

# **AMD Opteron™ Processor Benchmarking for Clustered Systems**

**Doug O'flaherty and Michael Goddard**

AMD  
One AMD Place  
Sunnyvale, CA 94088

## Objective

The selection of a High-Performance Computing (HPC) platform is a process of evaluation and analysis. The goal is to deliver the appropriate combination of SMP and cluster configuration, processor, memory, interconnect, and storage that optimizes performance. This is often done without being able to run the customer-specific application on the complete system prior to selection. Since the application specific metrics are absent, industry benchmarks provide a basis for comparison. Industry standard benchmarks are system or sub-system metrics for performing a particular computing task. There are as many benchmarks as there are computing tasks. This paper examines AMD Opteron™ processor performance across a suite of industry standard benchmarks. The metrics reflect the AMD Opteron processor's unique architecture, which balances bandwidth with outstanding computing performance to deliver true scalability for HPC applications.

## Introduction

Benchmarks seek to provide comparative results—a quantitative analysis of disparate systems performing similar computational tasks. The benchmarking effort typically attempts to set performance expectations of real world tasks. However, to achieve objective results, each benchmark is limited in the scope of its measurements. Some benchmarks emphasize raw floating-point operations per second, while others measure memory access latency and bandwidth, while others script a particular workload in a specific application. Because applications are diverse and more complex than a combination of metrics, no single benchmark is a metric of a processor's ability to perform real-world computation.

The challenge of benchmark analysis lies in discerning the appropriateness of a system or subsystem for user-specific applications. Integral to this process is an in-depth understanding of the particular applications for which the system will be used. Equally important is examining the suite of benchmarks to illustrate how the underlying processor architecture, subsystems, and compilers integrate to deliver application performance.

In this paper we examine a suite of benchmarks for AMD's new processor, the AMD Opteron processor, to illustrate its performance and scalability in single, multi-processor, and cluster configurations. The AMD Opteron processor integrates outstanding performance, low latency, high-bandwidth memory, and glueless multi-processor support to deliver improved application performance for High-Performance Technical Computing.

## The AMD Opteron™ Processor delivers

AMD64: a 64-bit computing platform that extends x86 for increased performance and scalability.

The advantages of 64-bit computing with scalable 32-bit computing power:

- AMD Opteron processor SPECint® performance leadership over Pentium® 4
- Significant AMD Opteron processor 2P and 4P SPECint\_rate performance leadership, dominating Xeon in 2P configurations by over 30%
- AMD Opteron processor SPECfp® performance leadership over Pentium 4
- Dominating 2P and 4P SPECfp\_rate over Xeon and Xeon MP (over 70% and 150% margin respectively)
- Highly efficient and scalable math libraries

## Low latency, scalable memory bandwidth with on-chip memory controller

Adding processors increases memory bandwidth:

- Over 5GB/s of available 128-bit wide DDR memory bandwidth per CPU
- Memory latency under 60ns; 70% less than Pentium 4
- Memory bandwidth and latency improve with frequency because the memory controller runs at processor frequencies

## Coherent HyperTransport™ Technology for glueless multiprocessor systems

Architecture eliminates Northbridge contention minimizing memory latency in multi-processor systems:

- Under 105ns for 2P platforms
- Under 140ns for 4P platforms; over 20% faster than Xeon MP
- Additional benefit for NUMA aware operating systems; over 48% for 2 processor and 35% for 4 processor AMD Opteron™-based platforms (with a worst-case 4P latency over 18% better than Xeon MP)
- We'll examine the performance benefits of the AMD Opteron processor as illustrated by leading benchmarks, beginning with the memory controller and growing to the cluster.

## Memory Bandwidth

In most systems, the memory controller is located on a separate chip, accessed via a shared bus operating at bus speed (typically a fraction of the processor's speed). While CPU power has increased apace with Moore's Law, memory bandwidth has not. In multiprocessor configurations and higher CPU clock speeds, shared bus bandwidth does not improve. As the processor outstrips the data delivery from the memory system, the percentage of CPU idle time increases. Increasing the on-board cache, commonly used to lessen the impact of slow memory, does not scale with multiprocessor systems and is only effective for smaller problems; but it does negatively impact manufacturing and sale costs because it requires an increased die size.

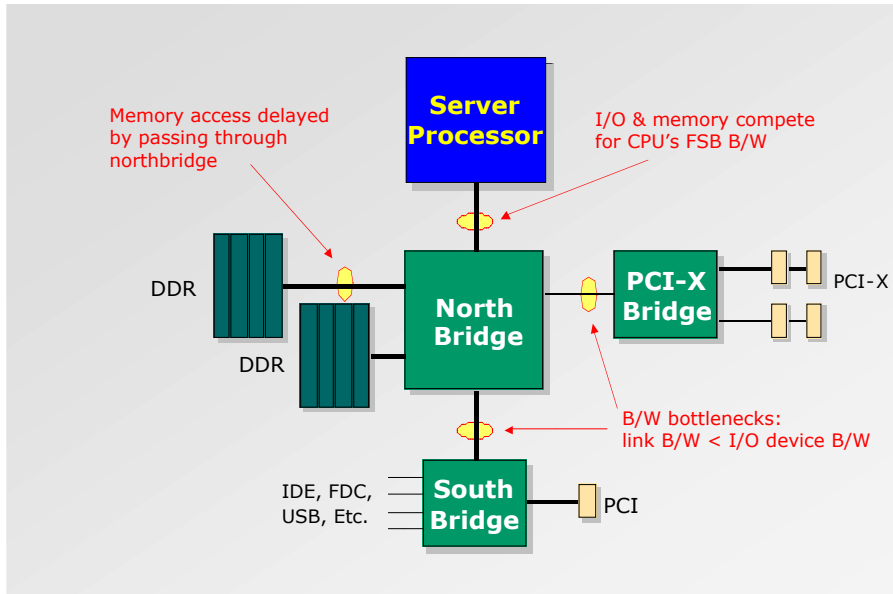


Figure 1: Previous Generation System Architecture: Northbridge Centric Design

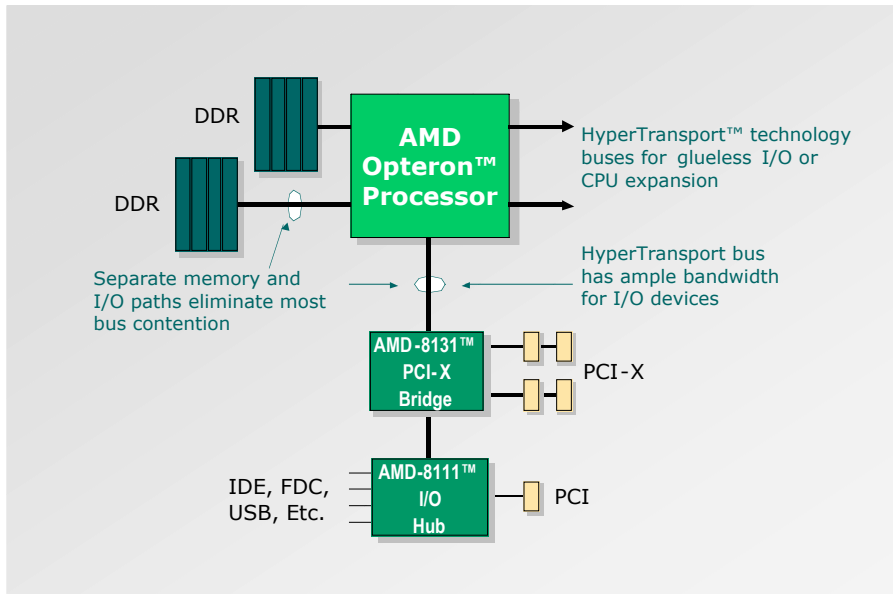


Figure 2: AMD Opteron Processor: Processor Centric Design

The AMD Opteron processor’s design rebalances the disparity between memory bandwidth and fast CPU operations with an on-chip memory controller. The AMD Opteron processor’s memory controller is 128-bits wide, operates at CPU clock speed, and supports up to 16GB (limited by the current maximum of 2GB/DIMM) of 200, 266, or 333MHz DDR RAM per CPU. This architecture provides true multiprocessor bandwidth scalability, dramatically increasing memory bandwidth with each AMD Opteron processor added to the system.

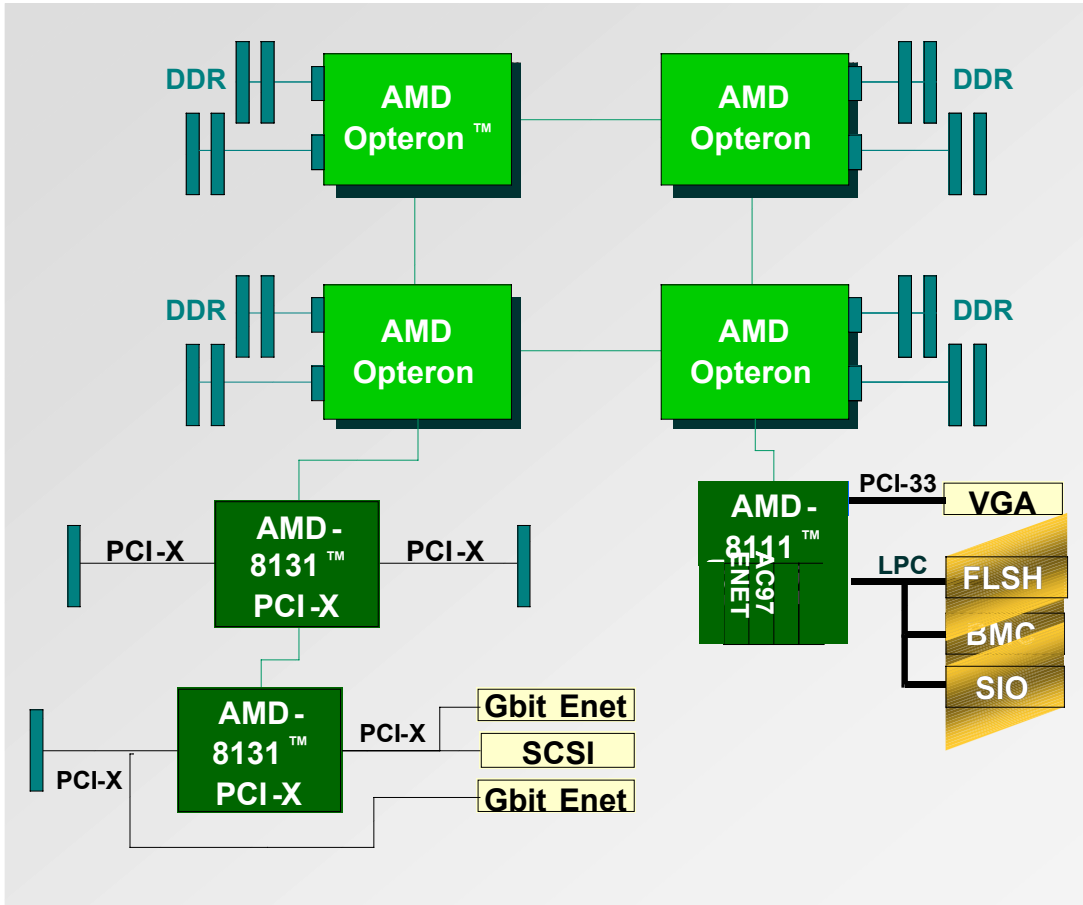



Figure 3: Glueless Multi-processing and Scaling Bandwidth with AMD Opteron processor 800 Series

## **SiSoftware Sandra Standard 2003/SP1 9.44 Memory Bandwidth (based upon the STREAMS benchmark)**

The results from SiSoftware's Sandra 2003 memory bandwidth benchmark (Figure 1a) clearly show the scalability of the AMD Opteron processor's architecture. The on-chip memory controller performance delivers over 4700MB/s when a single CPU is connected to 4 DIMMs, on par with the highest end Pentium 4 desktop platform. In a 2- and 4-CPU system, each CPU with 4 PC2700 DIMMs, the system delivers nearly 9GB/s and over 13GB/s of memory bandwidth respectively, significantly distancing the Xeon and Xeon MP ServerWorks-based chipset platforms.

Within the AMD Opteron processor 200 series, Figure 1b shows how the 2 CPU system memory bandwidth scales with different memory configurations. For example, a 4-4 configuration is a 2P platform with 4 DIMMs populated per processor and a 2-0 configuration is a 2P platform with 2 DIMMs populated only on 1 processor. The data clearly shows the memory bandwidth benefits of using multiple DIMMs off each 128-bits of each processor. The performance gain realized on any specific application will be dependent on the application sensitivity to memory bandwidth.


### Sisoftware Sandra Standard 2003/SP1 9.44



See Appendix A for system configuration information

Figure 4: SiSoftware Sandra Memory Bandwidth for AMD and Intel processors

**Memory Bandwidth by  
Memory Controller DIMM Population  
(AMD Opteron™ Model 246 2P Server)**



See Appendix A for system configuration information

Figure 5: SiSoftware Sandra Memory Bandwidth for AMD Opteron™ Processor Model 246 in Different Memory Configurations

## About the Benchmark

SiSoftware Sandra (<http://www.sisoftware.net/>, the **S**ystem **AN**alys**er**, **D**iagnostic and **R**eporting **AS**stant) is a modular suite of tests. The memory bandwidth test is adapted from the well-known STREAM (<http://www.cs.virginia.edu/stream>) memory bandwidth benchmark. There are several differences between STREAM and Sandra’s implementation, which precludes comparison of the two related benchmarks. Sandra aggressively optimizes the test to reduce possible degrading of the test by other subsystems. Aggressive instruction scheduling minimizes CPU impact. Sandra is multi-threaded, assigning a thread for each CPU for MP systems. Dynamic data of approximately 40–60% of the physical RAM as the data set is used as the problem size and alignment is optimized for throughput.

In contrast, STREAM uses static data of about 12M. STREAM alignment optimization is only available at compile time.


## Memory Latency

The AMD Opteron processor was designed to remove the inefficiencies introduced in legacy system architectures associated with relatively slow front-side buses and with memory controllers embedded in discrete chipsets. By integrating the memory controller, the AMD Opteron processor allows memory latency to more closely approach the limits of the memory technology and to improve with processor frequency. In addition, the overhead to ensure coherency is reduced by the glueless multiprocessing enabled by HyperTransport™ technology.

## ScienceMark 2.0 Beta MemBench Memory Latency

An immediate benefit of the on-board memory controller is reducing latency. Located within the processor and running at CPU clock speed, the memory controller is exceptionally responsive. The results of ScienceMark's MemBench clearly show this responsiveness. With a result of 58ns, the latency for an AMD Opteron processor Model 146 system is over 70% less than that of the highest-end Pentium 4 desktop platform. For AMD Opteron processor Model 246 systems, the memory latency is 70ns for local memory accesses and 104ns for remote memory accesses. With the introduction of NUMA support in Microsoft Windows Server 2003 and Linux Kernel, applications realize the benefit of optimizations for localized (0-hop) accesses.

## ScienceMark 2.0 Beta, 512-Byte Stride Latency (ns)



See Appendix A for system configuration information

Figure 6: ScienceMark 2.0 Beta Latency (ns) for AMD Opteron™ and Intel processors

## About the Benchmark

ScienceMark MemBench (<http://www.sciencemark.org>) runs over 15 different memory bandwidth algorithms and reports the highest performance for each platform. Each test is a variation on copying one stream or vector to another stream or vector.

## SPEC<sup>®</sup> cpu

The Standard Performance Evaluation Corporation ([www.spec.org](http://www.spec.org)) developed and administers the SPECcpu\_2000 industry standard benchmark of CPU performance. This test suite measures compute intensive tasks in two categories, SPECint for integer performance and SPECfp for floating-point performance. The tests are based on real applications using common compute-intensive tasks ranging from data compression to particle acceleration modeling. The tests are portable and can be run on any operating system and architecture using the SPEC code and problem set. Depending on the test, the SPEC code is C, C++, and Fortran (77 & 95). The SPECcpu\_2000 benchmark measures the processor performance, memory architecture, and the compiler.

### SPECint


There are 12 tests in the SPECint\_2000 suite based upon common productivity tasks including problems from digital rendering (ray tracing), CAD (circuit routing), and resource scheduling (network flow optimization).

The SPEC peak benchmark reflects speed at which a single task can be completed by the system when optimized for each of the tests in the suite. This reflects the HPC environment where each application is optimized to perform as well as possible on the platform. The SPEC benchmark result is computed as the geometric mean of the SPECratios of each problem. Across the range of integer tasks, the AMD Opteron processor Model 146 outperforms Xeon 3.0GHz.

SPEC rate measures the system's ability to perform multiple tasks, running multiple versions of the test simultaneously to show throughput. The metric reflects the system's ability to scale within a node.

The AMD Opteron processor's low-memory latency, high bandwidth, and glueless multiprocessing architecture are inherently scalable. In a 2P configuration, the AMD Opteron processor extends its advantage over the 2P Xeon systems. With HyperTransport technology and the integrated memory controller, even the AMD Opteron processor Model 242 2P integer outperforms the Xeon 3.06GHz system by 7%.

**SPECint®\_peak2000 Performance  
(Uniprocessor, Windows®)**



See Appendix A for system configuration information

Figure 7: SPECint\_peak2000 Performance

**SPECint®\_rate2000 Performance  
(Peak, 2P)**


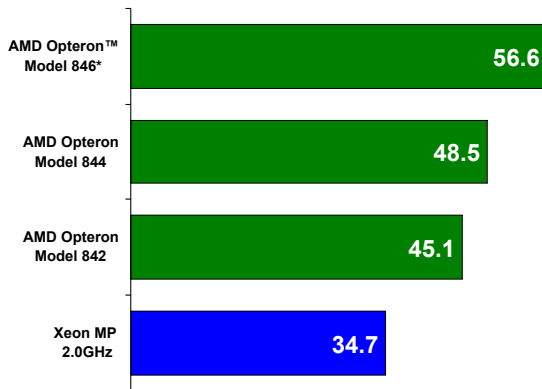


Figure 8: SPECint\_rate2000 Performance (2P)

**SPECint®\_rate2000 Performance  
(Peak, 4P, Windows®)**



See Appendix A for system configuration information

Figure 9: SPECint\_rate2000 Performance (4P)

**SPECint®\_rate2000 Performance (Peak,  
2-4P Scaling)**




Figure 10: SPECint\_rate2000 2P-4P Scaling

\*Denotes estimated score

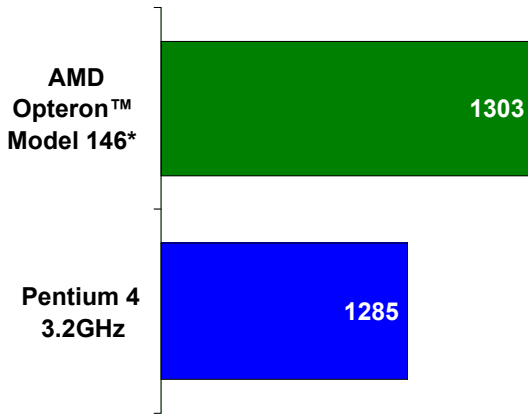
## SPECfp

There are 14 tests in the SPECfp\_2000 benchmark. The tests were derived from applications such as crash simulations, ocean modeling, and quantum chromodynamics. The majority of these tests are large number problems and written in Fortran. To reflect the real world, the problem sets in the SPECfp\_2000 will not fit in 32-bit CPU cache. SPECfp exercises the memory subsystem and exposes limitations in scalability to multiple processors and large problem sizes.

The AMD Opteron processor is a general-purpose x86 CPU, with a balance of floating-point and integer performance. The AMD Opteron processor performs well on the uniprocessor SPECfp\_2000 benchmark, well ahead of Xeon, Intel's x86 processor.

As the system scales, the advantage of balanced processor architecture is clear. The shared bus architecture of the Xeon MP 4P scales worse than the Xeon 2P, here the AMD Opteron processor Model 846 delivers over 20% better scaling. Using coherent HyperTransport to connect the processors, the AMD Opteron processor's architecture avoids bus contention. The glueless architecture and dedicated memory controller deliver outstanding 4P floating-point performance.

**SPECfp®\_peak2000 Performance (Uniprocessor)**



See Appendix A for system configuration information

Figure 11: SPECfp\_peak2000 Performance (1P) (2P)

**SPECfp®\_rate2000 Performance (Peak, 2P)**


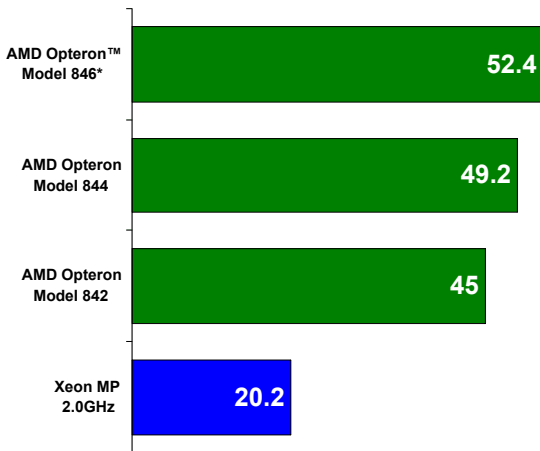


Figure 12: SPECfp\_rate2000 Performance

**SPECfp®\_rate2000 Performance (Peak, 4P)**



See Appendix A for system configuration information

Figure 13: SPECfp\_rate2000 Performance (4P)

**SPECfp®\_rate2000 Performance (Peak, 2-4P scaling)**

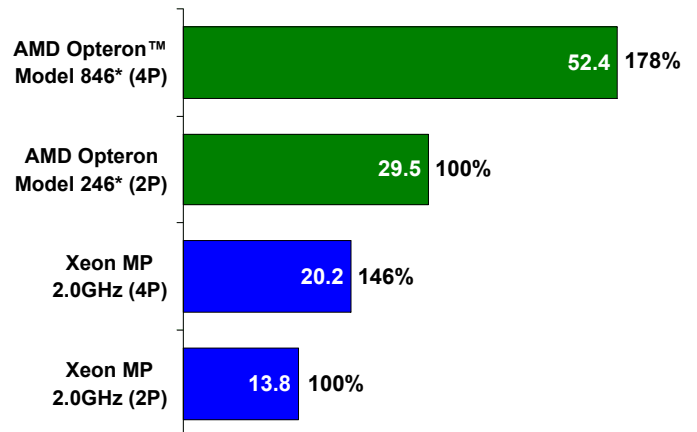


Figure 14: SPECfp\_rate2000 2P-4P Scaling

\*Denotes estimated score

## High-Performance Linpack (HPL)

Developed by Jack Dongarra, Linpack has been the metric for determining the world's top supercomputers since 1993. Linpack is a multiplication matrix solver for a dense linear system in double precision on distributed memory systems using Gaussian elimination. Different versions of the Linpack benchmark are distinguished by their problem size. Since Linpack uses MPI, it's also very well suited for evaluating cluster peak computing power.

Linpack results are a combination of the raw floating-point resources, the efficiency of the CPU architecture, the optimized math library, and the compiler; with the results measured in floating point operations per second or FLOPS. Many HPC applications, such as mechanical-analysis and computational fluid dynamics, are floating-point intensive.

The HPL (High-Performance Linpack) benchmark allows user-selection of the problem size yielding the highest performance. This is often the largest problem size that will fit in memory. The HPL benchmark provides the following information:

- Nmax: The matrix size. The memory required is  $N^2$  times 8 bytes.
- Rmax: The maximum number of FLOPS achieved for that problem size.
- N1/2: The problem size achieving 50% of Rmax. A low N1/2 shows a robust system delivering strong performance on a broad range of problem sizes.
- Rpeak: The theoretical maximum FLOPS for the system determined by multiplying the floating-point operations per clock cycle, the CPU clock, and the number of processors.
- Rmax/Rpeak: The calculated efficiency of the computer. This value is never 100%.

Several math libraries are currently available for the AMD Opteron processor including ACML (AMD Core Math Libraries), Atlas (Automatically Tuned Linear Algebra Software), and the specialized GOTO BLAS (Basic Linear Algebra Subroutines). Additional libraries are under development; see [developer.amd.com](http://developer.amd.com) for recent announcements. For the HPL benchmarks in this paper, AMD has used the GOTO libraries which are especially well tuned for HPL benchmarking.

Linpack benchmarks demonstrate the AMD Opteron processor's impressive efficiency, achieving 87.1% of the theoretical peak FLOPS on a single processor system. As reported by Supercomputing Online, a single CPU 2.4GHz Xeon using that platform's GOTO library reports only 81.2% efficiency. To put this in perspective, the Xeon CPU is not working on the problem almost 19% of the time, while the AMD Opteron CPU only spends less than 13% on non-result tasks.

The AMD Opteron processor's large memory bandwidth, improved prefetch, and outstanding branch prediction, support high levels of computational efficiency while minimizing idle time and wasted CPU cycles. While GOTO library results are outstanding, the AMD Opteron processor's architectural efficiency is not library dependant. Using the open source Atlas libraries on the same systems achieves over 84.5% efficiency, still well ahead of the reported Xeon results.

As indicated by the SPEC rate and memory benchmarks, the AMD Opteron processor scales exceptionally well. Multiprocessor AMD Opteron processor-based systems use coherent HyperTransport technology between the processors for dedicated low-latency and high-bandwidth communication. Traditional Northbridge bottlenecks have been eliminated in the multiprocessor AMD Opteron processor-based design by separate paths for memory access. The AMD Opteron processor's glueless MP architecture shows less than 1.5% efficiency loss moving from single processor to dual processors, an inherent inefficiency in multiprocessor scaling caused by increased kernel overhead required to manage multiple threads. Traditional Northbridge architecture with a shared memory bus, such as the Xeon, does not perform as well. The Dual CPU efficiency for the 2.4GHz Xeon, as reported by Supercomputing Online, is 71.4%—a loss of almost 10% relative to the single CPU.

AMD Opteron processors also remain highly efficient scaled to four processor systems. The HPL efficiency for a 4P system is almost 84%; 77.8% with the Atlas library on a large problem size.

Common practice builds HPC clusters from dual CPU nodes; architecture driven by price and performance. Price increase and computational efficiency decrease are not offset by the corresponding reduction in interconnect and administration costs resulting from using 4P nodes. Scaling with AMD Opteron processors challenges this assumption with a 1P to 4P efficiency penalty as low as 3.3%.

## GOTO Library Benchmarks

GOTO Library Results							
AMD Opteron™ system	#P	Rmax (Gflops)	Nmax (order)	N1/2 (order)	Rpeak (Gflops)	GFLOP/Proc	Rmax/Rpeak
4P AMD Opteron 1.8GHz 2GB/proc PC2700 8GB Total	4	12.06	28000	1008	14.4	3.02	83.8%
2P AMD Opteron 1.8GHz 2GB/proc PC2700 4GB Total	2	6.22	20617	672	7.2	3.11	86.4%
1P AMD Opteron 1.8GHz 2GB PC2700	1	3.14	15400	336	3.6	3.14	87.1%

See Appendix A for system configuration information

Table 1: High-Performance Linpack (HPL) Performance—GOTO r0.1 Libraries

## Atlas Library Benchmarks

ATLAS 3.5.1 Library Results							
AMD Opteron™ system	#P	Rmax (Gflops)	Nmax (order)	N1/2 (order)	Rpeak (Gflops)	GFLOP/Proc	Rmax/Rpeak
4P AMD Opteron 1.8GHz 8GB/proc PC2700 32GB Total	4	11.60	60114	1123	14.4	2.90	77.8%
2P AMD Opteron 1.8GHz 2GB/proc PC2700 4GB Total	2	6.009	19320	616	7.2	3.00	83.5%
1P AMD Opteron 1.8GHz 2GB PC2700	1	3.042	14000		3.6	3.04	84.5%

See Appendix A for system configuration information

Table 2: High-Performance Linpack (HPL) Performance—Atlas Libraries version 3.5.1

**Notes:****References:**

Top 500 Supercomputers:

<http://www.top500.org/>

Linpack:

<http://www.netlib.org/linpack/>

Supercomputing Online:

<http://www.supercomputingonline.com/article.php?sid=3427>

**Links for Libraries:**

ACML:

[http://www.amd.com/us-en/Processors/DevelopWithAMD/0,,30\\_2252\\_2282,00.html](http://www.amd.com/us-en/Processors/DevelopWithAMD/0,,30_2252_2282,00.html)

Atlas:

[http://www.amd.com/us-en/Processors/DevelopWithAMD/0,,30\\_2252\\_2272\\_8716,00.html](http://www.amd.com/us-en/Processors/DevelopWithAMD/0,,30_2252_2272_8716,00.html)  
<http://sourceforge.net/projects/math-atlas/>

GOTO:

<http://www.cs.utexas.edu/users/flame/goto/#obtain>

**Conclusion**

When choosing a processor for compute intensive applications maximizing the efficiency and scalability are critical. The AMD Opteron processor design reduces the bottlenecks of the previous generation x86 processors and delivers bandwidth for memory and I/O. Increased and scalable access to data in the AMD Opteron architecture balances the processor's ability to compute with the architecture's ability to deliver data to the processor core. Higher benchmark scores reflect a faster time to solution for compute intensive applications. In the 32-bit SPEC results, the AMD processor family consistently outperforms the Intel x86 family of processors.

In the traditional measurement of cluster performance, High Performance Linpack, the 64-bit AMD Opteron nodes deliver floating-point efficiency and larger problem sizes than is possible on the Xeon systems. Unlike the traditional Front Side Bus architecture, the AMD Opteron architecture maintains efficiency as the number of CPUs in the node increases.

In comparison to parallel products in the marketplace, these benchmarks mean that the AMD Opteron delivers a 64-bit computing platform with increased performance through MP scaling and an on-chip memory controller that reduces memory access latency. For both floating point and integer operations, the AMD Opteron consistently outpaced its Xeon competitors.

## Appendix A: System Configuration Information

### Memory Bandwidth

SiSoftware Sandra Standard 2003/SP1 9.44

- 1 AMD Opteron™ processor Model 146 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft® Windows® Server 2003 Enterprise Edition, BIOS PQDT006, Max6L040J2 36GB IDE hard disk, all drivers from standard OS install
- 2 AMD Opteron processors Model 246 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition, BIOS PQDT006, Max6L040J2 36GB IDE hard disk, all drivers from standard OS install
- 4 AMD Opteron processors Model 846 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition, BIOS PQDT006, Max6L040J2 36GB IDE hard disk, all drivers from standard OS install
- 1 Intel P4 processor 3.0GHz with 512KB L2 cache in Intel 875PBZ motherboard, Microsoft Windows Server 2003 Enterprise Edition, BIOS PQDT006, Max6L040J2 36GB IDE hard disk, all drivers from standard OS install
- 4 Intel Xeon MP processors 2.0GHz with 2MB L3 cache in Dell PowerEdge 6600 server, Microsoft Windows Server 2003 Enterprise Edition, BIOS A08, Ultra 3 SCSI ST 3367522C 36GB hard disk, all drivers from standard OS install

### Memory Latency

ScienceMark 2.0 Beta MemBench

- 1 AMD Opteron™ processor Model 146 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft® Windows® Server 2003 Enterprise Edition, BIOS PQDT006, Max6L040J2 36GB IDE hard disk, all drivers from standard OS install
- 2 AMD Opteron processors Model 246 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition, BIOS PQDT006, Max6L040J2 36GB IDE hard disk, all drivers from standard OS install
- 4 AMD Opteron processors Model 846 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition, BIOS PQDT006, Max6L040J2 36GB IDE hard disk, all drivers from standard OS install

- 1 Intel P4 processor 3.0GHz with 512KB L2 cache in Intel 875PBZ motherboard, Microsoft Windows Server 2003 Enterprise Edition, BIOS PQDT006, Max6L040J2 36GB IDE hard disk, all drivers from standard OS install
- 4 Intel Xeon MP processors 2.0GHz with 2MB L3 cache in Dell PowerEdge 6600 server, Microsoft Windows Server 2003 Enterprise Edition, BIOS A08, Ultra 3 SCSI ST 3367522C 36GB hard disk, all drivers from standard OS install

## SPEC<sup>®</sup> cpu2000

### SPECint Uniprocessor Benchmarks

- Estimated AMD Opteron processor Model 146 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition
- Intel Pentium 4 3.2GHz with 512K cache and 800 FSB in Dell Precision Workstation 360, Windows XP Professional SP1  
(<http://www.spec.org/osg/cpu2000/results/res2003q3/cpu2000-20030616-02265.html>)

### SPECint 2P Benchmarks

- Estimated 2 AMD Opteron processors Model 246 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition
- 2 AMD Opteron processors Model 244 with 1MB L2 cache in Einux A4800 server, Microsoft Windows Server 2003 Enterprise Edition  
(<http://www.spec.org/osg/cpu2000/results/res2003q2/cpu2000-20030421-02118.html>)
- 2 AMD Opteron processors Model 242 with 1MB L2 cache in Einux A4800 server, Microsoft Windows Server 2003 Enterprise Edition  
(<http://www.spec.org/osg/cpu2000/results/res2003q2/cpu2000-20030421-02119.html>)
- 2 Intel Xeon 3.06GHz with 512KB L2 cache in Dell Precision WorkStation 650, Windows 2000 Server  
(<http://www.spec.org/osg/cpu2000/results/res2003q2/cpu2000-20030407-02056.html>)

## SPECint 4P Benchmarks

- Estimated 4 AMD Opteron processors Model 846 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server Microsoft Windows Server 2003 Enterprise Edition
- 4 AMD Opteron processors Model 844 with 1MB L2 cache in Einix A4800 server, Microsoft Windows Server 2003 Enterprise Edition.  
(<http://www.spec.org/osg/cpu2000/results/res2003q2/cpu2000-20030421-02115.html>)
- 4 AMD Opteron processors Model 842 with 1MB L2 cache in Einix A4800 server, Microsoft Windows Server 2003 Enterprise Edition.  
(<http://www.spec.org/osg/cpu2000/results/res2003q2/cpu2000-20030421-02123.html>)
- 4 Intel Xeon MP 2.0GHz with 2MB L3 cache in Dell PowerEdge 6650, Windows 2000 Advanced Server (SP2).  
(<http://www.spec.org/osg/cpu2000/results/res2002q4/cpu2000-20021021-01745.html>)

## SPECint 2P-4P Benchmarks

- Estimated 4 AMD Opteron processors Model 846 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition
- Estimated 2 AMD Opteron processors Model 246 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition
- 4 Intel Xeon MP 2.0GHz with 2MB L3 cache in Dell PowerEdge 6650, Windows 2000 Advanced Server (SP2).  
(<http://www.spec.org/osg/cpu2000/results/res2002q4/cpu2000-20021021-01745.html>)
- 2 Intel Xeon MP 2.0GHz with 2MB L3 cache in Dell PowerEdge 6650, Windows 2000 Advanced Server (SP2)  
(<http://www.spec.org/osg/cpu2000/results/res2002q4/cpu2000-20021021-01744.html>)

## SPECfp Uniprocessor Benchmarks

- Estimated AMD Opteron processor Model 146 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition
- Intel Pentium 4 3.2GHz with 512K cache and 800 FSB in Dell Precision Workstation 360, Windows XP Professional SP1  
(<http://www.spec.org/osg/cpu2000/results/res2003q3/cpu2000-20030616-02266.html>)

## SPECfp 2P Benchmarks

- Estimated 2 AMD Opteron processors Model 246 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition
- 2 AMD Opteron processors Model 244 with 1MB L2 cache in Einux A4800 server, Microsoft Windows Server 2003 Enterprise Edition.  
(<http://www.spec.org/osg/cpu2000/results/res2003q2/cpu2000-20030421-02117.html>)
- 2 AMD Opteron processors Model 242 with 1MB L2 cache in Einux A4800 server, Microsoft Windows Server 2003 Enterprise Edition  
(<http://www.spec.org/osg/cpu2000/results/res2003q2/cpu2000-20030421-02120.html>)
- 2 Intel Xeon 3.06GHz with 512KB L2 cache in Dell PowerEdge 2650, Windows 2000 Server  
(<http://www.spec.org/osg/cpu2000/results/res2003q2/cpu2000-20030404-02020.html>)

## SPECfp 4P Benchmarks

- Estimated 4 AMD Opteron processors Model 846 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition
- 4 AMD Opteron processors Model 844 with 1MB L2 cache in Einux A4800 server, Microsoft Windows Server 2003 Enterprise Edition.  
(<http://www.spec.org/osg/cpu2000/results/res2003q2/cpu2000-20030421-02114.html>)
- 4 AMD Opteron processors Model 842 with 1MB L2 cache in Einux A4800 server, Microsoft Windows Server 2003 Enterprise Edition.  
(<http://www.spec.org/osg/cpu2000/results/res2003q2/cpu2000-20030421-02122.html>)
- 4 Intel Xeon MP 2.0GHz with 2MB L2 cache in Dell PowerEdge 6650, Windows 2000 Advanced Server (SP2).  
(<http://www.spec.org/osg/cpu2000/results/res2002q4/cpu2000-20021021-01741.html>)

## SPECfp 2P-4P Benchmarks

- Estimated 4 AMD Opteron processors Model 846 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition
- Estimated 2 AMD Opteron processors Model 246 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, Microsoft Windows Server 2003 Enterprise Edition
- 4 Intel Xeon MP 2.0GHz with 2MB L2 cache in Dell PowerEdge 6650, Microsoft Windows 2000 Advanced Server (SP2)  
(<http://www.spec.org/osg/cpu2000/results/res2002q4/cpu2000-20021021-01741.html>)
- 2 Intel Xeon MP 2.0GHz with 2MB L2 cache in Dell PowerEdge 6650, Microsoft Windows 2000 Advanced Server (SP2)  
(<http://www.spec.org/osg/cpu2000/results/res2002q4/cpu2000-20021021-01740.html>)

Competitive numbers shown reflect results published on [www.spec.org](http://www.spec.org) as of July 15, 2003. For the latest SPEC results visit <http://www.spec.org>.

## Linpack

### GOTO Library Benchmarks

- AMD Opteron processor Model 144 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, 64-bit SuSe 8.1 Linux Professional Edition with NUMA kernel and Myrinet MPIch-gm-1.2.5..10 message passing library.
- 2 AMD Opteron processors Model 244 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, 64-bit SuSe 8.1 Linux Professional Edition with NUMA kernel and Myrinet MPIch-gm-1.2.5..10 message passing library.
- 4 AMD Opteron processors Model 844 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, 64-bit SuSe 8.1 Linux Professional Edition with NUMA kernel and Myrinet MPIch-gm-1.2.5..10 message passing library.

### Atlas Library benchmarks

- AMD Opteron processor Model 144 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, 64-bit SuSe 8.1 Linux Professional Edition with NUMA kernel and Myrinet MPIch-gm-1.2.5..10 message passing library.
- 2 AMD Opteron processors Model 244 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, 64-bit SuSe 8.1 Linux Professional Edition with NUMA kernel and Myrinet MPIch-gm-1.2.5..10 message passing library.
- 4 AMD Opteron processors Model 844 with 1MB L2 cache in M&A Technology Patriot 64 Model 4400 server, 64-bit SuSe 8.1 Linux Professional Edition with NUMA kernel and Myrinet MPIch-gm-1.2.5..10 message passing library.

## AMD Overview

Founded in 1969 and based in Sunnyvale, California, AMD (NYSE: 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, Japan, and Asia. AMD, a Standard & Poor's 500 company, produces microprocessors, Flash memory devices, and silicon-based solutions for communications and networking applications.

© 2003 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow logo, AMD Opteron, and combinations thereof, AMD-8111 and AMD-8131 are trademarks of Advanced Micro Devices, Inc. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States and/or other jurisdictions. Pentium is a registered trademark of Intel Corporation in the United States and/or other jurisdictions. SPEC and the benchmark name SPECint are registered trademarks of the Standard Performance Evaluation Corp. Other product and company names in this paper are for identification purposes.