




Trends in High Performance Computing

Simon McIntosh-Smith
simonm@cs.bris.ac.uk
16th July 2009

Agenda



- Three trends in High Performance Computing
 - Processors
 - Interconnect
 - Storage
- Onwards and upwards – Exascale computing
- Wrap-up




Trends in processors

Trends in Processors

- Consolidation
 - Sun, SiCortex, ClearSpeed, ...
- Power consumption
 - Keeps going up!
 - HP predicting racks up to 50kW
 - Will require water cooling at multiple levels
- Quad core CPUs now the norm, six and eight cores now appearing
- Multiple heterogeneous systems now in Top500
 - Intel, AMD, IBM all heading in this direction





		Niagara-like core	Niagara-like core
Large core		Niagara-like core	Niagara-like core
Niagara-like core	Niagara-like core	Niagara-like core	Niagara-like core
Niagara-like core	Niagara-like core	Niagara-like core	Niagara-like core



Trends in Processors – cont.

- GPUs becoming more HPC-friendly
 - OpenCL, 64-bit capable, passive cooling, IEEE754
 - Systems vendors providing GPU-capable systems
- Trend towards 2U servers to reduce power consumption, cost and weight
- Form-fitting, mother-board sized heat-sinks with water cooling
- First signs of commodity computing being displaced by consumer computing
- Hardware faults will become more prevalent





Trends in interconnect

Trends in Interconnect

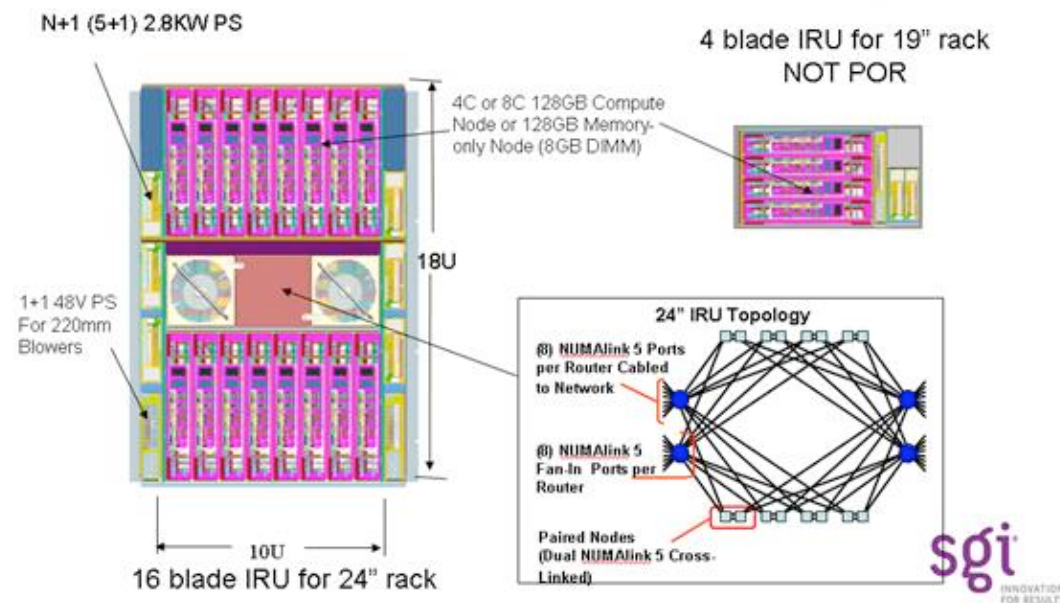
- Consolidation
 - Quadrics folding, proprietary interconnects disappearing
 - Top 100 mostly InfiniBand (40Gbps being deployed)
 - Bottom 400 mostly Gbit Ethernet
- Disruption?
 - Mass market is >> HPC and dominated by Ethernet
 - InfiniBand delivers better latency and bandwidth than Ethernet but at higher cost
 - Opportunity for disruption: better than IB latency and bandwidth via Ethernet?



www.gnodal.com

Trends in Interconnect

- Acceleration in the interconnect
 - E.g. sgi's UltraViolet technology to accelerate MPI



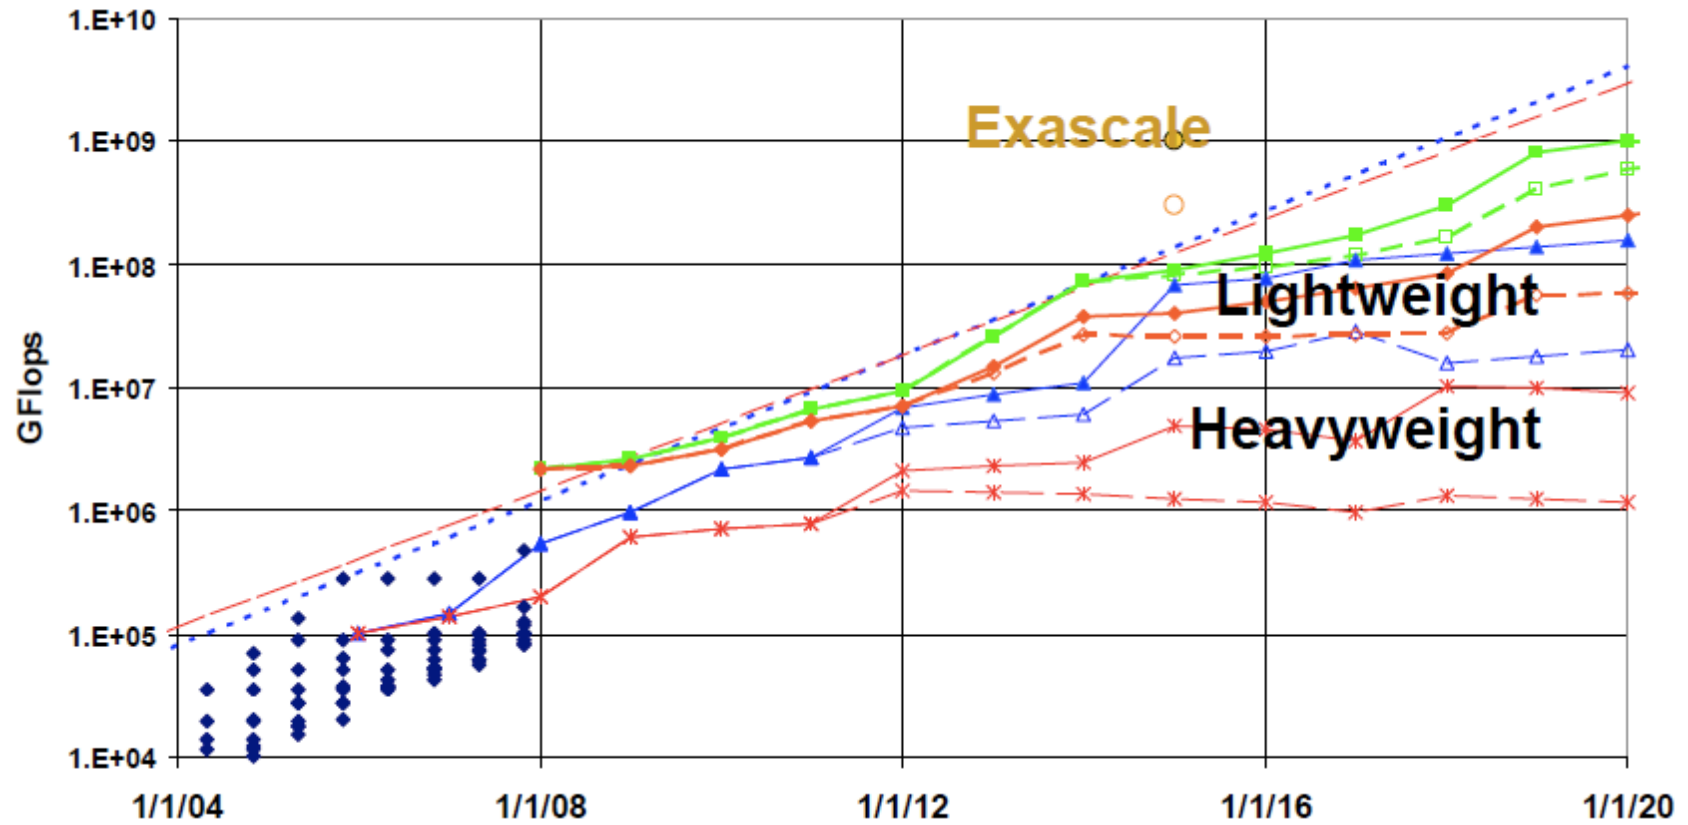
- But (current) MPI won't scale forever



Trends in storage


Trends in Storage

- Solid State Disks (SSDs) becoming ubiquitous
- Will supplement traditional hard drives in the short term
- Their greater reliability will be essential for future, large scale storage solutions
- Currently ~10X less capacity but ~3X lower power consumption and 4X smaller form factor
- ASPs on flash memory reducing by nearly 50% per year per GigaByte
- Random read performance already excellent – future SSDs may be more like DRAM than disks



The path to exaflops

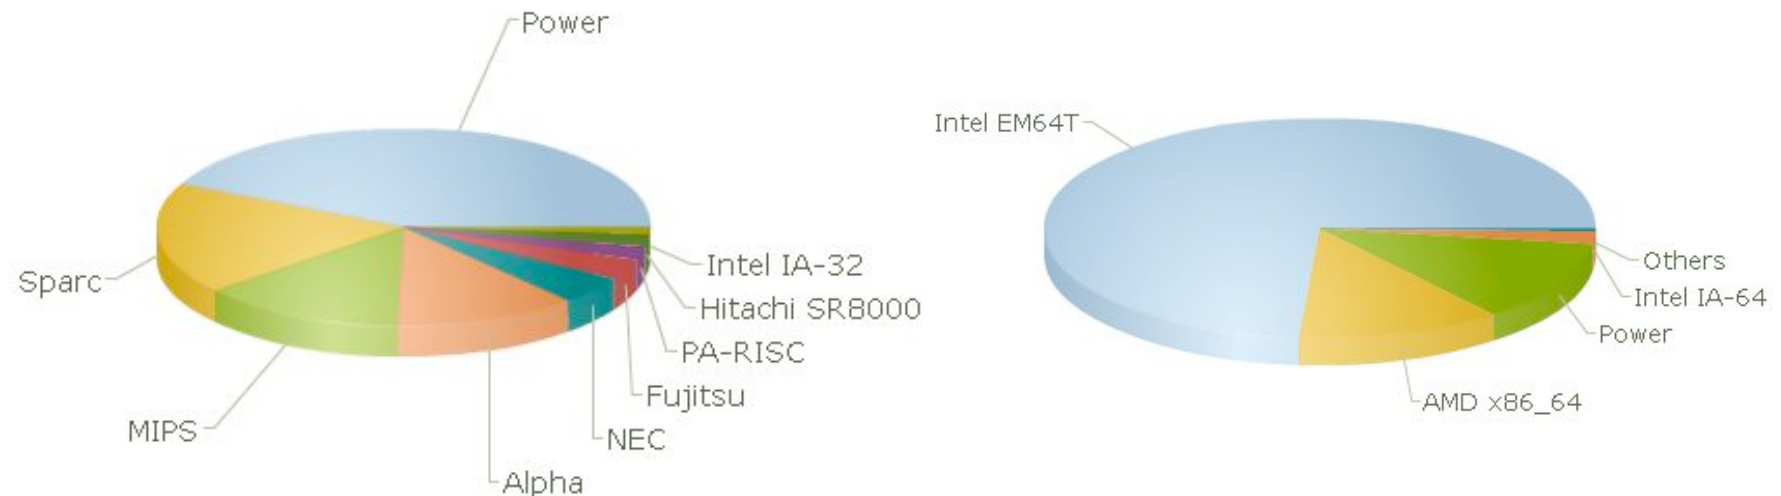
Performance Development



Commodity Clusters Revolutionised HPC

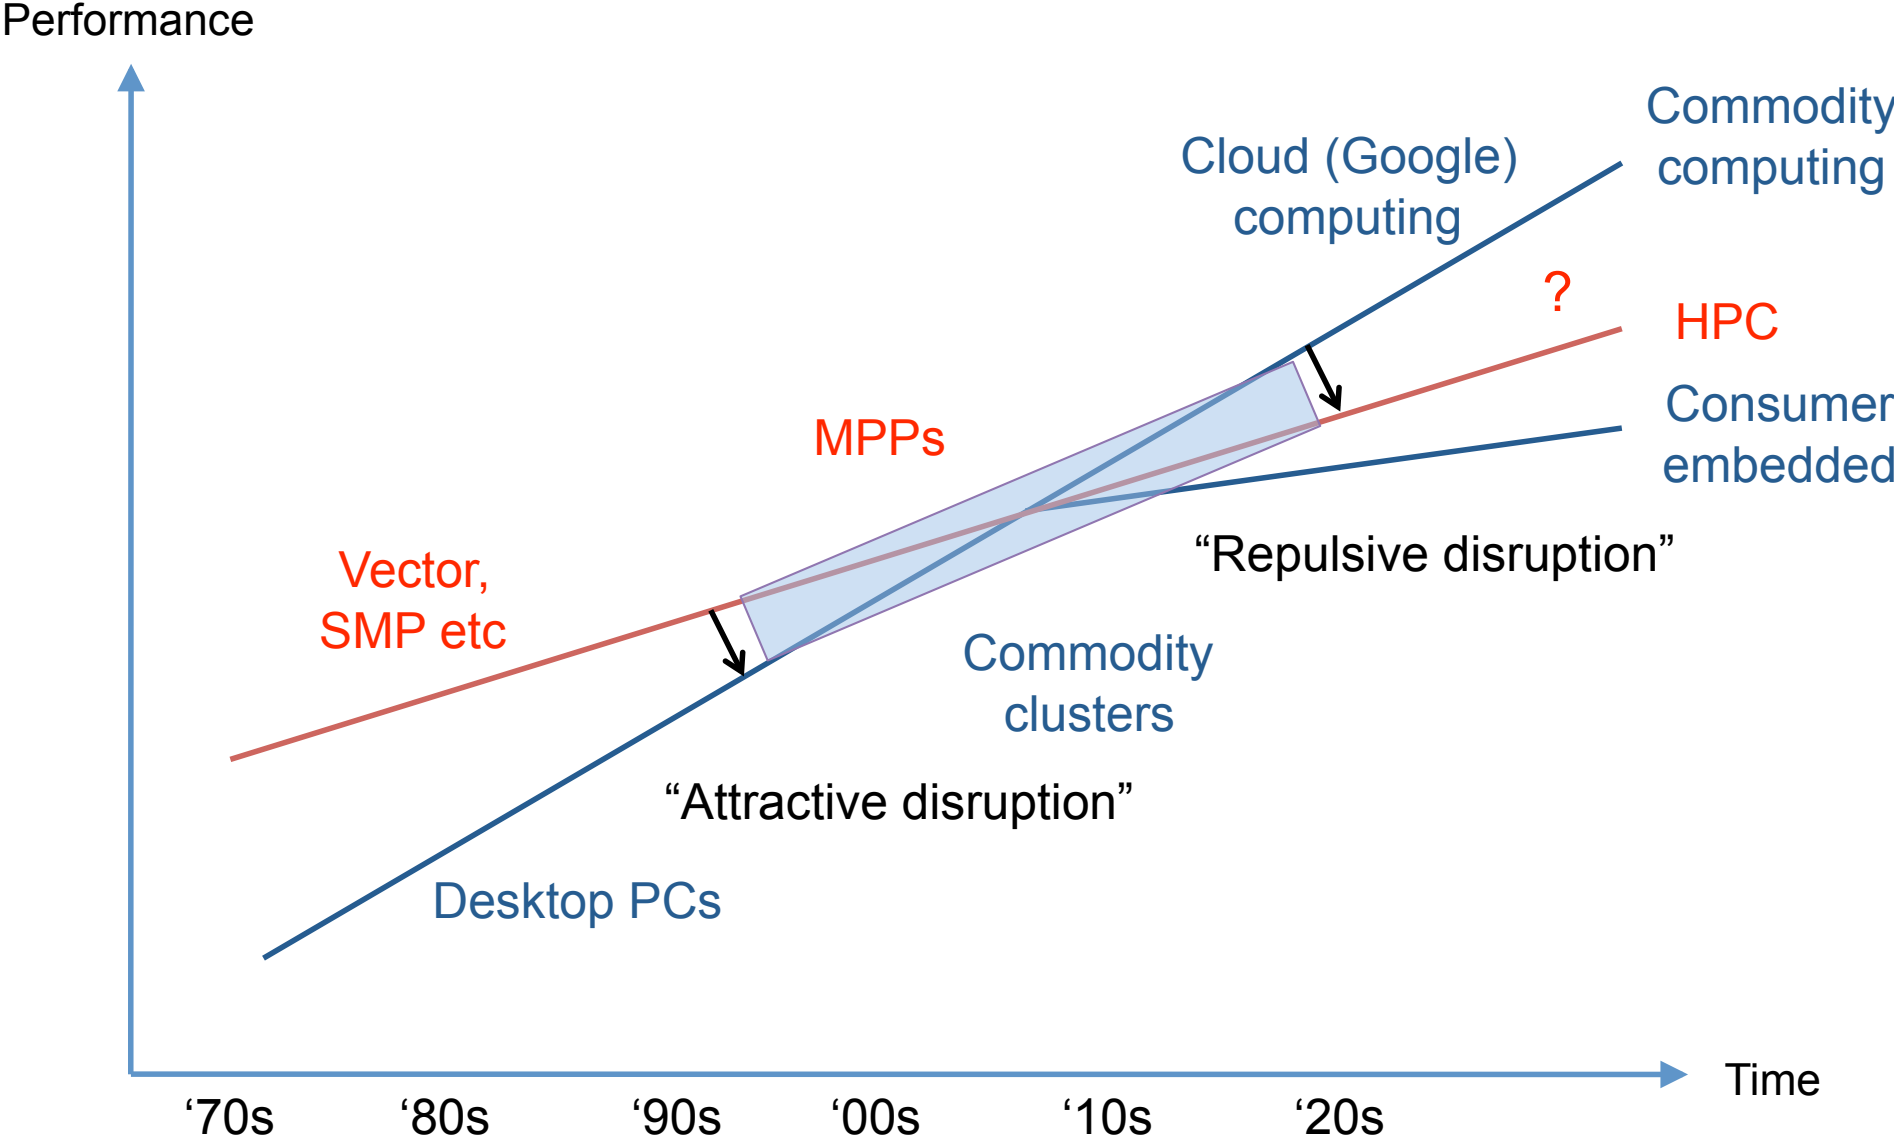
- HPC has been riding the commodity cluster gravy train since the start of the decade

Top500 systems by Processor Family



November 2000, x86 account for 1.2% of systems. By November 2008 this had grown to 85.8%.

A parting of the ways?




Onwards and upwards

The path to Exascale:

- General purpose to give way to specialist machines (again)?

- Grape-DR
- D.E. Shaw's Anton
- Cray, sgi, NEC, ...



- The predicted continued growth in data centre size and power consumption must surely hit a wall?
 - In the June 2009 Green Top500, average machine efficiency increased by 10% while average power consumption increased by 15%
 - Lowest power consumption prediction for first Exaflop machine is *50 MW!*

Thomas Sterling's canonical system

From the June 2009 Top500:

- Commodity Cluster
- Intel Xeon E54xx Harpertown
- Quad core
- 8192 cores
- 2048 sockets
- HP systems integrator
 - IBM a major second also dominates in overall performance
- 45.2 Teraflops performance Rmax
 - #91 on the Top-500 list
- 8 Terabytes main memory
- 1 Gigabytes/processor core
- Infiniband interconnect
 - Ethernet for system administration and maintenance
- MPICH-2
- OpenMP gaining in interest to address multicore
- Linux
- Power Consumption: 384 Kwatts
- Industry owned and run

Summary

- Processors: more cores, heterogeneity
- Interconnect: IB & Ethernet, time for disruption?
- Storage: SSDs, but where do they fit in the memory hierarchy?

HPC@Bristol

