

CompSci07

Computational Science 2007: Interdisciplinary Challenges and Perspectives, from the Grid to e-Science

25 - 26 June 2007 at the Royal Society, London. There is no registration fee for attending this conference.

Abstracts

(Download [PDF Version](#))

Computational Science in the Twenty-First Century

Peter Coveney, University College London

Abstract Awaited.

Simulating Complex Fluids, Both Living and Dead

Mike Cates, University of Edinburgh

The lattice Boltzmann method (LB) allows fast, efficient simulation of isothermal fluids by discretizing both real space and momentum space. By introducing additional distribution functions to describe various types of local order, and/or coupling the lattice fluid to the molecular dynamics of colloidal particles, efficient simulations of many complex fluids becomes possible. Application areas for LB are now undergoing an expansion, to include not only passive (or 'dead') systems in which the components have conventional physicochemical interactions, but also 'active' materials. The latter, generally of biological origin, are systems in which the constituent particles are self-propelled entities, such as bacteria or cytoskeletal assemblies, whose dynamics depends on a continuous throughput of chemical energy. Recent advances in using LB to study complex fluids, both living and dead, will be outlined. These include (i) a method of representing both active and passive colloidal particles as pointlike objects, speeding their simulation; and (ii) development of hybrid codes in which LB handles momentum transport only, with other order parameter fields treated by finite difference methods. To showcase these hybrid codes, recent results for both both active and passive liquid crystals will be presented.

From quantum mechanics to computational atomistic materials design: e-science challenges and opportunities from the CCP9 and Psi_k networks

James Annett, Bristol University and CSE, CCLRC Daresbury

The CCPs have been active for many years as collaborative projects for developing and applying high quality scientific software in support of UK science. CCP9 focuses on using computational quantum mechanics to study the electronic structure of materials, and hence to predict materials properties from 'first principles'. It is a network of over 30 academics, with members in almost all major UK research universities. Psi-k is a Europe wide network, with over 500 active researchers, and up to 1500 participants. These networks have survived and succeeded despite crossing institutional and national barriers, and through a variety of different funding mechanisms over time. How can e-science help this activity? E-science has been defined as supporting collaborative virtual organisations across institutional barriers. In this talk I will present a personal perspective of the challenges for maintaining networks such as ccp9 and Psi-k, and how the ideals and achievements of the e-science community can assist these large scale cooperative projects. Challenges include code development across institutions, improved code interoperability, interfaces, access to HPC resources, visualization, data access and storage, as well as training of postdocs and PhD students in the latest developments.

Managing scientific workflows with BPEL

Wolfgang Emmerich, University College London

Modern scientific applications often need to be distributed across Grids. Increasingly applications rely on services, such as job submission, data transfer or data portal services. We refer to such services as Grid services. While the invocation of Grid services could be hard coded in theory, scientific users want to orchestrate service invocations more flexibly and compose them into scientific workflows. In enterprise applications, the orchestration of web services is achieved using emerging orchestration standards, most notably the Business Process Execution Language (BPEL). We describe the OMII-BPEL environment that is distributed as part of the Open Middleware Infrastructure Institute grid middleware. The environment supports the visual modelling of scientific workflows, the automated deployment of these workflows into an execution environment, enactment of workflows and monitoring of long-running workflows. We report on experience in using the OMII-BPEL environment gained from an extensive case study that orchestrates Grid services for the automation of a polymorph prediction application.

Challenges and Opportunities for the Development of Biomolecular Simulation

Charles Laughton, University of Nottingham

Molecular Dynamics simulations of biological macromolecules (e.g. proteins, nucleic acids, and their complexes) can provide information about the properties and behaviour of such molecules that are unavailable by any experimental technique. Therefore MD simulation is a research tool that should be in the arsenal of the majority of biologists, chemists and biotechnologists - but at present this is not generally the case, being restricted to a relatively small number of specialized groups. The Collaborative Computational Project for Biomolecular Simulation (CCPB) was established last year with the aim of increasing the UK's capacity and capabilities in this area. In this talk I will outline some of the issues that need to be addressed to improve the uptake of these computational techniques by the wider research community, and illustrate some of the initiatives that CCPB is involved with for this purpose.

Overview of the NW-GRID Project - Progress to Date and Plans

Terry Hewitt, Director of Research Computing, Manchester Computing

The NW-GRID Project is a collaboration between CCLRC Daresbury Laboratory and the Universities at Lancaster, Liverpool and Manchester funded by the North-West Regional development Agency. The Project has established a computational Grid comprising commodity computing systems coupled by a high-speed private network and is seeking to evaluate middleware in a broad base of materials, lifescience, engineering and environmental applications. The talk will over view the aims and objectives of the Project which are structured around three testbeds:

- high-performance computing and parametric calculations;
 - a safe, bookable and accountable grid; and,
 - a collaborative real-time grid.
- and review progress to date and future plans.

Some Computational Problems in Astrophysical MHD

Sam Falle, University of Leeds

In astrophysical magnetohydrodynamics (MHD) the equations that need to be solved are those of compressible flow with the addition of the induction equation for the magnetic field and the Lorentz force in the momentum equation. There can also be complicated heating and cooling processes and self gravity is sometimes important. Since the flows are often supersonic, they contain shocks and these need to be captured via an upwind finite volume method. Explicit methods of this type on uniform grids can readily be parallelised on distributed memory machines using MPI, but there are cases where the ideal MHD approximation breaks down and one has to use an implicit approximation to the generalised Ohm's law. It is possible to implement such methods on distributed memory machines, but it is not trivial to do this without loss of efficiency. Many astrophysical problems involve a very large range of length and time scales, which means that a uniform grid is very inefficient. For this reason, Adaptive Mesh Refinement (AMR) is often used, but this is certainly not easy to implement on distributed memory machines. However, it can be done on clusters using either Myrinet or Infiniband communication. Although self-gravity is an essentially non-local effect, the Poisson equation can still be solved efficiently on a distributed memory machine using local algorithms, such as multigrids. For certain problems, such as the generation of stellar magnetic fields by dynamo action, the flow is essentially incompressible. For these problems it is usual to use spectral methods. Although these are very accurate, any implementation on a distributed memory machine requires a great deal of communication, which limits the number of processors that can be used to less than 100 with current technology.

Ab-initio modelling of surfaces

Steven Kenny, Loughborough University

This talk will focus on two areas surface reconstructions on (100) SrTiO₃ and fullerenes on Si (100). SrTiO₃ has been extensively studied experimentally as it is used as a support for the growth of high T_c superconductors. The (100) surface of SrTiO₃ undergoes a series of reconstructions when exposed to UHV and heated, including the formation of nanoscale islands. The driving force for the formation of these islands is not understood. We have studied a number of possible surface reconstructions in order to explain some of the observed experimental results. Fullerenes on Si (100) are of interest both from a fundamental science perspective and also as a possible system for building devices, including quantum computers. I will present work from our studies of these systems which explain the nature of the interaction between the fullerenes and the surface.

Improving User Interaction in Grid based Computational Environments

R. S. Kalawsky, University of Nottingham

Computational science is radically changing the way large science is performed today. This has partly been enabled by the evolution of distributed computing environments based on grid infrastructure. Consequently, the scale and complexity of scientific investigations has increased but this has meant harnessing the power of computational resources has not been without its challenges. Whilst in principle, scientific problems can be treated with computational modelling and simulation the analysis of data very often leads to unsatisfactory results due to the limited evaluation functionality of the tools available. On one hand, the reason is bound up in the visualization of 3D data, which is more frequently performed in 3D, and on the other hand is due to the physical relationship between parameters, which is usually not reproduced in the analysis. This paper looks at the steps that have been taken to tackle the user's interaction with large grid based computing resources. We will report on our use of architectural frameworks to help understand and define the scientist's interactive requirement and illustrate how this has been successfully applied to computational science investigations and medical visualization. The context of the paper will be based on the provision of a ubiquitous interactive real-time environment (high-end compute, desktop, mobile computing solutions (PDAs)) that provides an effective user interface to the scientists computational environment

and data space.

Grid-Based Domain Decomposition for the Simulation of Hydrodynamic Turbulence

Bruce Boghosian, Tufts University

We present recent results of the Vortronics and HYPO4D projects, aimed at applying Grid-based geographically distributed domain decomposition to the study of hydrodynamic turbulence. The vision, motivation, and broader implications of the HYPO4D project for the future of numerical simulation of complex spatiotemporal dynamical systems will also be discussed.

High Performance Computing for Interest Rate Derivatives Trading

Jonathan Chin, Morgan Stanley

Interest rate derivatives are financial derivative transactions whose underlying asset is an interest rate. They range from "vanilla" instruments whose price can be determined analytically, to path-dependent "exotics", and hybrids whose value derives from many underlying asset classes. We review some properties and common models of these transactions, and the computational challenges involved in using these models to manage a large trade portfolio; we describe a distributed computing environment which has evolved to meet these challenges in a highly competitive market.

Adaptive performance control of coupled scientific models executing in a distributed environment

John Gurd, University of Manchester

For certain forms of coupled scientific model, effective control of performance under dynamically varying load in a distributed execution environment can be achieved via adaptive redeployment of the component models based on simple execution and load monitoring schemes coupled with various kinds of performance predictor. Experiments based on various load-change scenarios illustrate how example coupled model codes, including HybridMD, react dynamically to their execution environment using such techniques.

Modern software engineering techniques applied to the simulation of materials and nanosystems : the ABINIT project

Xavier Gonze, Louvain-la-Neuve, Belgium

The field of first-principles simulation of materials and nanosystems has seen an amazing development in the past twenty years. In order to cope with the increasing software complexity, it became apparent, about a decade ago, that software engineering techniques and a group collaborative effort would be major ingredients of a successful first-principles project. Following the model of the Linux operating system development, the open source ABINIT project was launched in 1997. As of now, there are more than 1000 mailing addresses in the main mailing list, and about 40 active contributors. After a short presentation of the capabilities of the package, I will present and analyze the software engineering techniques that are at the heart of ABINIT, permitting an efficient group effort. The following issues will be covered : the coding rules, the reliability and quality control (self-testing), the portability, the self-documentation, the re-use of software, the file formats, the version control system. Some characteristics common to open software projects will be highlighted.

Computational Fluid Dynamics: an Engineering Perspective

Stewart Cant, Department of Engineering, University of Cambridge

Within the vast and diverse area of computational engineering, the topic of Computational Fluid Dynamics (CFD) emerges as by far the largest consumer of high-performance computing resources. CFD is now widely accepted as a key component of the design process within the aerospace and automotive industries, and yet there remain significant fundamental challenges which have yet to be faced. The central issue in scientific terms is the current lack of understanding of turbulent flow an adequate mathematical description of turbulence does not exist, and it is necessary to construct phenomenological models in order to solve industrial problems. Three levels of CFD are being pursued. Direct Numerical Simulation (DNS) aims to solve the governing Navier-Stokes equations without physical modelling, but the computational cost is prohibitive for all but the simplest research problems. Large Eddy Simulation (LES) resolves only the largest scales of the turbulent motion and is slowly gaining acceptance for industrial problems. Reynolds-Averaged Navier-Stokes (RANS) represents only the average behaviour of the flow and remains the method of choice for most industrial problems. The paper describes ways in which the different approaches complement each other and provides examples of current practice in a range of engineering application areas.

Large-scale simulations for WiFi-based wireless telecommunication networks

Dr. Maziar Nekovee, University College London and BT Research

Wireless Fidelity (WiFi), a digital communication technology operating in unlicensed frequency bands has

revolutionized wireless networking and is the fastest growing wireless technology today. Ad hoc wireless networks, a collection of mobile wireless nodes that form a temporally network without any infrastructure or centralized control, can greatly increase the coverage WiFi-based communication by enabling wireless devices route packets in a distributed manner either to the nearest Access Points or among themselves. The increasing complexity and very large scale of Wi-Fi-based wireless communication systems has created a need for high-fidelity simulation platforms that can help scientists and engineers accurately predict and optimize their performance prior to large-scale deployment. However, such large-scale simulations are computationally extremely intensive and are beyond the reach of standard simulation engines.

The National Grid Service and Computational Science

Stephen Pickles, University of Manchester

The mission of the National Grid Service (NGS) is to provide coherent electronic access for UK researchers to all computational and data based resources and facilities required to carry out their research, independent of resource or researcher location. In this talk, I shall review the services that the NGS offers to computational scientists, and report on recent and planned developments. I shall also discuss how several years experience of operating production grids in support of collaborative research has helped refine our understanding of the notion of Virtual Organisations, and I consider the implications for user communities and service providers.

Simple API for Grid Applications: Towards a Standardized Application Level Interface for Grids

Shantenu Jha, Louisiana State University

The aim of this talk is to give an overview of efforts to develop a standardised API for Grid applications -- A Simple API for Grid Applications (SAGA). We will discuss the scope of the API, present a road-map and through couple of examples present the power and importance of such an application level abstraction.

Steering and Visualization: Enabling Technologies for Computational Science

Helen Wright, University of Hull

Computational steering is a mechanism for scientific investigation in which the parameters of a simulation can be altered whilst the program is running. Already an important paradigm, its popularity is growing as access to high-end computing facilities improves. Successful steering also relies on the interactive visualization of results - an essential component but still a challenge for large-scale, distributed applications. This talk will review these important enablers for computational science.

Recent developments in coupled coastal and shelf sea modeling

Andrew J. Willmott, Proudman Oceanographic Laboratory

Human wealth, health and well-being are all affected by the sea. For example, we are heavily dependent fish stocks, many of which we exploit in the coastal seas. Many of the most rapidly growing ?mega-cities? are located on the coast and will be at increasing risk from flooding exacerbated by rising sea levels. Large coastal cities can create contamination of coastal waters and sediments with adverse impact on marine eco-systems. Climate change is also likely to lead to increased coastal erosion.

The combined impact of climate change and coastal developments on the health of our coastal seas is clearly a problem that requires urgent attention. Scientists at the Proudman Oceanographic Laboratory are at the forefront of developing coupled hydrodynamic, wind wave, sea ice, sediment, ecosystem models for predicting how coastal and shelf seas will function, in an integrated sense, as the climate warms. Our predictions typically cover the next 50 years and resolve horizontal length scales of the order of one kilometer.

In my presentation I will review our current coupled coastal and shelf sea models, including their strengths and weaknesses. The presentation will conclude with a review of the next generation UK ocean circulation model that POL is developing in partnership with other university and marine laboratories. This model will exploit unstructured grids which will enable it to resolve small-scale processes such as mixing over topography and dense water overflows. With increasing interest in small scales, particle tracking techniques are also being used for sediment transport modelling.

Enabling Scientific Collaboration in Computational Science

John Brooke and Michael Parkin, University of Manchester

It is commonly accepted that current extended computational infrastructures, such as computational Grids, are perceived as difficult to use with a high barrier to entry. The RealityGrid project developed the concept of "lightweight" Grids, which should be both easy to use and also be accessible by lightweight devices such as PDAs and mobile phones over wireless access. In order to do this, the protocols and methods used to connect together end users accessing computational resources for the purposes of scientific collaboration need to be different from those used to tie together large scale computing and data sources connected via cabled networks. This research proceeds from the premise that the lightweight and heavyweight Grids are complementary or dual to each other. Given this, it becomes much easier to engineer solutions for each type of infrastructure, but also to provide points

of connectivity between them. We illustrate this with a practical example of a group of scientists joining together to form a dynamic Virtual Organisation to steer a large-scale calculation running on a fixed computational Grid. We compare this method of providing collaborative access with the effort involved in joining a more conventional static Virtual Organisation (e.g. VOMS).

TeraGrid: Analysis of Use and Impact on Science

Charlie Catlett, University of Chicago and Argonne National Laboratory

TeraGrid is the US National Science Foundation's nation-wide high-performance computational science facility. Comprised of over a dozen institutions and several dozen major computational, analysis, and data management resources, TeraGrid is a federated infrastructure supporting many modes of computation. Today TeraGrid includes individual systems in the 50-90 Teraflops range, with an aggregate of nearly 300 Teraflops. By 2008 TeraGrid's aggregate computational power will exceed 1 Petaflops, and by 2009 we anticipate individual computing systems within TeraGrid operating at or above the Petaflops level. TeraGrid's software and services have been prioritized since inception through a series of interactions with the computational science community. We analyze the use of TeraGrid by examining the use of software and services and through interviews and surveys of the over 4,000 TeraGrid users. This analysis guides system improvements, services, and policies as the data identifies new and growing user communities with different needs relative to more established user communities. We also examine the scientific impact of large-scale "cyberinfrastructure" such as TeraGrid, factoring in indicators including publications, examples of scientific discovery not possible without TeraGrid, and surveys and interviews with computational scientists.

The AgentX Framework: automating the exchange of information between scientific applications

P. A. Couch, R. Tyer and R. J. Allan, CCLRC Daresbury Laboratory, Daresbury, Warrington, Cheshire WA4 4AD

It is becoming increasingly common for scientists to make use of several applications in close cooperation to address complex scientific problems. There are many examples in the biological domain where the use of hybrid methods, such as QM/MM, is not uncommon. Therefore, the advent of Grid computing and software for specifying and executing workflows promises exciting opportunities for the scientist. Unfortunately, it is very difficult to realise these opportunities; a critical lack of data standards hinders the automatic exchange of information between applications. Although workflows can be specified and executed, bespoke converters and wrappers need to be developed to deal with the transformation of any data that must be exchanged. The development of such converters is a time consuming and complex process, often involving several expert application developers. The same problem arises in the development of common tools for analysing and visualising the results from simulations. In the worst case, specific software components must be developed to handle the data generated from each application. Again, this requires a detailed understanding of the data models specific to each application. Recent efforts to address this problem focus on the development of common data models. This is exemplified by an increasing number of XML based standards for scientific data. However, there are a number of difficulties in finalising these standards, both technical and social. The AgentX framework aims to support the exchange of information between scientific applications without a need to complete the standardisation process. The framework allows information to be extracted from different data sources (such as XML and HDF5 documents and relational databases) through a series of queries based on terms in an ontology. AgentX includes a library that can be used directly by scientific applications and that presents a standard API for extracting data based on its meaning. This frees the scientific application developer from having to understand the interfaces and underlying data models of specific data sources. The framework also provides a suite of tools that can be used to visualise the results of scientific simulations through the generation of XHTML and SVG. These tools process the output of applications through the AgentX library and produce files that can be processed by standard web browsers. In summary, the AgentX framework aims to facilitate the inter-operability of applications that form part of complex workflows and to promote the development of common tools that can be used to analyse and post process scientific data.

Software development using a model driven architecture: the CCPN project

Waybe Boucher, Dan O' Donovan, Darima Lamazhapova, Rasmus Fogh, Wolfgang Rieping, Tim Stevens, Alan da Silva and Ernest Laue, Department of Biochemistry, University of Cambridge

Structural studies using NMR can be considered as a pipeline. The process begins with sample preparation and then various NMR experiments are performed, the data is processed and analysed (i.e. signals are assigned to particular nuclei), and following this assignment process structures are determined and validated. Finally, the data is deposited in a database. Some of the steps in this pipeline are shared with other structural biology disciplines, e.g. X-ray crystallography.

In high-throughput NMR studies, data management along this pipeline is a key bottleneck. In particular, it is desirable to have a formal description of the data and for the data to be consistently tracked all the way through the pipeline, from sample preparation through to deposition. This means that data is not lost, data integrity is maintained, and in the end, data mining can be comprehensively carried out, to the benefit of everyone in the community.

The CCPN (Collaborative Computing Project for NMR, <http://www.ccpn.ac.uk>) was founded in 1999 and one aim was to define a data model for NMR structural biology, and provide associated software libraries to use this data model. The core part of the data model includes descriptions of molecules, structures and NMR. In addition, the model has been extended to include descriptions of sample preparation and tracking, and more recently X-ray crystallography. The data model and libraries (currently implemented in Python, C and Java) together provide a comprehensive framework for data management.

The data model and libraries also allow data exchange between different software programs. Traditionally in NMR, different programs used proprietary data models and formats to store data. This meant that it was difficult for users to move between different programs. If applications use a standard data model and libraries, users can seamlessly move between different programs. Application developers also benefit by being able to code against a well-defined standard, and by not having to worry about the details of persistent data storage. Already there are several NMR programs that use the CCPN data model and we are actively collaborating with many groups (both in NMR and elsewhere) to promote this standard.

The Role of High Performance Computing in CFD within Rolls Royce

Peter Stow, Rolls Royce

Within Rolls Royce CFD or Engineering Simulation is already used extensively in the aerodynamic design of all major engine components. The number of areas of simulation continues to emerge and expand as people see what can be achieved. Even in the more traditional or established areas like Compressor and Turbine design, applications continue to expand with the reality of the simulations increasing as we are now starting to be able to analyse flows of the real engine configuration with the introduction of component interaction effects, introduction of leakage flows or cooling flows, introduction of casing treatment. We are also seeing the start of more design optimisation as people begin to have confidence in the fidelity of the CFD and start to adopt more parametric modelling of engine components. There are a number of reasons for this increase in usage of CFD. In part this is due to our increased physical understanding and ability to model this mathematically, coupled with advances in numerical methods to solve the resulting differential equations. But at the same time a very major impact has come from advances in computers both in the performance of machines but just as important in the price-performance of machines resulting in our ability to make computing power more extensively available to the engineering workforce. This presentation will give an indication of why Rolls Royce has an interest in CFD - what we are hoping from it - what we are achieving - what the future looks like - and the role of High Performance Computing in supporting our aspirations.

Optimising the Data - Information - Knowledge Transformation: solutions for astronomy and their wider relevance

N. A. Walton, University of Cambridge

Astronomy faces a number of major data discovery and interpretation challenges, resulting from the size and complexity of the data streams being generated by current and planned observational (e.g. the Hubble Space Telescope, ALMA, VISTA Infra Red survey telescope) and simulation missions (e.g. the Millennium simulation). This is coupled to the increased complexity of the scientific problems being addressed ranging from the deeper understanding of the early formation of the Universe to the search for extra-solar planets.

The astronomy Virtual Observatory (VO) initiative has emerged to provide solutions to many of these challenges. Interoperability standards have been developed (see <http://www.ivoa.net>) which ensure that all astronomy data and applications can be discovered and accessed by VO systems - e.g. that developed by AstroGrid in the UK (see <http://www.astrogrid.org>).

This presentation focuses on how these VO systems exploit generic computational infrastructures (such as those provided by the National Grid Service in the UK) to enable the mid-to-large scale data transformations supporting a wide range of science cases. The techniques developed have a wider relevance, for instance the use of analysis and access frameworks developed to enable the analysis of large multi-Tera Byte multi-colour image surveys as applied to the analysis of Tissue Microarray data supporting oncological research.

I will discuss areas of cross-domain activity, for example the use of Taverna to support astronomy work-flows, and how the astronomy community is in turn contributing to the future development of this system. Input to the computational community through standards bodies such as the Astronomical Applications Interest Group of the Open Grid Forum will also be highlighted.

Computational Science: Implications for Education and Training

David Gavaghan, University of Oxford

The increasingly interdisciplinary nature of scientific research at the end of the 20th and start of the 21st centuries has major implications for the education and training of the next generation of research scientists. Much of this research has been facilitated by breakthroughs in advanced information and high throughput technologies underpinned by computational science techniques. Drawing on the experience of providing research training at the interface between the physical and the life sciences within the Doctoral Training Centre in Oxford, and in particular in providing training to a cohort of students going on to work within the Integrative Biology e-Science Pilot Project, I will suggest that a new approach is required to computational science training for current and future graduate students moving into interdisciplinary research.

Software Interoperability in Computational Nano- and Materials Science

Michael S. Summers¹, Anton V. Kozhevnikov², Thomas C. Schulthess¹, ¹Oak Ridge National Lab, USA, ²Institute of Metal Physics of the Russian Academy of Sciences

Computational nano- and materials science are interdisciplinary fields that rely on a combination of computational techniques developed in different traditional science areas such as physics and chemistry. As many of these

techniques are implemented in legacy codes with incompatible interfaces, interoperating simulation codes in computational nanoscience poses a great challenge. It is widely believed that this challenge can be overcome with a conversion of the input/output subsystems of simulation codes to technology based on the Extensible Markup Language (XML). We will discuss how this conversion can be done in practice. We begin with a presentation of a simple and portable library to process XML files and Schemas from Fortran and C/C++ codes. We will discuss strategies for Schema development and show how the XML sub-system has been incorporated into electronic structure codes and atomistic simulation systems.

Work funded by the Laboratory Directed Research and Development program at ORNL as well as the Division for Scientific User Facilities of the U.S. Department of Energy.

Production Level Scientific Simulation Management on International Federated Grids

S. Zasada, University College London

Key to broadening participation in the grid is the provision of easy to use access mechanisms and user interfaces, to allow a wide range of users with different skill sets to access the computational and data resources on offer. The Application Hosting Environment is one such middleware tool, hiding much of the complexity of dealing with grid resources from the user, and allowing them to interact with applications rather than machines. The nature of the AHE means that it can be used as a single interface to a wide variety of resources, ranging from those provided at a departmental or institutional level to international federated grids of super computers. The number, range and size of resources made available by federating grids makes possible scientific investigations that would previously not have been feasible. In this paper we describe how we have deployed the AHE to offer access to federated resources provided by the TeraGrid, UK National Grid Service and EU DEISA grid. We also present three case studies where the AHE has been used to facilitate production level scientific simulation across these federated resources.

Software Services for Science from Science: the emergence of software as a facility

Malcolm Atkinson, Director of NeSC

E-Science depends on advances in software and generates new software and methods. A systematic and collaborative approach is needed to establish a virtuous circle of software creation and exploitation. To sustain advances in e-Science this software should be treated as a research facility. A community effort is needed to shape and support the services that will provide software facilities. The talk will review the progress in establishing such services through the UK e-Science Programme, OMII-UK and international collaborations. It will suggest there is a significant transition in research enabled by recent technical advances. This change provokes questions about the road map for software facilities.

