The new AMD 6200 series CPU and its relevance to HPC

Binesh Majithia
Technical Pre-Sales
binesh.majithia@amd.com

Lee-Martin King
UK&I Business Development
lee-martin.king@amd.com

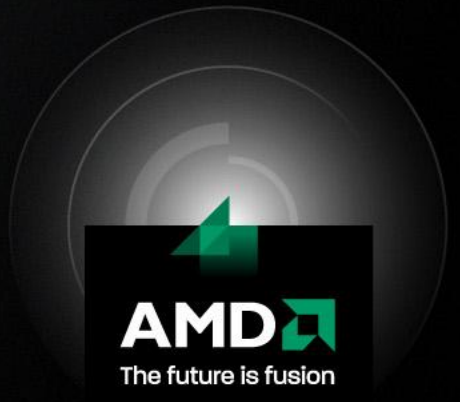

AMD in High-Performance Computing
CPUs and GPUs Drive Efficient Performance

TOP500.ORG list
CPUs
Oak Ridge "Jaguar"
DOE/NNSA/LANL/SNL – "Cielo"
The National Energy Research Scientific Computing Center (NERSC) – "Hopper"
GPUs
Nat'l Supercomputer Center, Tianjan "Tianhe-1"

Since 2005, AMD technology has powered more than 70 top-25 entries in the list of the world's most powerful supercomputers

Source: www.top500.org

AMD Server Platform Strategy

	Today	2011	2012
4P/8P Platforms ~5% of Market*	**Performance-per-watt and Expandability**	**AMD Opteron™** 6000 Series Platform	AMD Opteron 6100 Series processor 8 and 12 cores	AMD Opteron 6200 Series processor 4, 8, 12 and 16 cores	Future Product
		• 2/4 socket; 4 memory channels • For high core density			
2P Platforms ~75% of Market*	OPTERON AMD	"Hydra"	"Bulldozer"	Future Core
		SR5600 Series Chipsets		
	Highly Energy Efficient and Cost Optimized	**AMD Opteron™** 4000 Series Platform	AMD Opteron 4100 Series processor 4 and 6 cores	AMD Opteron 4200 Series processor 6 and 8 cores	Future Product
		• 1/2 socket; 2 memory channels • For low power per core			
1P Platforms ~20% of Market*	**Low cost for dedicated web hosting**	**AMD Opteron™** 3000 Series Platform	AMD Opteron 3200 Series processor 4 and 8 cores	Future Product
		• 1 socket; 2 memory channels • For low cost per core		

*AMD internal estimates of total server market as of Q3 2011

* Based on AMD estimates.

Bulldozer Module - Advanced Performance/Watt

Leadership Multi-Threaded Micro-Architecture

Full Performance From Each Core

- Dedicated execution units per core
- No shared execution units as with SMT

High Frequency / Low-Power Design

- Core Performance Boost
  - "Boosts" frequency of cores when available power allows
- Power efficiency enhancements
  - Deeper core sleep states

Virtualization Enhancements

- Faster switching between VMs
- AMD-V extended migration support

Shared Double-sized FPU

- Amortizes very powerful 256-bit unit across both cores

Improved IPC

- Micro-architecture and ISA enhancements
  - SSE4.1/4.2, AVX 1.0/1.1, SSSE3

Enhanced Systems Management

- Greater power management control via APML

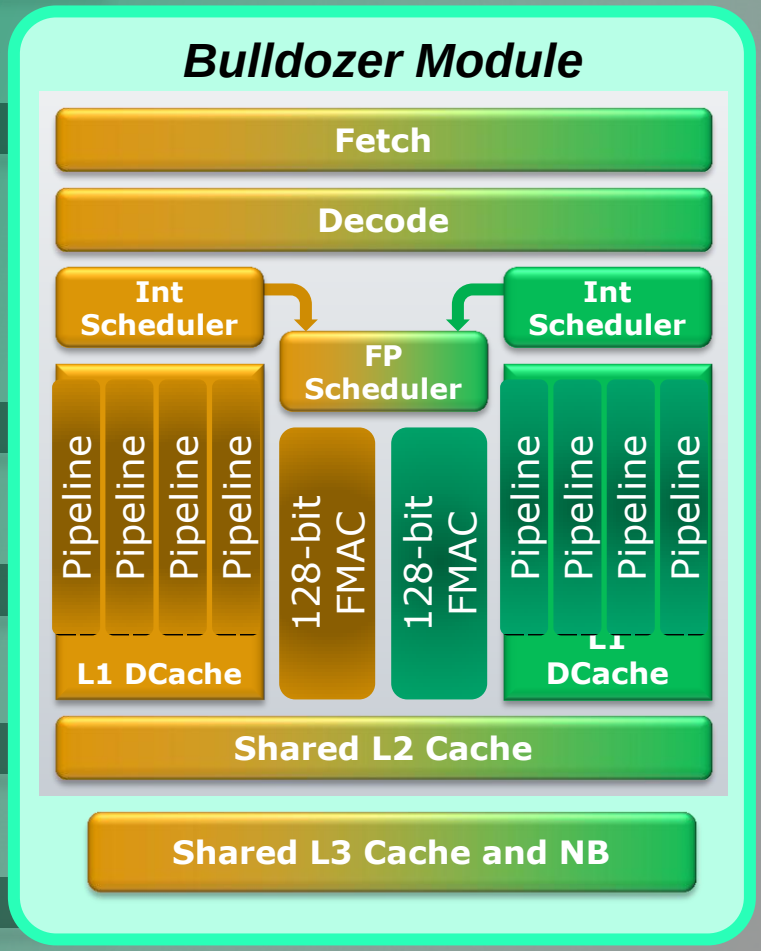

THE HPC LEADER

HPC

- Linux OS
- Open64
- GCC
- PGI Compilers

Customer Requirements:

- Scalable performance
- Strong floating point performance
- High memory throughput
- More cores for highly threaded apps
- Wide range of technical instructions

73GB/s memory throughput[3]

73% more memory bandwidth than Intel[3]

Maximum cores per rack[2]

More FLOPs per sq. foot[2]

33% lower cost per core[4]

Superior Performance[1]

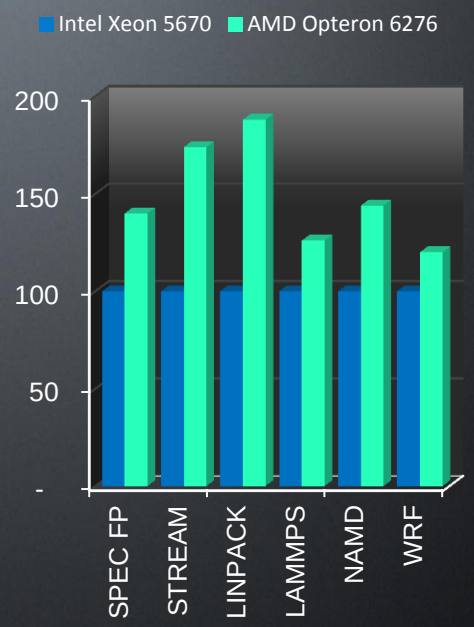

24-84% better performance at significantly lower price[1]

Greatest FLOPs per Sq. Foot

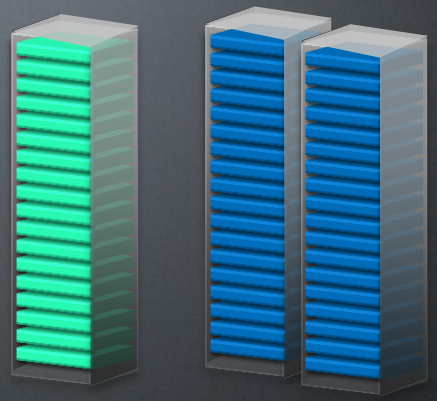

With almost twice the FLOPs per sq. ft. with AMD Opteron™ 6276 Series processors, it would take 2 racks of Intel Xeon 5670 racks to match AMD in density and performance[2]

[1-3] See complete benchmark data on slides 35-37.

Flex FP: More flexible technical processing

More performance and new instruction support

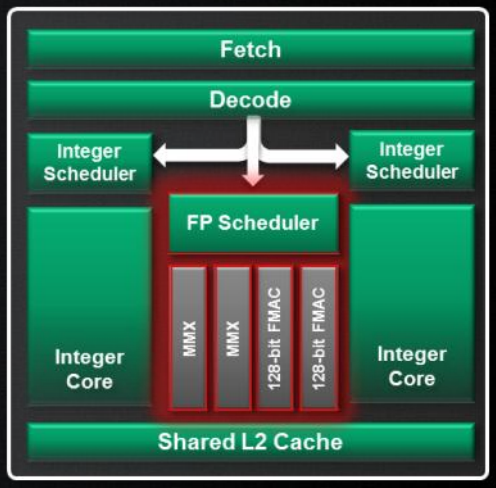

Runs SSE and AVX simultaneously	• No processing penalty for simultaneous execution
Executes two SSE or AVX (128-bit) instructions simultaneously or one AVX (256-bit) instruction per Bulldozer module	• Wider range of FP processing capabilities than competition
Processes calculations in a single cycle using FMA4* and XOP instructions	• Executes more instructions in fewer cycles than competition
Uses dedicated floating point scheduler	• No waiting on integer scheduler to run instructions • Designed to be always available to schedule floating point operations

*FMAC can execute an FMA4 execution (a=b+c*d) in one cycle vs. 2 cycles that would be required for FMA3 or standard SSE floating point calculation.

AMD OPTERON™ 4200 AND 6200 SERIES CPU
OS AND HYPERVISOR SUPPORT SUMMARY

ASSUMES latest updates/patches are installed*

Enabled
Optimized to support some or all of "Bulldozer's" new features

Includes new instruction support:
- Hyper-V Nex Gen (in development)
- Linux kernel 2.6.37 +
- Novell SLES 11 SP2 Beta (includes Xen)
- RHEL 6.2 with KVM (in development)
- Windows Server 2008 R2 SP1
- Windows 8 Server (in development)
- Xen 4.1
- Ubuntu 11.04 (w/ KVM)
- VMware vSphere 5.0

Versions in this category also include latest software advances

Compatible
Will boot and run but not take advantage of "Bulldozer's" new features outside of new instructions

Incudes new instruction support:
- Linux kernel 2.6.32 – 2.6.36
- Novell SLES 11 SP1
- RHEL 6.1
- Ubuntu 10.10

Does not support new instructions for either Bulldozer or Sandy Bridge:
- Hyper-V R1
- Hyper-V R2, Hyper-V R2 SP1
- Novell SLES 10 SP4 and higher
- RHEL 5.7 (included KVM)
- Solaris 10u9, 11
- VMware vSphere 4.1u2 (in development)
- Windows Server 2003 R2 SP2
- Windows Server 2008 R2
- Windows Server 2008 SP2
- Xen 3.4.2

Will run but not necessarily provide performance uplift

Not Supported
Will not run on "Bulldozer" platforms and/or will not be supported by OSV

- Linux kernel 2.6.31 or earlier
- Novell SLES 10 thru SP3
- Novell SLES 11
- RHEL 4.x
- RHEL 5.0 – 5.5
- RHEL 5.6 (can run with patches but is not supported by Red Hat)
- RHEL 6.0
- Solaris 10 – 10u8
- VMware ESX 3.5
- VMware ESX 4.0 – 4.1u1
- Windows Server 2003 versions prior to R2 SP2

* Please note: For proper support of available features/processors, the latest updates/patches always needs to be installed

AMD OPTERON™ 4200 AND 6200 SERIES PROCESSORS

Compiler	Status	SSSE3 SSE4.1-.2 AVX	FMA4 XOP	Auto Generates Code	Comments
GCC 4.6.1	Available	✓	✓	✓	GCC 4.4 is included in RHEL 6.0 distribution and should be updated to GCC 4.6.1 for optimized support
Microsoft Visual Studio 2010 SP1	Available	✓	✓	No	Supports new instructions but does not auto generate code
Open64 4.2.5	Available	✓	✓	✓	http://developer.amd.com/open64
Open64 4.5	Planned for Dec 2011	✓	✓	✓	Will provide incremental performance and functionality improvements
PGI 11.9	Available	✓	✓	✓	PGI Unified Binary™ technology combines into a single executable or object file code optimized for multiple AMD and Intel processors
ICC 12	Available	✓ (-mAVX flag)	No	✓	–mAVX is designed to run on any x86 processor, however the ICC runtime makes assumptions about cache line sizes and other parameters that causes code to run on AMD CPUs

Compiler Optimization Quick Guide: http://developer.amd.com/Assets/CompilerOptQuickRef-62004200.pdf

THE NEW "BULLDOZER" INSTRUCTIONS A CLOSER LOOK

Instructions	Applications/Use Cases
SSSE3, SSE4.1, SSE4.2 (AMD and Intel)	• Video encoding and transcoding • Biometrics algorithms • Text-intensive applications
AESNI PCLMULQDQ (AMD and Intel)	• Application using AES encryption • Secure network transactions • Disk encryption (MSFT BitLocker) • Database encryption • Cloud security
AVX (AMD and Intel)	Floating point intensive applications: • Signal processing / Seismic • Multimedia • Scientific simulations • Financial analytics • 3D modeling
FMA4 (AMD Unique)*	• Vector and matrix multiplications • Polynomial evaluations • Chemistry, physics, quantum mechanics and digital signal processing
XOP (AMD Unique)*	• Numeric applications • Multimedia applications • Algorithms used for audio/radio

XOP and FMA4 instruction set extensions are AMD unique 128-bit and 256-bit instructions designed to:

- Improve performance by increasing the work per instruction
- Reduce the need to copy and move around register operands
- Allow for some new cases of automatic vectorization by compilers

* http://blogs.amd.com/developer/2009/05/06/striking-a-balance/

XOP AND FMA4 | A CLOSER LOOK

FMA4 Overview (AMD Unique)

Performs fused multiply–add (FMA) operations. The FMA operation has the form d = a + b x c. FMA4 allows a, b, c and d to be four different registers, providing programming flexibility.

- A fast FMA can speed up computations which involve the accumulation of products
- FMA capabilities are also available in IBM Power, SPARC, and Itanium CPUs.
- Intel is anticipated to introduce FMA3, a more limited implementation of FMA (where d is the same register as either a, b, or c) to Xeon in 2013 timeframe*

XOP Overview (AMD Unique)

Provides three- and four-operand non-destructive destination encoding, an expansive new opcode space, and extension of SIMD floating point operations to 256 bits.

- Horizontal integer add/subtract
- Integer multiply/accumulate
- Shift/rotate with per-element counts
- Integer compare
- Byte permute
- Bit-wise conditional move
- Fraction extract
- Half-precision convert

For more details: *AMD64 Architecture Programmer's Manual Volume 6: 128-Bit and 256-Bit XOP and FMA4 Instructions*
http://support.amd.com/us/Embedded_TechDocs/43479.pdf

* http://software.intel.com/en-us/blogs/2011/06/13/haswell-new-instruction-descriptions-now-available/

NEW "BULLDOZER" INSTRUCTIONS USAGE RECOMMENDATIONS

Software using SSE instructions should be recompiled with AVX 128 and FMA4 compiler options (see Compiler Optimization Guide*) & linked to ACML 5.x libraries.

If software currently supports the new instructions that are common with Intel (SSSE3, SSE4.1/4.2, AES-NI, AVX):

- No recompile of code needed if the software only checks ISA feature bits
- Recompile needed if software also checks for CPU VENDOR *(For Example: ICC generates code that checks for CPU VENDOR)*

For software to support FMA4 or XOP (AMD-specific instructions):

- Rewritten to call new instructions

-OR-

- Recompiled with options to automatically generate code that uses these instructions

-OR-

- Linked to a library that offers support for these instructions

* http://developer.amd.com/Assets/CompilerOptQuickRef-62004200.pdf

WHY DOES AMD SUPPORT A RANGE OF COMPILERS?

No one compiler services all of our target markets

Compilers	Languages Supported	Processors Supported	Operation Systems Supported	Comments
GCC	C,C++, Fortran, Objective-C, Java, Ada, Go	Wide variety including: x86, AIX, SPARC, ARM	Wide variety including: Linux, Windows, Mac OS, Android, Solaris	Default compiler for Linux
Intel	C, C++, Fortran	Intel x86, Itanium	Linux, Windows, Mac OS	Performance compiler for Intel
Open64	C, C++, Fortran	AMD and Intel x86	Linux	Performance compiler for AMD
PGI	C, C++, Fortran	AMD, Intel x86, NVIDIA CUDA	Linux, Mac OS, Windows	Performance compiler for HPC
MSFT Visual Studio	C, C++, C#, Basic	AMD and Intel x86	Windows	Default compiler for Windows

- Default compilers are used to compile the kernel, some of the system software, and libraries for the OS
- Customers are often reluctant to change compilers
- Compilers used to generated high performance code are not necessarily the ones used for mainstream server applications

OPEN64 COMPILER | A CLOSER LOOK

Setting the "–march" (microarchitecture) flag will automatically optimize code for the target processor's instruction set

Open64 Settings	Processor Type
-march=bdver1	AMD Opteron™ 4200 and 6200 Series
-march=barcelona	AMD Opteron™ 13xx, 14xx, 23xx, 24xx, 83xx, 84xx, 4100, and 6200 Series
-march=any86	Any x86 processor

"Bulldozer" compiler optimizations enabled by –march=bdver1*

- Support for all new instructions (SSSE3, SSE4.1, SSE4.2, AVX, FMA, and XOP)
- Automatically selects instructions to improve performance (intrinsics and inline)
- Automatic calls to libM (math library) functions that use these new instructions
- Code generation tuned for microarchitecture, e.g. instruction latencies, cache sizes
- Adjusted to take advantage of the improved hardware prefetcher
- Improvements in code layout and alignment to take advantage of shared compute unit, e.g. "dispatch scheduling"

* Additional information: http://developer.amd.com/tools/open64/Documents/open64.html

GCC COMPILER | A CLOSER LOOK

Setting the "–march" (microarchitecture) flag will automatically optimize code for the target processor's instruction set

Open64 Settings	Processor Type
-march=bdver1	AMD Opteron™ 4200 and 6200 Series
-march=amdfam10	AMD Opteron™ 13xx, 14xx, 23xx, 24xx, 83xx, 84xx, 4100, and 6200 Series
-march=generic	Any x86 processor

"Bulldozer" compiler optimizations enabled by –march=bdver1

- Support for all new instructions (SSSE3, SSE4.1, SSE4.2, AVX, FMA, and XOP)
- Automatically selects instructions to improve performance (intrinsics and inline)
- Scalar and vector libm calls available with AMD Libm
- Code generation tuned for microarchitecture, e.g. instruction latencies, cache sizes
- Memset/Memcpy inliner heuristics
- Defaults to 128-bit vectorization
- Improvements in code layout and alignment

Additional information: http://developer.amd.com/tools/gnu/pages/default.aspx

"AMD OPTERON™ 4200 AND 6200 SERIES PROCESSORS LIBRARY SUPPORT

A library is a collection of pre-written code and subroutines

	Description	Comments
ACML (AMD Core Math Library)	Set of optimized and threaded math routines for HPC, scientific, engineering and related compute-intensive applications	ACML 4.x is compatible with "Bulldozer" ACML 5.x is optimized for "Bulldozer"
AMD LibM	C library containing a collection of basic math functions optimized for x86-64 processors	AMD LibM 3.0 is optimized for "Bulldozer"
IPP (Intel Performance Primitives)	Library of multicore-ready, optimized software functions for multimedia, data processing, and communications applications	"For AMD 64-bit processors that support SSE3 the "m7" version of the IPP library will be dispatched automatically. Otherwise "mx" library will be used"*

For more information on ACML, go to: http://developer.amd.com/libraries/acml/pages/default.aspx
For more information on AMD LIbM, go to: http://developer.amd.com/libraries/libm/pages/default.asp

* http://software.intel.com/en-us/articles/use-ipp-on-amd-processor/

ACML (AMD CORE MATH LIBRARY) | A CLOSER LOOK

- A full implementation of Level 1, 2 and 3 Basic Linear Algebra Subroutines (BLAS), with key routines optimized for high performance on AMD Opteron™ processors.
- A full suite of Linear Algebra (LAPACK) routines. As well as taking advantage of the highly-tuned BLAS kernels, a key set of LAPACK routines has been further optimized to achieve considerably higher performance than standard LAPACK implementations.
- A comprehensive suite of Fast Fourier Transforms (FFTs) in both single-, double-, single-complex and double-complex data types.
- Random Number Generators in both single- and double-precision.

Compiler Support

- Absoft Pro Fortran
- GFORTRAN
- Intel Fortran (Linux, Windows)
- NAG Fortran
- Open64
- PGI Fortran (Linux, Windows)

For more information on ACML, go to: http://developer.amd.com/libraries/acml/pages/default.aspx

ACML SUPPORT | A CLOSER LOOK

	Linear Algebra	Fast Fourier Transforms (FFT)	Others	Compiler Support
ACML 5.0 (Aug 2011)	• SGEMM (single precision) • DGEMM (double precision) • L1 BLAS	• Complex-to-Complex (C-C) single precision FFTs	• Random Number Generators • AVX compiler switch for Fortran	• Absoft • GCC 4.6 • Open64 4.2.5 • PGI 11.8, 11.9 • ICC 12 • Cray to begin deployment of ACML with their compiler with ACML 5.0
ACML 5.1 (Dec 2011)	• CGEMM (complex single decision) • ZGEMM (complex double precision)	• Real-to-complex (R-C) single precision FFTs • Double precision C-C and R-C FFTs		All compilers listed for ACML 5.0 will be supported

For additional information on ACML, go to:
http://developer.amd.com/libraries/acml/pages/default.aspx

STARTING POINTS FOR APPLICATION OPTIMIZATION

	Operating System	Compiler	Library
Recommended for SPECCPU, LINPACK, HPC Challenge	Novell SLES 11 SP1 or RHEL 6.1	Open64 4.2.5	ACML 5.0
Recommended for application development and benchmarks with gcc	Novell SLES 11 SP1 or RHEL 6.1	GCC 4.6	ACML 5.0 and/or libM 3.0
Recommended for HPC application code development	Novell SLES 11 SP1 or RHEL 6.1	Open64 4.25 or PGI 11.9	ACML 5.0
Recommend for integer code development for Windows	Windows Server 2008 RS SP1	Microsoft Visual Studio 2010 SP1	AMD libM 3.0

Recommendations are based on AMD evaluations, please evaluate for your specific workload

Three Eras of Processor Performance

Single-Core Era

Enabled by:
- ✓ Moore's Law
- ✓ Voltage Scaling
- ✓ Micro-Architecture

Constrained by:
- ✗ Power
- ✗ Complexity

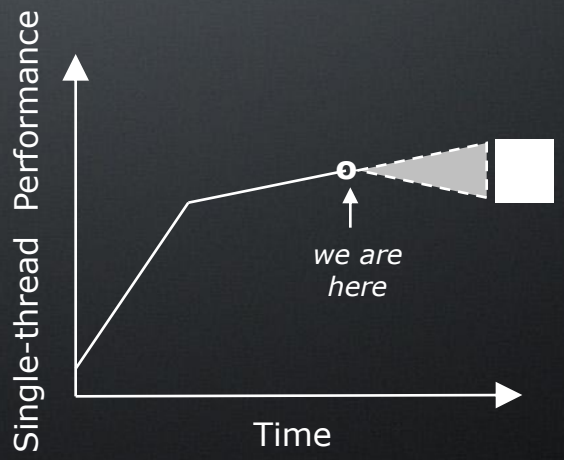

Multi-Core Era

Enabled by:
- ✓ Moore's Law
- ✓ Desire for Throughput
- ✓ 20 years of SMP arch

Constrained by:
- ✗ Power
- ✗ Parallel SW availability
- ✗ Scalability

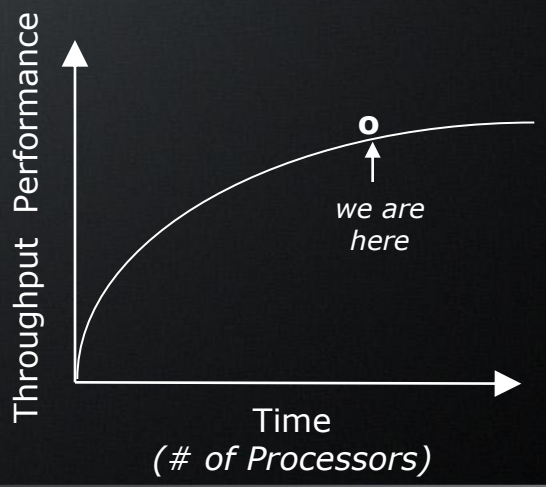

Heterogeneous Systems Era

Enabled by:
- ✓ Moore's Law
- ✓ Abundant data parallelism
- ✓ Power efficient GPUs

Temporarily constrained by:
- ✗ Programming models
- ✗ Communication overheads

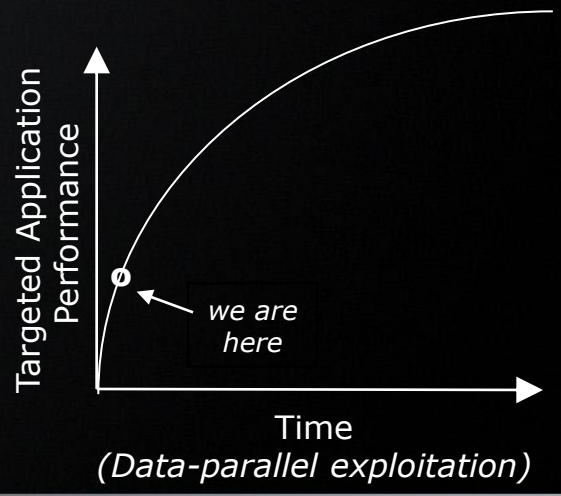

GPU Compute Offload – 3 Phases

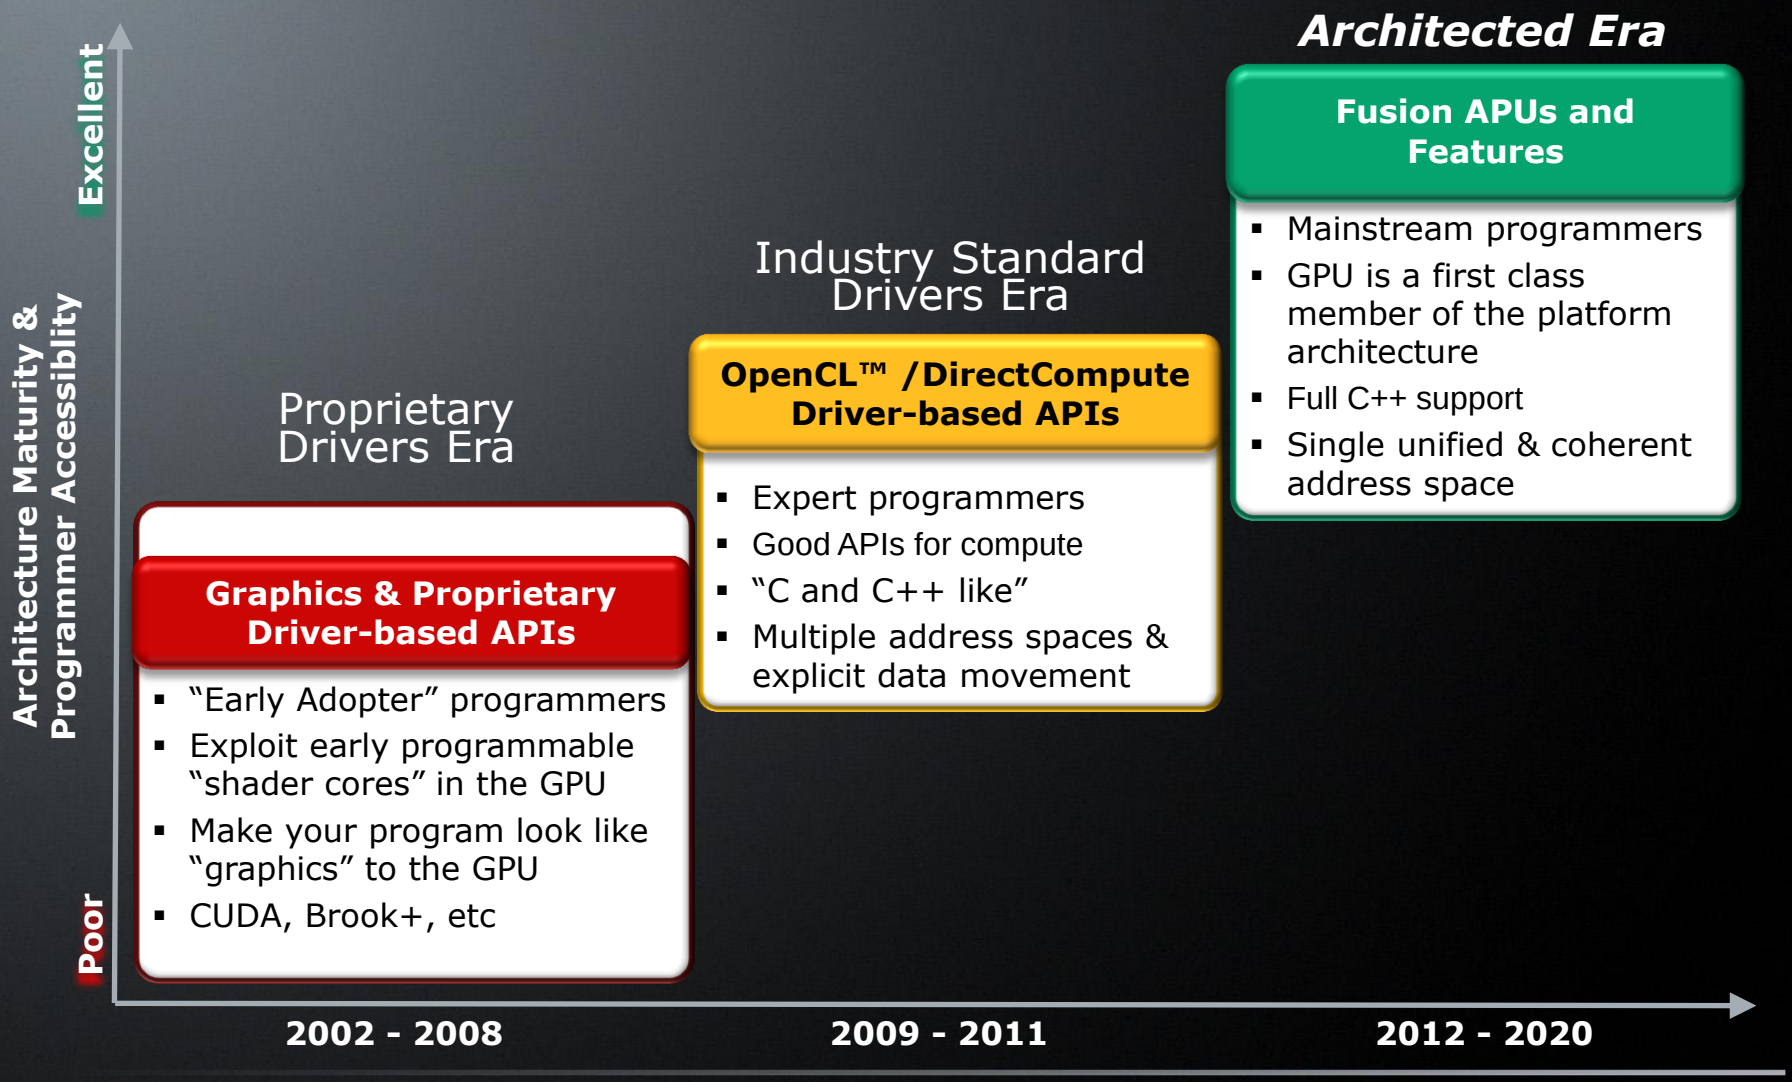

A New Era of Processor Performance

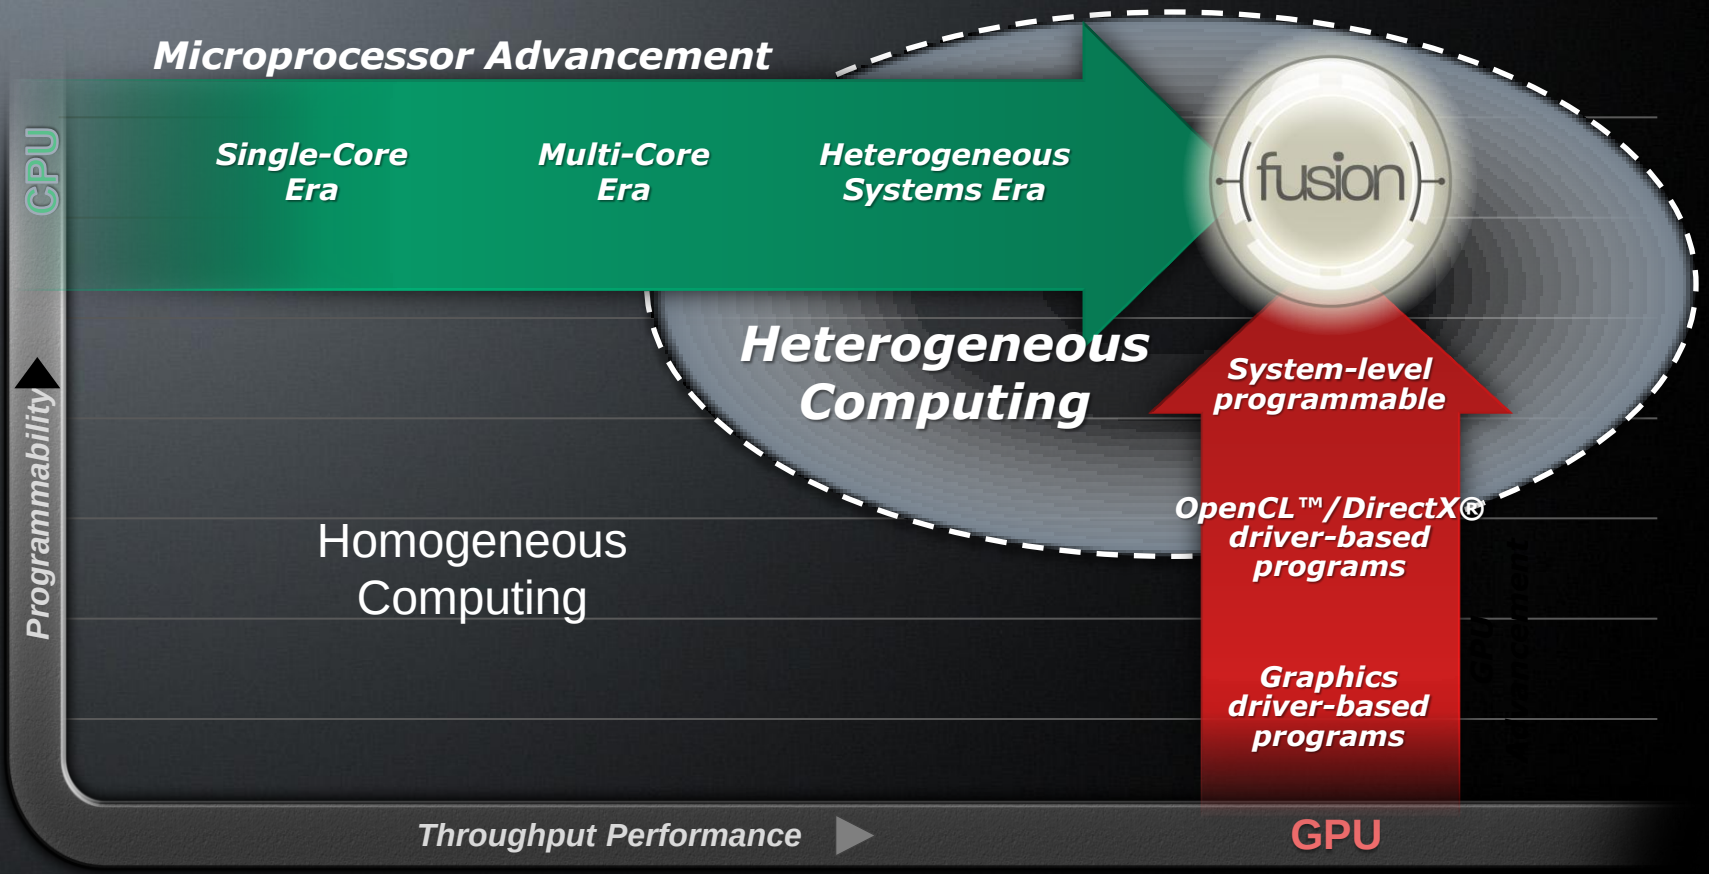

AMD Fusion APUs Fill the Need

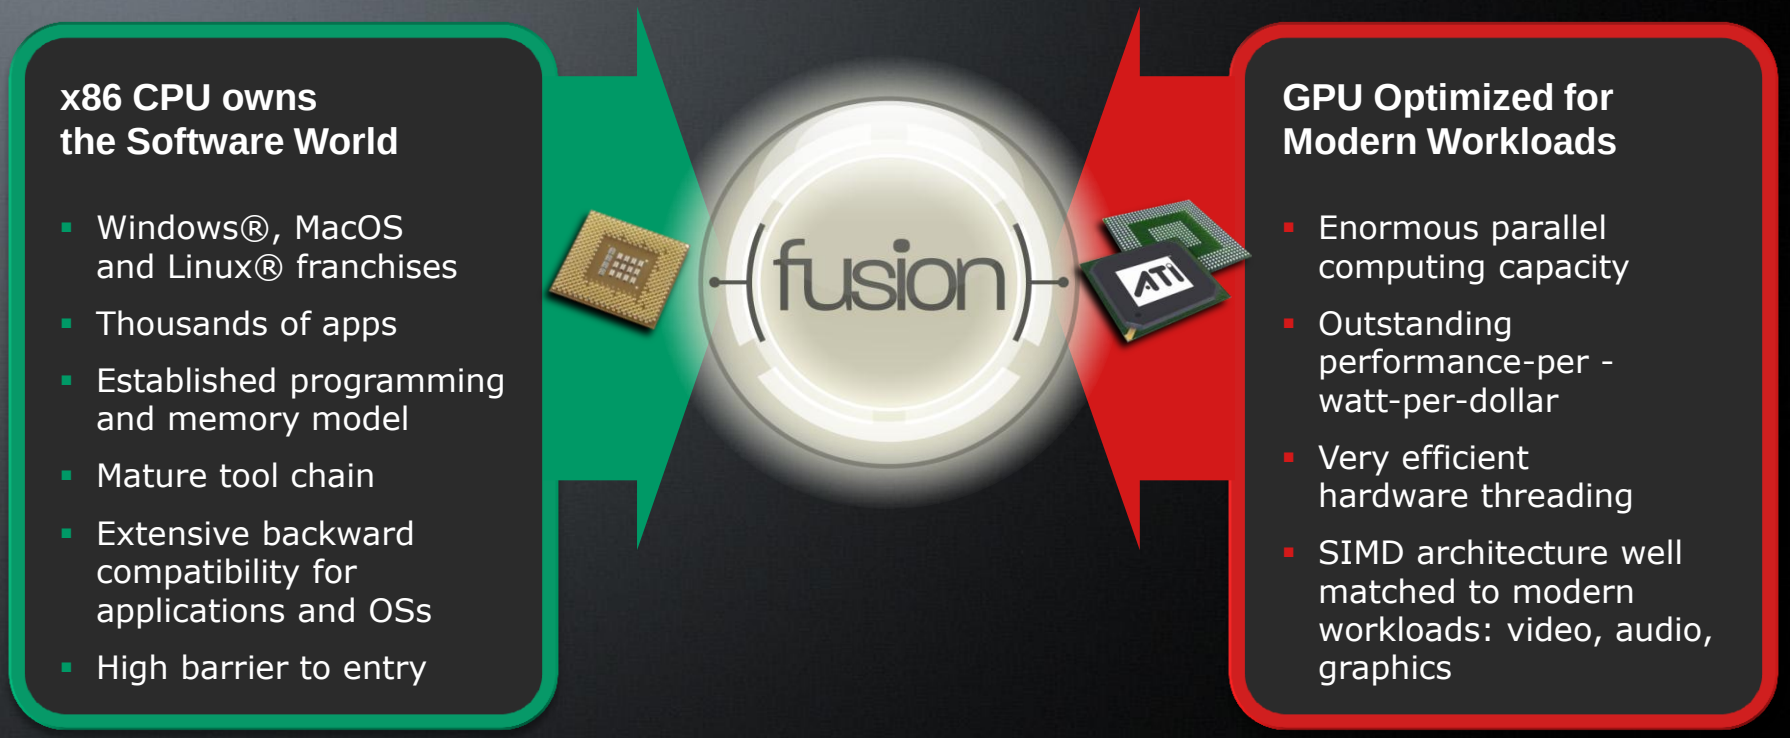

AMD Fusion: Enabling Heterogeneous Computing in A Broader Set of Applications

Discrete architectures can deliver performance acceleration…

CPUs

- ✓ Single-threaded performance
- ✓ Efficient flow control

Add GPUs

- ✓ Parallel Data performance
- ✓ High performance per watt

But…

Costly data movement means complex programming (optimal kernel sizes, hand-tuning, etc.)

AMD Fusion

All of the CPU and GPU advantages, plus:

+ No data movement bottleneck
+ No limits on memory size
+ Easier to program; broader application availability
  - Scientific and engineering
  - Sorting/searching
  - Real time Audio/video
  - Face/object recognition
  - Image processing
  - Physics and AI

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

This presentation contains forward-looking statements concerning AMD and technology partner product offerings which are made pursuant to the safe harbor provisions of the Private Securities Litigation Reform Act of 1995. Forward-looking statements are commonly identified by words such as "would," "may," "expects," "believes," "plans," "intends," "strategy," "roadmaps," "projects" and other terms with similar meaning. Investors are cautioned that the forward-looking statements in this presentation are based on current beliefs, assumptions and expectations, speak only as of the date of this presentation and involve risks and uncertainties that could cause actual results to differ materially from current expectations.

ATTRIBUTION