

Oxford Supercomputing Centre: the Next Generation

Prof. Mike Giles

`giles@comlab.ox.ac.uk`

Chairman, OSC Management Committee

Brief History

OSC was established in 1998:

- inter-departmental consortium
- £1.4M external funding from JREI
- OSCAR: 96-processor shared-memory SGI
- provided a unique facility in Oxford for large computations

Brief History

Second generation arrived in 2002:

- £1.5M external funding from Wellcome Trust and HEFCE
- OSWALD: an IBM cluster with 128 processors and high-performance networking
- OSWELL: 4 shared-memory SUN's with $24+24+24+12=84$ processors
- provided a very useful high-end computing resource for large computations

Current Status

- OSWALD and OSWALD replaced early/mid 2005 by a “stop-gap” solution with remnants of Wellcome funding
- ZUSE: 96-processor Dell cluster
- KONRAD: three shared-memory machines each with 16 Opteron “cores”

2.5 centre staff (bare minimum!):

- Jon Lockley: centre manager / senior systems manager
- Jo Shields: systems manager
- Joe Pitt-Francis: 50% scientific computing advisor (other 50% working on an eScience project)

Current Usage

- Research groups from almost every department in MPS, and many others across the university;
- Major users include:
 - Dr. Barnard (Materials)
 - Dr. Bradley (Physiology)
 - Prof. Green (Chemistry)
 - Prof. Noble (Social Policy and Social Work)
 - Dr. Petrinic (Engineering)
 - Prof. Sansom (Biophysics)
 - Prof. Yeomans (Physics)
- Over 250 journal and conference publications have acknowledged use of OSC facilities in the past 7 years.

Management

- Decisions on operational policy and major purchases are made by a Management Committee of users;
- OSC currently reports to the IT Research Committee chaired by the pro-VC (Research);
- Updates also given to some divisional ICT panels/committees;
- Operationally, OSC and IeRC are separate entities at present, but will probably become more linked in the future.

Rationale for Continuation

The benefits of centralised HPC are:

- exploits economies of scale in operating a very large system;
- provides rapid-start computing facilities to new academics;
- relieves established academics of burden of managing their own systems;
- relieves post-docs/students of burdens of running systems instead of doing research;
- promotes a community of computational scientists, engineers and others.

The need for HPC was also identified in the Strategic Plan.

The Next Generation

University has endorsed a new long-term vision for OSC, backed initially by SRIF3 funding:

- £0.6M for new computer room
- £1.6M for new hardware/software in 2006
- £0.8M for new hardware/software in 2007
- £0.8M/yr subsequently to replace 1/3 of the system each year (subject to ICT Strategy approval)
- staff costs currently underwritten until 2008/9

New Computer Room

- Estates Department is responsible for creating a new computer room in the Denys Wilkinson building;
- cost is now estimated at £1.3M with shortfall to be covered by university;
- entry date has slipped – now aiming for June 1st this year, but could slip further.

New Hardware

- 1600-2000 cores in distributed-memory clusters
 - probably 1U “pizza boxes” with 2 dual-core Opterons and 8GB memory
 - possibility of blade servers
 - probably gigE networking for most, but maybe some Myrinet/Infiniband for highest performance
 - Linux o/s
- 160-200 cores in shared-memory systems
 - probably 8 dual-core Opterons with 32GB memory
 - probably Linux o/s but not certain
- 20-40TB of disk space
- an extra 50% in 2007

Future Budget

- £800k for new equipment
- £400k for running costs (esp. electricity)
- £160k for staff

This is being included in the ICT strategy review.

Unresolved issues: I

How do we pay for all of this?

- departments should recover full costs through fEC accounting, but not necessarily pass all of the money to OSC;
- one suggestion is for user charges to pay for running costs and staff, but SRIF/QA money pays for equipment
- there may also be a hybrid model for fEC charges
 - direct charge for major users (“capability”)
 - indirect charge for minor users (“capacity”)
- the university is now keen to get this resolved.

Unresolved issues: II

How do we support a greatly expanded user community?

- we expect many more users than we have had in the past when it was largely restricted to those who secured the funding
- OSC staff will not be able to cope without more staff
- it is maybe better to encourage support from local IT staff?

Unresolved issues: III

What is the relationship with leRC/OxGrid?

- leRC is primarily a research entity; OxGrid is providing a service element
- OxGrid's cycle-scavenging is good for certain applications (“trivially parallel”)
- off-loading these from OSC will improve OSC's service for those needing its high-performance networking
- OSC is being integrated into OxGrid so that registered OSC users have the option of using OSC via OxGrid

Unresolved issues: IV

Should OSC take the lead on scientific software licensing for the university?

- OUCS runs a license server for Matlab, with charge-back to departments
- OSC runs a license server for Intel/PGI compilers (paid out of our budget) and certain application software (charged to users)
- various departments manage site-licensed software with charge-back to users
- maybe we should coordinate all of this, and pay for a lot of it to avoid charge-back complexities?
- have discussed this with MPS division – need to discuss with MSD

Unresolved issues: V

How is OSC kept responsive to the needs of the university community?

- should the Management Committee be enlarged to have representatives from more departments?
- would this make it unwieldy?
- should we have a few people, each representing a division rather than a department?
- would this give us both management and technical expertise/advice?
- which body will set our budget? how will users say if it's too high or too low?

Unresolved issues: VI

Should academics be discouraged from setting up their own small clusters?

- no central directive – academic autonomy is important
- but why would anyone want to set up their own cluster?
- key is that OSC must demonstrate its ability to provide a good responsive service, and university must demonstrate its commitment to maintaining it.

Concluding remarks

- OSC is going through a major change
- vast expansion of our computing capability and our user base
- significant shift from external to internal funding . . .
- . . . and from a “rogue” self-funding operation to a central university resource.

Any questions?