

---

## Applications/Algorithms Roadmapping Activity

---

### Workshop 3 Report

---

January 2009

---

## Contents

Preamble.....	2
Background.....	2
Overview of Workshop 3.....	2
Overview of Presentations.....	3
Prof Stan Scott (Queen’s University, Belfast) Emerging HPC Technologies: back to the future?.....	3
Dr Steven Kenny (Loughborough University) Accelerating Simulations using Computational Steering..	4
Prof Bruce Boghosian (Tufts University) Spacetime computing: A Dynamical Systems Approach to Turbulence.....	4
Dr Massimiliano Fatica (NVidia) CUDA for High Performance Computing.....	4
Dr Ram Rajamony (IBM Research, Austin) Productive Petascale systems and the challenges in getting to Exascale systems.....	5
Dr Maziar Nekovee (BT Research & UCL) High-Performance Computing for Wireless Telecom Research.....	5
Prof Philip Treleaven (UCL) UK PhD Centre for Financial Computing.....	5
Dr Charlie Laughton (Nottingham University) Biomolecular Simulation: Where do we go from here?..	6
Dr Sabine Roller (HLRS, Stuttgart) Challenges and opportunities in hybrid systems.....	6
Dr Kevin Stratford (University of Edinburgh, EPCC) HPC for Soft Matter Physics: Present and future....	7
HPC-NA Roadmap Presentation & Discussion.....	7
Algorithms.....	8
Major issues for the future.....	8
Prioritization Axes.....	8
APACE Website.....	8
Issues Identified.....	9
Summary of Breakout Group Discussions.....	9
Breakout Group 1: Numerical Aspects of HPC-NA.....	9
Breakout Group 2: Applications and Algorithms.....	10
Breakout Group 3: Infrastructure & Architectures.....	10
Final Discussion.....	11
Next Steps.....	12
Contacts and further information.....	12
Acknowledgements.....	12
Annex 1: Workshop Attendees.....	13
Annex 2: HPC/NA Workshop 3 Agenda.....	14

## **Preamble**

This report provides an overview of the third and final workshop of the current HPC/NA Roadmapping project. It includes all information pertaining to the workshop. The report is organized such that all outputs and likely outcomes are reported in the body of the report; all supporting material is attached in annexes.

The main purpose of the HPC/NA activity is to get community input into the roadmap development. This document therefore should be seen as a report of a specific event in that activity and not as a final statement of any kind– we welcome constructive input of any sort whether to support the findings or indeed to question them.

Contributions to the discussion can happen by emailing the contacts provided below or by engaging through the project website.

## **Background**

The applications/algorithms roadmapping activity has developed the first instantiation of a high performance numerical algorithm roadmap. The roadmap identifies areas of research and development focus for the next five years including specific algorithmic areas required by applications as well as new architectural issues requiring consideration. It provides a co-ordinated approach for numerical algorithm and library development.

Many applications from different fields share a common numerical algorithmic base. We have captured the elements of this common base, identified the status of those elements. The next step, in conjunction with the EPSRC Technology and Applications roadmapping activity, is to determine areas in which the UK should invest in algorithm development.

A significant sample of applications, from a range of research areas, is included in the roadmapping activity. The applications investigated include those in the EPSRC Technology and Applications roadmap, and others that represent upcoming and potentially new HPC areas.

The applications provided the basis to understand:

- The role and limits of a common algorithmic base
- How this common algorithmic base is currently delivered and how should it be delivered in the future
- The current requirements and limitations of the applications, and how these should be expanded
- The “road-blocks” that limit the scope of the future exploitation of these applications.
- A better comprehension of the “knowledge gap” between algorithmic developments and scientific deployment
- How significant computing language as well as other “practical” issues weigh in the delivery of algorithmic content

## **Overview of Workshop 3**

This workshop focused on discussing and reviewing the first instantiation of the National Roadmap for HPC Algorithms and Applications, published just before the workshop in January 2009.

V1.1 Wednesday, February 18, 2009

Additional aims of the workshop were to:

- Gain national and international perspectives from a range of speakers
- Review the roadmap document
- Present a community view to the Research Councils via EPSRC
- Think about the next steps and continuing the dialogue

There were 34 attendees including four international participants at the workshop, held at the Royal Society, London over 26<sup>th</sup>-27<sup>th</sup> January 2009. A full list of attendees is provided in annex 1. There was a good mix of application developers, computer scientists, numerical analysts and vendors including system architects.

The agenda for the meeting is provided in annex 2.

## Overview of Presentations

Copies of the presentations may be downloaded from the HPC-NA project website.

### Prof Stan Scott (Queen's University, Belfast) Emerging HPC Technologies: back to the future?

Prof Scott gave a comprehensive overview of the current hardware developments having an impact on the immediate future of HPC: GPUs, FPGAs, heterogeneous chips (e.g. Intel Larrabee), floating point accelerators (e.g. Clearspeed), and heterogeneous systems (e.g. IBM Roadrunner). He commented that in the future hybrid computing, i.e. collections of heterogeneous components, would become more common. However, he sounded a warning as novel technologies often appeared unstable and might well be discontinued. Stan highlighted a study by Sandia, showing the impact of the "memory wall", i.e. the inability to "feed" data fast enough to hungry processors due to bandwidth constraints, and how this resulted in decreasing parallel efficiency as the number of cores grew. He added that there was the possibility of a rebirth of SIMD algorithms, at least in some specific cases. Commenting on GPUs, he expressed some concern about their potential non-compliance to IEEE 754 standards, and wondered whether that could have a detrimental effect on numerical stability. Stan highlighted some efforts to ease the situation: PLASMA, a multi-core aware follow-up to the LAPACK genealogy; mixed-precision numerical computation, to achieve best performance from mostly 32-bit engines such as GPUs; HONEI, a collection of libraries targeting multi-processing architectures; Sequoia, a memory layout aware programming language that could lead to self-adapting algorithms.

Finally, he described some recent work in the UK.

- A recently formed consortium, a joint effort by the Universities of Belfast, Imperial, Cardiff, Bristol and Southampton and Daresbury Laboratory, to study a collection of large codes from the CCP community. The aim being not only to provide high performance versions of these codes but also to abstract general principles and guidelines for the design of applications and algorithms for emerging and future HPC systems
- The recently formed Multi-core and Reconfigurable Supercomputing Network (MRSN), an initiative led by Oxford, Imperial, Belfast and Manchester.

He reported to the meeting that a conference dedicated to emerging HPC technology, MRSC 2009, would be taking place in Berlin on March 25-26 2009.

### **Dr Steven Kenny (Loughborough University)** **Accelerating Simulations using Computational Steering**

Dr Kenny reported on his group's investigation on materials required for energy demand reduction, particularly on glass plates. Steven reported that the simulation of these glasses would require a three-fold iteration, each iteration been determined by a host of parameters that were very difficult to compute automatically. Convergence criteria and requirements were, in particular, difficult to assess and tended to vary considerably even across closely related problems. Hence, their solution consisted in allowing computational steering, in the sense of altering manually on-the-fly the parameters responsible for the computation at each iteration level. He commented that this allowed researchers to make best use of their own understanding of the problems studied and resulted in increased speed-up and better utilisation of resources.

He found that current set-ups, particularly at National centres, inhibited the possibility of running of a wide adoption of computational steering and that further developments would be required for complicated coupled simulations.

### **Prof Bruce Boghosian (Tufts University)** **Spacetime computing: A Dynamical Systems Approach to Turbulence**

Prof B.Boghosian reported on a class of algorithms being developed to tackle the very computationally intensive problem of turbulence at high Reynolds' numbers, such as in flow past aircraft. This spacetime computing, exemplified by so-called Parareal algorithms, employs coarse and fine time grids. The coarse grid, with purely sequential evolution, can be used as a predictor, and then a number of time slices (the fine grid) can be computed in parallel for each coarse time step, and acting as a corrector for the coarse grid results. In other words, these methods can be viewed as achieving domain decomposition in time. Bruce commented that these or similar methods could provide the only way of making turbulence simulation faster than real time.

Additionally, with periodic boundary conditions in time, it would be possible to use this method to generate the discrete set of unstable periodic orbits (UPOs) of a given flow. The enumeration of these, a project that requires petascale computing, would be of enormous help in extracting averages of observables over the turbulent flow. The so-called dynamical zeta function formalism reduces such averages to combinations of those over the UPOs.

Bruce added that other these ideas, particularly parareal algorithms, could affect other areas by providing means of parallelising evolution equations and efficiently extracting statistical results.

### **Dr Massimiliano Fatica (NVidia)** **CUDA for High Performance Computing**

Dr M.Fatica, from NVidia, first introduced the newest NVidia GPUs (Tesla T10), their architecture and their performance capabilities. Of particular relevance was the introduction of a number of dedicated Double Precision threads (cores) within the GPUs and their (the DP units, that is) compliance to IEEE 754 standards.

CUDA, based on standard C, is the language used for programming these GPUs. CUDA encapsulates a thread-based approach to parallelism and allows the mapping of threads to the GPU thread arrays. Massimiliano reported that through CUDA a number of applications had benefited by either being

V1.1 Wednesday, February 18, 2009

ported directly to GPUs or by employing GPUs as accelerators for specific computationally intensive portions, with considerable faster performances than achievable by conventional hardware.

**Dr Ram Rajamony (IBM Research, Austin)**  
**Productive Petascale systems and the challenges in getting to Exascale systems**

Dr R.Rajamony reviewed the current contribution of IBM to advanced computing. In particular, he singled out several issues that affect scalable performance, but were not always evident from headlines performance figures such as overall network bandwidth and access to memory. Ram pointed out that the new Blue Gene/Q would be addressing these issues and would preturn sustainable performance figures beyond what could be achievable elsewhere.

Ram also reported on IBM efforts to create a new model of parallelism, based on direct access to remote data (PGAS), in a way akin to virtual shared memory model. This approach to parallelism would depend on the application code to guarantee memory integrity in the course of computation, using a number of provided locking primitives.

**Dr Maziar Nekovee (BT Research & UCL)**  
**High-Performance Computing for Wireless Telecom Research**

Dr M.Nekovee gave an overview of computational modelling and optimisation for modern wireless telecom systems. In particular, he concentrated on V2V (Vehicle-to-Vehicle) networks for future intelligent transports. In the future, he argued, such capabilities would be used for provision of broadband access to millions of vehicles, for traffic monitoring and optimisation, to convey and relay information to vehicles, to allow intelligent decision making, for example, for routing and reducing traffic congestion.

The simulation poses extreme difficulties as phenomena of widely differing time scales need be considered: vehicles on a slow time scale, of the order of seconds, the wireless network and its requirements on a time scale of microseconds. Maziar said that such coupled simulation played an increasing role in industry as well as defined the state-of-the-art and posed extreme challenges to HPC, possibly requiring altogether novel approaches and algorithms, e.g, for scalable parallel discrete event simulations.

**Prof Philip Treleaven (UCL)**  
**UK PhD Centre for Financial Computing**

Prof P.Treleaven described the new Doctoral Training Centre for Financial Computing at UCL. The new Centre aroused considerable interest and was supported by a large number of major financial outfits. He commented that, despite contingent worldwide economic difficulties, HPC was viewed by all major banks and financial institutions as a key technology. The UCL-lead PhD programme had great appeal to them and aimed to facilitate the already strong links between UK banks and UK Universities.

Philip then gave an overview of the various aspects of Financial Computing: financial modelling, mathematical finance, financial engineering and computational finance. A wide range of algorithm was employed and more would need to be developed to cater for the needs of the financial world: e.g. automated learning, sophisticated statistical analysis, probabilistic methods, many flavours and techniques of optimisation, dynamic programming, Monte Carlo simulations, etc. On the computational end of the spectrum he saw the increasing importance of automatic and semiautomatic trading systems.

He added that a number of key banking applications required HPC approaches, such as derivatives pricing, portfolio management and risk management. The importance of risk management was increasing also due to the expanding role of financial regulators.

### **Dr Charlie Laughton (Nottingham University) Biomolecular Simulation: Where do we go from here?**

Dr C.Laughton gave a comprehensive overview of the challenges facing MD (Molecular Dynamics) approaches to biomolecular simulations. He said that the focus of interest is shifting towards the study of complex systems over the millisecond-to-second timescale. This was unachievable by current technologies and algorithms as it was many orders of magnitude beyond their capabilities. The evaluation of the forces between interacting particles represented a serious bottleneck. In particular, various approximations used to compute long-range forces had not proved to scale sufficiently well to very large numbers of cores. A second bottleneck is the short simulation time step needed to keep the algorithms stable.

Charlie concluded that radically new approaches were necessary. Coarse graining, i.e. aggregating a number of particles into larger objects (e.g. several atoms in one molecular group) could speed-up things considerably, but still not adequately for the long-term requirements of the field. Larger still objects, such individual biomolecules, should be representable. However, Charlie said, these objects would have internal structure and flexibility and developing efficient methodologies to represent them, and the interactions between such complex objects, would pose a formidable challenge. Many of the properties of these larger objects with internal structure could be inferred from studies at a smaller scale of them and their components. This approach would then make best use of a body of knowledge already accumulated. Charlie surmised that should this approach prove feasible, then grand challenges of computational biology, such as the simulation of a whole bacterial cell, could be realistically tackled.

### **Dr Sabine Roller (HLRS, Stuttgart) Challenges and opportunities in hybrid systems**

Dr S.Roller described how HPC was currently organised in Germany and the role of HLRS (Stuttgart). She said that funding was divided into three, roughly identical, portions: the first for the three National centres (Stuttgart, Munich and Juelich); the second for the ten Regional Centres with specific domain focus (Aachen, Berlin, Hannover, etc); the third portion for University-based HPC-servers. This pyramidal structure, Sabine commented, served a number of purposes, from allowing applications to “scale up” to the large National platforms, to allow the “trickling down” of know-how and algorithms from the high-end to smaller systems. This last point, she added, was seen as having great importance, particularly in view of the strong ties between the research and industrial communities.

Sabine then reviewed the work carried out at HLRS to employ different architectures for different portions of a specific application. She explained that that was the meaning of “hybrid computing”: creating a computing environment made up of different technologies and optimise applications on this. In the HLRS case, traditional cache-based as well as vector processors were made available to an aero-acoustic application and the grid and numerical methods were mapped to the two architectures employed. She then proceeded to show that much higher performance could be achieved by a hybrid system than by a purely cache-based or vector-based system.

## **Dr Kevin Stratford (University of Edinburgh, EPCC) HPC for Soft Matter Physics: Present and future**

Dr K.Stratford first explained that by “soft matter” he meant the study and simulation of liquid, gels, foams etc., such as the study of liquid crystals “Blue Phases”, binary fluid under strain, suspensions (ferrofluids).

Kevin showed that the study of blue phases of liquid crystals is acquiring great technological importance, for example for next generation of fast-switching, low-power displays. He reported that the phase transition could be simulated by solving the Navier-Stokes equations via a lattice Boltzmann method. A similar computational approach could be used to simulate binary fluids under strains as well as colloidal suspensions of particles subject to long-range forces (e.g. magnetic particles).

Kevin reported that their main code for lattice Boltzmann computations employed MPI parallelism and had been ported successfully to a number of HPC platforms. The code was not publicly available, and was unlikely to become so in the immediate future for a number of reasons. Work was underway to include better kernels (BLAS, PLASMA, etc), algorithmic enhancements and, possibly, to port part of the computation to novel architectures such as FPGAs, GPUs, etc. He also said that the computation of long-range electromagnetic forces inhibited scalability to large number of processors.

## **HPC-NA Roadmap Presentation & Discussion**

The full National HPC-NA Roadmap for Applications and Algorithms is published in a separate document; this report provides a summary of the presentation and previous work:

- Workshop 1: Oxford, Nov 2008
- Workshop 2: Manchester, Dec 2008
- Background work considering DOE/DARPA/NSF workshops
- Discussions with applications outside of workshops

The first version of the roadmap document is the outcome of the two community meetings together with input from similar activities elsewhere. The roadmap activity aims to provide a number of recommendations that together will drive the agenda toward the provision of:

- Algorithms and software that application developers can reuse in the form of high-quality, high performance, sustained software components, libraries and modules
- a community environment that allows the sharing of software, communication of interdisciplinary knowledge, and the development of appropriate skills.

The first version of the roadmap is built around five themes that have evolved during the discussion within the community.

- Theme 1: Cultural Issues
- Theme 2: Applications and Algorithms
- Theme 3: Software Engineering
- Theme 4: Sustainability
- Theme 5: Knowledge Base

Each of these is represented in the roadmap. As the roadmap activity goes forward we expect that these initial actions to develop into a detailed map of priorities across a sensible timeframe.

## Algorithms

- Optimisation
- Scalable FFT
- Adaptive mesh refinement
- Eigenvalue/eigenvector (all or few)
- Iterative & direct solvers
- Monte Carlo
- Out of core algorithms to enable larger problems

## Major issues for the future

The roadmap identifies a number of major issues of high importance that future work should be focussing on:

- Load balancing
  - meshes
  - particle dynamics and computation of interactions
- Better software environments for complex application development
- Adaptive software to automatically meet architectural needs
- Use of novel architectures (in the immediate future)
  - FPGAs
  - GPUs
  - IBM Cell
  - Other....
- Coupling between different models and codes
- Error propagation
- Scalable I/O
- Visualisation

## Prioritization Axes

- Key applications
- Algorithms
- New approaches due to architectural issues
- Software development issues
- Skills
- time frame for each

## APACE Website

The APACE website is planned as a solution to support the development of a community environment that allows the sharing of software and communication of interdisciplinary knowledge.

- **AP**plication **A**dvanced **C**omputing **E**xchange
- Community site built on same lines as myExperiment<sup>1</sup>, a collaborative environment where scientists can publish their workflows and experiment plans, share them with groups and find

---

<sup>1</sup> <http://www.myexperiment.org/>

those of others. Workflows, other digital objects and collections (called Packs) can be swapped, sorted and searched like photos and videos on the Web. myExperiment enables scientists to contribute to a pool of scientific workflows, build communities and form relationships. It enables them to share, reuse and repurpose workflows and reduce time-to-experiment, share expertise and avoid reinvention.

- APACE will facilitate collection of information around
  - numerical analysis algorithms,
  - definition of applications in terms of algorithms
  - Expertise in applications and algorithms
  - Global activity in development etc
  - Build community groups and sharing ideas, information and software

## Issues Identified

Arising out of the discussions over the two days, a number of issues were identified that would improve and focus the initial draft of the HPC-NA roadmap

- Need to identify exemplar “baseline” projects
- Develop scenarios & timelines
- Prioritisation of themes and/or algorithms
- NA specific “actions” for roadmap
- Getting & retaining engagement from the various communities
- “Sustainability” as one of the themes or as a cross-cutting issue?
- Next step – EPSRC Network application – participation & ideas

These issues were taken up through three breakout groups focusing initially on issue 2 – the development of scenarios and timelines.

## Summary of Breakout Group Discussions

### Breakout Group 1: Numerical Aspects of HPC-NA

Prof. N.Higham reported on behalf of the breakout group. He highlighted a number of key points:

#### 1. Numerical precision aspects

This arises arising from the non-IEEE compliant (single or double precision) arithmetics on GPUs and FPGAs, along with variable and fixed precision on FPGAs. Its importance is enhanced by the large number of time steps required by integrators (order  $10^5$  or higher), which magnifies rounding errors. This issue has arisen only recently and its importance has become increasingly apparent during the 3 workshops (as well as at the Jan 09 MRS Network workshop in Belfast).

Urgency: high

Timescale: short. Good progress can be made over the course of a 3 year project. Work is already underway in Manchester (Jan - Mar 09) as part of the MRS network to survey the literature and identify key applications where precision problems arise.

## 2. Error propagation in coupled models

In particular, this includes error control in adaptive PDE solvers, a topic mentioned in previous workshop reports.

## 3. Input from numerical analysts to applications scientists

This could take the form of advice on choice of algorithms. While it will be facilitated by APACE., numerical analysts would find difficult to find the time to provide "free consultancy" - their time will need to be costed.

## 4. Study Groups

The annual Smith Institute "study groups with industry"<sup>2</sup> in applied mathematics have been very successful. An analogous activity could be undertaken here in the form of numerical analysts and computer scientists working with applications scientists in intensive EPSRC-funded workshops focused on a small number of key applications. These are a necessary follow-on to the 3 workshops so far in order to delve deeper into technical aspects. Experience from the workshops suggests there are willing participants, subject to their availability.

### Breakout Group 2: Applications and Algorithms

Dr S.Salvini reported on behalf of the group. The group discussed a number of numerical aspects common to a range of applications that could provide exemplar "baseline" PROJECTS

- Multiscale problems/simulations: encapsulation, manipulation of complex physical objects (i.e. with an internal structure, e.g. molecule, cellular structures etc) and their interactions. Timescale: long term.
- Long range interactions for particle models. Several speakers from different fields (molecular dynamics, plasma physics, astrophysics, material sciences, etc.) reported that this constituted a serious bottleneck that inhibited scalability to large numbers of cores. Current algorithms, mostly based on FFT, have proved inadequate and new ideas and solutions need to be sought. Time scale: short to medium term.
- Generalised Hermitian/symmetric eigenproblems arise in many fields (quantum chemistry, material sciences, etc). Standard LAPACK/ScalAPACK provisions do not scale satisfactorily with increasing number of cores; in many cases most of the computation time is spent in solving these eigenproblems.

There was some discussion also about the delivery vehicles for algorithmic content, possibly beyond simple libraries, automatic code generation to achieve optimal performance on specific target systems, high level abstractions and their suitability and use (possibly along the lines of PLASMA). Unfortunately, the time allocated was not sufficient to explore these themes in sufficient depth.

### Breakout Group 3: Infrastructure & Architectures

Dr John Brooke reported on behalf of the third breakout group. The group discussed on a number of issues of general import relating to infrastructure, system architectures and suggested some practical steps that should be undertaken.

The group noticed that the UK HPC landscape is currently dominated by the National Centres. At the same time, porting codes across architectures has proved a particular bottleneck because of the lack of a systematic approach.

---

<sup>2</sup> [http://www.smithinst.ac.uk/Mechanisms/StudyGroups/index\\_html](http://www.smithinst.ac.uk/Mechanisms/StudyGroups/index_html)

V1.1 Wednesday, February 18, 2009

In order to improve the situation, it would be essential to ensure future proofing of codes, to avoid any much time consuming re-engineering. Future UK funding and purchasing decision could be based on actual delivered performance.

A number of practical steps should be undertaken:

- In similar fashion to Germany, effort should be spent to make sure that technologies and know-how can “trickle down” from the National Centres to smaller installations. This would be of benefit not just to the academic research world but to industrial and commercial concerns.
- The NA community should provide insight and support to application developers
- The APACE website could be a good starting point, but it should provide more than just a library of algorithms and knowledge on how to employ them
- Coordination with those involved in supporting applications on high-end systems, e.g. CCPS, NAG, etc, would also be essential

## Final Discussion

It was agreed that the existing version 1 of the HPC-NA Roadmap should be circulated widely at once, beyond the immediate circle of the Workshop participants. In the light of this third Workshop, a redrafted HPC-NA Roadmap should be put on the Website at once and circulated among the participants for comments and corrections. It was agreed that the final revised version would be completed by the 20<sup>th</sup> February.

There was also agreement that the HPC-NA Roadmap presented to EPSRC would contain the following “exemplar” applications:

- Numerical precision issues, raised by breakout group 1.
- Coupled problems, and error propagations in mixed models, as described by breakout group 1.
- Scalable algorithms for the modelling of long range forces, as described by breakout group 2.

It was also agreed that others should be solicited and could be proposed from outside the Workshop. These would be added to the basket of “exemplar” issues/applications in due course.

Computational chemistry was highlighted as another possible exemplar, focussing on the bottleneck of the Hermitian generalized eigenvalue problem (GEP), as identified e.g. by Dr Kenny and Dr Sutherland. This would have the advantage of UK NA expertise, as well as wide interest in its application.

Establishing a Network activity as a follow-up to this series of Workshop was unanimously supported. It was generally perceived as an important step in bringing together numerical analysts, computer scientists and application researchers/developers.

APACE was supported by all present; it was also felt that a prototype should be set up as soon as possible.

The Meeting also recognised the importance of an international dimension to all UK efforts and funding towards this should be actively sought. European funding for projects within this general remit should also be applied for.

Short duration study groups as in applied maths were seen as a very good idea, provided funding could be secured

## Next Steps

- Submit a proposal for Network funding to EPSRC to keep the HPC-NA community activity alive after the end of this initial project.
  - Suggestions for activities and membership can be sent to the project contacts below
- The roadmap is being further developed from the outputs of this workshop and is now in a separate document.
- Development of APACE website & initial user testing

## Contacts and further information

### Issues and input to this Report

Dr Mark Hylton:

[mark.hylton@oerc.ox.ac.uk](mailto:mark.hylton@oerc.ox.ac.uk)

### General input to Activity

Prof. A. E. Trefethen, OeRC, University of Oxford

[anne.trefethen@oerc.ox.ac.uk](mailto:anne.trefethen@oerc.ox.ac.uk)

Prof P. V. Coveney, University College London

[p.v.coveney@ucl.ac.uk](mailto:p.v.coveney@ucl.ac.uk)

Prof N. J. Higham, University of Manchester

[nicholas.j.higham@manchester.ac.uk](mailto:nicholas.j.higham@manchester.ac.uk)

Prof I. S. Duff, STFC, Rutherford-Appleton Laboratory

[iain.duff@stfc.ac.uk](mailto:iain.duff@stfc.ac.uk)

### Project website

[www.oerc.ox.ac.uk/research/hpc-na](http://www.oerc.ox.ac.uk/research/hpc-na)

## Acknowledgements

We are grateful for the support provided by EPSRC for the algorithm/application roadmapping activity. Thanks also to Peter Coveney for organizing the speaker programme and Nilufer Betik of UCL for local arrangements.

## Annex 1: Workshop Attendees

Jamil Appa	BAe Systems
Tony Arber	Warwick University
Bruce Boghosian	Tufts University
John Brooke	University of Manchester
Anthony Ching Ho Ng	Smith Institute/NAG
George Constantinides	Imperial College London
Peter Coveney	UCL
Iain Duff	STFC, RAL
Massimiliano Fatica	NVIDIA
Grzegorz Gawron	HSBC IB
Mike Gillan	UCL
Gerard Gorman	Imperial College London
John Gurd	University of Manchester
Nick Higham	University of Manchester
Mark Hylton	Oxford University
Emma Jones	EPSRC
Crispin Keable	IBM
Steve Kenny	Loughborough University
Igor Kozin	STFC Daresbury
Charlie Laughton	University of Nottingham
Maziar Nekovee	BT Research & UCL
Stephen Pickles	STFC Daresbury
Ram Rajamony	IBM Research, Austin
Graham Riley	University of Manchester
Sabine Roller	HLRS, Stuttgart
Radhika Saksena	UCL
Stef Salvini	Oxford University
Stan Scott	Queen's University Belfast
David Silvester	University of Manchester
Edward Smyth	NAG
Kevin Stratford	University of Edinburgh (EPCC)
Andrew Sunderland	STFC Daresbury
Anne Trefethen	Oxford University
Philip Treleaven	UCL

## Annex 2: HPC/NA Workshop 3 Agenda

Royal Society, 6-9 Carlton Terrace, London, SW1Y 5AG

### Day 1: Monday 26<sup>th</sup> January 2009, Council Room (First Floor)

- 12:00-13:00 Lunch
- 13:00-13:15 Welcome & Introduction  
*First Speaker Session Chaired by Peter Coveney*
- 13:15-13:45 • Prof Stan Scott, Emerging HPC technologies: back to the future?
- 13:45-14:15 • Dr Steven Kenny, Accelerating Simulations using Computational Steering
- 14:15-14:45 • Prof Bruce Boghosian, Spacetime Computing: A Dynamical Systems Approach to Turbulence
- 14:45-15:15 Refreshments
- 15:15-16:30 *Presentation of the HPC-NA Roadmap*  
Lead: Anne Trefethen  
Discussion around roadmap and feedback  
*Second Speaker Session Chaired by Iain Duff*
- 16:30-17:00 • Dr. Massimiliano Fatica, NVIDIA, CUDA for High Performance Computing
- 17:00-17:30 • Dr Ramakrishnan Rajamony, IBM Research, Austin, Productive Petascale systems and the challenges in getting to Exascale systems
- 17:30-18:00 • Dr Maziar Nekovee, BT Research & UCL, *High-Performance Computing for Wireless Telecom Research*
- 18:00 Drinks Reception – ‘City of London’ Room 2 (Ground Floor)
- 19:00 Workshop Dinner – ‘City of London’ Room 1 (Ground Floor)

### Day 2: Tuesday 27<sup>th</sup> January 2009, Kohn Centre and Marble Hall (Ground Floor)

- 09:00-09:15 Welcome & Introduction  
*Third Speaker Session Chaired by Nick Higham*
- 09:15-09:45 • Prof Philip Treleaven, New UK PhD Centre in Financial Computing
- 09:45-10:15 • Dr Charlie Loughton, Biomolecular Simulation: Where Do We Go From Here?
- 10:15-10:45 • Dr Sabine Roller, Challenges and opportunities in hybrid systems
- 10:45-11:15 Refreshments
- 11:15-11:45 • Dr Kevin Stratford, *HPC for Soft Matter Physics: Present and future*
- 11:45-12:30 Facilitated discussions around roadmap and feedback
- 12:30-13:30 Lunch
- 13:30-14:30 Further development of roadmap and plans for taking it forward
- 14:30-15:00 Next steps, final comments & close