

HPC Future Look

Exascale and Challenges

Dr Mark Bull, EPCC
markb@epcc.ed.ac.uk



Outline

- Future architectures
 - Exascale initiatives
 - Processors
 - Memory
 - Impacts on performance
- Software challenges
 - Parallelism and scaling
 - New algorithms
 - What about software that does not scale?
- Impact for standard computing

Future architectures

What will HPC machines look like?

What will future systems look like?

	2013	2016	2020
System Perf.	20 PFlops	100-200 PFlops	1 EFlops
Memory	1 PB	5 PB	10 PB
Node Perf.	200 GFlops	400 GFlops	1-10 TFlops
Concurrency	32	O(100)	O(1000)
Interconnect BW	40 GB/s	100 GB/s	200-400 GB/s
Nodes	100,000	500,000	O(Million)
I/O	2 TB/s	10 TB/s	20 TB/s
MTTI	Days	Days	O(1 Day)
Power	10 MW	10 MW	20 MW

Processors

- More Floating-Point compute power per processor
 - Only exploit this power via parallelism
 - Lots of low power compute elements combined in some way

Memory

- Will be packaged with processor
 - Increases power efficiency, speed and bandwidth...
 - ...at the cost of smaller memory per core

System on a chip

- Instead of separate:
 - Processor
 - Memory
 - Network interface
- Combined system package where all these things are included in one manufactured part
 - This is the only way to improve power efficiency
 - Less scope for customisation
 - If you need more memory than in package you will have to have levels of memory hierarchies

Software challenges

What does software need to do to exploit future HPC?

What does this mean for applications?

- The future of HPC (as for everyone else):
 - Lots of cores per node (CPU + co-processor)
 - Little memory per core
 - Lots of compute power per network interface
- The balance of compute to communication power and compute to memory are both radically different to now
- Must exploit parallelism at all levels
- Must exploit memory hierarchy efficiently

Algorithms

- For many problems new algorithms will be needed
- May not be optimal but contain more scope for parallelisation
- Mixed-precision will become more important

Applications that do not scale

- The good news is that if you do not need to be able to treat larger/more-complex problems then you can access more of current resource size
 - May be caught out by decrease in memory per core
 - Options to scale in trivial-parallel way: increase sampling, use more sophisticated statistical techniques
 - This may well be the best route for many simulations

Impact on standard computing

What does this mean for my workstation?

Parallel everywhere

- All current computers are parallel
 - From supercomputers all the way down to mobile phones
 - Most parallelism is task-based on 4-8 cores – each application (task) runs on an individual core.
- In the future:
 - More parallelism per device – 10s to 100s cores running at lower clock speeds
 - All applications will have to be parallel
 - Parallel programming skills will be required for all application development.
- More system on a chip – more things will be packaged together

Summary