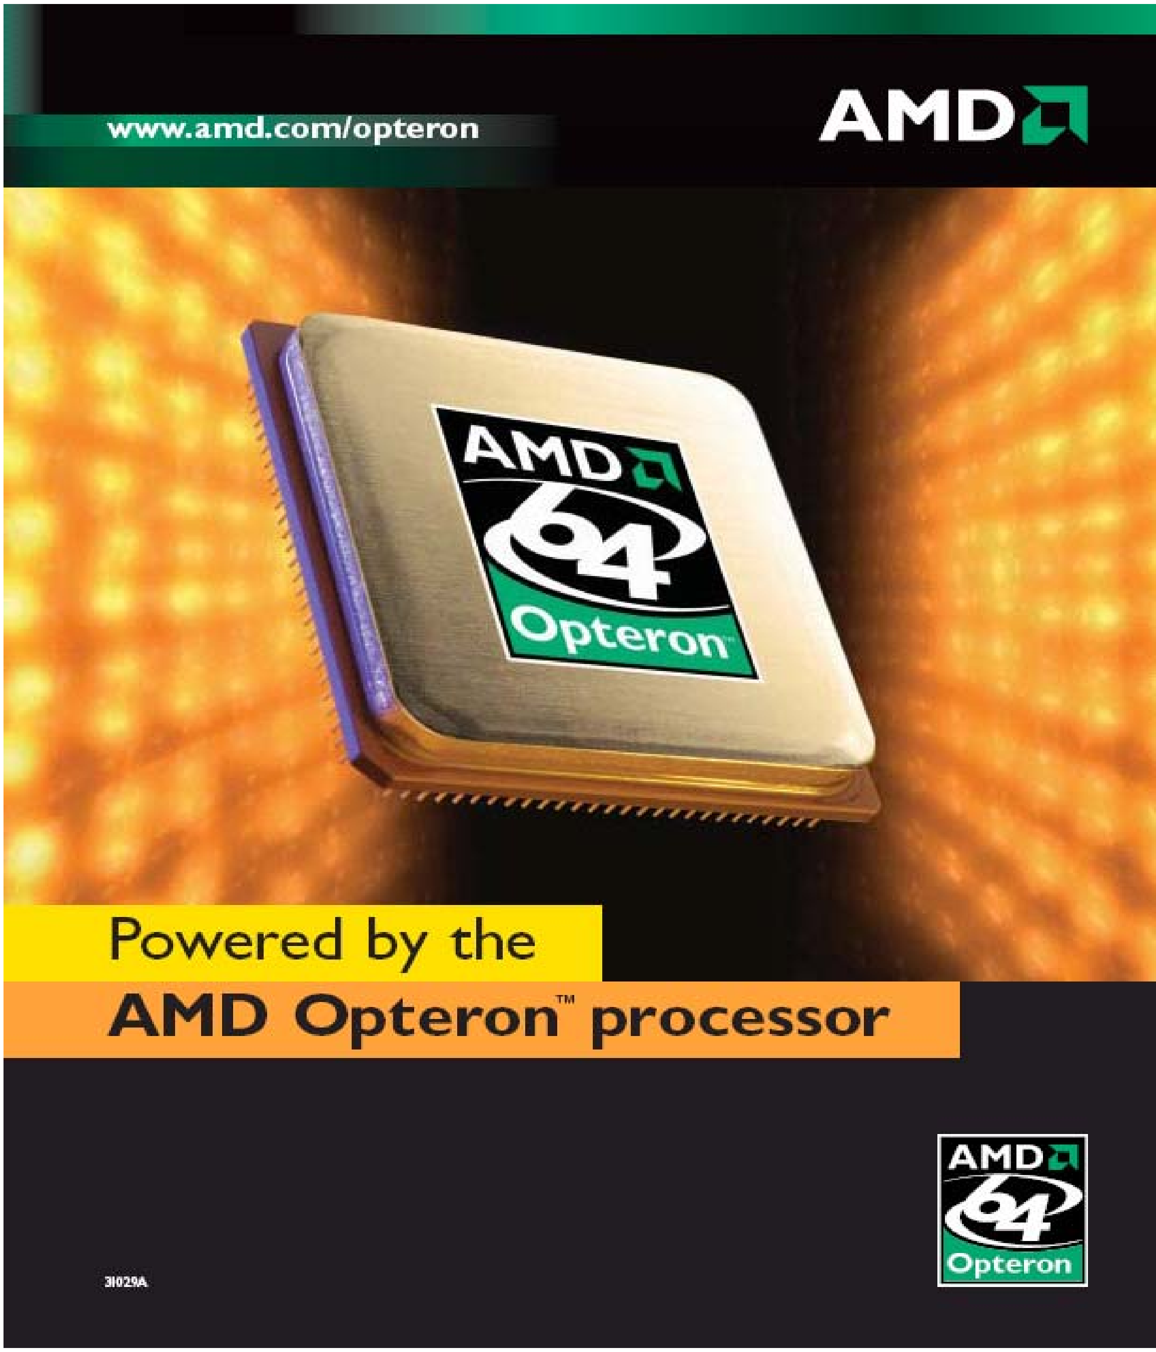

AMD on AMD

Enabling High-Performance Computing

Written By

Dean Marvin, Director, CAD Engineering, AMD

Ramesh Joginpalli, Strategic Marketing Manager, AMD

Quentin Fennessy, Senior Member Technical Staff, AMD

AMD on AMD

AMD continues to produce leading-edge microprocessors that provide customers with compatible and compelling solutions. Customers can rely on AMD for mission-critical business applications or high performance computing (HPC) solutions such as Engineering Design Automation (EDA). This paper will focus on the EDA methodology that was used to design AMD Athlon™ 64 and AMD Opteron™ processors. These processors were designed using the open source operating system Linux, EDA software tools from leading ISVs for design, synthesis, layout and verification, and systems based on AMD processors to develop new AMD processors – "AMD on AMD."

AMD initially developed its own specific software EDA tools to design AMD Athlon processors. Over time, commercial EDA tools have been developed specifically for systems based on AMD Athlon processors, and more recently AMD Athlon 64 and AMD Opteron processors that are based on AMD64 technology. The best tools run on AMD processor-based systems, which AMD uses to develop the next-generation leading-edge processors.

Introduction

The ultimate goal in EDA is to develop, test, revise and complete a processor design as quickly as possible in order to achieve the shortest time-to-market. AMD is an innovator in the EDA market by developing the best price/performance systems for chip development. Because of the work AMD is doing with EDA, other companies can benefit directly from AMD's own experience.

Today's systems using AMD Opteron processors and running 64-bit Linux and EDA tools from leading software vendors, provide a very effective solution. This combination resolves the challenges of managing the design complexity of shrinking device geometries, large datasets, and verification and regression simulations to produce the best possible product as quickly as possible. These simulations require a balanced 64-bit architecture that extends beyond the 4GB memory limit and maintains 32-bit compatibility. Only AMD64 processors provide the best of both scale-up and scale-out approaches to meet high-performance scientific computing needs while preserving users' investments in 32-bit legacy software.

Accelerating the adoption of 64-bit processing

Engineering design automation tools are a perfect example of applications whose requirements extend beyond the 4GB memory limit that the 32-bit x86 architecture imposes. EDA applications tend to have large databases, large design files and large simulation files that all need to extend beyond 4GB. These large datasets can be forced to work within 32-bit limitations at the expense of performance, depending on virtual memory and relatively slow disk I/O. The AMD Opteron processor architecture diagram shows the key architectural features that make it a perfect computing engine for meeting advanced high-performance computing needs while preserving users' investments in 32-bit legacy software.

AMD responded to customer needs with AMD64 technology, which is a new standard to replace the legacy 32-bit x86 architecture. AMD64 technology allows the AMD Opteron processor to run both 32- and 64-bit software as *native* applications. OEMs and other solution providers are able to offer advanced industry-standard server solutions that fit their customer needs built around the 32- and 64-bit AMD Opteron processor.

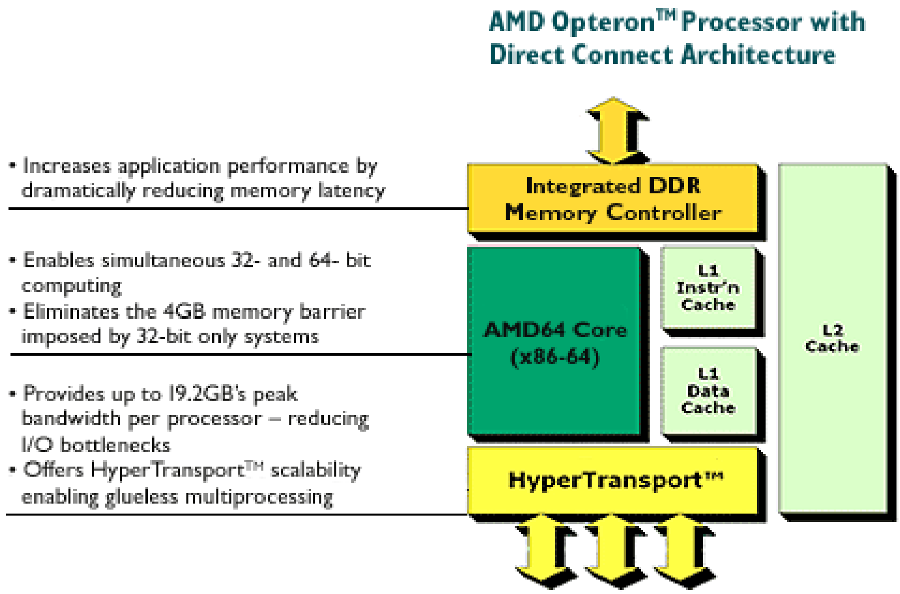

Workstations and servers are also able to take full advantage of the AMD Opteron processor with Direct Connect Architecture. Direct Connect Architecture directly connects the memory controller and I/O to the processor through HyperTransport™ technology, eliminating the bottlenecks and inefficiencies associated with a front side bus. Customer-centric innovation such as the AMD Opteron processor with Direct Connect Architecture is what enables AMD to offer exceptional price/performance for a broad spectrum of high-performance applications such as EDA.

AMD Chronology

UNIX to Linux Transition: AMD first used Linux to test AMD-K6® processor-based systems in 1996. We started the transition to Linux for microprocessor design and verification in 1997. The transition to Linux was rate-limited by the fact that the Linux operating system was an emerging but promising technology at that time, which resulted in the lack of available EDA software for Linux and the lack of industry-standard 64-bit hardware. AMD led the way with commercial Linux applications. AMD worked with Platform Computing to port their Platform LSF (Load Sharing Facility) cluster software to Linux long before the market moved to Linux.

In 1996 when AMD was designing the AMD-K6 processor, virtually all commercial EDA tools were running on proprietary UNIX workstations/servers resulting in a very high cost per seat and therefore a very expensive design process for AMD-K6. Though price was an issue, performance remained a primary consideration because certain tasks needed the highest performance.

Compute-intensive tasks such as:
- Functional verification
- Physical verification/mask generation
- Physical design/routing
- Extraction/timing analysis

Memory-intensive tasks such as:
- Physical verification/mask generation

Functional verification on AMD Athlon™ processors led the way: Generally, simulation is a problem bound only by time and the number of resources one has available. The quality of results is largely a function of how much simulation can be performed. The x86 architecture is complex and notoriously difficult to verify as evidenced by the AMD Athlon processor design where more CPU power was needed than was available using the traditional simulation model.

The solution was for AMD to develop its own simulator so that software cost per seat became almost zero. AMD ran the functional verification on AMD-K6 processor-based PCs so that hardware cost per seat was greatly reduced with a dramatic increase in available simulation cycles.

Recent History – AMD Opteron™ Processor Design: AMD used the AMD Athlon processor-based Linux systems for many parts of the AMD Athlon 64 and AMD Opteron processor design project (those that would fit into 32-bit address space). The other parts of the processor design required a 64-bit environment. The best available UNIX hardware was used for the following design steps that required 64-bit processing:

- Full-chip physical verification
- Clock extraction and skew analysis
- Mask generation

Future Processor Design - Real Time: AMD Opteron processor-based platforms are now available from both Tier 1 vendors and system builders along with mature, stable and scalable enterprise versions of AMD64 Linux. AMD uses many commercial EDA applications on several Linux distributions including RedHat and SUSE (now part of Novell). Most of the EDA vendors now support their memory intensive applications on AMD Opteron processors. A substantial number of EDA applications are ported to the AMD64 architecture every week. The future processor design effort uses both 64- and 32-bit tools and 32-bit Linux applications continue to work well in 32- and 64-bit mode on AMD Opteron processors. This is the one of the greatest benefits AMD64 technology has brought to the HPC and EDA markets.

Currently, all AMD internal EDA tool development is done on Linux for AMD64, eliminating the need for specialized UNIX hardware. The benefit to AMD is tremendous - the entire chip design process is on a homogeneous platform with the future processor design completely done and taped out on AMD CPUs.

AMD Opteron Processor-based Clusters

AMD has created a large and ever-increasing set of directed regression tests and random test procedures to verify AMD64 technology for each step of its design process. These compute-intensive tests number in the thousands. In addition, the final step of mask generation is also compute intensive, has very large memory requirements and is the *last step* in the silicon design process. It occurs when there is no longer any additional flexibility in the design schedule

In order to meet the challenge of successful and timely design and manufacturing, AMD has invested in several very large compute clusters. AMD has successfully run most of the necessary compute cycles on AMD Athlon processor-based systems over the last few years. In the time since AMD Opteron processor-based systems became commercially available, AMD engineering has invested in AMD Opteron processor-based systems, both dual- and quad-processor configurations. Over the last few years the percentage of the compute power provided by AMD processors has grown from 30% in 2000 to over 90% in 2004. This huge increase over the past year has been provided by AMD Opteron processor-based systems. The remaining 10% of non-AMD processor-based compute capability are the legacy systems that are no longer in the critical path for future processor design.

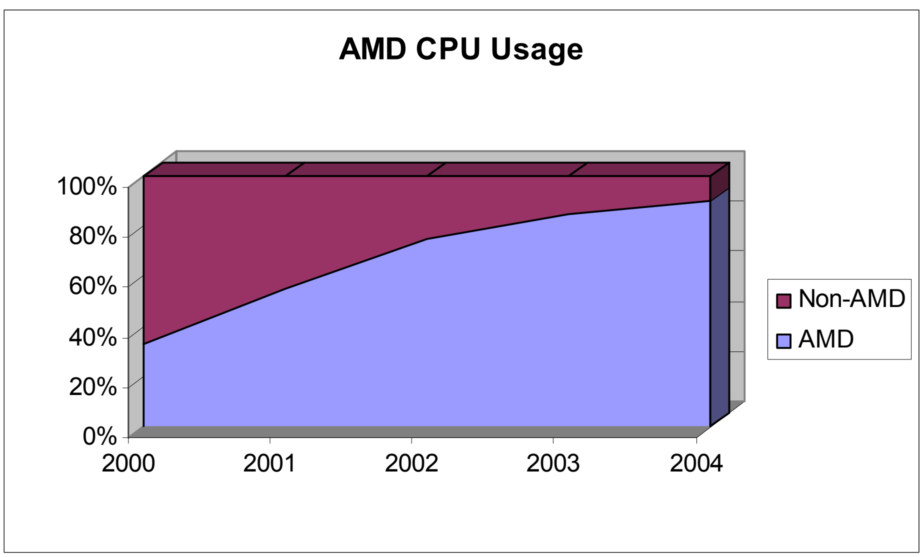

AMD Opteron processor-based systems provide engineering with an unbeatable combination of high performance with commodity system price and operational cost. AMD Opteron processor-based systems have an architectural advantage with Direct

Connect Architecture, and AMD engineers would not be able to duplicate the individual system performance using any competitive processor currently on the market. Although compute needs are growing quickly, and are fueled by the added complexity of designs, the ongoing refinement of the design process and the number of concurrent products in the pipeline, the unbeatable performance of AMD Opteron processor-based systems has allowed AMD to increase capacity in a timely fashion.

AMD Opteron™ Processor-based Platform LSF Clusters: AMD compute clusters are managed by Platform LSF. One significant challenge in managing large pools of compute servers is the scalability. Can a Platform LSF-based cluster manager system be provisioned to handle tens of thousands of jobs simultaneously? High performance in cluster management is required for low latency job dispatch and support of sophisticated resource allocation policy. If the cluster manager systems are too slow then users might have to wait 10 minutes or more for a job to start executing. This delay reduces the effectiveness of design engineers and increases their frustration level significantly. Historically, AMD has partitioned its compute clusters to reduce the compute needs for the cluster manager systems. AMD Opteron processor-based systems (dual- and quad-processor) have been proven in production as capable of running world-class compute clusters. AMD's Platform LSF-based cluster manager systems are commodity-based systems, using off-the-shelf hardware and network connectivity. The cluster manager system for AMD's largest cluster could be duplicated for less than $10,000 with an AMD Opteron processor-based system from a Tier I vendor.

AMD was able to migrate very smoothly from 32-bit Linux systems using AMD Athlon processors to 64-bit Linux systems using AMD Opteron processors. It has historically been a challenge for large design environments to convert from one system vendor to another, or from one operating system to another. The changes required for these conversions range from replacing physical infrastructure to completely revamping system operations to a complete re-education of end users. The conversion from AMD Athlon processor-based systems to AMD Opteron processor-based systems at AMD was a much easier task.

The system support structure for hardware has remained essentially unchanged. The AMD Opteron dual-processor systems fit in the same racks and use the same power, network and cooling facilities that have been used for years with AMD Athlon processor-based systems. Coincidental to the AMD Opteron processor installation, the network infrastructure was upgraded to Gigabit Ethernet. This conversion was also seamless due to the existing high-quality network cables and patch panels.

AMD supports Linux and UNIX systems with a corporate toolbox of freely available software (such as the GNU C compiler gcc and the scripting language Perl). Integrating the first AMD Opteron processor-based systems into production, AMD was able to re-use over **99%** of the existing software installed for 32-bit Linux. AMD's current practice for non-performance critical software is to use a 32-bit version of the software for both AMD Athlon and AMD Opteron processor-based systems. This practice reduces the incremental software support cost of our AMD Opteron processor to almost zero. There are a few programs that require recompilation for the highest performance of the Direct Connect Architecture or for access to the 64-bit address space or for access to the 64-bit OS interface. These were easily recompiled and put into production.

Conclusion

The simplicity of AMD's conversion from AMD Athlon processor-based systems to AMD Opteron processor-based systems demonstrates that the first and subsequent steps of a design environment migration can be smooth and quick. If you are already supporting 32-bit Linux you are ready to support 64-bit Linux on AMD64 systems.

AMD has seen outstanding performance increases with the migration to AMD Opteron processor-based systems. The typical pattern is that a 32-bit application when run on an AMD Opteron processor in 32-bit Linux runs up to 2.5x faster, comparing the fastest 2P AMD Athlon MP servers with AMD Opteron processor Model 246 dual systems. AMD engineers have noticed another 5-8% increase when the same 32-bit application was run on the same AMD Opteron processor Model 246 dual system—running 64-bit Linux. Preliminary results suggest that a performance gain of about 5% may be expected for the recompiled application.. This performance increase for each step of modernization made the AMD Opteron processor conversion a huge success.

###