STFC's Scientific Computing Department and The Hartree Centre

Paul Sherwood

paul.sherwood@stfc.ac.uk

Overview

- Background to STFC and its Facilities
- STFC's Scientific Computing Department
- STFC Hartree Centre: future developments to meet future challenges

Science & Technology Facilities Council

Organisation

HM Government (& HM Treasury)

↓

BIS Department for Business Innovation & Skills

↓

RESEARCH COUNCILS UK — RCUK Executive Group

↓

Arts & Humanities Research Council

BBSRC bioscience for the future

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL

EPSRC Pioneering research and skills

MRC Medical Research Council

NATURAL ENVIRONMENT RESEARCH COUNCIL

Science & Technology Facilities Council

STFC's Sites

UK Astronomy Technology Centre, Edinburgh, Scotland

Polaris House Swindon, Wiltshire

Chilbolton Observatory
Stockbridge, Hampshire

Daresbury Laboratory
Daresbury Science and Innovation Campus
Warrington, Cheshire

Rutherford Appleton Laboratory
Harwell Science and Innovation Campus
Didcot, Oxfordshire

Isaac Newton Group of Telescopes
La Palma

Joint Astronomy Centre Hawaii

Science & Technology Facilities Council

STFC's Facilities

Neutron Sources

ISIS - pulsed neutron and muon source/ and Institute Laue-Langevin (ILL), Grenoble
Providing powerful insights into key areas of energy, biomedical research, climate, environment and security.

High Power Lasers

Central Laser Facility - providing applications on bioscience and nanotechnology
HiPER
Demonstrating laser driven fusion as a future source of sustainable, clean energy

Light Sources

Diamond Light Source Limited (86%) - providing new breakthroughs in medicine, environmental and materials science, engineering, electronics and cultural heritage
European Synchrotron Radiation Facility (ESRF), Grenoble

Science & Technology Facilities Council

Understanding our Universe
STFC's Science Programme

Particle Physics

Large Hadron Collider (LHC), CERN - the structure and forces of nature

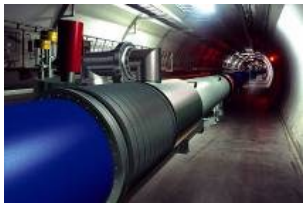

Ground based Astronomy

European Southern Observatory (ESO), Chile
Very Large Telescope (VLT), Atacama Large Millimeter Array (ALMA), European Extremely Large Telescope (E-ELT), Square Kilometre Array (SKA)

Space based Astronomy

European Space Agency (ESA)
Herschel/Planck/GAIA/James Webb Space Telescope (JWST)
Bi-laterals – NASA, JAXA, etc.
STFC Space Science Technology Department

Nuclear Physics

Facility for anti-proton and Ion research (FAIR), Germany
Nuclear Skills for - medicine (Isotopes and Radiation applications), energy (Nuclear Power Plants) and environment (Nuclear Waste Disposal)

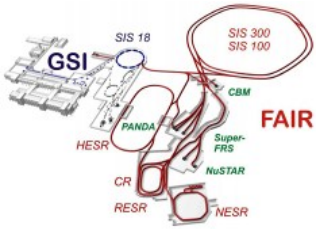

Science & Technology Facilities Council

Scientific Computing Department

NEW

Director: Adrian Wander
Appointed 24th July 2012

Major funded activities

- 180 staff supporting over 7500 users
- Applications development and support
- Compute and data facilities and services
- Research: over 100 publications per annum
- Deliver over 3500 training days per annum
- Systems administration, data services, high-performance computing, numerical analysis & software engineering.

Major science themes and capabilities

- Expertise across the length and time scales from processes occurring inside atoms to environmental modelling

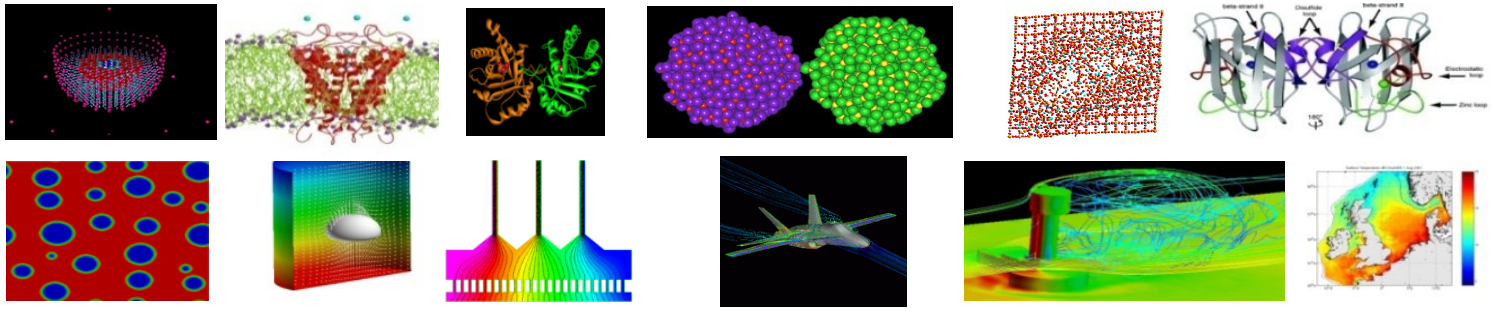

Science & Technology Facilities Council

SCD – Divisional Structure

- Established following merger of Computational Science and Engineering and e-Science Departments

- Data
  - Large data storage (e.g. LHC)
  - Middleware development
  - Data management for UK and international facilities
- Systems
  - Operation of HPC and cluster systems including STFC, national, regional and developmental systems
- Technology
  - New architectures, HPC optimisation and porting, software engineering
- Applications
  - Science based groups including numerical analysis, grant projects, Service Level Agreement with EPSRC for UK academic support

Applications Groups (i)

- Computational Chemistry
  - QM and QM/MM
  - MD, DPD, LBE
  - MC, aKMC
- Engineering
  - CFD, microfluidics
  - environmental simulation

Applications Groups (ii)

- Theoretical and Computational Physics
  - Periodic HF, DFT (CRYSTAL, CASTEP, ONETEP)
  - Atomic and molecular Physics (R-Matrix)
- Computational Biology
  - Protein and membrane simulation
  - Bioinformatics (new hire)
  - Sofware and data infrastructures for structural biology (X-ray, NMR, EM)
- Numerical Analysis
  - Sparse matrix solvers
  - Harwell subroutine library (HSL)

Science & Technology Facilities Council

Scientific Highlights

Journal of Materials Chemistry 16 no. 20 (May 2006) - issue devoted to HPC in materials chemistry (esp. use of HPCx);

Phys. Stat. Sol.(b) 243 no. 11 (Sept 2006) - issue featuring scientific highlights of the Psi-k Network (the European network on the electronic structure of condensed matter coordinated by our Band Theory Group);

Molecular Simulation 32 no. 12-13 (Oct, Nov 2006) - special issue on applications of the DL_POLY MD program written & developed by Bill Smith (the 2nd special edition of Mol Sim on DL_POLY - the 1st was about 5 years ago);

Acta Crystallographica Section D 63 part 1 (Jan 2007) - proceedings of the CCP4 Study Weekend on protein crystallography.

The Aeronautical Journal, Volume 111, Number 1117 (March 2007), UK Applied Aerodynamics Consortium, Special Edition.

Proc Roy Soc A Volume 467, Number 2131 (July 2011), HPC in the Chemistry and Physics of Materials.

Last 5 years metrics:

- 67 grants of order £13M
- 422 refereed papers and 275 presentations
- Three senior staff have joint appointments with Universities
- Seven staff have visiting professorships
- Six members of staff awarded Senior Fellowships or Fellowships by Research Councils' individual merit scheme
- Five staff are Fellows of senior learned societies

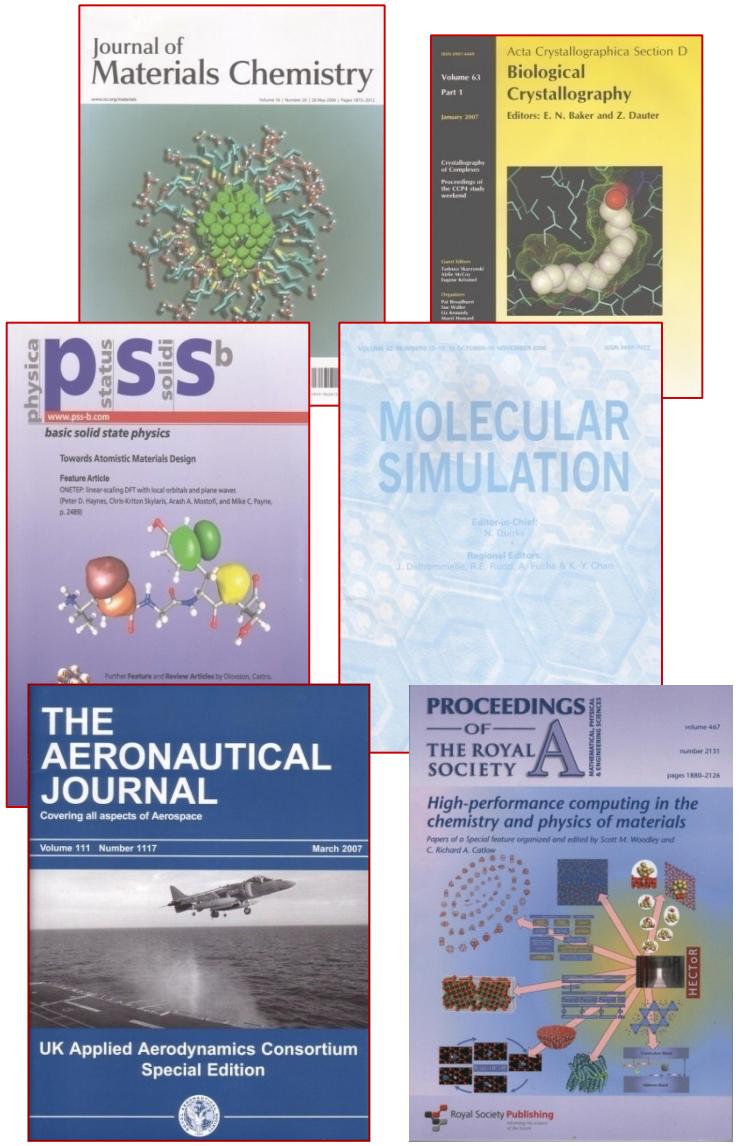

The CCPs

- UK national consortia
- Bring together leading theoretical and experimental research groups in thematic areas
- Tackle large-scale scientific software development
  - Flagship Projects
- Maintenance (long-term)
- Distribution
- Training

Science & Technology Facilities Council

Collaborative Computational Projects

- CCP4 – Macromolecular crystallography
- CCP5 – Computer simulations of condensed phases
- CCP9 – Computational Electronic Structure of Condensed Matter
- CCP12 – High performance computing in engineering
- CCP-ASEArch – Algorithms and Software for Emerging Architectures
- CCPBioSim – Biomolecular simulation at the life sciences interface
- CCPN – NMR spectroscopy
- CCP-EM – Electron cryo-Microscopy
- CCPi – Tomographic Imaging
- CCPN – NMR
- CCP-NC – NMR Crystallography
- CCPQ – Quantum dynamics in Atomic, Molecular and Optical Physics

Supported by
EPSRC
BBSRC
MRC

Science & Technology Facilities Council

High End Computing Consortia

- Awarded large amounts of national HPC time (to allocate internally) and staff support including DL staff for HPC code development and support
- Last round of awards included
  - Materials Chemistry
  - UK Car-Parinello
  - Biomolecular simulation
  - Mesoscale Engineering Science
  - UK Turbulence Consortium
  - Plasma Physics

Science & Technology
Facilities Council

The Hartree Centre

Science & Technology Facilities Council

Science and Innovation Campuses

- We will develop the Science and Innovation Campuses at Harwell and Daresbury as focal points for collaboration and knowledge exchange with industry and academic researchers
- Work to establish gateways to our in-house expertise and that of the communities we support
- We will increasingly focus our technology competencies on an outward facing collaborative role

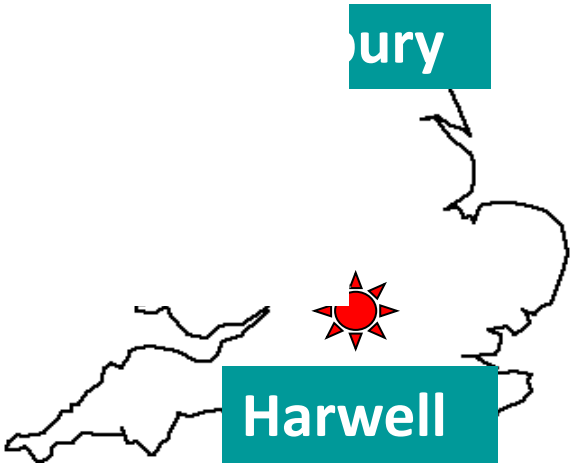

Science & Technology Facilities Council

Science and Technology Gateway Centres

Focus our technology competencies increasingly on an outward facing collaborative role

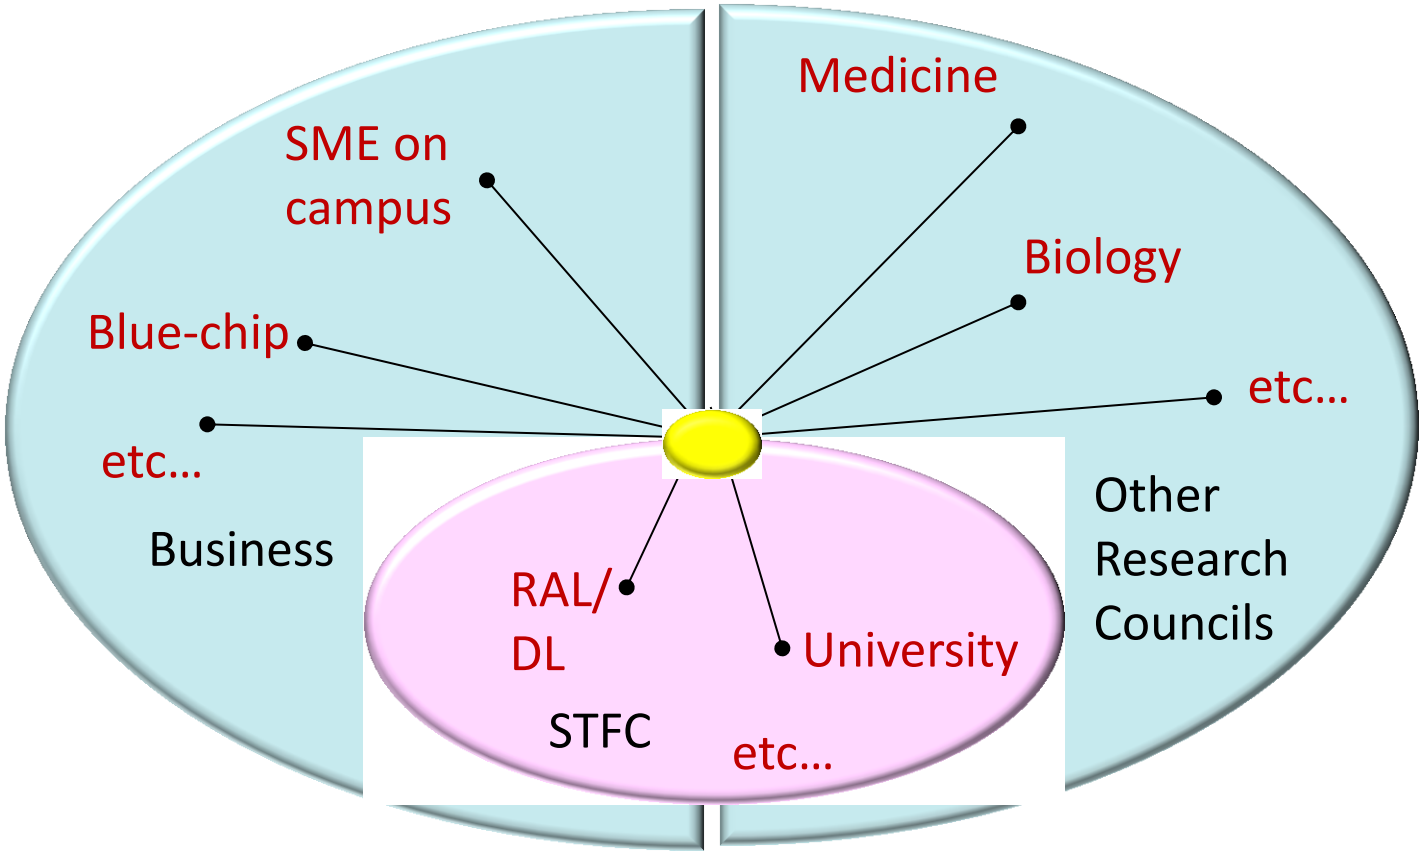

Science & Technology Facilities Council

Government Investment in e-infrastructure - 2011

17th Aug 2011: Prime Minister David Cameron confirmed £10M investment into STFC's Daresbury Laboratory. £7.5M for computing infrastructure

3rd Oct 2011: Chancellor George Osborne announced £145M for e-infrastructure at the Conservative Party Conference

4th Oct 2011: Science Minister David Willetts indicated £30M investment in Hartree Centre

30th Mar 2012: John Womersley CEO STFC and Simon Pendlebury IBM signed major collaboration at the Hartree Centre

Clockwise from top left

Science & Technology Facilities Council

Tildesley Report

- BIS commissioned a report on the strategic vision for a UK e-Infrastructure for Science and Business.

- Prof Dominic Tildesley led the team including representatives from Universities, Research Councils, industry and JANET. The scope included compute, software, data, networks, training and security.

- Mike Ashworth, Richard Blake and John Bancroft from STFC provided input.

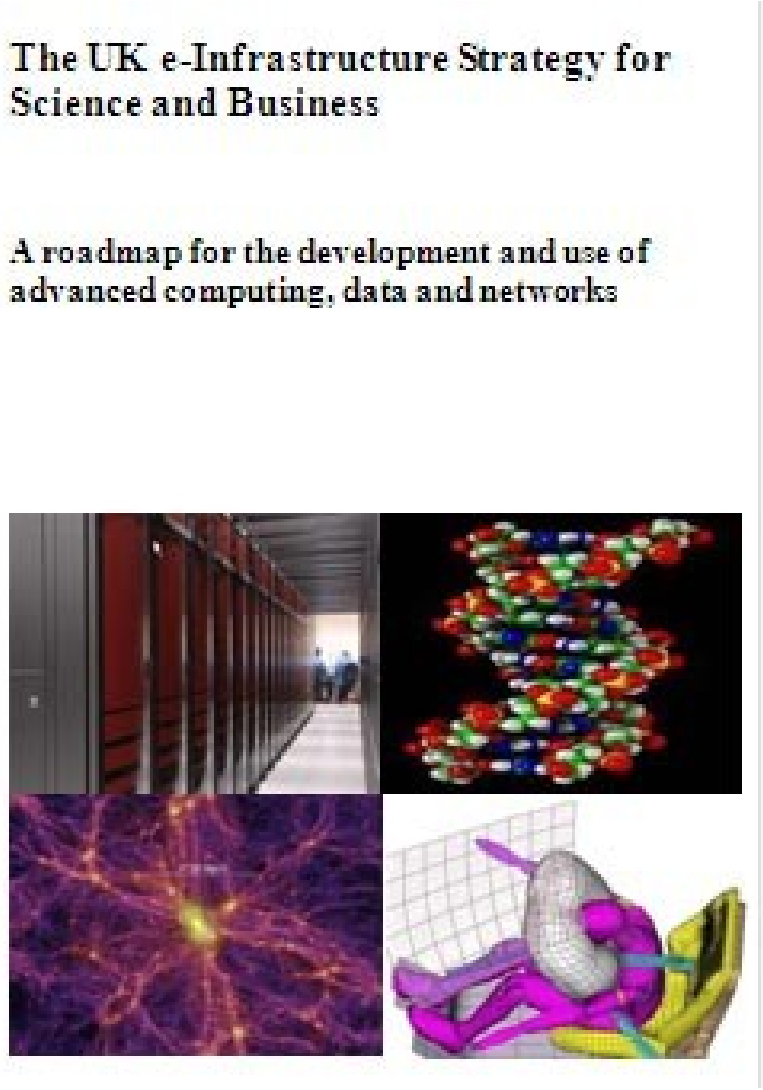

Published in December 2011. Google the title to download from the BIS website

Science & Technology Facilities Council

Responding to the Challenges

- Expert optimisation of existing software
  - Profiling and identification of hotspots, use of libraries, tackling issues associated with large core counts, serial bottlenecks, redesign of I/O etc
- Application Co-design
  - Software and hardware must evolve together
  - Requires close collaboration between hardware architects, computer scientists, and application software experts
- Re-engineering software requires a specialised development platform
  - Highest possible core count
  - Configured and operated for software development (interactive use, profiling, debugging etc)
- Ease of Use
  - To achieve economic impact, HPC must be made easier to adopt throughout organisations, not just a specialist capability

The Hartree Centre

a Research Collaboratory with IBM

Science & Technology Facilities Council

Hartree Centre capital spend 2011/12

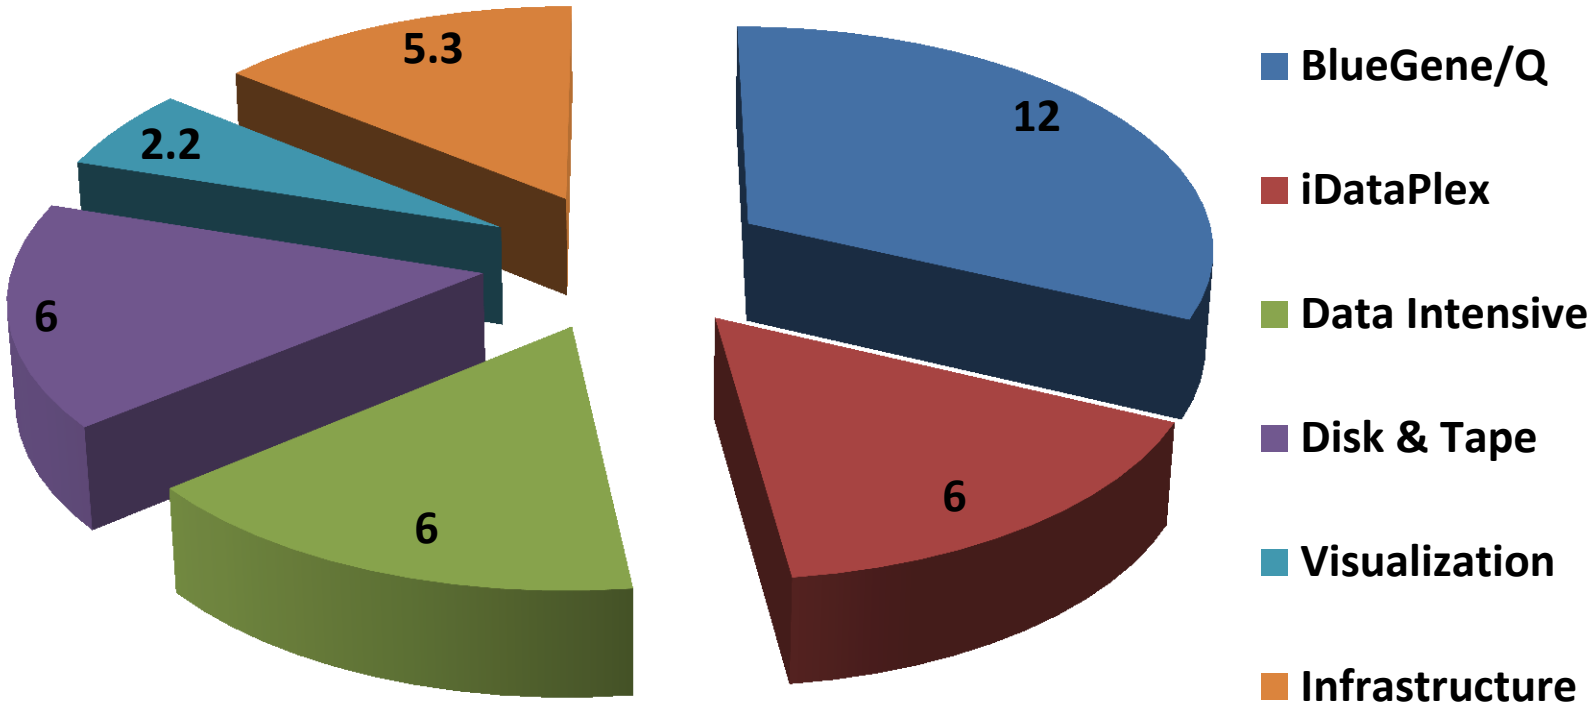

approximate capital spend £M

Total £37.5M

Science & Technology Facilities Council

Hartree Centre
IBM BG/Q Blue Joule

TOP500
#16 in the Nov 2012 list
#5 in Europe
#1 system in UK

6 racks
- 98,304 cores
- 6144 nodes
- 16 cores & 16 GB per node
- 1.26 Pflop/s peak

1 rack of BGAS (Blue Gene Advanced Storage)
- 16,384 cores
- Up to 1PB Flash memory

Science & Technology Facilities Council

Hartree Centre
IBM iDataPlex Blue Wonder

TOP500
#158 in the Nov 2012 list

8192 cores, 196 Tflop/s peak
node has 16 cores, 2 sockets
Intel Sandy Bridge (AVX etc.)

252 nodes with 32 GB
4 nodes with 256 GB
12 nodes with X3090 GPUs

256 nodes with 128 GB
ScaleMP virtualization software up to 4TB virtual shared memory

Science & Technology Facilities Council

Hartree Centre Datastore

Storage:

5.76 PB usable disk storage
15 PB tape store

Science & Technology Facilities Council

Hartree Centre Visualization

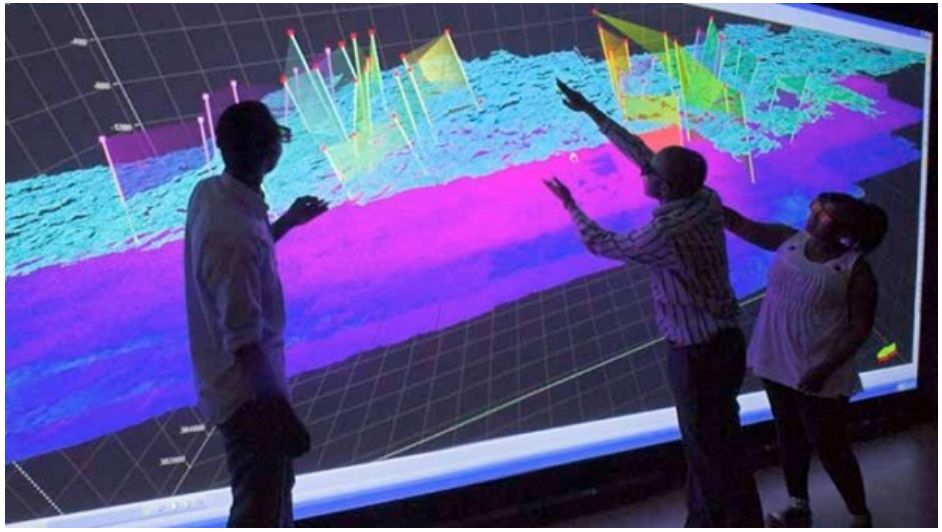

Four major facilities:

Hartree Vis-1: a large visualization "wall" supporting stereo
Hartree Vis-2: a large surround and immersive visualization system
Hartree ISIC: a large visualization "wall" supporting stereo at ISIC
Hartree Atlas: a large visualization "wall" supporting stereo in the Atlas Building at RAL, part of the Harwell Imaging Partnership (HIP)

Virtalis is the hardware supplier

Science & Technology Facilities Council

Hartree Centre Objectives

- Hartree is establishing a new model for innovation and interdisciplinary collaboration around HPC, Big Data and Cloud technology bringing together STFC, Industry and Academia

- We are creating an Innovation Ecosystem for science and research linked strongly to industry.

Science & Technology Facilities Council

Hartree Centre Implementation

- Established as a 5th division of Scientific Computing Department
- Cliff Brereton appointed director April 2013
- IBM BD and technical staff on secondment
- HC BD team responsible for building projects based on Centre facilities and expertise from SCD and other partners (BD lead Michael Gleaves)
- The centre uses staff from other divisions on projects (systems, technology and applications)

Science & Technology Facilities Council

Government Investment in e-infrastructure - 2013

1st Feb 2013: Chancellor George Osborne and Science Minister David Willetts opened the Hartree Centre and announced a further £185M of funding for e-Infrastructure

£19M for the Hartree Centre for power-efficient computing technologies

£11M for the UK's participation in the Square Kilometre Array

This investment forms part of the £600 million investment for science announced by the Chancellor at the Autumn Statement 2012.

"By putting out money into science we are supporting the economy of tomorrow."

George Osborne opens the Hartree Centre, 1st February 2013

Science & Technology Facilities Council

Power-efficient Technologies 'Shopping List'

£19M investment in power-efficient technologies

- System with latest NVIDIA Kepler GPUs
- System based on Intel Xeon Phi
- Active storage project using IBM BGAS
- Dataflow architecture based on FPGAs
- Instrumented machine room

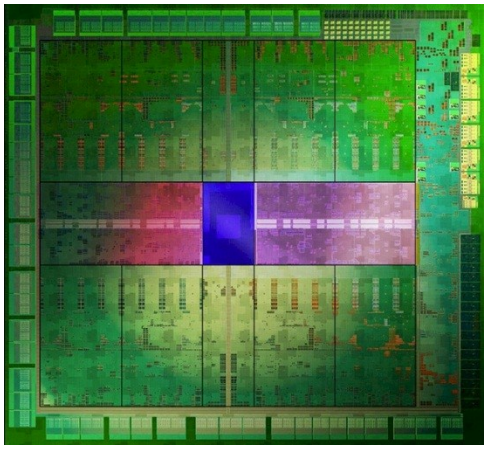

Systems will be made available for development and evaluation projects with Hartree Centre partners from industry, government and academia

Science & Technology Facilities Council

Collaboration with Unilever

1st Feb 2013: Also announced was a key partnership with Unilever in the development of Computer Aided Formulation (CAF)

Months of laboratory bench work can be completed within minutes by a tool designed to run as an 'App' on a tablet or laptop which is connected remotely to the Blue Joule supercomputer at Daresbury.

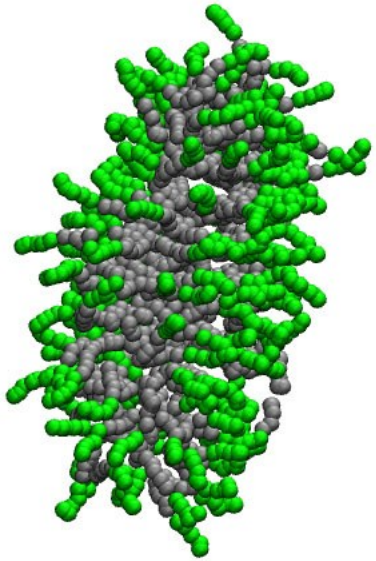

The aggregation of surfactant molecules into micelles is an important process in product formulation

This tool predicts the behaviour and structure of different concentrations of liquid compounds, both in the bottle and in-use, and helps researchers plan fewer and more focussed experiments.

John Womersley, CEO STFC, and Jim Crilly, Senior Vice President, Strategic Science Group at Unilever

Science & Technology Facilities Council

The Hartree Centre and HECToR/ ARCHER

Hartree Centre is complementary to the UK National Super-computing Services, managed by EPSRC on behalf of the UK academic communities

HECToR is the current system

HECToR Phase3 90,112 cores Cray XT6 (800 Tflop/s)

ARCHER is the new service, due for installation around Nov/Dec 2013, service starting in early 2014

HECToR/ARCHER provide cycles with limited science support

Hartree focuses on commercial applications and next generation software development

Summary

- STFC is one of the UK's research councils
- Scientific Computing Department has been formed from merger of Computational Science and Engineering and e-Science
- The Hartree Centre, established as a "Gateway Centre" following capital investment
  - a Collaboratory with IBM
  - focussed on the economic impact of HPC