

HPC for Scientific Computing

Dr Matt Ismail
Centre for Scientific Computing



Talk overview

- Introduction to CSC
- Overview of HPC facilities
- Minerva

What is the CSC?

- nurtures computationally-driven research in scientific computing
- national teaching and training centre
- provides a central HPC service


CSC as a research centre

- Core research
 - fundamentals
 - molecular dynamics and modelling
 - computational engineering and fluids
 - Monte Carlo and stochastic simulation
 - quantum simulations
 - computation of living systems
- 20 core academic staff from 7 departments



Molecular Dynamics and Modelling


- MD concerns particle motion inherent to natural processes such as molecular vibration or bond stretching and bending
- Disciplines include Chemistry, Physics and Mathematics
- Example projects are e.g. The biological interface with materials or differential equations with random initial data



Periodic boundaries for molecular simulations

Computational Systems Biology


- Systems biology studies complexity in biological systems e.g. gene expression, cell signalling or biological image analysis
- Disciplines include Mathematics, Computer Science, Physics
- Example projects are electronic transport in DNA and molecular simulation and modelling in the life sciences



Transport and reaction
simulation in cells

Computational Fluid Dynamics


- CFD concerns the simulation of fluid flow usually in the presence of some complication such as e.g. turbulence, chemical reaction or mass transfer
- Physics, Mathematics and Engineering
- Example projects include simulation of magnetohydrodynamics (MHD) and turbulence simulation




Adaptive mesh high resolution
image of Rayleigh-Taylor
unstable flows

Teaching and Training

- National EPSRC Training Centre
- Taught MSc in Scientific Computing
 - High Performance Scientific Computing
 - Computational Linear Algebra and Optimisation
 - Data structure and algorithms
 - Monte Carlo Methods
 - Computational PDEs
- MSc by research or PhD in Scientific Computing




MSc & PhD Degrees
Centre for Scientific Computing
Molecular Dynamics and Modelling
Monte Carlo Methods
Computational PDEs
Quantum Simulations
Computational Fluid Dynamics
Computational Systems Biology



Underlying the research goals of the Centre for Scientific Computing (CSC) is a graduate level educational and training program including a taught MSc in Scientific Computing, as well as research programs at both the PhD and MSc level. The taught MSc course is part of an EPSRC-funded Training Centre in High-End Computing. Students will work a state-of-the-art High-Performance Computing environment throughout the course. The CSC is built around a core of researchers from several Science Departments and fosters collaborative and interdisciplinary research across the whole University.

Centre for Scientific Computing, University of Warwick,
Gibbet Hill Road, Coventry, CV4 7AL, UK
Phone: +44 (0) 2476 732411 | Fax: +44 (0) 2476 737123
E-mail: csc-admissions@warwick.ac.uk



www.warwick.ac.uk/go/CSC

Interacting with CSC

- CSC@Lunch
 - Mondays, 13:00 – 14:00
- Coffee and Cookies
 - Fridays, 14:00 – 15:00



Talk overview

- Introduction to CSC
- **Overview of HPC facilities**
- Minerva

Branscomb Pyramid¹




¹International Review of Research Using HPC in the UK, December 2005, EPSRC

HPC kit

- HPC environments
 - Distributed memory cluster for parallel computing
 - includes many core processors
 - includes SMP nodes
 - Cluster of Workstations for task farming


Distributed Memory Cluster



Distributed Memory Cluster

- “Minerva” installed early 2011
- “Francesca” installed 2007
 - 960 Intel “Woodcrest” 3GHz cores / 1.92 TB RAM / low latency InfiniBand / Linux / water-cooled / GPFS parallel file system
 - now end of life

SMP




SMP

- “skua” installed December 2004
 - SGI Altix
 - 56-way single system image / 112GB RAM / Intel Itanium2 1.6 GHz processors / Linux
 - now reached end of life
- Minerva has a SMP component



Cluster of Workstations



Cluster of Workstations

- Intended for Task Farming
- Layered on top of Linux workstations
- Utilises “free” CPU cycles for computation
- ~600 heterogeneous cores

Architecture Comparison

- Distributed memory cluster
 - harder to program for because of need to use MPI
 - scalable and portable
 - Commodity clusters are cost effective
- SMP
 - OpenMP easier and compiler can do some of the work
 - less scalable and less portable than MPI
 - SMP machines are expensive
- COW
 - data parallel applications
 - visualization
 - code development and testing

Talk overview


- Introduction to CSC
- Overview of HPC facilities
- **Minerva**

Minerva

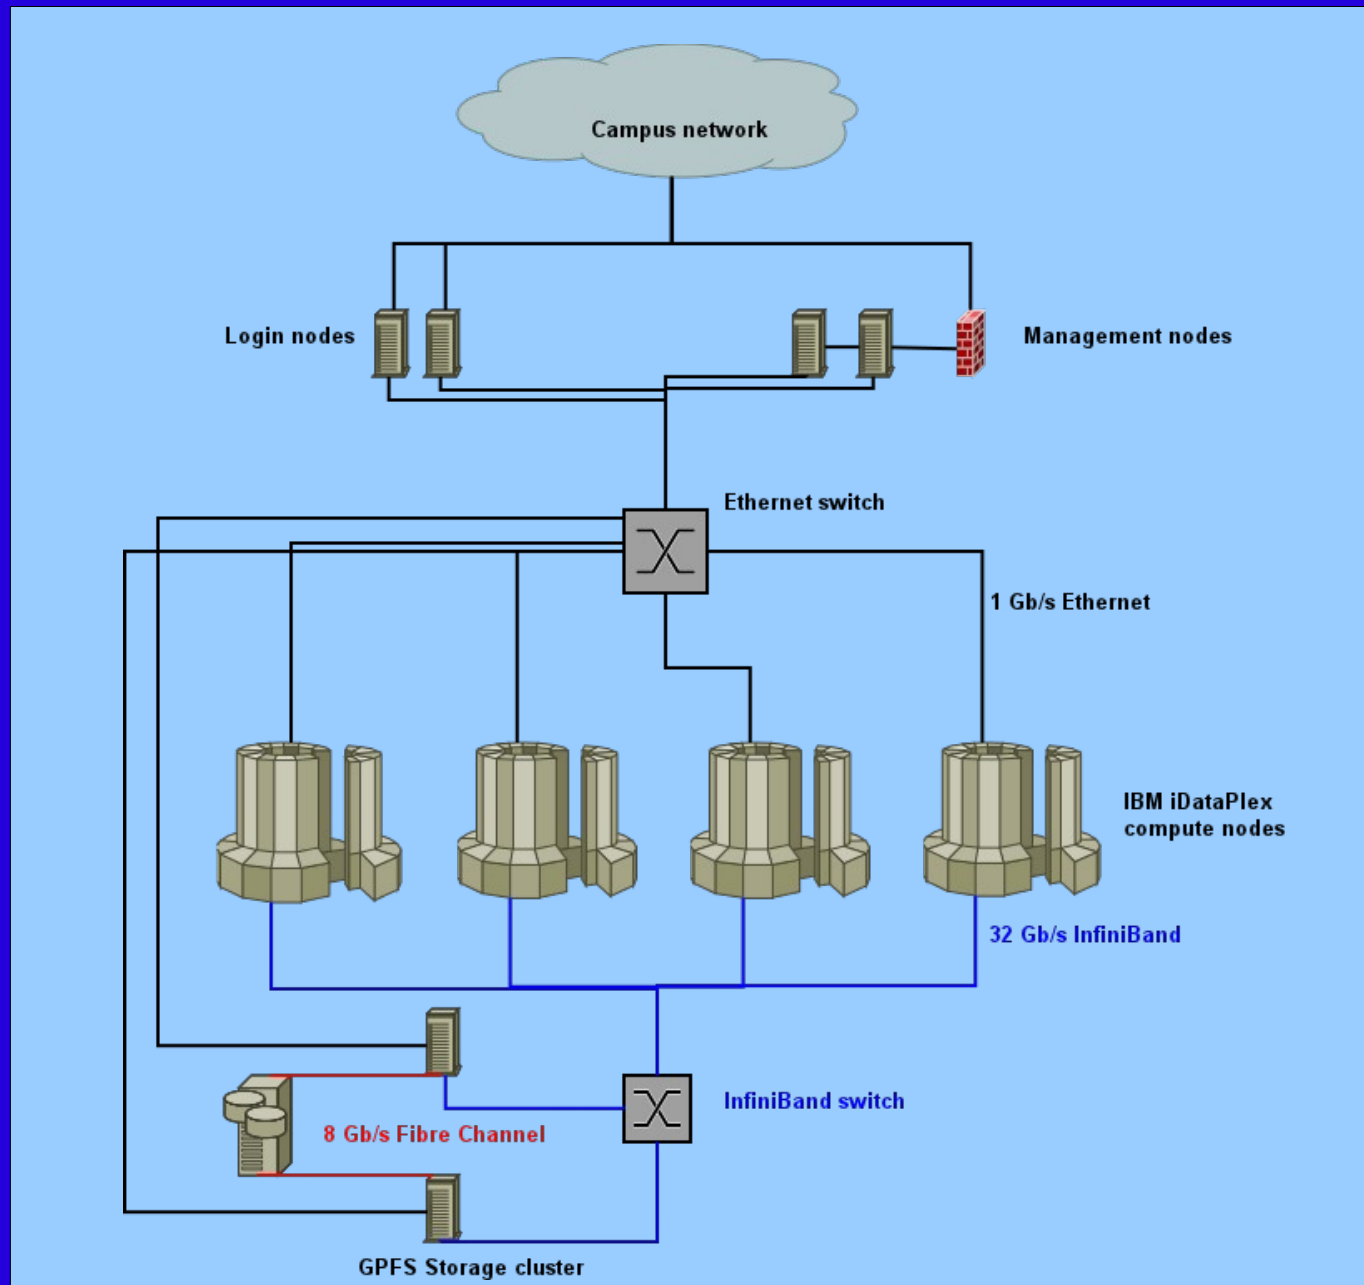
- Shared HPC facility between University and UK MHD consortium
- Entered production in early 2011
- IBM system based on “iDataPlex” hardware
- Collaboration with IBM Research

iDataPlex

- Energy efficient design



Minerva




Minerva

- Technical specifications
 - 2.66 GHz Intel “Westmere” 6-core processors
 - 258 compute nodes with 12 cores per node = 3096 cores
 - 6 GPU nodes with NVIDIA Tesla M2050 GPUs
 - 2 SMP nodes with 144 GB memory
 - ~35 TFLOPS
 - QLogic QDR InfiniBand interconnect

Minerva

- Software
 - SuSE Linux Enterprise server
 - Open MPI
 - Intel Cluster Tools
 - IBM GPFS
 - MOAB/Torque
 - xCAT



Minerva

- xCAT
 - “Extreme Cloud Administration Toolkit”
 - provision OS on physical or virtual machine
 - remote “lights out” management
 - open source
- MOAB
 - powerful meta-scheduler
 - can drive xCAT to provision resources on the fly
- xCAT + MOAB

MOAB/Torque

- CLI interface plus scripting

```
#!/bin/bash
```

```
#PBS -l nodes=256:ppn=4,pvmem=2gb,walltime=08:00:00
```

```
mpirun ./a.out
```

```
msub myscript.pbs
```

Minerva

- GPU nodes
 - NVIDIA M2050 GPU Computing Modules
 - GPUs offer potential large performance gains for the right application
 - Worlds fastest supercomputer “Tianhe-1” is GPU-based
 - CUDA or OpenCL
- SMP nodes
 - Commodity nodes with large RAM
 - cost-effective alternative to dedicated SMP

Minerva

- Collaborative relationship with IBM Research
 - cloud computing
 - pilot projects in cloud computing for research
 - potential application to regional HPC
 - hardware donation to kick-start the project
 - access to Blue Gene technology and expertise at e..g. IBM Watson
 - alignment with IBM's “Smarter Planet” themes

Conclusions

- CSC is a research centre with expertise in scientific and high performance computing
- CSC supports teaching and training in scientific computing
- CSC provides substantial HPC services both nationally and locally
- Why not consider HPC or scientific computing for further study?

