HPC – The Next 5 Years

Towards a Fusion™ Accelerated Processing Unit (APU)

William C. Brantley, Ph.D.
Manager HPTC Performance
Bill.Brantley@AMD.COM

March 4, 2010

A New Era of Processor Performance

Single-Core Era

Constrained by:

- Power
- Complexity

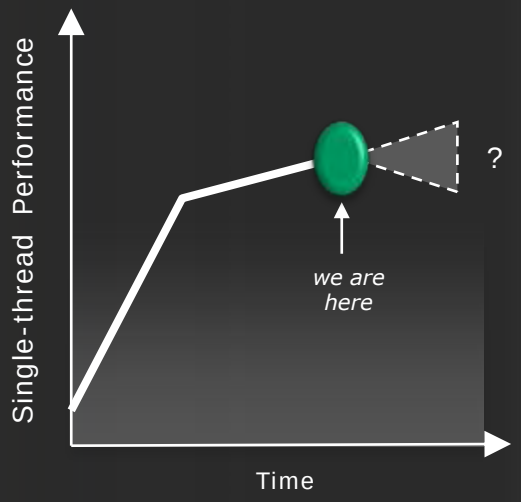

Multi-Core Era

Constrained by:

- Power
- Parallel SW availability
- Scalability

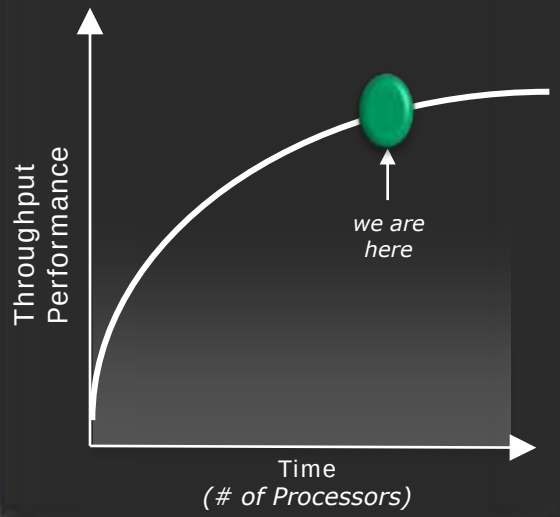

Direct Connect Architecture 1.0

Balanced and Scalable Design to Support up to 6 Cores

- ✓ No front side bus
- ✓ Integrated memory controller
- ✓ HyperTransport™ technology
- ✓ NUMA memory architecture

Direct Connect Architecture 2.0

Balanced and Scalable Design to Support up to 16 Cores* per CPU

- ✓ 1-hop between processors
- ✓ Up to 50% more DIMMs
- ✓ Four memory controllers
- ✓ Up to 33% increase in CPU to CPU communication speed**

*16-core configuration planned for upcoming AMD processor core codenamed "Bulldozer."
**Based on HyperTransport™ technology links @ up to 4.8 GT/s for Six-Core AMD Opteron™ processor vs. 6.4 GT/s for AMD Opteron™ 6100 Series processor.

"Bulldozer" x86 Architecture: AMD's Latest Leap Forward

- Two tightly linked cores share resources to increase efficiency
  - A dedicated integer core for each parallel thread
  - Flexible floating point unit that can be dedicated or shared
- ISA extensions, including FP "FMAC"
- Extensive new power mgmt innovations
- 32nm SOI with high-K metal gate
- 2011 / desktop and server

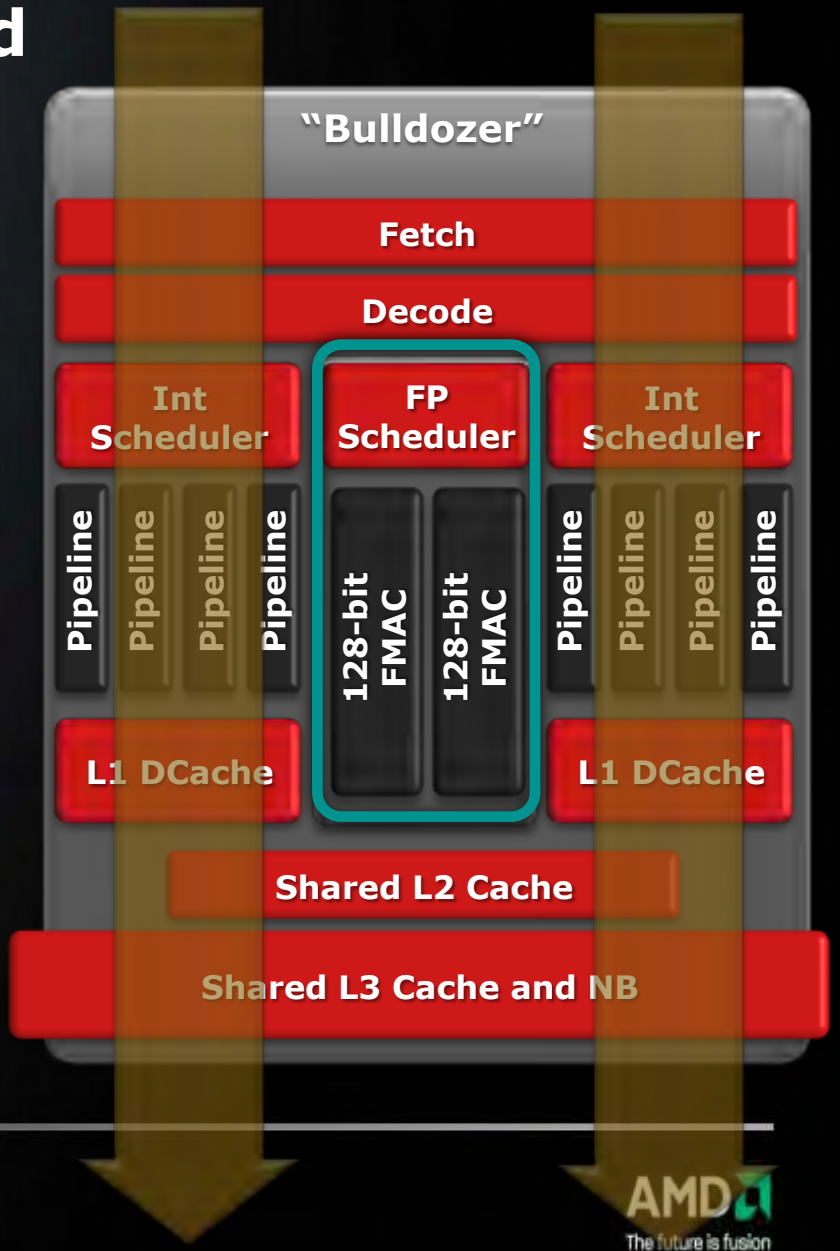

A New Era of Processor Performance

Single-Core Era

Constrained by:
- ✘ Power
- ✘ Complexity

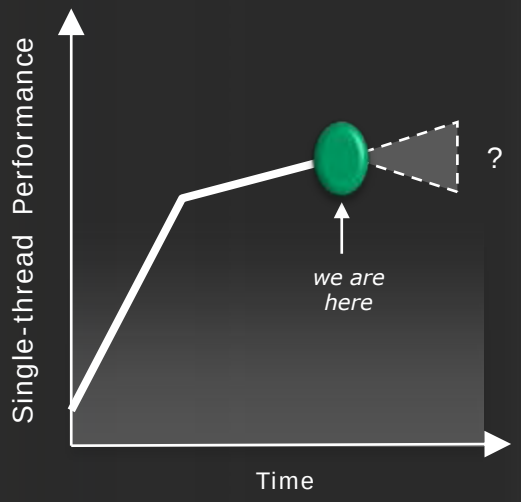

Multi-Core Era

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

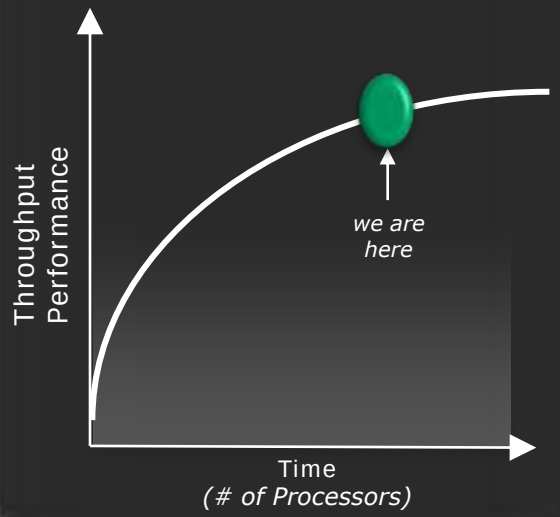

Heterogeneous Systems Era

Enabled by:
- ✓ Abundant data parallelism
- ✓ Power efficient GPUs

Constrained by:
- ✘ Programming models
- ✘ Communication Overheads

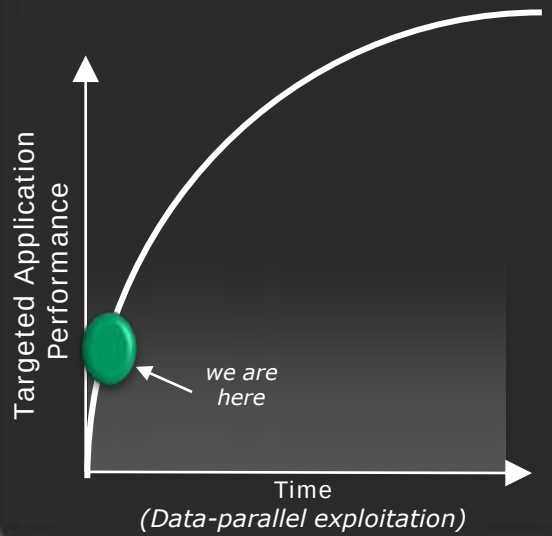

AMD
The future is fusion

AMD Fusion™ APUs: Delivering Heterogeneous Computing

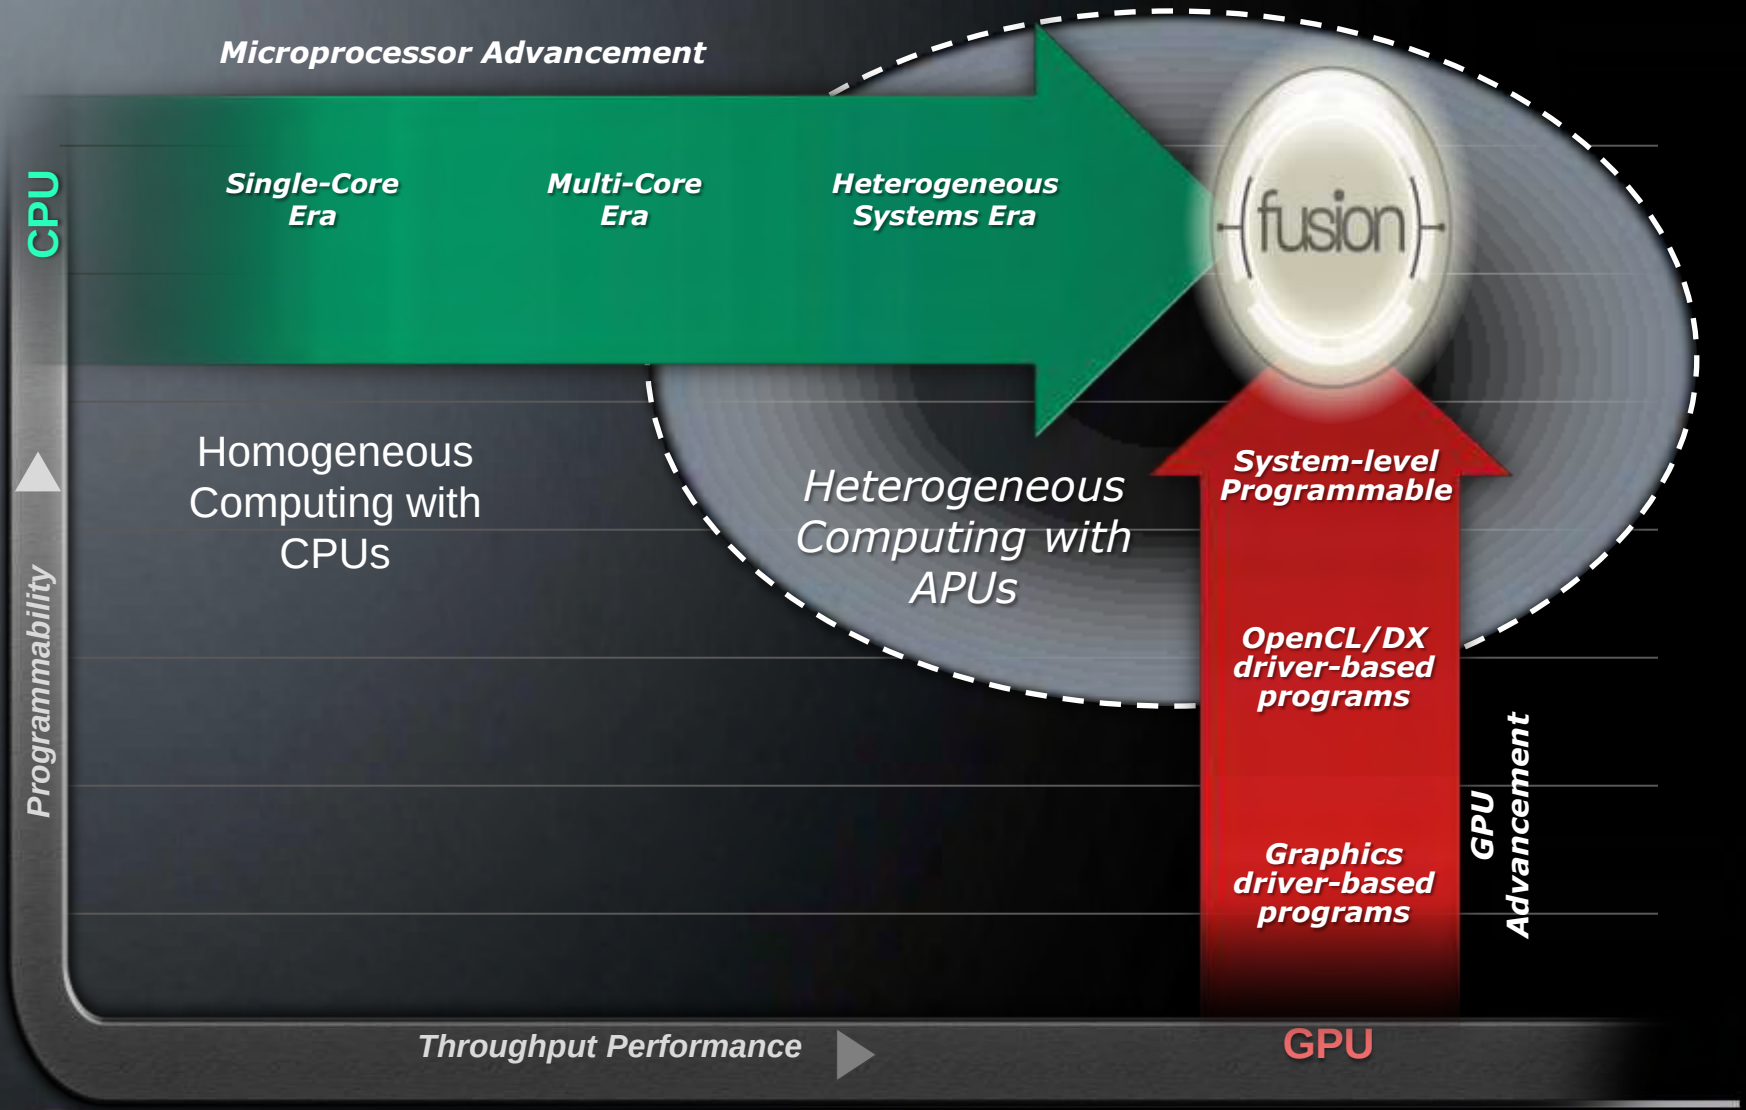

The Future For Servers Is Fusion

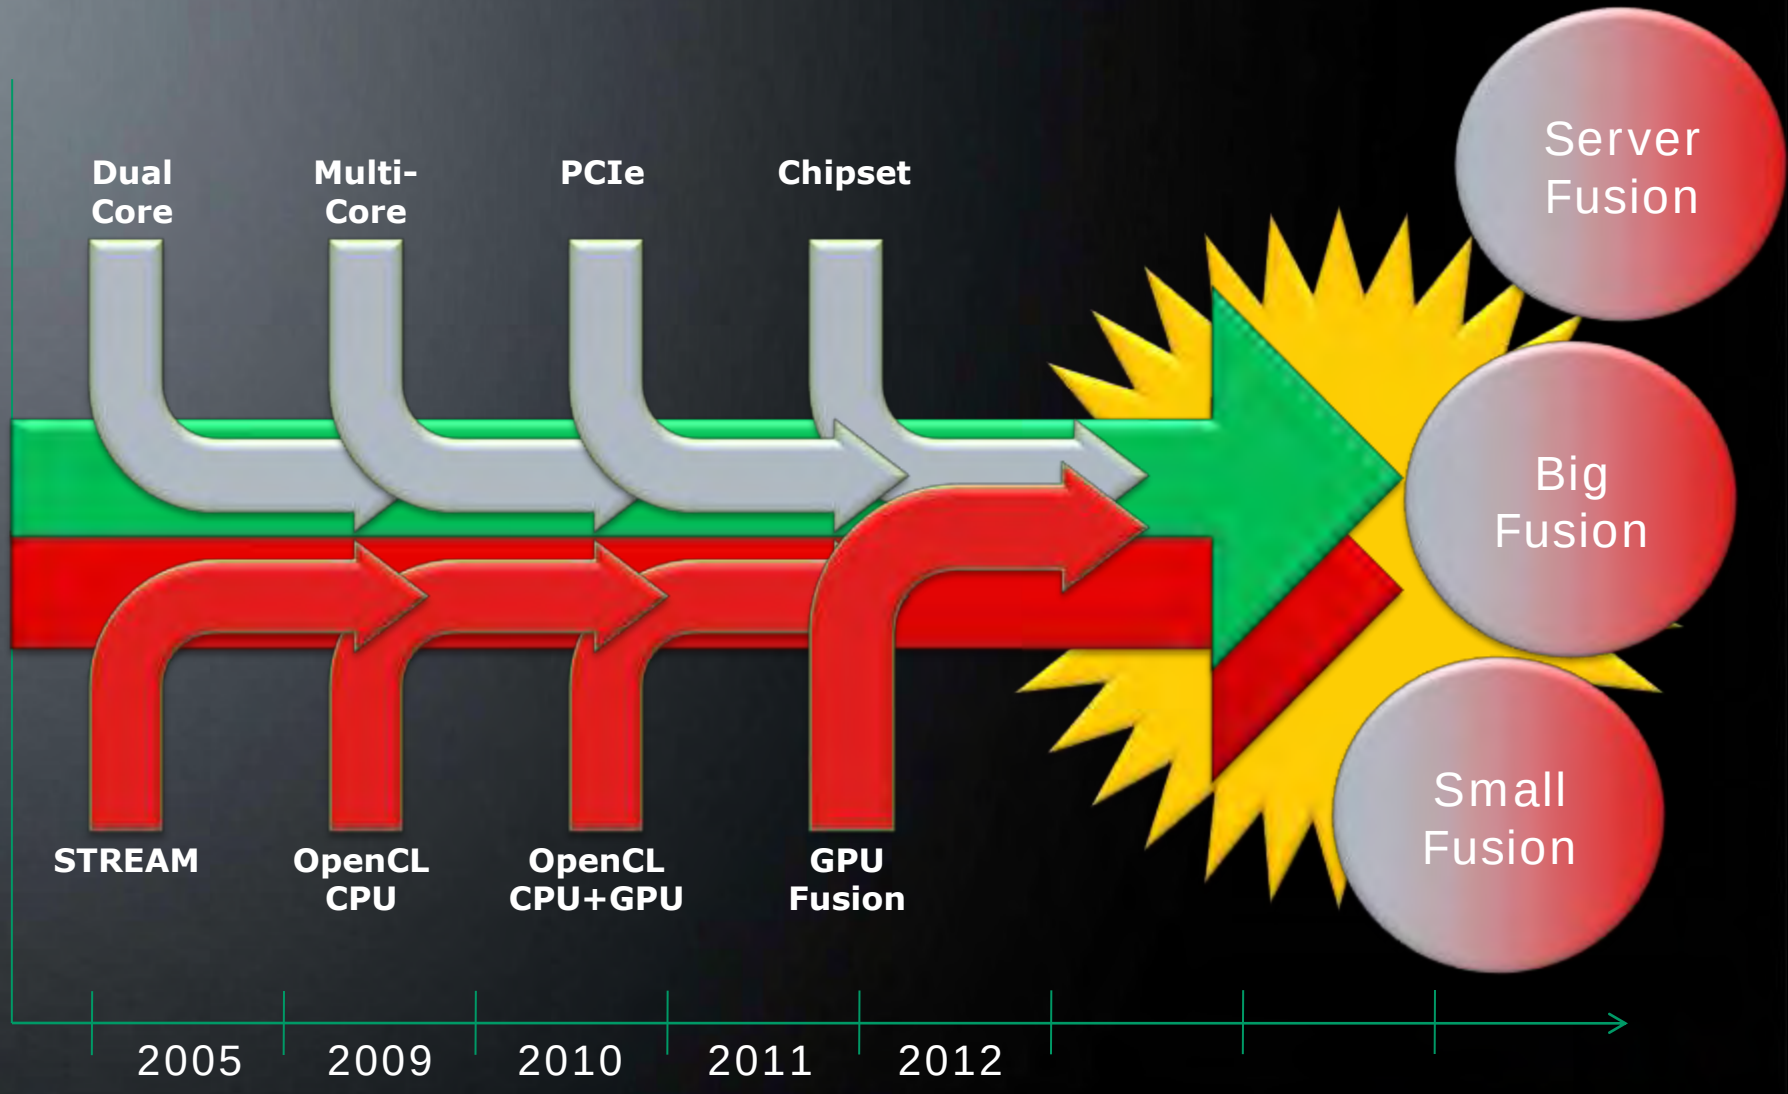

Thank You!

Disclaimer & Attribution

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

Software Seminar – Late April 2010
(West Houston area)

Converting CUDA code to industry-standard OpenCL

Dr. Udeepta Bordoloi will present a technical comparison of CUDA and OpenCL, highlighting similarities and differences along with hints on how best to convert common CUDA constructs into OpenCL. Then the team will walk through conversion of one or more code examples taken from among algorithms frequently used in Oil & Gas industry applications. After a technical Q&A session, Dr. Bordoloi will be available to discuss and provide guidance on individual CUDA conversion projects.

For details:

tommy.toles@amd.com

512-327-5389

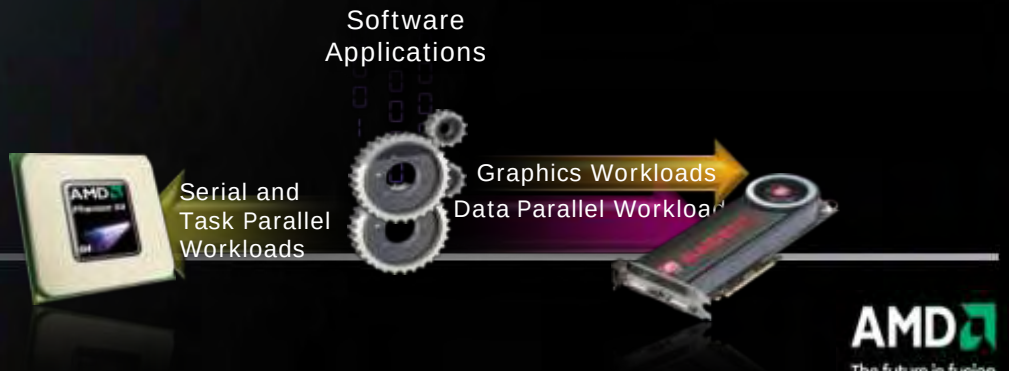

AMD The future is fusion