Trends in Processors, Memory and IO:

How Integration will Impact High Performance Computing

OPENFABRICS ALLIANCE

Kevin McGrath, Sr. Fellow
AMD Research and Advanced Development Labs

Caveats…

- The forward thinking ideas collected and discussed in this talk represent the varied thinking of the members of AMD Research and Advanced Development Labs (RAD Labs) and others in the company.

- This is NOT necessarily a reflection of the AMD roadmap!!!

The *other* Moore's Law: Semiconductor Economics

- Gordon Moore's Law relates to the ***economics*** of the semiconductor industry just as much as it does to technology progression

Chuck Moore, Corporate Fellow, AMD

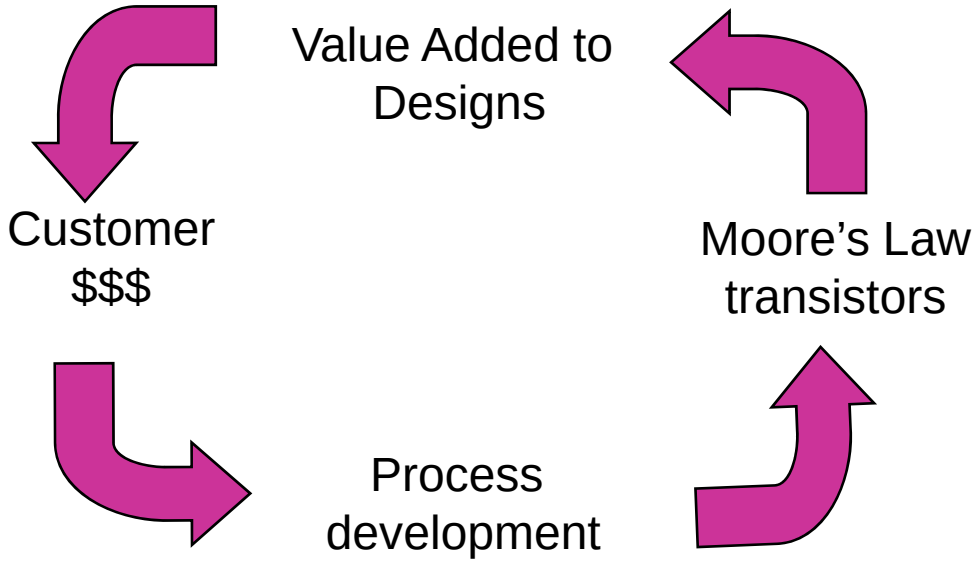

Evolution of HPC Systems

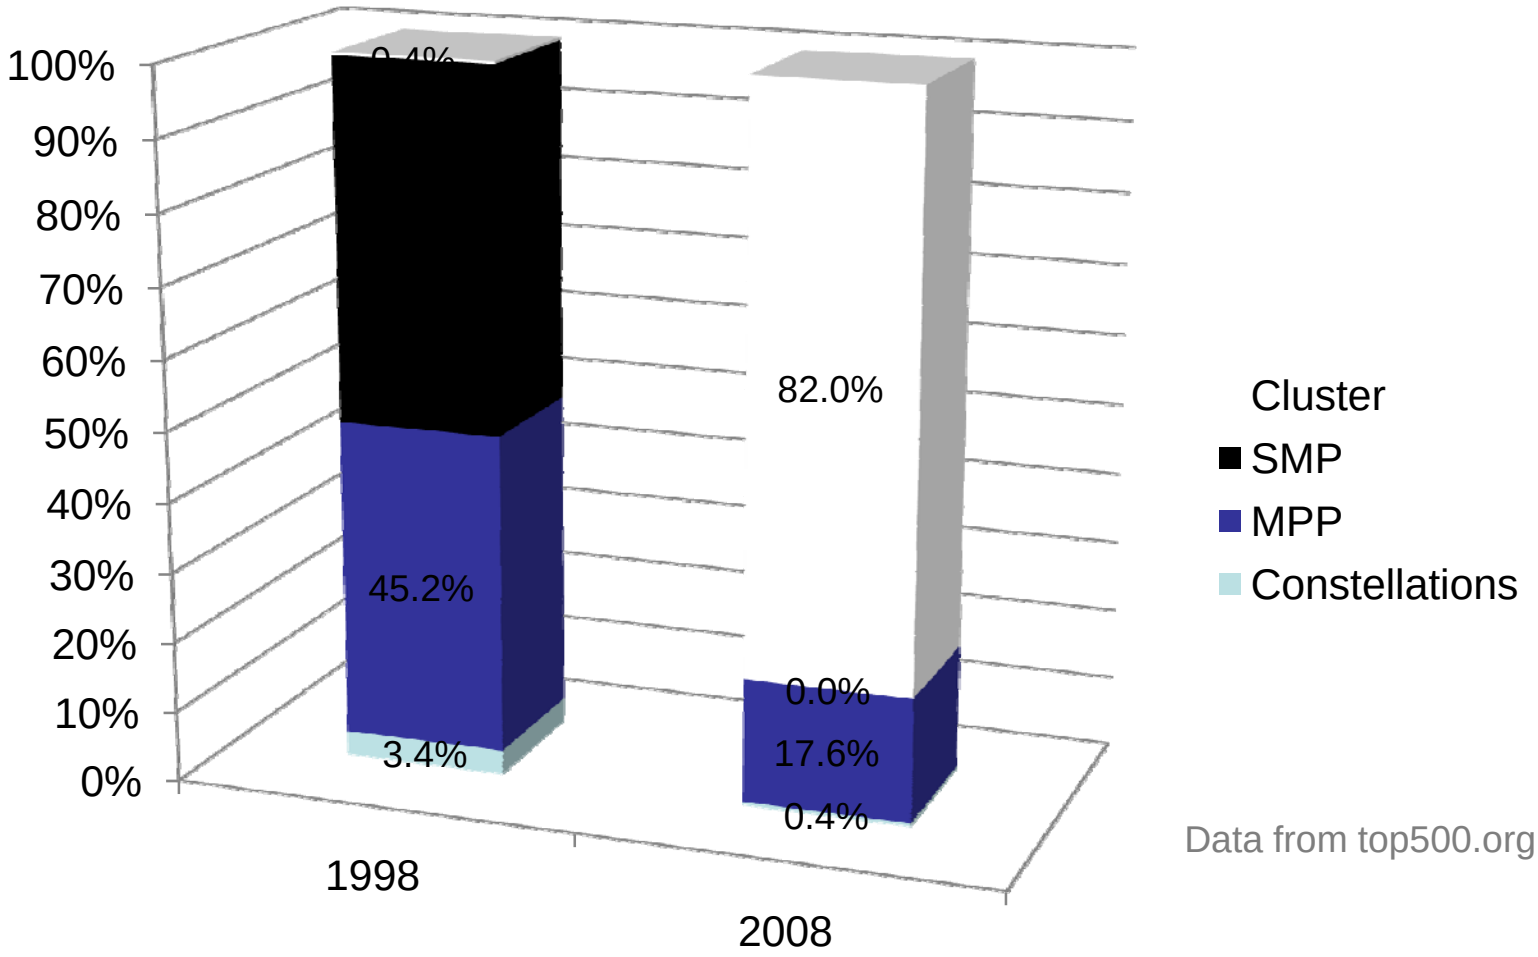

Data from top500.org

Evolution of HPC Processors

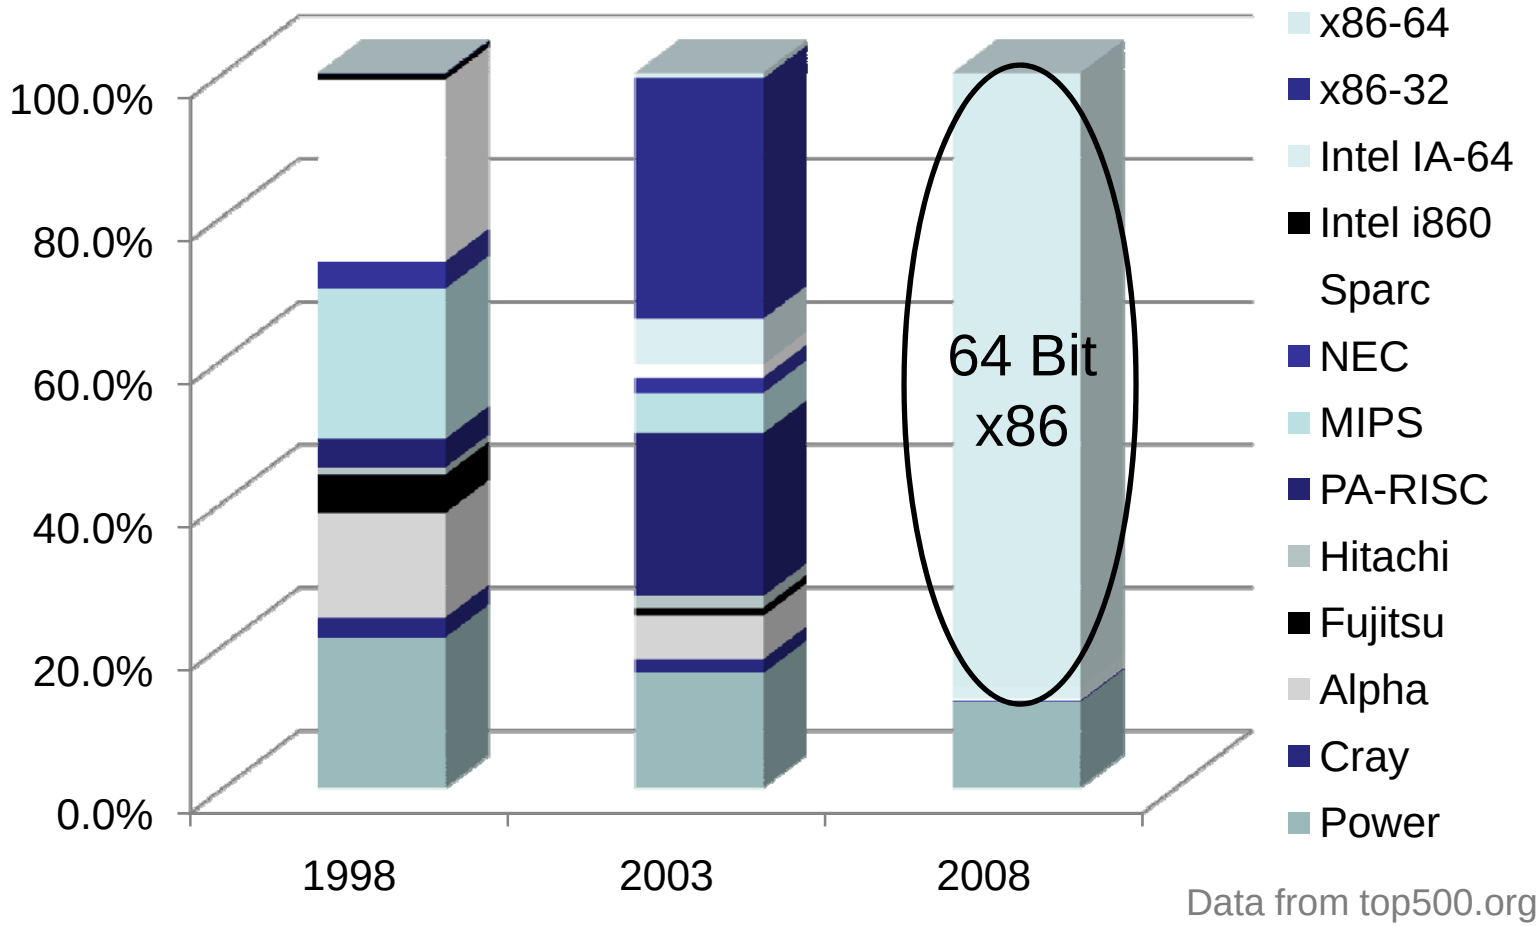

Data from top500.org

So… what does this mean?

The low end tends to eat the high end

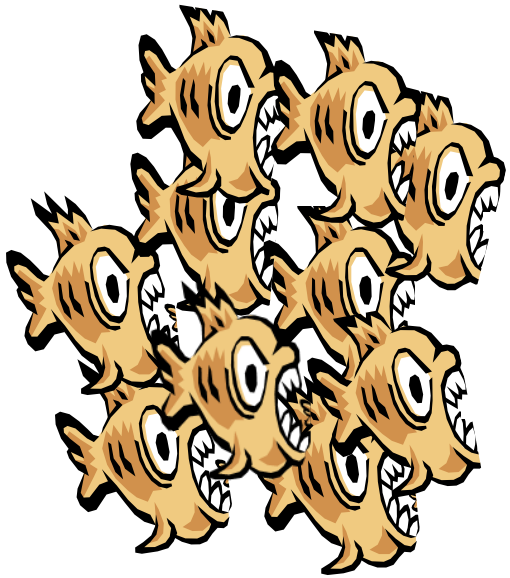

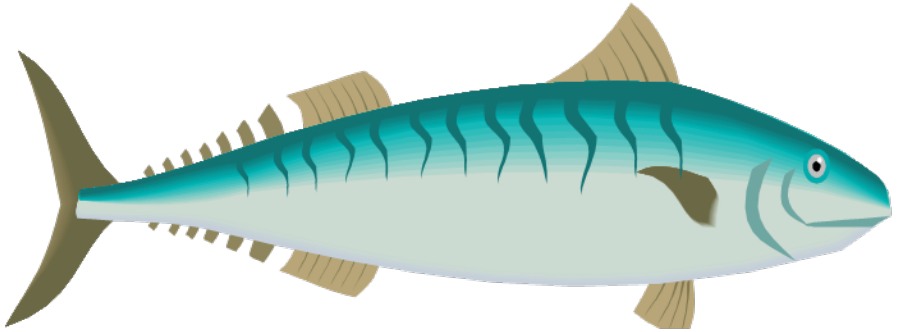

“good enough” can be a Disruptive Technology

- When combined with low TCO

Trends

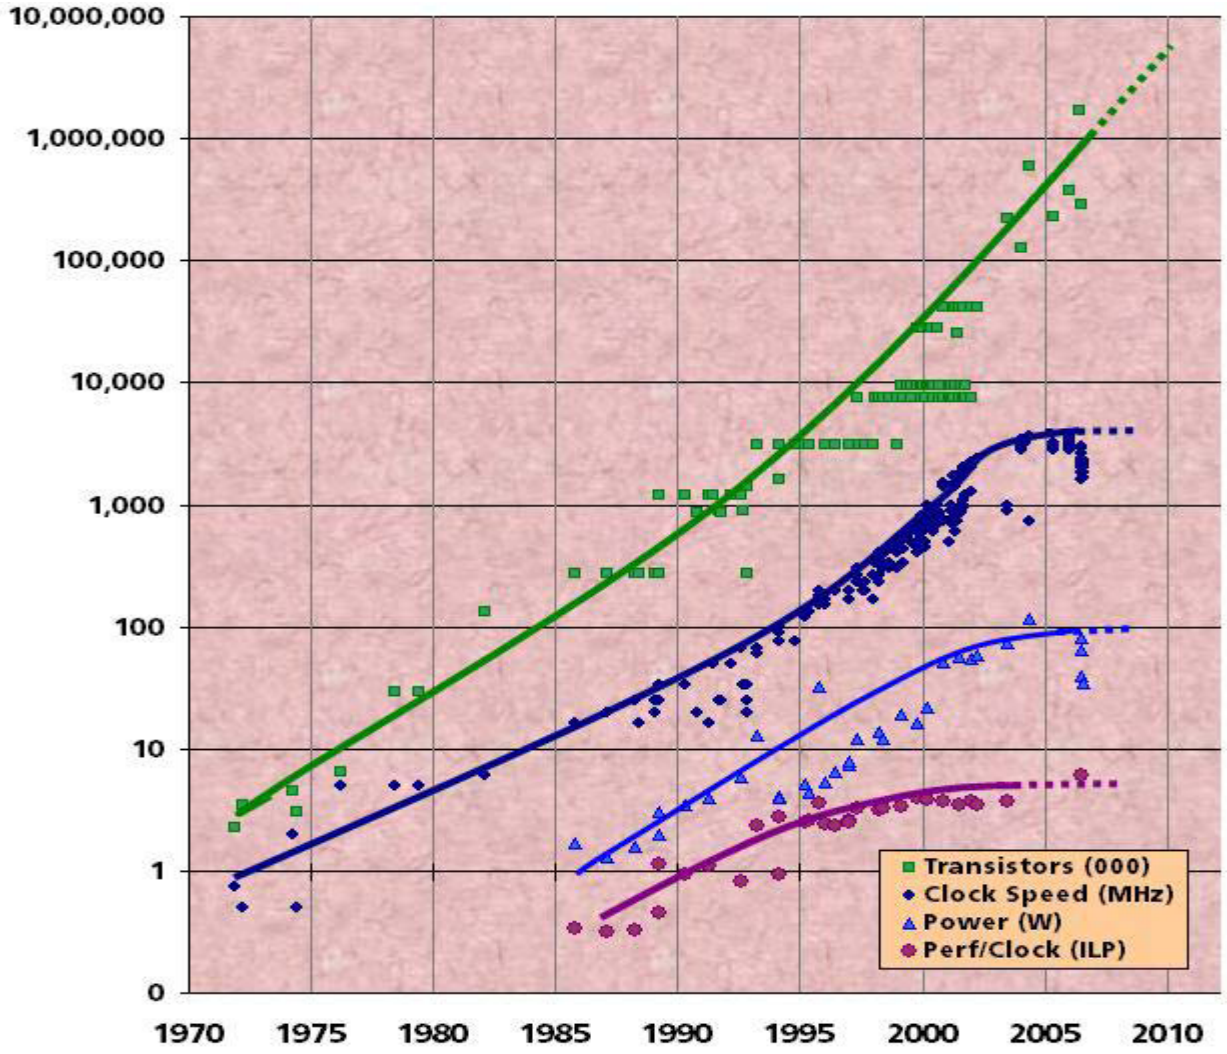
10,000,000
1,000,000
100,000
10,000
1,000
100
10
1
0
1970
1975
1980
1985
1990
1995
2000
2005
2010
Transistors (000)
Clock Speed (MHz)
Power (W)
Perf/Clock (ILP)

Parallel Programs and Amdahl's Law

$$\text{Speed-up} = \frac{1}{S_W + (1 - S_W) / N}$$

S_W: % Serial Work

N: Number of processors

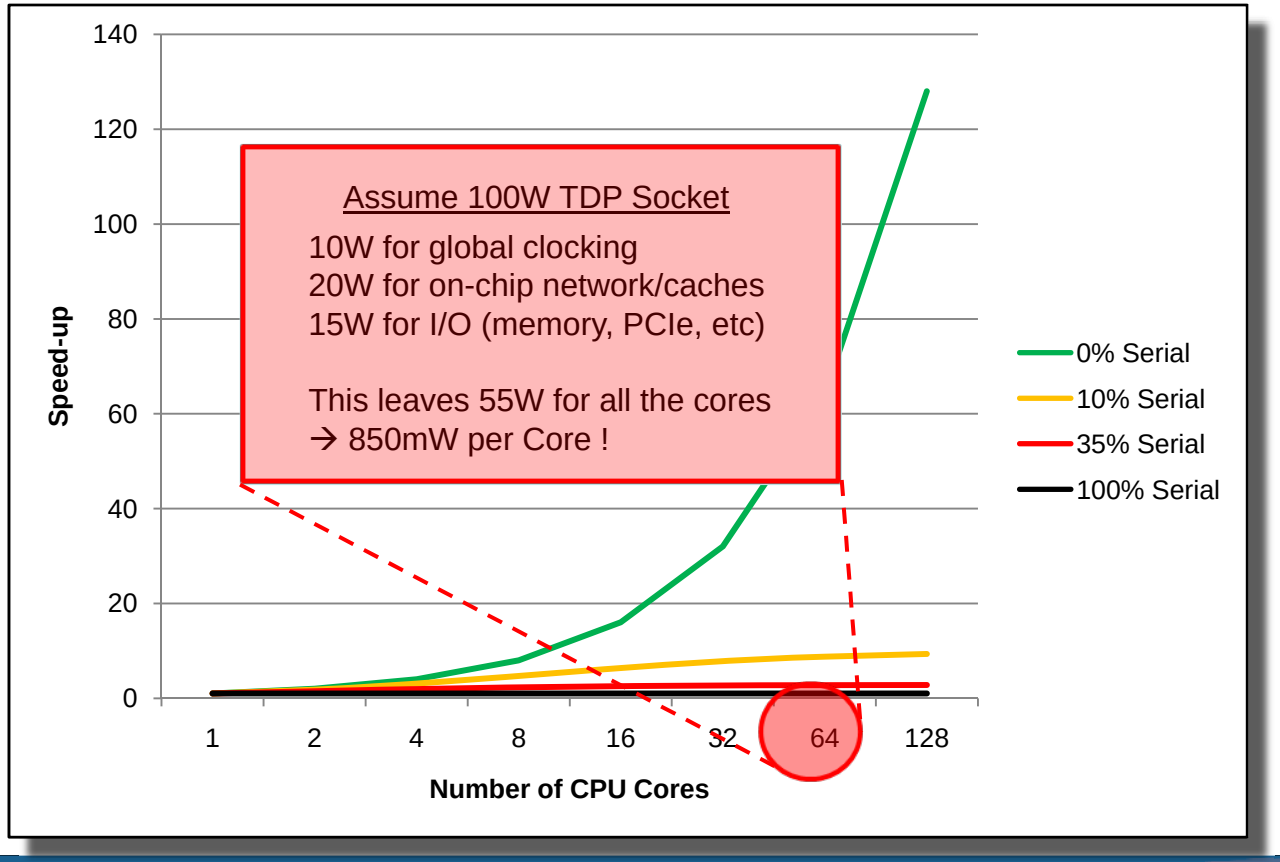

The Power Wall -- Implications

- Chip Multiprocessors (CMPs) will evolve to be heterogeneous
  - First, integration of cores with different capabilities
  - Second, integration of alternative programmable devices with superior performance/watt characteristics when running workloads of interest
  - Third, integration of extremely power efficient dedicated hardware assists for very commonly used functions
- Most meaningful metrics are (or will be) a ratio with power
  - Processor Cores and Chips: **Perf/Watt** and **Perf/Watt/$$**
  - GPUs & other data parallel throughput solutions: **FLOPS/Watt**
  - I/O SERDES PHYs: **mW/Gbit/sec**
- Very sophisticated next generation power management
  - Provision power based on real time monitoring and/or explicit requests
  - SOC components all on separable voltage and clock domains
  - Programmable uController for the base power management controller

Application Acceleration

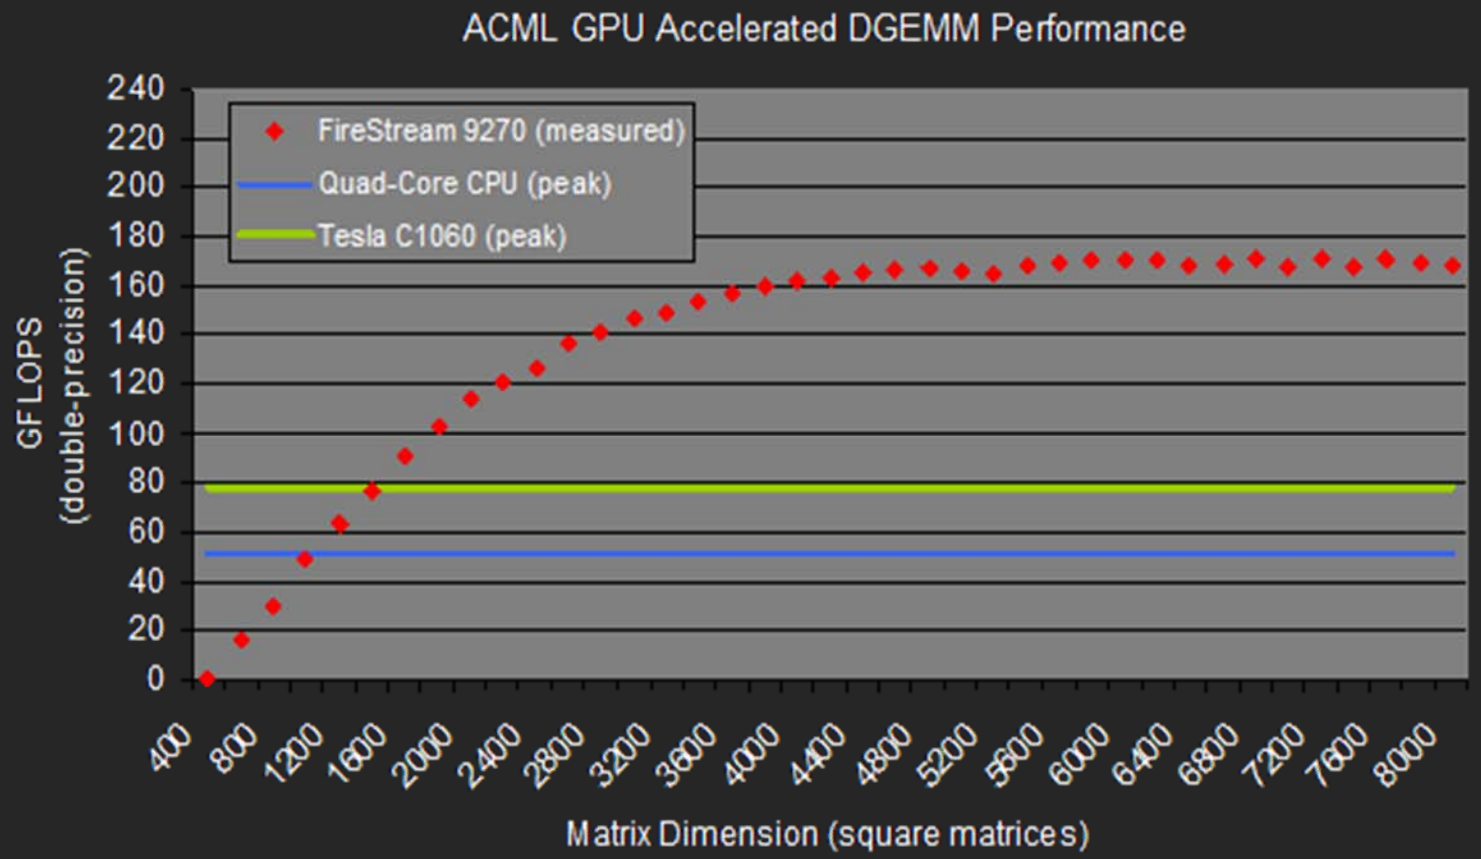

- AMD FireStream 9270 on AMD Phenom X4 9950/790FX/4GB DDR2 running RHEL 5.1 x86_64
- FireStream measured performance includes transfer of operand and result matrices
- Quad-Core peak theoretical performance quoted for 3.2GHz Nehalem processor
- C1060 peak performance derived from published specifications

GPU Performance

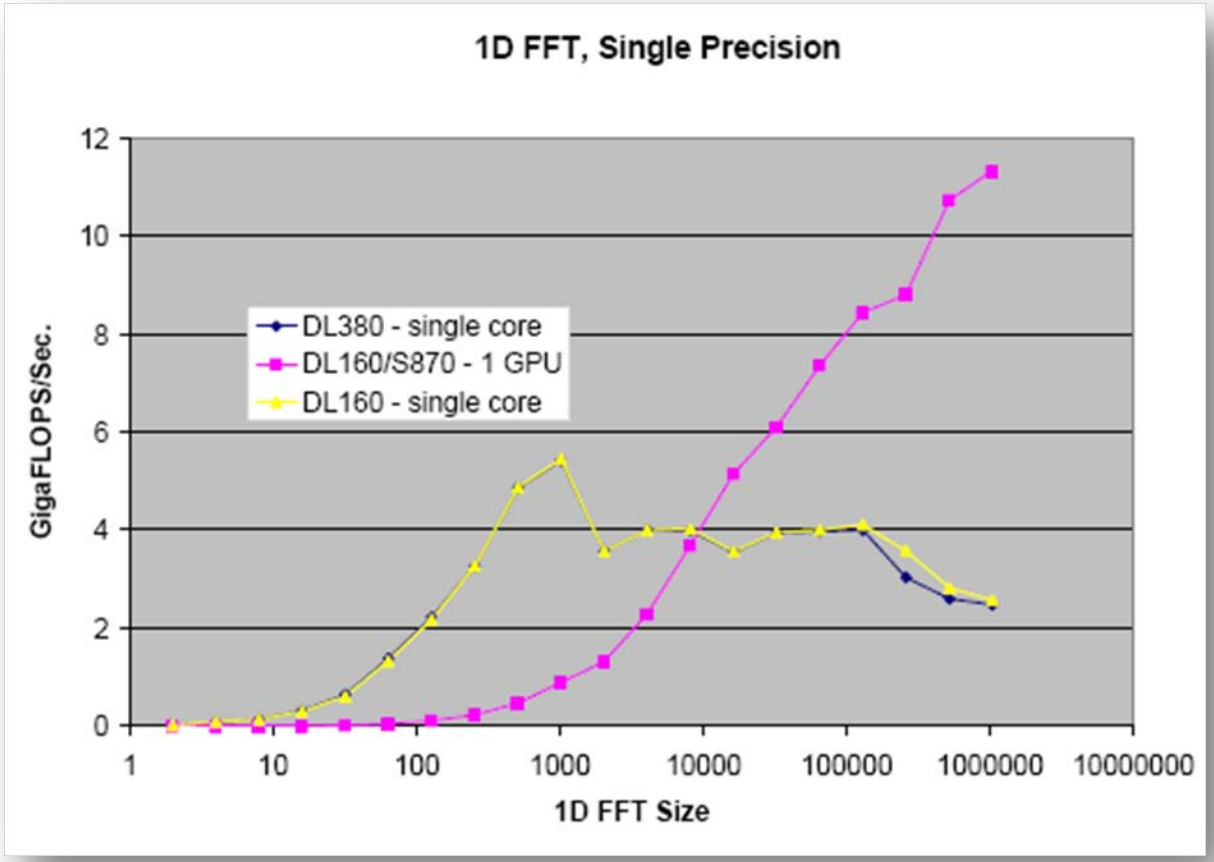

Accelerating HPC Using GPUs
G. Lupton, D. Thulin, HP

Benefits of many-core

- Moore's law points to more integration (more cores)
- Decent scalar performance
- Fine-grain parallelism (SSE)
- ~~Easy~~ Familiar programming model
- Fast intra-socket communication
- So what's the problem?

The Beast Is Hungry

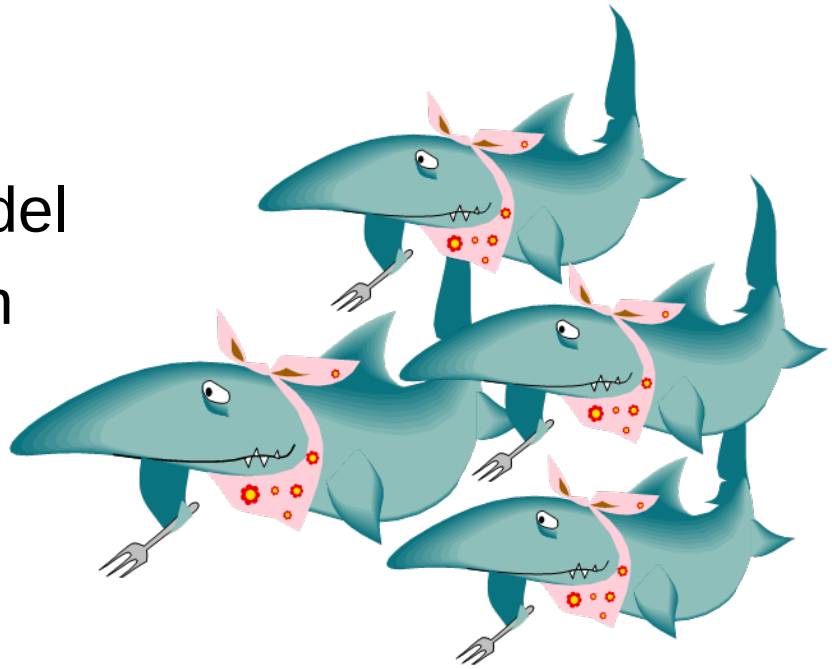

Feeding the Multicore Beast, Michael Perrone, IBM Master Inventor

DRAM & CPU Access Time Evolution

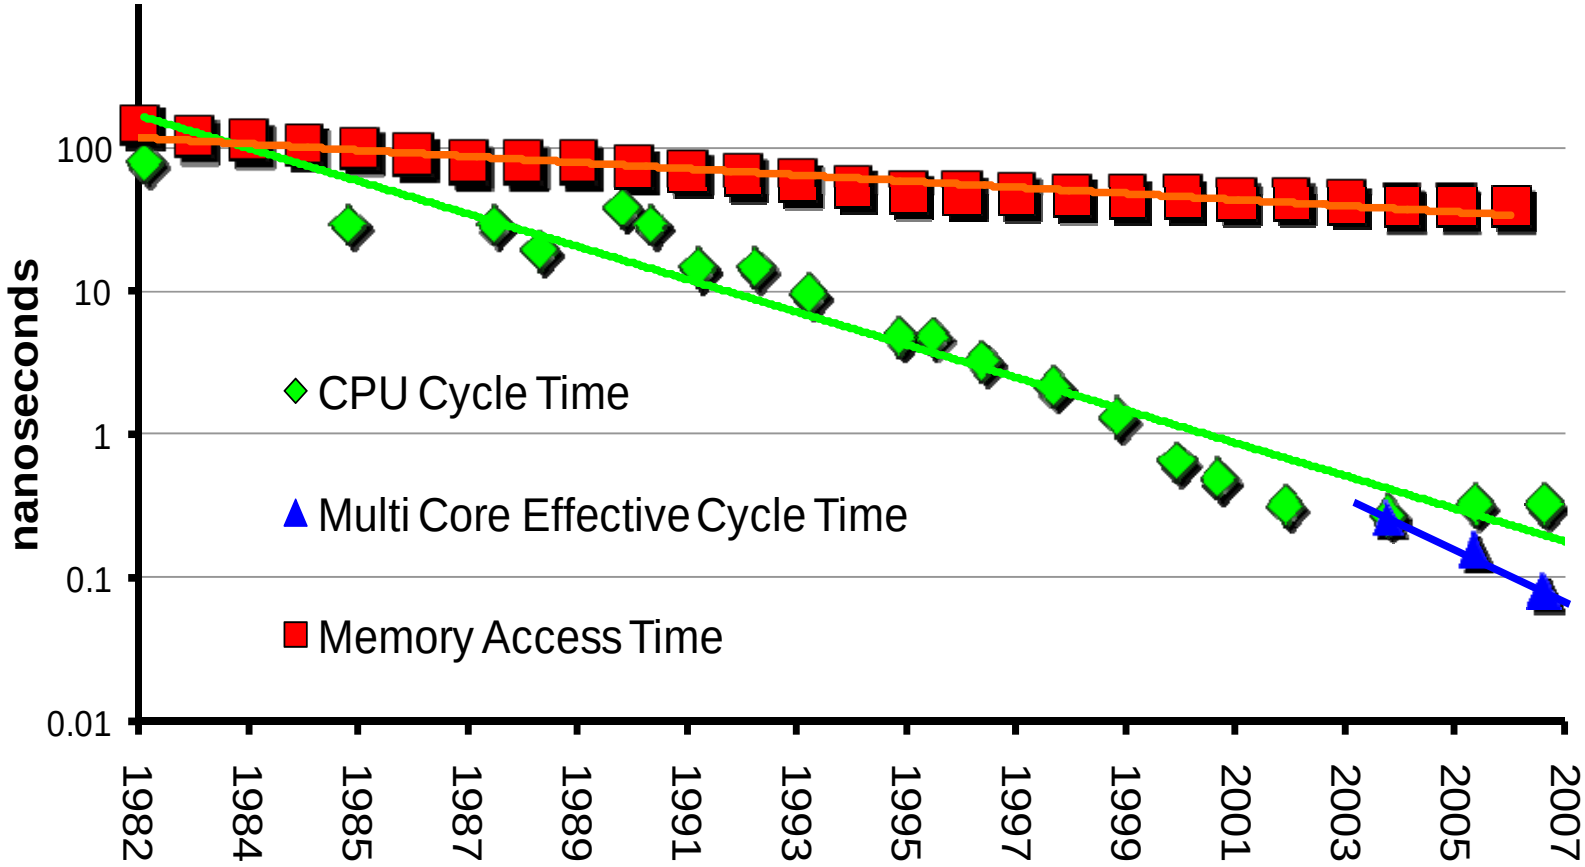

Source: Micron

Stacked Memory

- Several memory dies 1-16 can be placed under the CPU
- TSVs in the memory stack deliver power and external I/O to the CPU

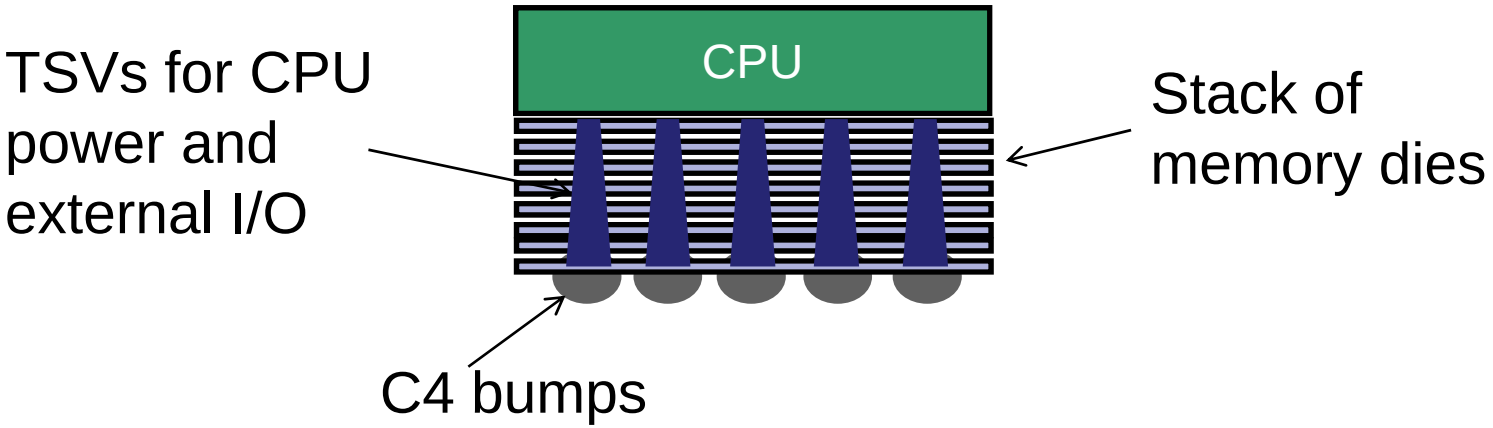

Interposer Top View

- Interposers look like MCMs but utilize Si with dense RDL to increase die to die BW from 100's of interconnect to 1000's
- Very large interposers could include 2 to 4 stacks of memory

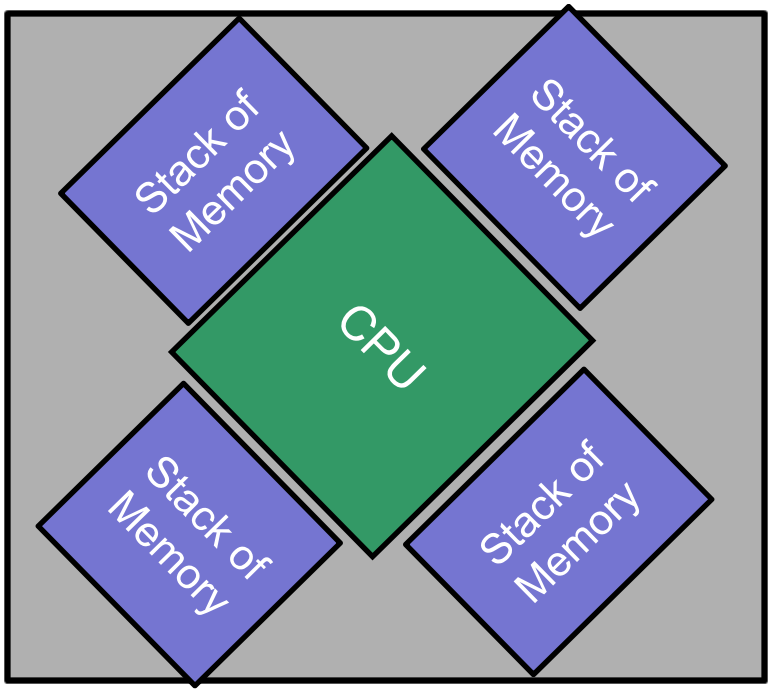

Si Interposer with several layers of RDL for dense signal routing

NIC & Integration

- Benefits:
  - Power savings
  - Latency reduction

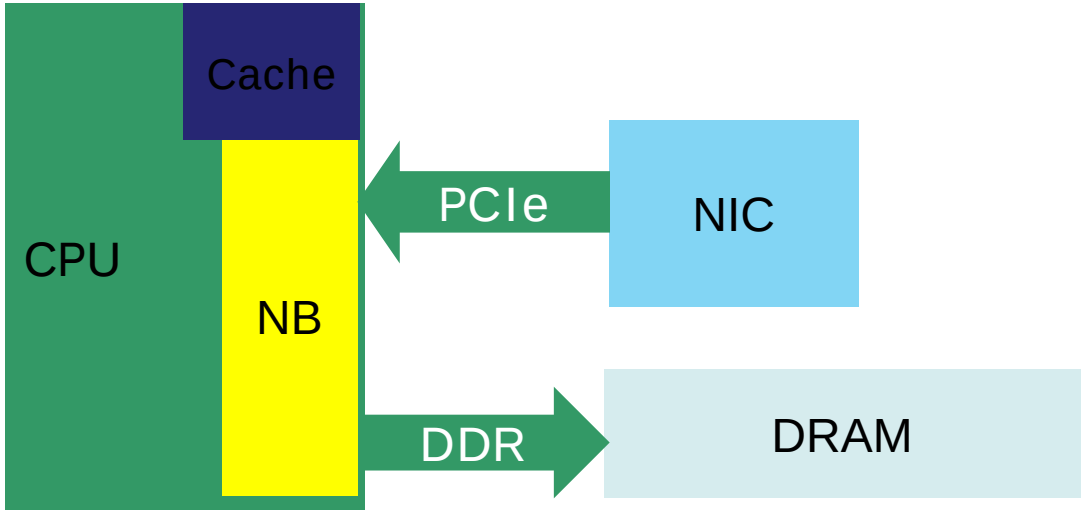

Interconnect Trends of Top500

- Ethernet is dominate today

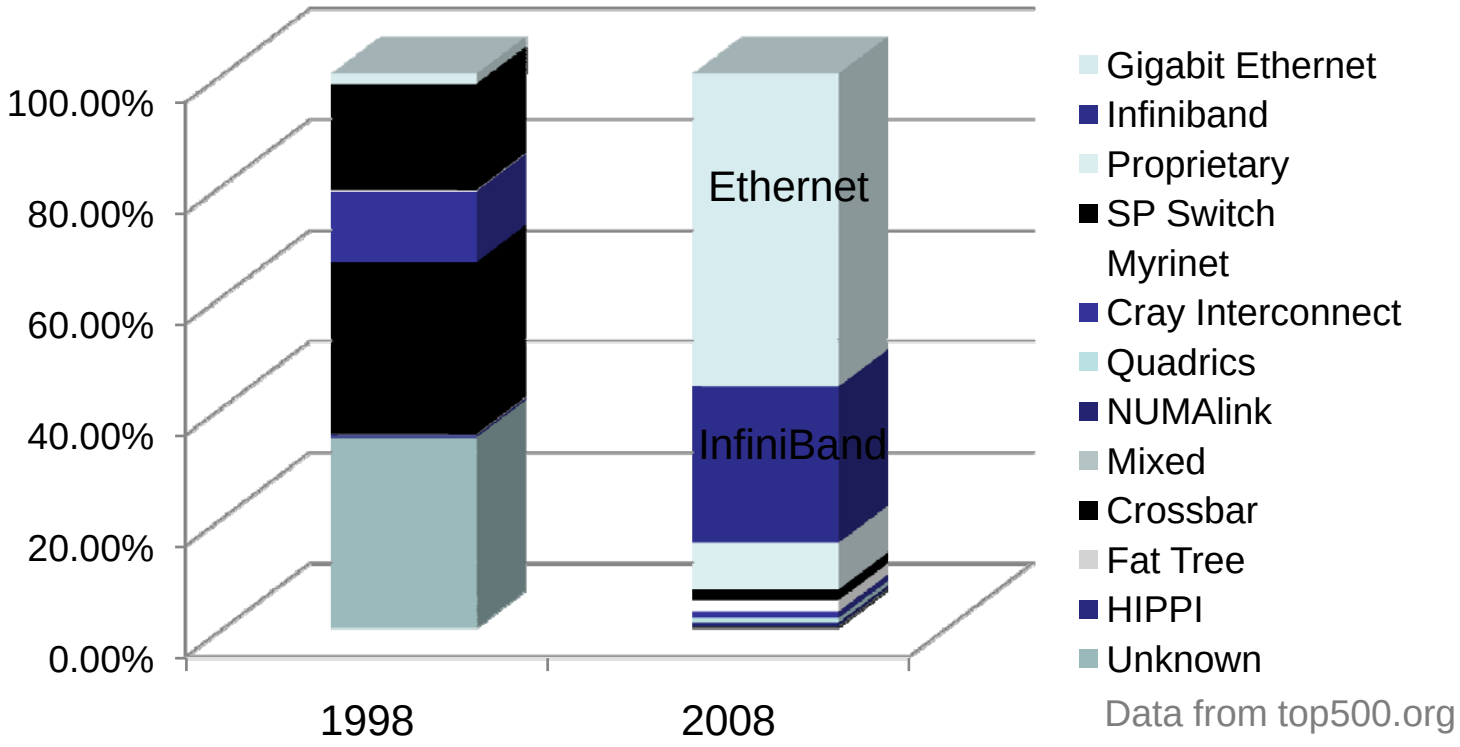

Ethernet Connection Trends

x86 Server Ethernet Connection Speeds with 40GbE & 100GbE

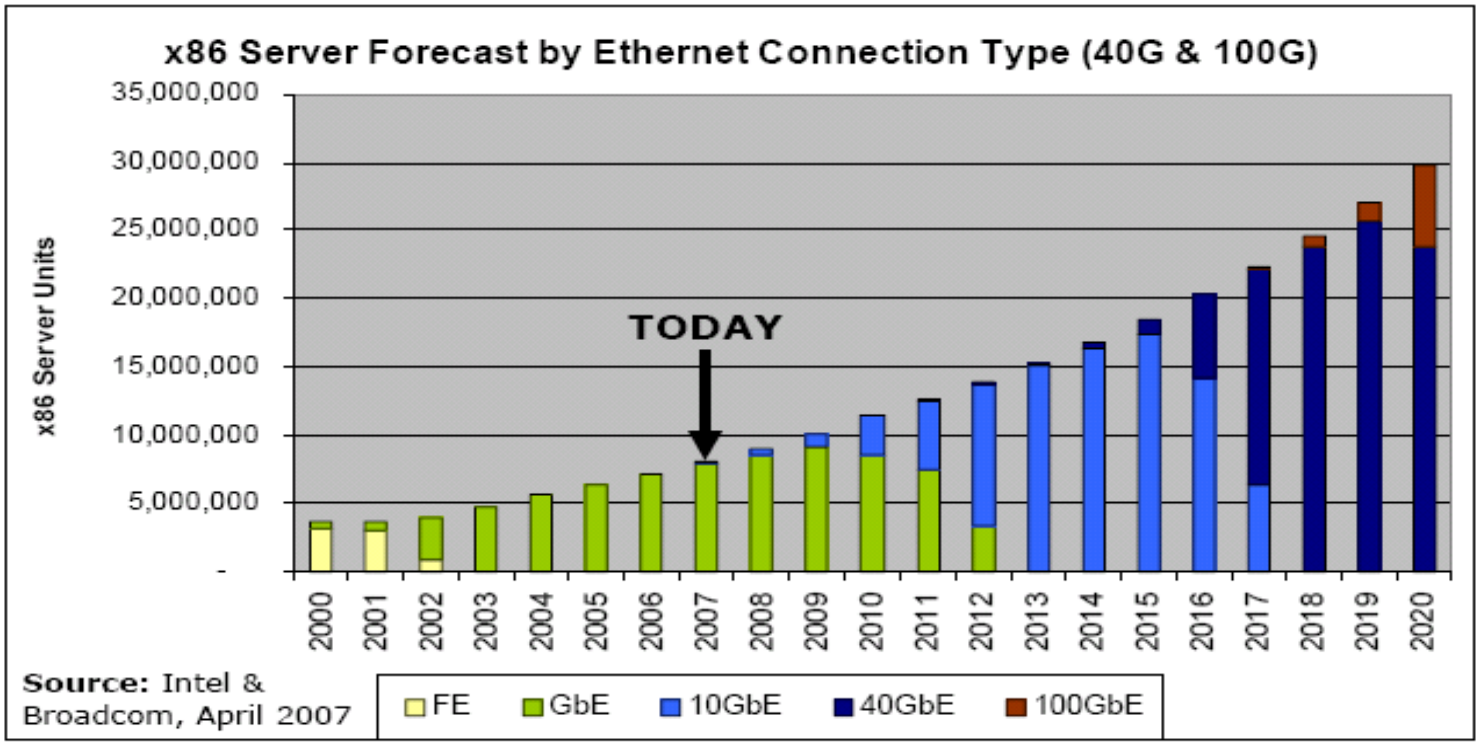

InfiniBand Trends

InfiniBand Link Speed Roadmap

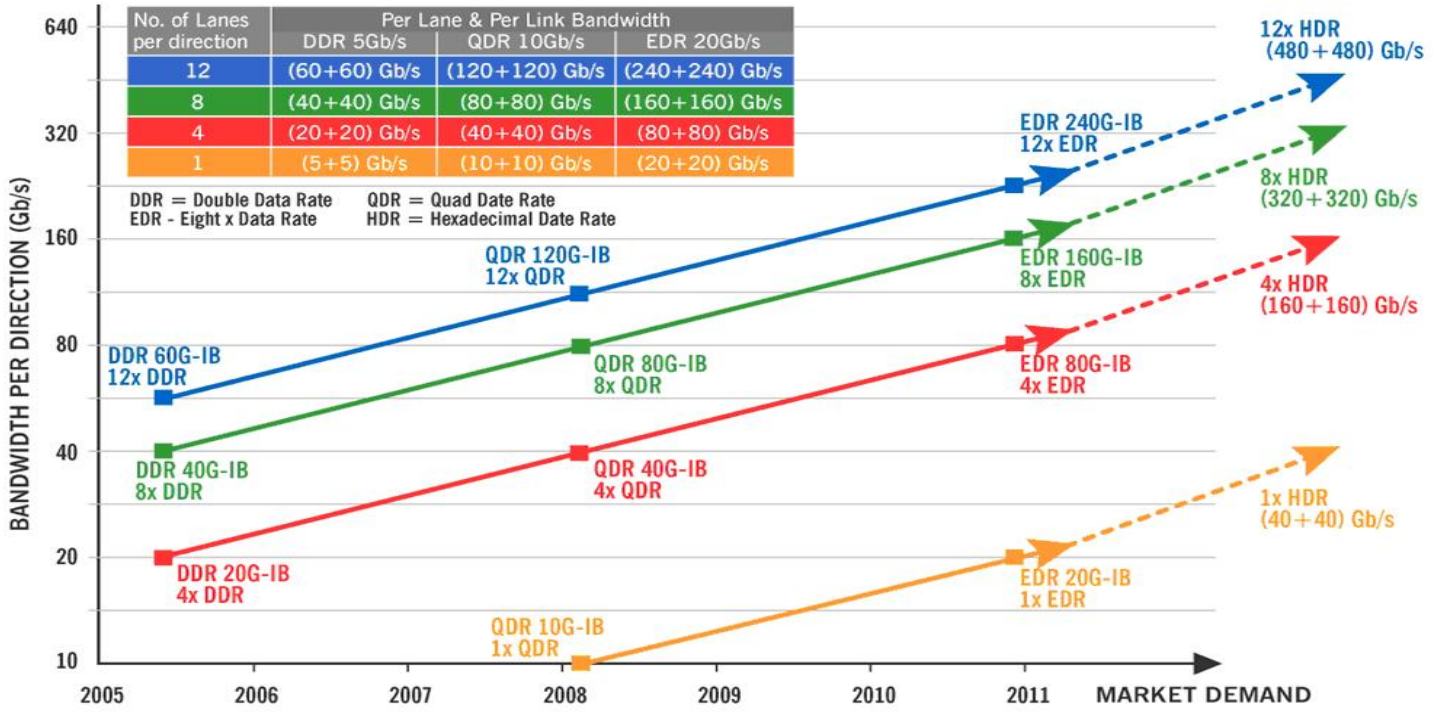

Future Single-Socket Server

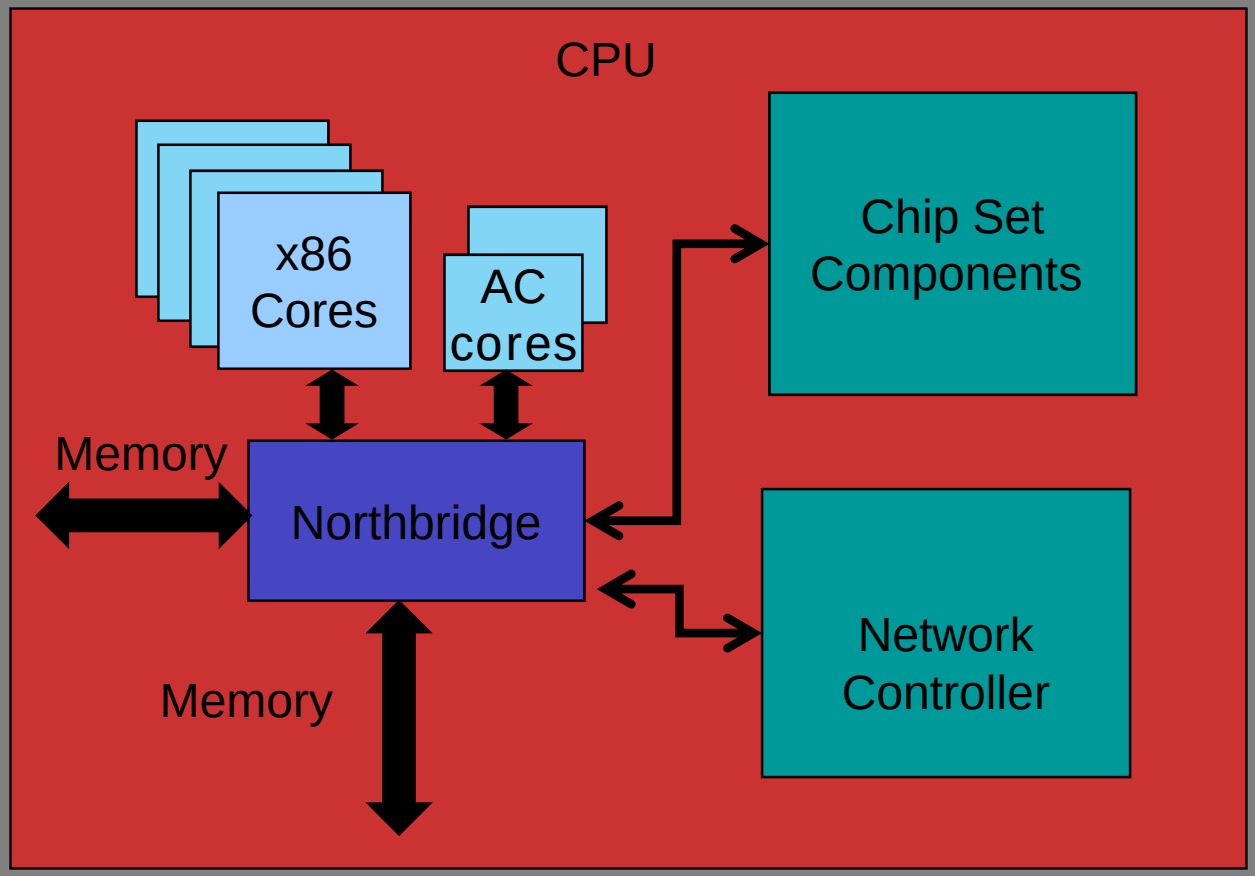

In Summary

- Commodity CPUs are driving down TCO
- Advances in the low end are relentless
  - “good enough” can be a disruptive technology
- Power efficiency is key
- Future CPUs will include small, power-efficient, domain optimized compute offload engines
- Integration such as stacked memory and NIC offer tremendous opportunities

Acknowledgements

- Chuck Moore
- Mike Uhler
- Mike O'Connor
- Lucas Kreger-Stickles
- Andy Kegel
- Steve Reinhardt
- Dave Christie
- Brad Beckmann
- Alan Lee
- Bobbie Manne
- Mike Houston
- Krishna Anne
- Dave Mayhew
- Patti Harrell
- Scot Schultz

Thank You!

Have a great Workshop!

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.

Trademark Attribution

AMD, the AMD Arrow logo, AMD CoolCore, AMD Opteron, AMD Virtualization, AMD-V, AMD PowerNow! , Dual Dynamic Power Management and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States and/or other jurisdictions. Other names used in this presentation are for identification purposes only and may be trademarks of their respective owners.