

The background features a collage of digital and technical imagery. On the left, there's a blue vertical bar. The main area is a mix of textures: a world map, a hand typing on a keyboard, a circuit board, and various alphanumeric strings and dates like '18/03', '22/03', '2KV6JTK', '0VZEK48', '2196S2T', '19DR40', '68BG', and '9DBM'.

nag

Numerical Algorithms Group

Mathematics and technology for optimized performance

Early Experiences with HECToR

Mike Dewar

Results Matter. Trust NAG.

HECToR

- Latest high-end computing service for UK
 - managed by EPSRC on behalf of RC-UK
 - funded by EPSRC, NERC & BBSRC
 - will run from 2007-2013
- Objective:
 - to provide a service to the academic community enabling it to do true capability science
- Partners:
 - Service Provision: UoE HPCx Ltd
 - hardware hosting and maintenance
 - user services, help desk etc
 - Hardware: Cray Inc
 - Computational Science & Engineering Support (CSE): NAG Ltd

Overview of the CSE Service

- Partnership with HECToR user community to assist in deriving maximum benefit from the hardware:
 - Initial assessment of project suitability
 - Training
 - Documentation, case studies, FAQs
 - Help desk
 - Assistance with porting, performance tuning and optimisation of codes
 - 60% of resources dedicated to specific users, research groups or consortia
- NAG brings expertise in
 - Numerical Analysis and statistics
 - Software engineering
- General brief to help to improve scientific software in the UK

Hardware Solution

- Phase 1 (September 2007)
 - Cray XT4 (63 Tflops)
 - 5664 dual core Opterons
 - 33.2 TB memory
 - SeaStar2 Interconnect
 - BlackWidow Vector (2 Tflops)
 - after 9 months
 - 576 TB direct storage (rising to 934 TB)
- Phase 2 (2009)
 - Cray ??? (c 250 Tflops)
 - quad core 'Barcelona'
 - faster interconnect
- Phase 3 (2011)
 - > 500 Tflops

Results Matter. Trust NAG.

nag



HECToR Early Users

- Early user service started September 2007
 - Material Science
 - AIMPRO, VASP
 - Engineering
 - PCHAN
 - Chemistry/Biochemistry
 - AMBER, NAMD, PMEND, DLPOLY
 - Environmental Science
 - UM, NEMO
 - Physics
 - HELIUM, CENTORI, GS2, ORB5, QCD codes
- Next major group: “HECToR Capability Challenge”
December 2007

Issue 1: Data Storage / Transport / Analysis

- 576 TB of storage initially
 - of which 40 TB can be backed-up
 - maximum 9 GB/s read-write bandwidth
 - disk-to-disk copy of 250TB would take at least 51 hours
- HECToR connected to SuperJanet
 - 10 GB/s backbone
- Example: Unified Model
 - typical output data $O(10GB)$
- Some scope for post-processing on HECToR
 - mainly batch environment
 - no support for visualisation, computational steering etc.

Visual Analytics

- Combine statistical & visualisation techniques
 - identify key characteristics of data set
 - understand those characteristics through visualisation
- ADVISE: UK TSB-funded research project
 - NAG / VSN / University of Leeds
- User community
 - pharmaceutical / environmental science / engineering / medicine
 - data sets up to 1.5GB
- Service-oriented technology
 - web services, distributed visualisation
 - 'intelligent' datatypes supporting distributed, out-of-core storage
- First prototype implemented

Issue 2: Code Management

- Typical user code
 - Implemented in Fortran
 - Uses MPI for parallelism
 - Written by generations of graduate students and PDRAs
- NAG's role in HECToR
 - Advise on good software-engineering practices
 - Help with optimisation, tuning and scaling
 - Preserve performance across hardware upgrades
 - Make better use of standard software components and libraries

Hardware Upgrades

- HECToR Phase 1: 2.8GHz dual-core opteron
 - SSE via 64bit FPU
- HECToR Phase 2: 2.6GHz (?) quad-core “barcelona”
 - SSE via improved 128bit FPU
- AMD Core Math Library (ACML)
 - developed in collaboration with NAG
 - includes LAPACK, FFTs, RNGs, ...
 - many kernels coded in assembler and tuned for particular chips
- To maximise performance many assembler kernels in ACML have been (laboriously!) re-written
 - better choice of instruction to reduce latencies
 - re-ordering code to ensure instructions done in parallel

Results Matter. Trust NAG.

nag

Issue 3: Developments in Hardware

- Not HECToR-specific but relevant to anybody developing scientific codes
- Multi-core / many-core / GPGPUs / FPGAs etc
 - performance per watt becoming important
 - disruptive technology, codes will need to be rewritten
 - need for software that adapts to host environment
 - golden opportunity to introduce new algorithms and computational techniques

Conclusions

- HECToR will be a major tool for “high end” UK Science for the next six years
- The ability to store and transport data is becoming a bigger restriction than the availability of computational resources
- The wide availability of good-quality components can help insulate users against changes in hardware in short term
- Many codes may soon need to be re-written to take account of architectural changes ...
- ... making this a perfect time to incorporate radically new algorithms and computational techniques



HECTOR

HIGH END COMPUTING TERASCALE RESOURCE

Results Matter. Trust NAG.

nag