

Imperial College
London

Imperial College London

Simon Burbidge 13 April 2016

Imperial College London

- Premier UK University and research institution ranked #2= (with Cambridge) in QS World University rankings (MIT #1)
- #9 in worldwide Universities (THES) and #3 in Europe
- 15,000 students, 6,200 of them postgrad
- 7,200 staff (2,600 of which dedicated to research)
- Science based – engineering/technology/medicine
- Lots of industrial research collaboration
- £350M per year research income

Main Campus South Kensington London

- Located on the site of the 1851 Great Exhibition in London (pic wikipedia Dickinsons)



HPC Service

- University wide service
- £2M annual budget for hardware and software
- Additional income directly from research groups
- 5 dedicated HPC staff
- Largest UK University HPC system (HPC-SIG)

3 HPC systems

- cx1 – low end cluster mostly ethernet
- cx2 – high end parallel MPI cluster (SGI ICE-X)
- ax3/4 – large shared memory dedicated to genomics (UV)

ax3/4

- Large memory SMP
- Two SGI UV
 - 384 processors 5 TB memory
 - 1024 processors 16 TB memory
- Main user is Bioscience – Next Generation Gene Sequencing
- PBS topology aware scheduling, cpusets

•CX2

- High end MPP commissioned Autumn 09, upgraded twice, system refresh Winter 2015, supporting large highly parallel jobs
- SGI Altix ICE 8400 EX + ICE-X 12288cores
- Dedicated to large MPI parallel jobs
- PBS topology aware scheduling – simple config
- Focus on capability, running jobs on thousands of cores

cx1

- Low end cluster
- Mostly gigabit ethernet, some infiniband islands
- Upgrading all the time, parts of the system owned by particular research groups
- 23026 cores, 1352 nodes, 12 distinct hardware types plus GPUs
- Serial and small parallel workload many shapes and sizes
- Tricky to schedule! 3 million jobs per year ...
- Focus on low cost/throughput

Why *cx1 and cx2?*

- Cost
- Flexibility
- Users able to contribute funds/hardware to *cx1*

10 years of Service

- Celebrating 10 years of successful service this summer
- Where to next?

More of the same

- For sure – users like what they get!
- They want more
- We have rolling replacement programmes
- Finance and space constraints
 - Users' grants may last for 5 years, which is a bit long in the tooth for a server, but they have no more funds in the 5 year period to refresh
 - Central funds are fixed and can't replace older hardware to a 3 year lifespan
- Space
 - Computer room is full and limits expansion

Expand?

- Demand is there for expansion, more resources for current users plus newer areas, like bioscience and social science are growing fast
- Additional demand for higher fidelity simulations from traditional computational science – larger molecules, more degrees of freedom, smaller timesteps, higher Reynolds number, more complex algorithms etc
- National resources constrained
- Good cases for expansion and likely funding – but how?

Possible Options?

- New computer room
 - Where, how big
 - Running costs (actually could be quite reasonable)
- Co-Lo's
 - Not really well suited for HPC hardware
 - Prices not fixed – market driven?
 - Readily available at the moment
- Cloud
 - For some work cloud gives good results, for large scale parallel simulations it doesn't
 - Can be pricey and costs are variable – hard to fund
- Stay tuned

Software Challenges

- Its not just hardware though, should we invest in software?
- New processors have more and more cores/threads running more slowly
- These processors need carefully optimised code to reach their rated FLOP rates
 - 3 levels of optimisation needed
 - At the thread level (vectorisation)
 - At the socket level (on chip communications, eg OpenMP)
 - Between nodes (via the interconnect, eg MP)
- Memory access times becoming more problematic
 - More NUMA levels:
 - Registers/vector registers/cache
 - On socket cache
 - Off socket cache and user controlled buffers
 - Main memory

Phew!

- Although some of this optimisation can be automated, most needs careful programming
- There is a real shortage of computational scientists who can do this work
- Its not fashionable to be a computational scientist in a research group
- Programming skills not being taught at lower levels
- Computing skills not regarded as worthwhile (not taught) in mainstream science and engineering degrees
- Skills shortage is a huge issue
- More training and recognition of programming and computation skills is badly needed across the board – the next generation of HPC won't succeed without it.

HPC is fragile!

- HPC is
 - Disruptive
 - Challenging
 - Expensive
 - Changeable
- It needs a nurturing ecosystem in order to succeed
- Universities need to take care not to break it
 - Long term sustainability needed
 - Computer room infrastructure (15-30 year)
 - It is not IT, it is not commodity, HPC is special and needs a special place of its own in the organisation. Essentially it's a research instrument
- Industrial collaboration and exchange is vital
 - Working with the real world focusses on real challenges and issues.
 - HPC could help in more areas than now by providing a safe place to conduct digital experiments, cheaply and effectively. It needs to be driven!
- Specialist Programming skills need to be fostered, encouraged and rewarded.

Thanks

Questions?

Simon Burbidge

s.burbidge@imperial.ac.uk