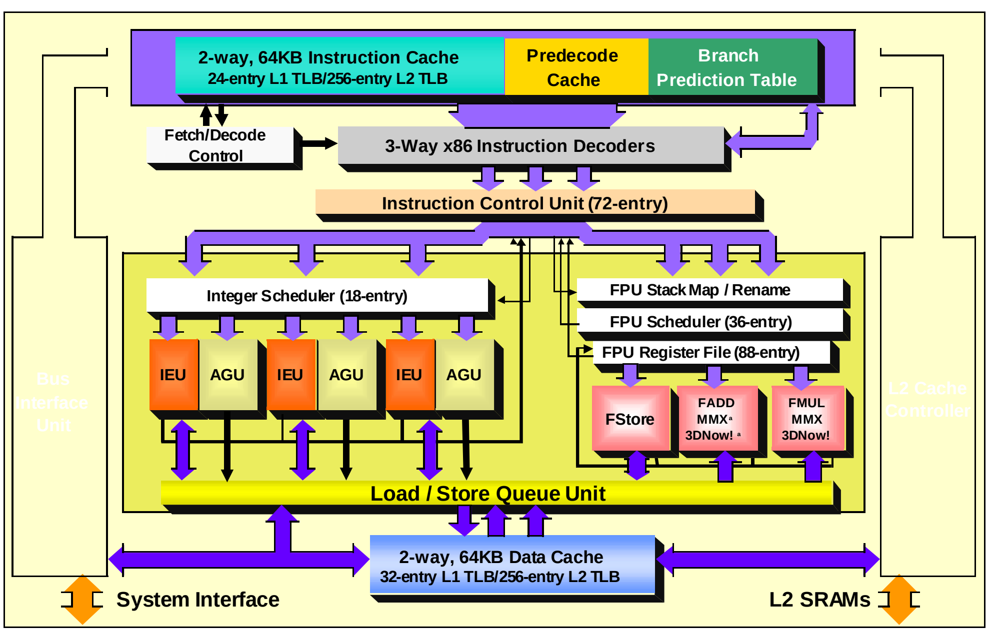

Figure 1: AMD Athlon™ Processor Block Diagram

When compared to other floating point architectures, the design and performance of the AMD Athlon processor's floating point engine set it apart as true next-generation implementation. For example, the AMD Athlon processor delivers twice the peak x87 floating point execution rate of the Intel Pentium® III processor and rivals the FP performance of many RISC processors currently used in workstations and servers.

The floating point engine of the AMD Athlon processor is fully x87 compatible and conforms to the IEEE 754 and 854 standards for single-, double-, and extended-precision data types. The AMD Athlon processor's floating point engine also can handle all MMX™, 3DNow!, and x87 floating point instructions.

The AMD Athlon processor's floating point engine uses three separate execution pipelines for superscalar x87, 3DNow!, and MMX operations:

- The floating point load/store (FSTORE) pipeline handles FP constant loads, stores, and miscellaneous operations.
- The adder pipeline (FADD) contains 3DNow! add, MMX ALU/shifter, and FP add execution units.
- The multiplier (FMUL) pipeline contains an MMX ALU, MMX multiplier, reciprocal unit, FP/3DNow! instruction multiplier, and support for FDIV instructions.

In addition to its superscalar design, the AMD Athlon processor's floating point engine is superpipelined. This technique supports higher clock frequencies and enables the floating point engine to process complex FP instructions more quickly and deliver higher overall FP instruction throughput. Table 1 compares the number of FP execution pipelines, as well as FP throughput, of the AMD Athlon processor versus the Intel Pentium III.

Table 1: AMD Athlon™ Processor Floating Point Engine Features/Latency/Throughput Comparison

	AMD Seventh Generation	Intel Previous Generation
Feature	**AMD Athlon™ Processor**	**Intel Pentium® III and Pentium III Xeon**
No. of FP execution pipelines	3	1
Latency / Throughput Rates for Classes of FP Instructions*		
❑ FADD reg	4/1	3/1
❑ FMUL reg	4/1	5/2
❑ FDIV reg (single)	16/13	17/17
❑ FDIV reg (double)	20/17	32/32
❑ FDIV reg (extended)	24/21	37/37
❑ FSQRT (single)	19/16	28/28
❑ FSQRT (double)	27/24	57/57
❑ FSQRT (extended)	35/31	68/68
❑ FCOM	2/1	1/1

***** Latency/throughput rates are measured in processor clocks; lower numbers mean higher throughput.

A pipeline is analogous to an automobile assembly line, which is divided into many functional parts, each dedicated to building a particular piece of the entire car. This efficient assembly line model increases the factory's overall throughput. Similarly, using a pipeline (or instruction "assembly line") enhances the processor's overall throughput.

Table 1 compares the floating point throughput of the AMD Athlon processor relative to the Pentium III processor for various FP instructions. The AMD Athlon processor is clearly superior to the Pentium III on the important multiplication, division, and square-root instructions.

The seventh-generation AMD Athlon processor's floating point performance is significantly better than that of any previous-generation x86 processor at the same clock speed, including the Pentium III. This higher floating point performance provides end users with a more productive computing experience, with such benefits as faster video frame rates, greater detail, superior realism for 3D graphics and computer-aided design applications, faster image rendering and 3D modeling for digital content creation, and faster mathematical and statistical calculations for business and workstation applications.

Summary: Unprecedented FP Performance for x86 Platforms

Cutting-edge, floating point-intensive software applications for next-generation x86 computing platforms require exceptional, new levels of floating point performance. AMD delivers this new level of performance in the next-generation AMD Athlon processor, which is designed to offer unprecedented floating point performance for high-end x86 desktop systems, workstations, and servers. The AMD Athlon processor's floating point engine is the most powerful floating point solution ever implemented in an x86 processor and is planned to power new and exciting visual computing applications for computing enthusiasts, small business users, and enterprise users for years to come.

AMD Overview

AMD (NYSE: AMD) is a global supplier of integrated circuits for the personal and networked computer and communications markets. AMD produces processors, flash memories, and integrated circuits for telecommunications and networking applications. The world's second-leading supplier of Windows® compatible processors, AMD has shipped more than 120 million x86 microprocessors, including more than 90 million Windows compatible CPUs. Founded in 1969 and based in Sunnyvale, California, AMD has sales and marketing offices worldwide and manufacturing facilities in Sunnyvale;

Austin, Texas; Dresden, Germany; Bangkok, Thailand; Penang, Malaysia; Singapore; and Aizu-Wakamatsu, Japan. AMD had revenues of $2.8 billion in 1999.

Cautionary Statement

This White Paper includes forward-looking statements that are made pursuant to the safe harbor provisions of the Private Securities Litigation Reform Act of 1995. Forward-looking statements are generally preceded by words such as “expects,” “plans,” “believes,” “anticipates,” or “intends.” Investors are cautioned that all forward-looking statements in this white paper involve risks and uncertainties that could cause actual results to differ from current expectations. Forward-looking statements in this white paper about the AMD Athlon processor involve the risk that the AMD Athlon processor will not achieve customer and market acceptance; and that software applications will not be optimized for use with the processor. We urge investors to review in detail the risk and uncertainties in the company’s Securities and Exchange Commission filings, including the most recently filed Form-10K.