

# AMDの最新テクノロジーと ヘテロニアスコンピューティング への取り組み.

(2012年度 理研シンポジウム)

日本AMD株式会社  
エンタープライズ事業本部 兼 コーポレートマーケティング部

林 淳二 (junji.hayashi@amd.com)

2012年 2月16日



# AGENDA

**Product Strategy**

**Roadmaps and Directions**

**HSA Strategy & Roadmap**

**Q&A**



# AMD PRODUCT FOCUS: FULLY ALIGNED TO MARKET TRENDS



**Exceptional user experience  
across device categories**

**Leadership Compute & Visualization  
at lower power and cost**

**Flexible SoCs, differentiated APUs  
across a complete range of platforms**



# FUTURE DIRECTIONS



## *Leadership IP*



## *HSA*



## *Low Power*



## *Ambidextrous Architectures*



## *Product Growth Vectors*

- Extend AMD's Product Roadmap into Ultra-Low Power (sub 2 watts)
- Enable new capabilities and User Interfaces (HSA)
- Exploit APUs in every market client-server-embedded
- Create innovative cloud and mega data center solutions



# EXECUTION: AGILITY AND FASTER TIME TO MARKET



## 1. SoC design methodology

- Modular design
- Consistent and **best-practice tools** and design flows
- IP re-use

## 2. Hardware / software co-design

- **Parallel hardware, software design**
- Leveraging emulation to get silicon right first time
- Enabling all software features

## 3. Driving low power

- **Power efficient SOCs optimized for workloads**
  - Performance per watt optimized cores
  - Ultra-low power APUs
- Dynamic platform power management



# AMD PRODUCT STRATEGY



## Graphics

- Create leadership IP and products that provide an unmatched visualization experience
- Leverage the highly parallel GPU architecture in new applications and markets

## Client

- Deliver a compelling end user experience across device categories with our disruptive APU technology
- Propel APUs into ultra-low power and new markets

## Server

- Focus on creating differentiated solutions in the next-generation of the cloud infrastructure
- Lead in performance & performance/watt in new workloads

## Execution

### Speed and Agility

IP SoC reuse

Improved design methodology

Time-to-market



# AGENDA

**Product Strategy**

**Roadmaps and Directions**

**HSA Strategy & Roadmap**

**Q&A**



# AMD OPTERON™ 6200 SERIES PROCESSORS



## Greater Performance

- World's first 16-core x86 processor<sup>1</sup>
- First processor with up to 1GHz boost over base frequency<sup>2</sup> using AMD Turbo CORE Technology
- 20% faster memory access<sup>3</sup>
- Up to 84% more throughput<sup>4</sup> on HPC workloads



**World's first truly modular x86 processor core design**

## Greater Efficiency

- First and only processor with TDP Power Capping for maximize power budgets and floor space<sup>5</sup>
- C6 power state enables ultra low power by gating power to idle cores
- As low as 5.3 W/core<sup>6</sup>, reduced processor power at idle by up to 46%<sup>7</sup>

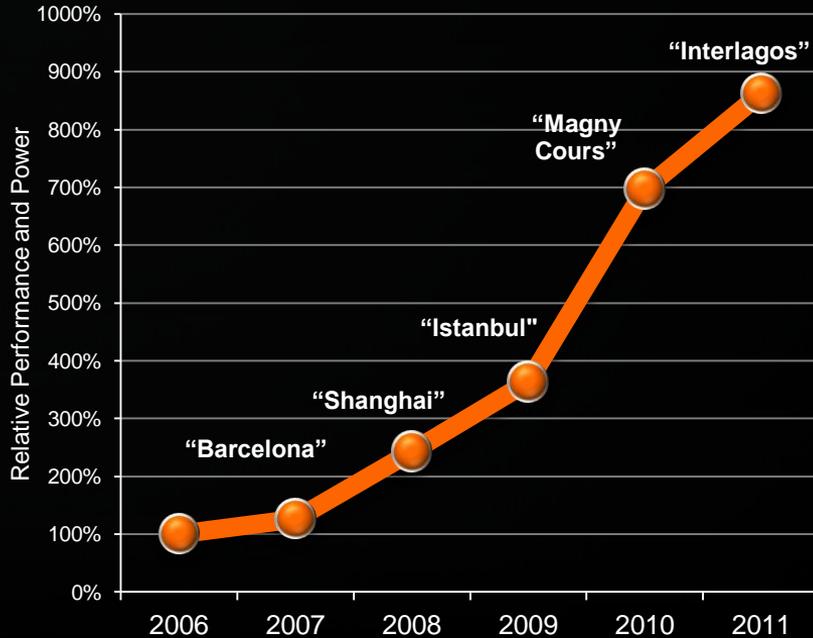
See Appendix B for footnotes.



# DRIVING LEADERSHIP PERFORMANCE-PER-WATT IN SERVER



## Performance/Watt



Source: AMD Internal

## Major Architectural Investment

- “Bulldozer” core is the foundation for continued performance enhancements for server workloads

## Revolutionary Power Design

- Shared resources in the core for maximum power efficiency
- Full power gating and power management at the system level

## Focused on TCO optimization

- Platform longevity and consistency



# 2012 – 2013 SERVER ROADMAP

32nm



**2P and 4P enterprise, mainstream platform**

## “Interlagos” CPU

- 4/8/12/16 “Bulldozer” cores
- 4x HT-3 (6.4GT)
- DDR3 (quad-channel)

## “Abu Dhabi” CPU

- 4/8/12/16 “Piledriver” cores
- DDR3 (quad-channel)

Socket G34

**1P and 2P cost-optimized, energy-efficient platform**

## “Valencia” CPU

- 6/8 “Bulldozer” cores
- 2x HT-3 (6.4GT)
- DDR3 (dual-channel)

## “Seoul” CPU

- 6/8 “Piledriver” cores
- DDR3 (dual-channel)

Socket C32

**1P Web Hosting/ Web Serving and Microserver platform**

## “Zurich” CPU (Q1)

- 4/8 “Bulldozer” cores
- 1x HT-3 (5.2GT)
- DDR3 (dual-channel)

## “Delhi” CPU

- 4/8 “Piledriver” cores
- DDR3 (dual-channel)

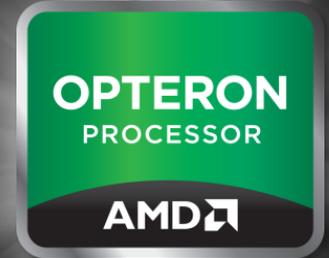
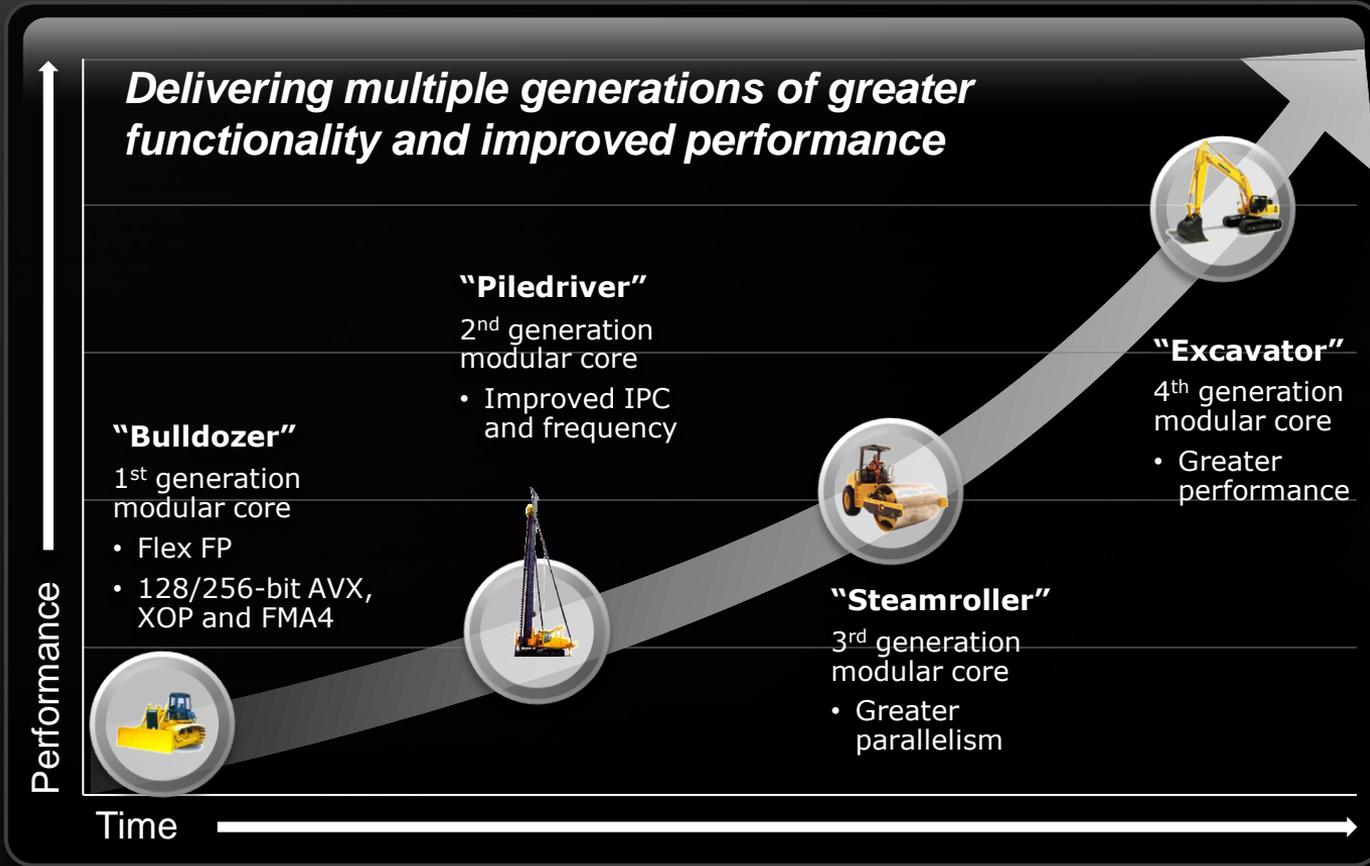
Socket AM3+

- “Piledriver” cores and process enhancements deliver more performance at the same power
- New lineup delivers better performance in the same infrastructure; replaces “Terramar” and “Sepang”

AMD roadmaps are subject to change without notice.



# AMD OPTERON™ FUTURE TECHNOLOGY



# CLIENT AND GRAPHICS ROADMAP



40nm



32nm



28nm



Discrete  
Graphics

“Northern  
Islands”

“Southern Islands”  
Graphics Core Next  
and DirectX® 11.1

Performance  
APU

“Llano”  
1<sup>st</sup> Gen APU

“Trinity”  
2<sup>nd</sup> Gen APU

Low Power  
APU

“Brazos”  
1<sup>st</sup> Gen Low Power APU

“Brazos 2.0”  
Turbo Core, USB 3.0

Ultra Low  
Power APU

“Hondo”  
1<sup>st</sup> Gen ULP APU

2011

2012

- 1<sup>st</sup> to Market with Next-Gen top-to-bottom discrete lineup
- 2<sup>nd</sup> Gen APU with significant performance and power improvements
- Extending APUs to the premium ultrathin and low power tablet market
- “Brazos 2.0” delivers better performance and features as upgrade path to “Brazos” platform; replaces “Krishna”

**INDUSTRY-LEADING GRAPHICS, COMPUTE IP RAPIDLY LEVERAGED IN  
LOW POWER PLATFORMS VIA APUs**

AMD roadmaps are subject to change without notice.



# CLIENT AND GRAPHICS ROADMAP



40nm



32nm



28nm



**Discrete Graphics**

**“Southern Islands”**  
Graphics Core Next and  
DirectX® 11.1

**“Sea Islands”**  
New GPU Architecture  
and HSA Features

**Performance APU**

**“Trinity”**  
2<sup>nd</sup> Gen APU

**“Kaveri”**  
3<sup>rd</sup> Gen APU  
“Steamroller” cores  
New HSA Features

**Low Power APU**

**“Brazos 2.0”**  
Turbo Core, USB 3.0

**“Kabini”**  
2<sup>nd</sup> Gen Low-Power APU  
“Jaguar” cores  
New HSA Features

**Ultra Low Power APU**

**“Hondo”**  
1<sup>st</sup> Gen ULP APU

**“Temash”**  
2<sup>nd</sup> Gen ULP APU  
“Jaguar” Cores

2012

2013

- Major GPU architecture enhancements for graphics, compute, HSA
- New 3<sup>rd</sup> gen APU with new x86 cores for IPC and power enhancements; Graphics Core Next and HSA enhancements
- 2<sup>nd</sup> gen low power APU with new low power x86 cores for IPC and power enhancements
- 1<sup>st</sup> gen SoC with integrated FCH

**INDUSTRY-LEADING GRAPHICS, COMPUTE IP RAPIDLY LEVERAGED IN LOW POWER PLATFORMS VIA APUs**

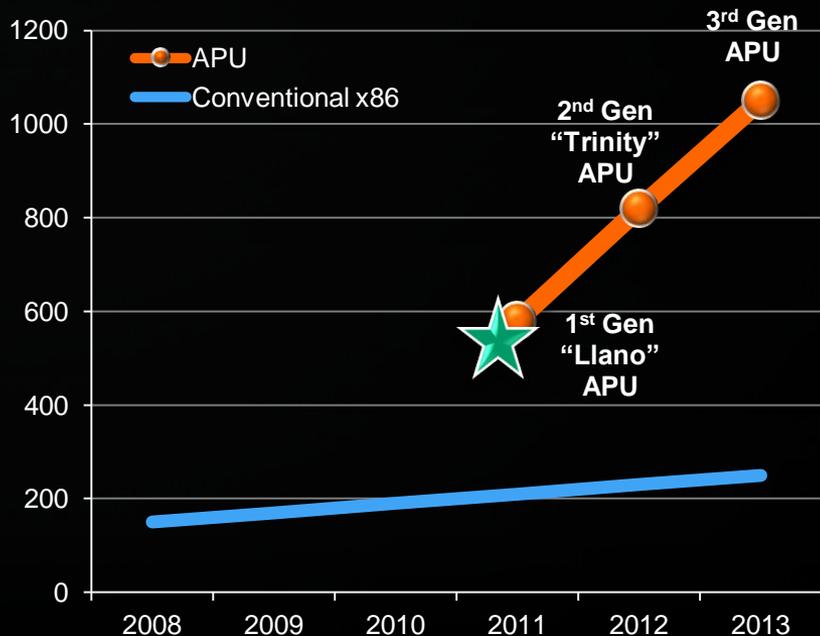
AMD roadmaps are subject to change without notice.



# APUs BRING LEADERSHIP GRAPHICS/COMPUTE IP TO MAINSTREAM



## Compute performance/GFLOPS<sup>1</sup>



- 2011: AMD first to introduce heterogeneous computing to mainstream applications
- "Llano" APU offers nearly 3X the performance in the same power envelope over conventional CPUs<sup>2</sup>
- Fully leverages the growing ecosystem of GPU-accelerated apps

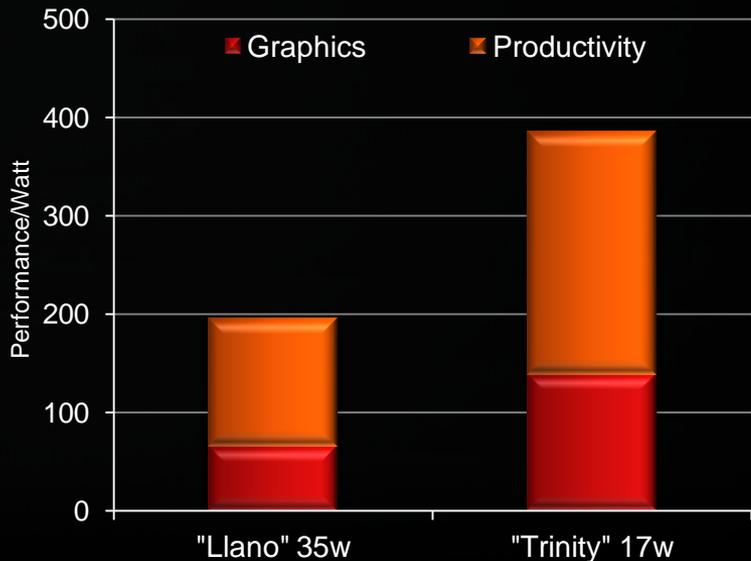
Source: AMD Performance labs. See Appendix A for footnotes.



# 2nd GEN APU "TRINITY"

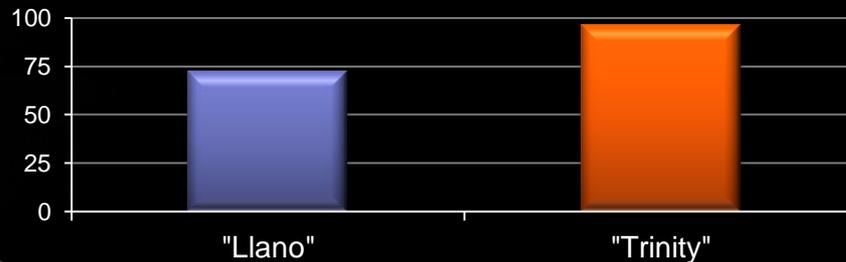


## Double the performance per watt of "Llano"



- "Piledriver" x86 cores: up to 25 percent performance increase over "Llano"<sup>3</sup>
- New Radeon™ graphic cores enable up to 50 percent graphics and compute uplift<sup>4</sup>
- "All day" battery life with over 12+ hours<sup>5</sup>
- Design wins tracking ahead of 2011 "Llano" platform

## OEM design wins at-launch



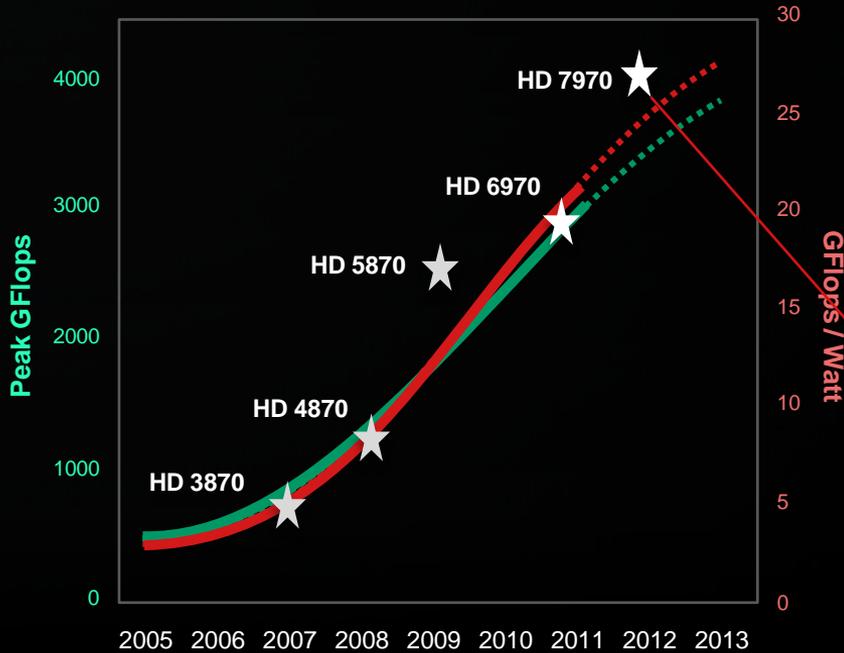
Source: AMD Internal  
See Appendix A for footnotes.



# LEADERSHIP IN GRAPHICS AND VISUALIZATION



## Leadership IP in performance-per-watt



- AMD Graphics: Leading the GPU industry inflection points
- AMD Radeon™ HD 7970 GPU (“Tahiti”) is the fastest in the world
- AMD Eyefinity: An immersive, exclusive multi-display technology, now with 3D
- 200+ applications are now accelerated by AMD GPUs and APUs

- 2,048 stream processors
- 3.79 TFLOPS of compute performance
- Support for 6 displays
- New Graphics Core Next Architecture (GCN)
- 4.3 Billion Transistors

Source: AMD Internal



# AGENDA



**Product Strategy**

**Roadmaps and Directions**

**HSA Strategy & Roadmap**

**Q&A**



# HETEROGENEOUS SYSTEM ARCHITECTURE (HSA)

## AN OPEN STANDARD



### *Why open standards matter*

- Open standards win over proprietary systems in the long term
- Enable large markets and ecosystems
- Single architecture target for OS and Apps

### *Establishing HSA as an open standard*

- Creating a consortium to own the architecture
- Specifications under review by technology partners now
- Specifications will be public at AMD Fusion Developer Summit 2012 (June 11<sup>th</sup> -14<sup>th</sup>)

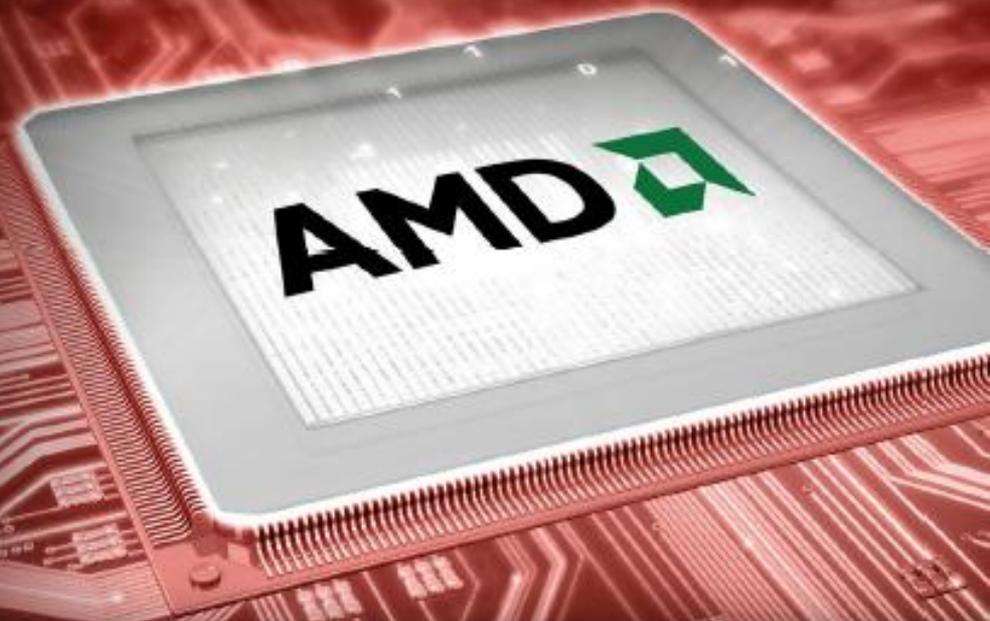


# THE FUTURE OF HETEROGENEOUS COMPUTING

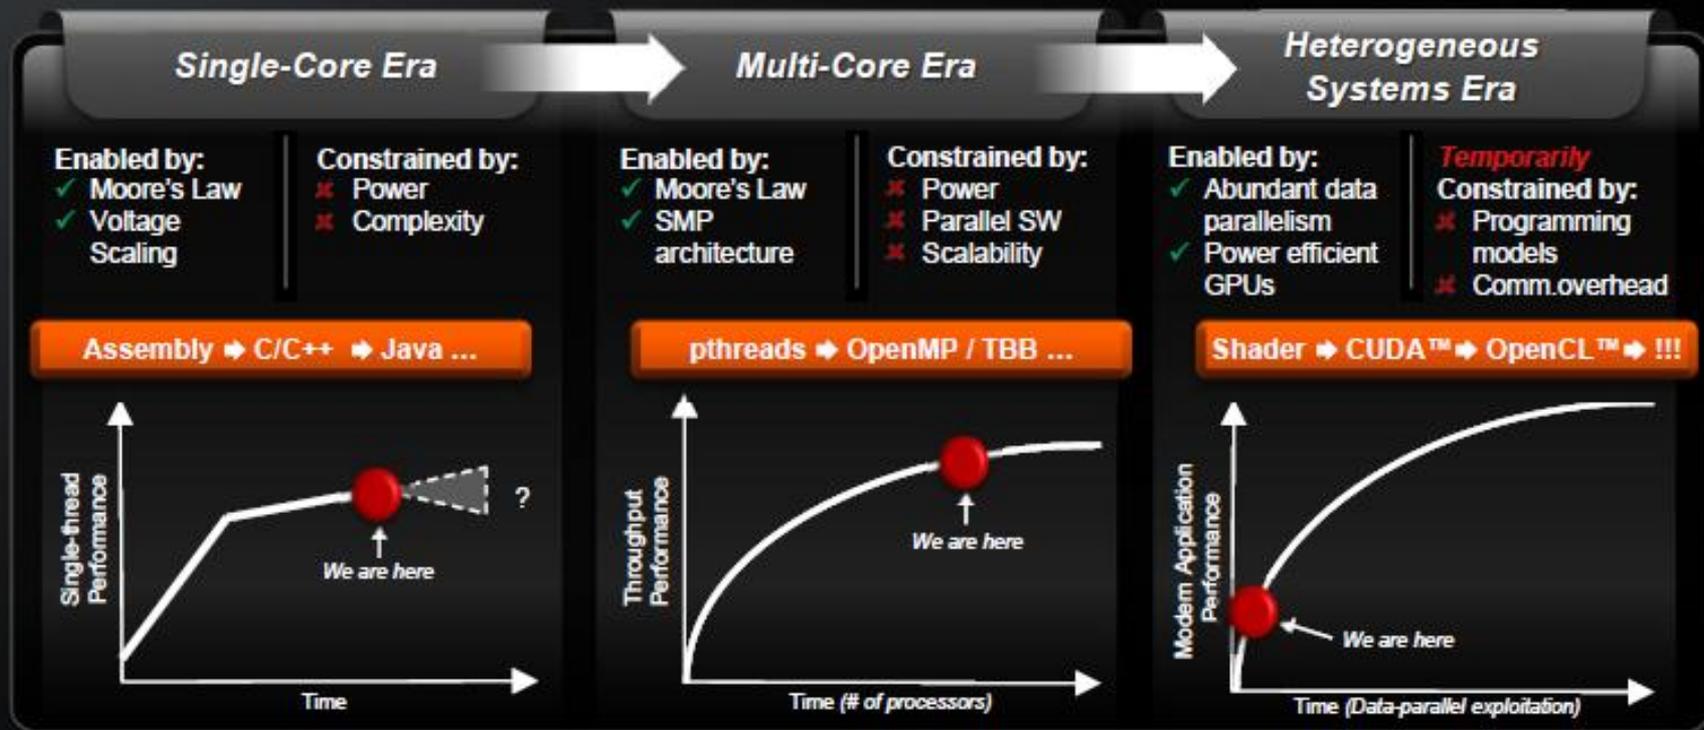


## *The architectural path for the future is clear*

- Programming patterns established on Symmetric Multi-Processor (SMP) systems migrate to the heterogeneous world
- An open architecture, with published specifications and an open source execution software stack
- Heterogeneous cores working together seamlessly in coherent memory
- Low latency dispatch
- No software fault lines



# A NEW ERA OF PROCESSOR PERFORMANCE





## Benefits

2011

Unified power efficiency



2012

Improved compute power utilization



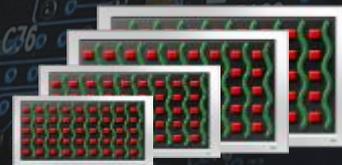
2013

Simplified data sharing



2014

HSA compatible GPUs



## Capabilities

Integrate CPU and GPU in silicon

GPU can access CPU memory

Unified memory for CPU and GPU

GPU compute context switching



# HSA FEATURE ROADMAP



2011

2012

2013

2014

Physical  
Integration

Optimized  
Platforms

Architectural  
Integration

System  
Integration

Integrate CPU & GPU  
in silicon

GPU Compute C++  
support

Unified Address  
Space for CPU and  
GPU

GPU compute  
context switch

Unified  
Memory Controller

User mode  
scheduling

GPU uses pageable  
system memory via  
CPU pointers

GPU graphics  
pre-emption

Common  
Manufacturing  
Technology

Bi-Directional Power  
Mgmt between CPU  
and GPU

Fully coherent  
memory between  
CPU & GPU

Quality of Service

Extend to  
Discrete GPU

# HETEROGENEOUS SYSTEM ARCHITECTURE

A KEY ENABLER TO OUR APU VALUE PROPOSITION

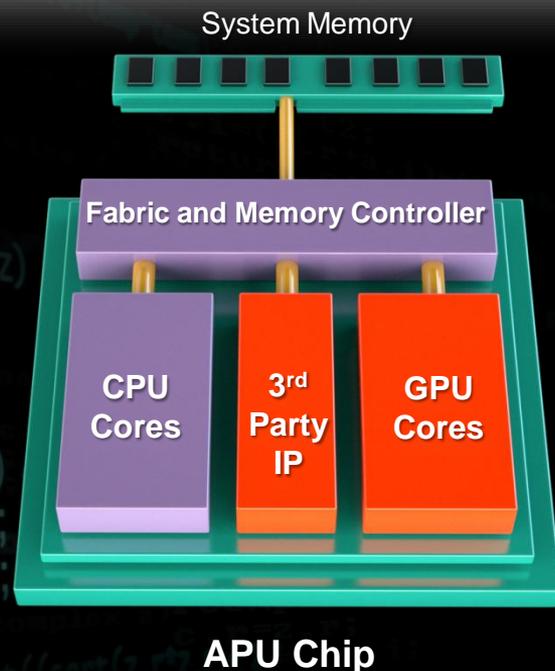


## HSA is an enabler for APU efficiency and differentiation

- Unleash our industry leading GPU cores on a broad range of applications beyond graphics
- CPU and GPU work cooperatively together directly in system memory
- Makes programming the GPU as easy as C++
- Up-to 125%\* OpenCL benchmark advantage vs. competition

## Key value propositions

- Lower power for modern applications!
- Easy for application developers to use
- Drives new class of applications
  - e.g., analytics, search, facial recognition



\*Actual performance will vary by application and machine configuration



# AGENDA

**Product Strategy**

**Roadmaps and Directions**

**HSA Strategy & Roadmap**

**Q&A**



# APPENDIX A



1. Testing performed by AMD Performance Labs. Calculated compute performance or Theoretical Maximum GFLOPS score for 2013 Kaveri (4C, 8CU) 100w APU, use standard formula of (CPU Cores x freq x 8 FLOPS) + (GPU Cores x freq x 2 FLOPS). The calculated GFLOPS for the 2013 Kaveri (4C, 8CU) 100w APU was 1050. GFLOPs scores for 2011 A-Series “Llano” was 580 and the 2013 A-Series “Trinity” was 819. Scores rounded to the nearest whole number.
2. Testing performed by AMD Performance Labs. Calculated compute performance or Theoretical Maximum GFLOPS score (use standard formula of CPU Cores x freq x 8 FLOPS) for conventional CPU alone in 2011 was 210 GFLOPs while the calculated GFLOPs for the 1<sup>st</sup> Gen APU using standard formula (CPU Cores x freq x 8 FLOPS) + (GPU Cores x freq x 2 FLOPS) was 580 or 2.8 times greater compute performance.
3. Projections and testing developed by AMD Performance Labs. Projected score for 2012 AMD Mainstream Notebook Platform “Comal” on the “Pumori” reference design for PC Mark Vantage Productivity benchmark is projected to increase by up to 25% over actual scores from the 2011 AMD Mainstream Notebook Platform “Sabine”. Projections were based on AMD A8/A6/A4 35w APUs for both platforms.
4. Projections and testing developed by AMD Performance Labs. Projected score for the 2012 AMD Mainstream Notebook Platform “Comal” the “Pumori” reference design for 3D Mark Vantage Performance benchmark is projected to increase by up to 50% over actual scores from the 2011 AMD Mainstream Notebook Platform “Sabine”. Projections were based on AMD A8/A6/A4 35w APUs for both platforms.
5. Testing performed by AMD Performance Labs. Battery life calculations using the “Pumori” reference design based on average power draw based on multiple benchmarks and usage scenarios. For Windows Idle calculations indicate 732 minutes (12:12 hours) as a resting metric; 421 minutes (7:01 hours) of DVD playback on Hollywood movie, 236 minutes (3:56 hours) of Blu-ray playback on Hollywood movie, and 205 minutes (3:25 hours) using 3D Mark '06 as an active metric. Projections for the 2012 AMD Mainstream Platform Codename “Comal” assume a configuration of “Pumori” reference board, Trinity A8 35W 4C – highest performance GPU, AMD A70M FCH, 2 x 2G DDR3 1600, 1366 x 768 eDP Panel / LED Backlight, HDD (SATA) – 250GB 5400rpm, 62Whr Battery Pack and Windows 7 Home Premium.
6. Testing done by AMD Performance Labs based on a 2012 Comal Reference Design Pumori. Results show 3D Mark Vantage for the A6 ULV 17W “Trinity” to score 2355 3D marks. Testing on a Core i5 ULV 2537M (17W) measured 1158 3D marks. With an assumed 30% increase for the Ivy Bridge architecture, the projected competitive score would be 1505 3D Marks. This provides the A6 ULV a 56% performance advantage over the projected Intel Ivory Bridge score. The 3D Mark Vantage score for the A10 LV 25W APU is 3600. This is 139% better than the projected Ivy Bridge score.



# APPENDIX B



1. Intel Xeon E7 Series has up to 10 cores per processor as of 2/1/11 at <http://www.intel.com/products/server/processor/xeonE7/index.htm>.
2. Intel's turbo boost is limited to 400MHz as of 2/1/11 on pg 147 at <http://www.intel.com/Assets/PDF/datasheet/325119.pdf>. AMD Turbo CORE technology claims based on internal AMD engineering projections of AMD Opteron 6200 Series processors with up to 500 MHz in P1 boost state and up to 1.4 GHz in P0 boost state over base P2 clock frequency.
3. Based on a maximum 1600Mhz DDR3 memory speed support for AMD Opteron 6200 Series processors vs. Intel Xeon 5600 Series maximum 1333Mhz DDR3 memory speed support: <http://ark.intel.com/products/series/47915>
4. 84% higher performance: LINPACK (2P) AMD Opteron processor Model 6276 generates 84% more FLOPS than Intel Xeon processor Model X5670  
239.1 FLOPS, 2 x AMD Opteron™ processors Model 6276 in Supermicro H8DGT server, 64GB (8 x 8GB DDR3-1600) memory, SuSE Linux® Enterprise Server 11 SP1 64-bit, gfortran compiler v4.6, OMPI 1.5.3, AMD Core Math Library 5.0.0.0  
Compiler Flags: -fomit-frame-pointer -O3 -funroll-loops -W -Wall -mavx -mfma4 -fopenmp  
  
130.1 FLOPS, 2 x Intel Xeon processors Model X5670 in Supermicro 6026TT-BIBQF server, 24GB (6 x 4GB DDR3-1333) memory, SuSE Linux® Enterprise Server 11 SP1 64-bit, Intel Professional Compiler v11.1, OMPI 1.5.1, Intel Math Kernel Library 10.3, Hyper-Threading disabled, Turbo Boost Technology enabled  
Compiler Flags: -O3 -w -ansi-alias -i-static -openmp -nocompchk
5. Not listed as a feature in Intel Xeon 5600 product brief.
6. AMD Opteron 6200 Series “HE” processors will be as low as 85W with 16 cores for 5.3 W/core.
7. (46%) - Based on testing in AMD Performance Labs as of August, 2011, an AMD Opteron™ processor model 6174 (12-core 2.2GHz) consumes 11.7W in the active idle C1E power state while an AMD Opteron™ processor model 6276 (16-core 2.3GHz) consumes only 6.4W in the active idle C1E power state with new C6 power gating employed. System configuration: “Drachma” reference design kit, 32GB (8 x 4GB DDR3-1333) memory, 500GB SATA disk drive, Microsoft® Windows Server® 2008 x64 Enterprise Edition R2. SVR-60





## 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.

## ATTRIBUTION

© 2012 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. All other products names and logos are for reference only and may be trademarks of their respective owners.

