
AMD “Bulldozer” Core Technology

INTRODUCTION

When AMD introduced the first AMD Opteron™ processor in 2003, it featured an innovative design that became the standard for the x86 server industry, with an integrated memory controller, scalable core architecture and innovative power efficiency features. For the following 7 years, AMD has improved on this design, increasing from a single core, all the way up to 12 cores. Multiple dies now occupy the high end processors and the performance has scaled significantly higher than what was originally introduced in 2003 – yet it has essentially remained in the same power and thermal envelope, with few socket/platform changes to help maintain the long-term stability of the servers.

BENEFITS FOR CUSTOMERS

In 2011, AMD plans to introduce the next generation of AMD Opteron processors, based on a new core architecture that is designed specifically to enhance the scalability of enterprise applications. Sixteen-core “Bulldozer”-based processors will deliver up to 50% greater throughput¹ than existing 12-core AMD Opteron processors, while maintaining the same power and thermal ranges. Virtualization will be greatly enhanced with new features that will deliver greater scalability and manageability. Database applications will enjoy more simultaneous threads than ever before with up to 33% more cores², and HPC customers will find they can reach a high density of cores with far fewer nodes than ever before, allowing for the maximization of data center space. Finally, cloud customers will find that the high core count and massive memory scalability are an ideal complement for highly scalable public or private cloud deployments.

THE PROCESS TECHNOLOGY

This new processor features the latest technology from AMD’s foundry partner, GLOBALFOUNDRIES. Based on a 32nm high-K metal gate design, the processor features a die size for an 8-core die that is slightly smaller than AMD’s current 6-core server dies. This means 33% more cores in essentially the same silicon budget as before. GLOBALFOUNDRIES is utilizing a gate-first approach in the fabrication of these products. The gate-first approach requires less-restrictive layout design rules, provides for a smaller die size, and provides other manufacturing benefits, while delivering on the power and performance needs of customers.

THE “BULLDOZER” DESIGN PHILOSOPHY

The design goal for the “Bulldozer”-based CPUs was simple – drive greater scalability and higher density of cores through a design methodology that shares some components of the processor die where die space can be maximized and yet keeps other parts discrete to avoid bottlenecks. Many of today’s modern processor designs include redundant components because of their multi-core design, and in many cases these redundancies merely consume more die space, which can increase cost and increase power consumption – without necessarily adding incremental value to the processor.



By moving to a new design, where two integer cores share a common front end (fetch and decode) along with a large shared L2 cache and an FPU complex, AMD has created a modular processor design that allows for a “building block” approach. These modules allow AMD to have a flexible design that serves as the blueprint for the upcoming 16-core AMD Opteron 6200 Series processor (code named “Interlagos”) as well as the 8-core AMD Opteron 4200 Series processor (code named “Valencia”). Because of the performance and scalability, “Bulldozer”-based designs are also scheduled to appear in the high end of the desktop processor market for AMD. AMD can use this modular approach to design the next generation of 20-core and 10-core products that are expected to reach the market in 2012.

HEAVY THREADING

Each module will feature two integer cores which are full cores, not half cores. Each core can execute a single thread, so each module can handle two threads per cycle across each cores’ dedicated integer pipelines. Previous AMD designs could also handle one thread per core and had only 3 pipelines, which shared ALU and AGU functions. Each “Bulldozer” integer core has four pipelines, 2 for ALU and 2 for AGU, allowing for greater throughput for processing. The platform, the operating system and the applications will not see modules (these are mainly designations for the designers) but will instead see only a pool of integer cores that can be utilized for processing.

THE FLEX FP

Floating point operations are essentially about 10% of the processing in a typical server workload, with the other 90% being integer. The challenge with advancements in floating point technology is that they require increasingly greater amounts of silicon and power to solve problems that most workloads are not taking advantage of today. The Flex FP is designed to help address this inequity, by bringing next generation FP capabilities in the way of support for 256-bit AVX instructions, yet still help minimize the die space and power impact for those that do not need AVX. By putting a Flex FP complex in each module, the “Bulldozer” design allows for the capability to handle a 256-bit AVX execution with either thread being able to dispatch. Since most applications will probably only need the 128-bit FP, the Flex FP can also function as a pair of FP units. The Flex FP can thus operate in one of 2 modes – as a single 256-bit AVX unit shared between the two cores or as two 128-bit units, with one dedicated to each core.

In addition, the Flex FP features FMAC units that can handle an FMA (Fused Multiply Accumulate) instruction as well as the standard FMUL (multiply) and FADD (add). An FMA operation is far more powerful because it allows a calculation like $A = B \times C + D$. In a standard FP unit, this would require 2 cycles, one for the multiply and one for the add instruction. In addition, the FMAC can handle both an FMUL and FADD, whereas the traditional designs have dedicated paths for those operations. Having to handle multiple FADDs or FMULs can be optimized by allowing the work to be split across both FMACs instead of having all of the operations stack up behind one, waiting for execution. With “Bulldozer”, AMD is utilizing the more powerful FMA4 operation which has four operands.

AMD TURBO CORE Technology

A new feature being added to the “Bulldozer” architecture is AMD Turbo CORE, a technology that allows processors to run above their base clock frequency provided there is additional power headroom available. When processors are rated for clock speed, it is typically done based on a worst case scenario for workloads, resulting in a speed rating that is very conservative for the vast number of workloads that will be run, potentially leading to a lost opportunity in clock speed for less strenuous workloads. This is one reason why processors generally consume far less power than their rated TDP. AMD Turbo CORE technology captures this additional power headroom and turns it into higher clock frequency, allowing the processor to run up to 500MHz higher than the rated base frequency with all cores active and utilized. In environments where all cores are not being utilized the frequency increase can be even greater than 500MHz.

AMD Turbo CORE technology is designed to monitor the workload as well and will not automatically drive to the highest clock frequency unless the workload demands it. So, just like AMD PowerNow!™ technology that can modulate the core frequency — driving it down in periods of low utilization – AMD Turbo CORE does the same, but in the opposite direction with the goal of squeezing extra clock speed out of the processor when loads are reaching their peak. AMD Turbo CORE works in conjunction with AMD CoolSpeed™ technology to help ensure that even when boosting frequencies, processors stay within safe operating ranges. Obviously, when increasing clock frequency, power consumption can go up, so customers must carefully consider their need for raw performance in conjunction with their power needs.

IMPROVED MEMORY CONTROLLER

When designing this new processor, AMD took the opportunity to redesign the memory controller, increasing the throughput by up to 50% over existing AMD Opteron processors³. Having an integrated memory controller for over seven years has helped AMD to learn much about the optimization of the memory controller and how to drive better efficiency for enterprise applications. Part of that throughput gain comes from optimizations to the memory controller – new algorithms and new ways to address reading and writing data that speed up access to information. In addition, AMD will support DDR3 1600MHz memory. Between the enhancements in memory controller circuitry and the faster clock speed memory we anticipate significantly higher throughput for memory-intensive environments like virtualization, HPC, database and business applications.

Two new memory standards that will be supported as well are LR DIMM (load reduced DIMM) which should allow for greater capacities of memory to be installed in the server as well as the emerging 1.25V low voltage DIMMs. Today's LV DIMMs are 1.35V so the new 1.25V should lead to even more power efficiency for customers.

CACHE STRUCTURE

Each module includes two levels of cache, an L1 cache that is focused on execution and an L2 cache that is focused on being the “working area” for data being processed. There are two components in the L1 cache, instruction and data. The L1 instruction cache is 64K and it is shared by the two integer cores as there is a high probability that instructions needed by one core will also be needed by the other, increasing the efficiency by allowing that instruction to only have to reside in one place in the cache. Each core also has its own L1 data cache that is designed to hold the data tied directly to processing those instructions for the cycle. Beyond the L1 cache, each module has its own L2 cache that is shared by the integer 2 cores. This is a much larger cache; it is 2MB. If only one core is active and storing data, it can have access to all 2MB of the L2 cache. Outside of the modules, at the die level, an 8MB L3 cache is employed. For AMD Opteron 6200 Series processors (code named “Interlagos”), there are two dies present inside of the processor, which means a total of 16MB of L3 cache. AMD Opteron 4200 Series processors (code named “Valencia”) will have only one die for a total of 8MB of L3 cache.

DEDICATED SCHEDULERS

Inside of each module are three dedicated schedulers, one for each integer core and one to feed the Flex FP. The integer schedulers are 40-entry and the Flex FP scheduler is 60-entry. By having a dedicated scheduler for each integer core, the “Bulldozer” architecture helps ensure that the 4 integer pipelines are being kept continually filled with instructions for maximum efficiency. Each integer core has control over its own scheduling so that there is virtually no bottleneck between the two dedicated threads that are executing in the module simultaneously. The Flex FP scheduler is a single entity because in AVX mode it needs to be able to send a single stream of 256-bit AVX operations through the FP pipes as well as handle dual 128-bit executions. In 128-bit mode, the extra entries in the Flex FP scheduler help ensure that the two 128-bit FMACs are receiving a constant stream of math instructions to execute. The key for great FP throughput is having a deeper scheduler, allowing for more out-of-order execution – this helps keep the FP pipelines full. FP code is more latency sensitive than integer code, so by being dedicated “Bulldozer” can give the full scheduler depth to the FPU for greater amounts of re-ordering. The scheduler depth on the “Bulldozer” FP scheduler is deeper than today's AMD Opteron processors, allowing for maximum performance potential.

VIRTUALIZATION FEATURES

With more and more environments moving from dedicated servers to virtualization, integrated virtualization features are becoming an essential part of any processor design. AMD continues the tradition of the AMDVirtualization™ (AMD-V™) technology, an integrated set of virtualization features that help to boost performance and efficiency. Bit-level enhancements like Flush by ASID, Larger ASID space, and VMCB clean bits, along with a host of other hypervisor-level enhancements have been made to beef up the AMD-V support. In addition, Virtual Cache Partitioning will be available, which allows a portion of the L3 cache to be dedicated to a single module (thus creating a self-contained “system in a system” for better manageability).

But the real star of “Bulldozer” in a virtualized environment is the greater number of cores and greater memory scalability as core density and memory addressability are key performance drivers for a virtualized platform. For those customers that utilize the “1 VM per core” methodology for deployment, “Bulldozer”-based platforms will allow them to deploy from 60-100% more VMs than on a comparable Intel-based platform⁴. And of course all of AMD's work with leading virtualization providers like VMware, Microsoft, Xen, Citrix, Parallels and the Linux KVM community help to ensure a highly supportable virtualization solution.

HPC ENVIRONMENTS

Customers that utilize high performance compute clusters in their business will find that the new “Bulldozer” – based platforms are a superb complement to their current cluster strategies. With 16 cores per processor and the ability to run 64 single precision operations (or 32 double precision operations) per processor per cycle, “Bulldozer” is designed to take on technical computing workloads. The memory controller enhancements will allow for high memory bandwidth³ and combined with 192GB/socket memory addressability; these platforms can handle significantly large memory footprints for large problems and map/reduce activities. In addition, some of the new instruction sets mentioned in this paper will allow those that are tuning code specifically for their applications to really harness the power of the “Bulldozer” architecture.

POWER EFFICIENCY FEATURES

AMD’s continued focus on power efficiency brings new power enhancements to the “Bulldozer” architecture. At the core level the new C6 power state is designed to power down a complete module when it has been in idle for a pre-determined amount of time. I/O C-stages and shared logic combine with TDP Capping and Dynamic Power Capping to help drive down core power consumption. The enhanced memory controller now features C-states (power states) that allow it to fluctuate in power consumption just like the P-states for the processor clocks.

All of these power features are wrapped up in 32nm Silicon on Insulator architecture that has been tuned for power efficiency. These features help modulate power in such a way that, even though the new design features 33% more cores, higher core frequencies, 20% higher memory clock speed and the inclusion of AMD Turbo CORE boosting technology, the processors still operate in the same TDP ranges as current AMD Opteron 4100 and 6100 Series processors.

NEW SOFTWARE INSTRUCTIONS

AMD plans to support a series of new instructions and operations with the “Bulldozer” core that are designed to help drive greater performance and compatibility for high powered applications:

INSTRUCTION	FUNCTION/BENEFIT
SSSE3	For media handling like video encoding and transcoding.
SSE4.1	Covers a range of applications with several new packed data operations including various specialized data movement instructions; dot product, min/max, compare and rounding operations for numeric processing, and further video encoding support.
SSE4.2	POPCNT (Population Count) instruction is applicable to bioinformatics algorithms; CRC32 is for accelerating CRC-based integrity checking of network or disk data transfers; string instructions are applicable to any text-intensive applications such as XML and HTML parsing.
AES, PCLMULQDQ	Accelerates any application that uses AES encryption, key uses being secure network transactions (internet and LANs), disk encryption such as Microsoft’s BitLocker, and database encryption.
AVX	Provides 256-bit floating point as well as a significant performance boost for vector floating point applications, and a lesser boost for multimedia apps and non-vector FP-intensive apps.
FMA4	A further boost for numeric applications, particularly HPC-type applications, providing up to a 2x increase beyond AVX on AMD processors.
XOP	Provides enhancements for pack data operations, applicable to a variety of numeric and multimedia applications, including DSP-type applications.

PLATFORM COMPATIBILITY

When AMD designed the AMD Opteron 4000 & 6000 Series platforms, we specifically looked at the needs of the next generation “Bulldozer” to help ensure that once this next generation processor is launched it could easily integrate into these platforms. AMD anticipates that most, if not all, servers will have only a BIOS update in order to take advantage of this new processor. As few customers actually ever update their processors once the systems have been deployed, the real benefit is not from the socket compatibility, but rather from the fact that current generation and next generation platforms can live side by side, being managed the same way with the same BIOS, drivers and interfaces. Migrating software and virtual machines should be a snap as the platform level components, like chipsets and peripherals are the same, making the process simple.

PRODUCT TIMING

“Bulldozer”-based products began sampling to our OEM partners in Q4 2010. AMD plans to have end user sampling in late Q1 or early Q2 of 2011. AMD plans to be in production of “Bulldozer”-based server products before the end of Q2 2011 with widespread availability and the actual product launch planned for Q3 2011. Once production commences customers will have the ability to purchase processors prior to launch, based on opportunity, through AMD’s OEM partners. For more information on availability and evaluation units, please contact your AMD commercial account executive.

SUMMARY

AMD’s next generation processors, based on the “Bulldozer” core, present a new way of addressing the need for high performing, highly scalable processors. The innovative modular architecture helps pack in more cores, making it the natural choice for enterprise applications that demand more cores and more memory. And as an added benefit, they’re designed to save costs through efficient design and power efficiency, something that most products in the data center don’t focus on nearly enough.

¹ Based on AMD internal performance estimates of top-bin “Interlagos” processor vs. 12-core AMD Opteron 6176 SE processors on a variety of server-based workload measurements.

² Comparison of 12-core AMD Opteron 6100 Series processors vs. 16-core “Interlagos” processors.

³ Based on AMD engineering estimates; up to 30% increase from memory controller enhancements and up to 20% increase from higher clock speed memory.

⁴ Comparison of 16-core AMD Opteron 6200 Series processor with expected 8-core Intel Xeon “Sandy Bridge EP” processor and expected 10-core Intel Xeon “Westmere EX” product.

