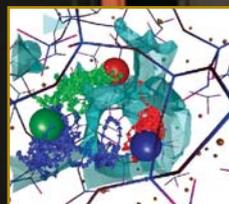
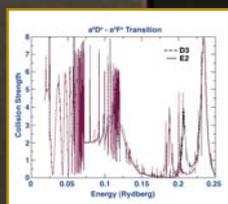
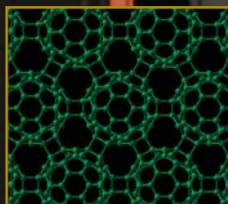
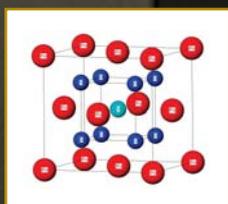
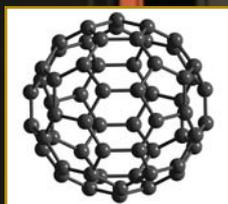
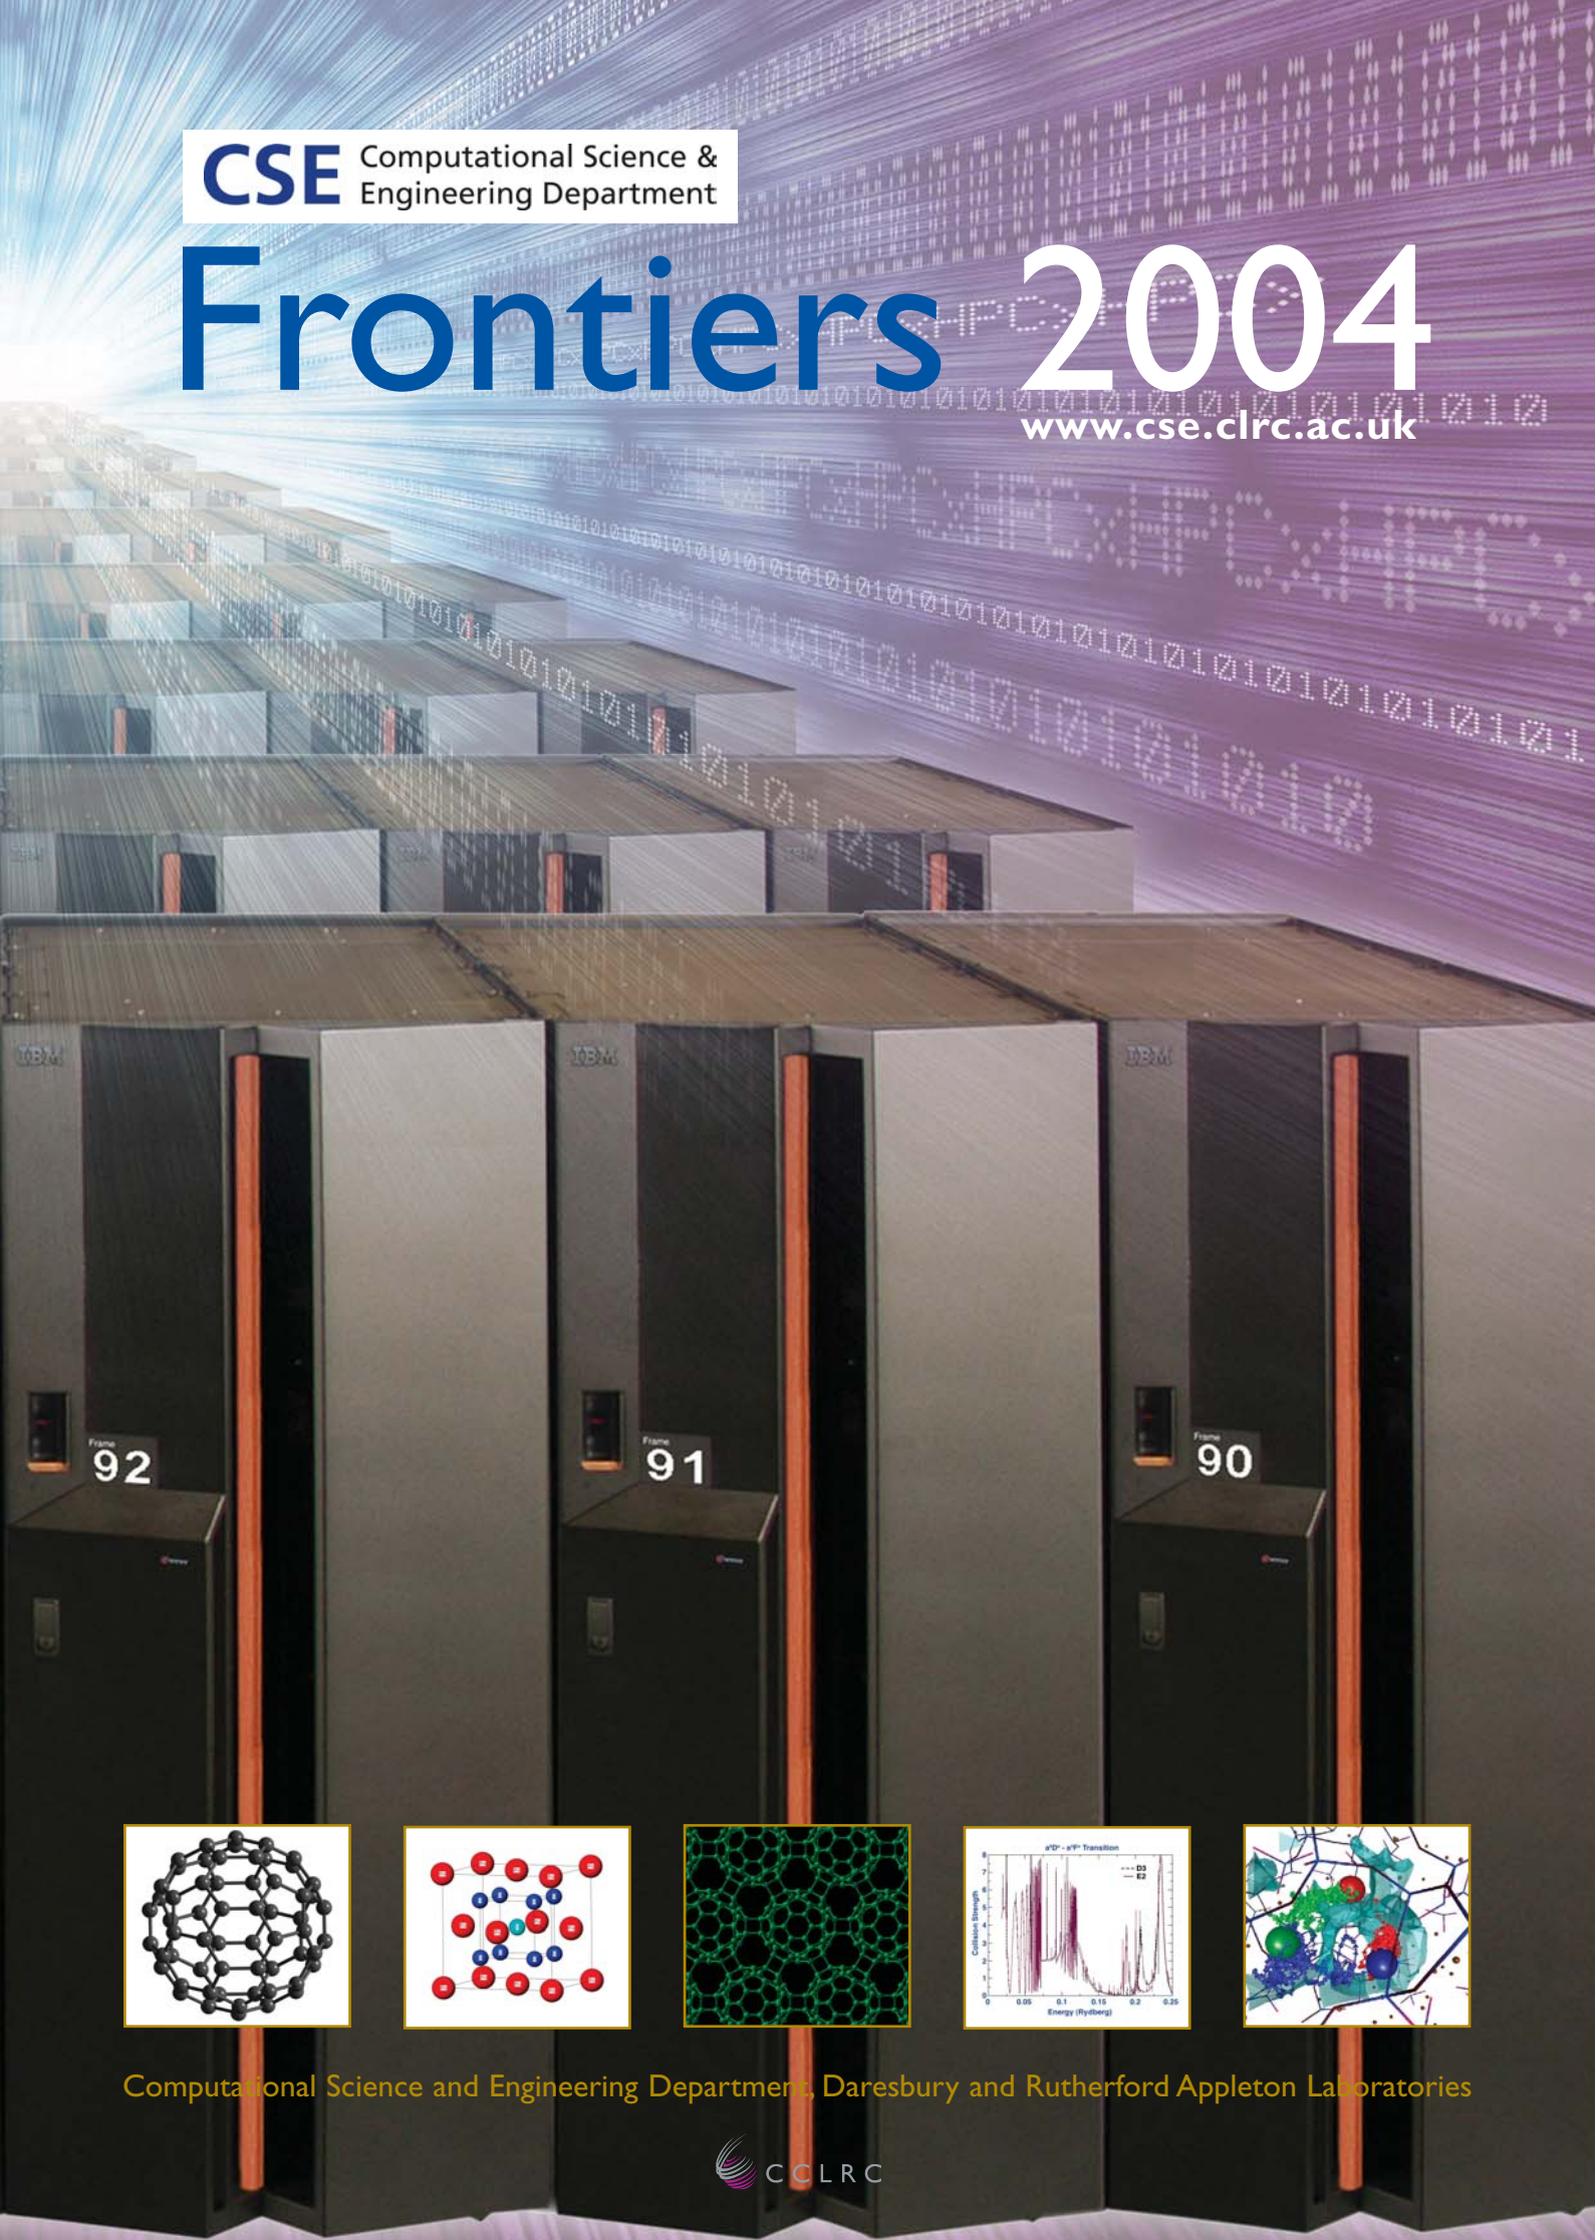


**CSE** Computational Science & Engineering Department

# Frontiers 2004

[www.cse.clrc.ac.uk](http://www.cse.clrc.ac.uk)



Computational Science and Engineering Department, Daresbury and Rutherford Appleton Laboratories



## Contacts and Further Information

Prof Paul Durham	Director, Computational Science and Engineering Department	p.j.durham@dl.ac.uk
Prof Martyn Guest	Associate Director, Computational Science and Engineering Department	m.f.guest@dl.ac.uk
Dr Richard Blake	Associate Director, Computational Science and Engineering Department	r.j.blake@dl.ac.uk
Mrs Shirley Miller	Personal Assistant	s.miller@dl.ac.uk
Mr Damian Jones	Administrative Assistant	d.r.jones@dl.ac.uk
Miss Nina Woodall	Administrative Assistant	n.m.woodall@dl.ac.uk

### Group leaders

Dr Steve Andrews	HPCx Team	s.j.andrews@dl.ac.uk
Dr Mike Ashworth	Advanced Research Computing	m.ashworth@dl.ac.uk
Dr Miles Deegan	Distributed Computing Systems	m.j.deegan@dl.ac.uk
Prof Iain Duff	Numerical Analysis	i.s.duff@rl.ac.uk
Dr Dave Emerson	Computational Engineering	d.r.emerson@dl.ac.uk
Prof Chris Greenough	Software Engineering	c.greenough@rl.ac.uk
Prof Nic Harrison	Computational Materials Science	n.m.harrison@dl.ac.uk
Mr David Laff	Village Systems	d.j.laff@dl.ac.uk
Dr Bob McMeeking	Chemical Database Service	r.f.mcmeeking@dl.ac.uk
Prof Cliff Noble	Atomic and Molecular Physics	c.j.noble@dl.ac.uk
Dr Paul Sherwood	Quantum Chemistry	p.sherwood@dl.ac.uk
Dr Bill Smith	Molecular Simulation	w.smith@dl.ac.uk
Prof Walter Temmerman	Band Theory	w.m.temmerman@dl.ac.uk

### Websites

Computational Science and Engineering Department	www.cse.clrc.ac.uk
HPCx Service	www.hpcx.ac.uk
EPSRC HPC Programme	www.epsrc.ac.uk/hpc
CCLRC	www.clrc.ac.uk

## Addresses

### CCLRC Rutherford Appleton Laboratory

Chilton  
Didcot  
OX11 0QX  
Tel: 01235 445000  
Fax: 01235 445808

### CCLRC Daresbury Laboratory

Keckwick Lane  
Daresbury  
Warrington  
WA4 4AD  
Tel: 01925 603000  
Fax: 01925 603100

<b>Contents</b>	<b>Page</b>
Introduction	1
Delivering Scientific Productivity on HPCx	2
Atomic Physics, Astrophysics and HPCx	4
Building an Infrastructure for Computational Quantum Chemistry	6
The Use of Partitioning in Solving Sparse Linear Systems	8
Non-intuitive Flow Phenomena in Gas-phase Microflows	10
Magnetic Carbon	12
First Principles Study of the Stability of PuO <sub>2</sub> Under Oxidisation	14
Virtual Reality Molecular Dynamics of Ion Diffusion in Glass	16
Dynamics of Crystals in Rigid Organic Molecules	18
High Performance Computing Molecular Dynamics Simulations	20
Studies on Manganese Oxide Perovskites in Conjunction with CDS	22
The TeraGyroid Experiment Supercomputing 2003	24

# Introduction

It is a pleasure to introduce the 2004 issue of CSE Frontiers, an annual report from CCLRC's Computational Science and Engineering Department.

Our Department's role is to ensure that UK researchers benefit from the best computational methods and support, and we can only do this by collaborating closely with University and Research Council colleagues to combine fundamental knowledge of relevant scientific or engineering domains with powerful numerical methods and modern software engineering techniques.

As usual, this particular document makes no attempt at a comprehensive account of all these activities; you can find lots of details of our projects, services, facilities, events etc on our web site <http://www.cse.clrc.ac.uk>, which I hope you'll visit and bookmark. Instead, this booklet gives a selection of highlights from our current work and interests. "Frontiers" gives a glimpse of the range of our work, from atomic physics through molecular science and novel materials to fluid flows, and the mathematical and computational techniques that underpin it. It also hints at the practical impact of this work on astrophysics, biochemical sensors, pharmaceutical processing, radiation damage, nuclear recycling and more.

Of course, the HPCx service, the UK's first tera-scale research computing facility, continues to be a strong focus for us and our partners in the Edinburgh Parallel Computing Centre. HPCx has had a very successful first year of service, with higher than expected capability usage and some excellent results on code performance. Mike Ashworth's article touches on this. At *Supercomputing 2003* in Phoenix last November, HPCx collaborated with the RealityGrid Project, CSAR at Manchester and the US TeraGrid centers (NCSA, San Diego, Pittsburgh, Argonne and Caltech) in implementing a Grid-based application of lattice Boltzmann simulations, and won a Global Challenge award for the most innovative data-intensive application. The first half of 2004 saw a lot of hard work in planning and executing the phase 2 upgrade to a system with a peak performance of about 11 Tflops, and an interconnect with much higher performance.

Important and headline-grabbing as these hardware developments are, the key to real scientific impact is our R&D programme on applications software - its functionality and its performance. This will always be at the heart of CSED's activities, as the UK computational research community begins to approach its future scientific and engineering goals at the peta-scale.

Paul Durham

# Delivering Scientific Productivity on HPCx

New scales of modelling and new generations of simulation remain critically dependent on advances in high-end computing. The functionality and performance of two key application codes - CASTEP in the area of materials modelling, and DL\_POLY in studies of the condensed phase - continue to play a central role in the scientific exploitation of high-end machines by members of CCLRC's Computational Science and Engineering Department and their academic collaborators. HPCx, the UK's flagship high-performance computing service provides a major vehicle for these studies.

The HPCx service is operated by the Department in collaboration with Edinburgh Parallel Computing Centre (EPCC) and IBM (UK) Ltd. In line with the service agreement with EPSRC, the HPCx system was enhanced this year to provide a doubling in capacity, achieved by more and faster processors, and a significant upgrade to the interconnect. Here we illustrate the work of the HPCx Terascaling team in improving the effective delivery of performance for scientific productivity by highlighting recent developments around both the CASTEP and DL\_POLY codes. In particular, we focus on the performance achieved on the IBM 1.7 GHz p690+ cluster with the new HPS switch (HPCx Phase2) compared to that on the previous 1.3 GHz p690 system with the Colony switch (HPCx Phase1). We also include results from competing systems: the SGI Altix 3700/1300 (newton) operated by CSAR at the University of Manchester, UK and the SGI Altix 3700/1500 (ram) at Oak Ridge National Laboratory, USA.

## CASTEP

CASTEP is a software package which uses density functional theory to provide a good atomic level description of all manner of materials and molecules. It performs quantum mechanical calculations of bulk materials, surfaces and interacting molecules. CASTEP was initially developed by Professor M.C. Payne and subsequently within the UKCP Consortium. It has recently been re-written as a fully modular code in Fortran 90 by the CASTEP Developers Group. It is distributed commercially by Accelrys Inc. The code makes extensive use of 'three-dimensional Fast Fourier Transforms' (3D-FFTs) on a distributed grid, performed by combining MPI\_AllToAllV with serial FFTs. Substantial optimizations have been made that have improved performance of this code on the IBM p690 systems by ensuring that the communications involved in the data redistribution surrounding the FFTs is aware of the architecture of the system (HPCx comprises a cluster of 32-way shared-memory nodes). This means that, where possible and advantageous, the communications within the node is separated from the inter-node communications.

Figure 1 shows performance for two benchmarks. TiN-8k is an 8 k-point spin-polarized total energy (SCF) calculation of TiN with a hydrogen defect. The cell contains 33 atoms and the 3d-FFT grid size is 108x36x36. Timings are presented for 34 SCF cycles. The second benchmark, Al2O3-G, is a total energy calculation of Al<sub>2</sub>O<sub>3</sub>. There are 5 k-points, the cell contains 120 atoms and the 3d-FFT grid size is 60x60x160. This benchmark is designed specifically to test communications and the parallelism is set-up so that the MPI\_AllToAllV is distributed across all the processors. This test is indicative of other systems which only have a single k-point.

We find that scalability is generally better on the p690+ than the p690 owing to the superior performance (especially latency) of the HPS interconnect. The 8-way k-point parallelism of the TiN-8k benchmark allows the collective communications to be limited to within the shared-memory nodes of the p690+ system. The Altix performs less well than the p690+. This is particularly evident in the Al2O3-G case, where the performance of the code is critically dependent on the performance of MPI\_AllToAllV across all the processors, and the Altix sees no performance gain in moving from 64 to 128 processors. There is still scope for optimisation of this code for the Altix.

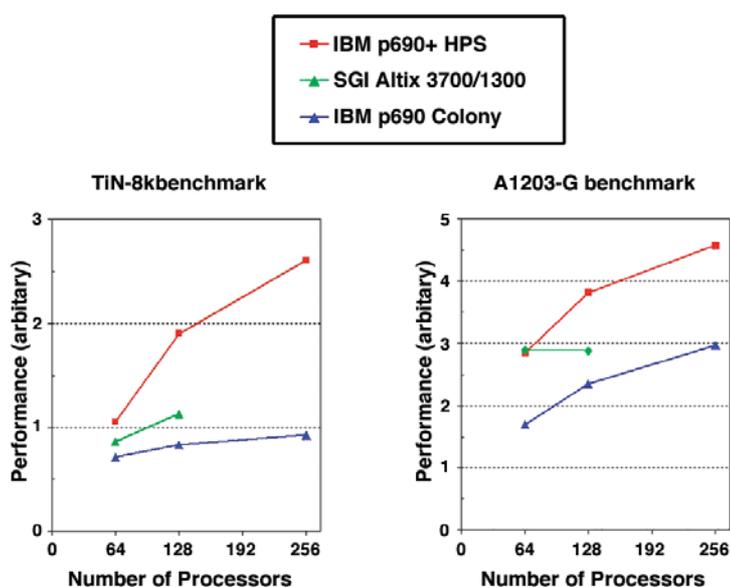


Figure 1: Performance of the CASTEP TiN-8k and Al2O3-G benchmarks on the two IBM clusters and the SGI Altix 3700/1300 system.

## DL\_POLY

DL\_POLY is a general-purpose molecular dynamics simulation package designed to cater for a wide range of possible scientific applications and computer platforms. DL\_POLY supports a wide range of application areas, including ionic solids, solutions, metals, zeolites, surfaces and interfaces, complex systems (e.g. liquid crystals), minerals, bio-systems, and those in spectroscopy. Significant enhancements to the code's capabilities have arisen from the recent release of the distributed data (or domain decomposition) version, DL\_POLY 3, and developments have been accelerated with the arrival of the HPCx system. Evaluation of the Coulomb potential and forces in DL\_POLY requires the calculation to be split into two parts, one performed in real space and one in Fourier space and the latter requires 3D-FFTs. A parallel 3D-FFT has been written which maps directly onto DL\_POLY's data distribution. This is very attractive at large processor counts. Although moving more data in total, the method requires much fewer messages, so that in the latency dominated regime it should perform better. Global operations (such as the *all-to-all* operations used for CASTEP above) are totally avoided. More generally the method is extremely flexible, allowing a much more general data distribution than those of other FFTs, and as such should be useful in other codes which do not map directly onto a distribution "by planes".

In figure 2 we present recent results which highlight the drastic improvements in both system size and performance made possible through these developments. The simulation is for a system of eight Gramicidin-A species; a total of some 792,960 atoms.

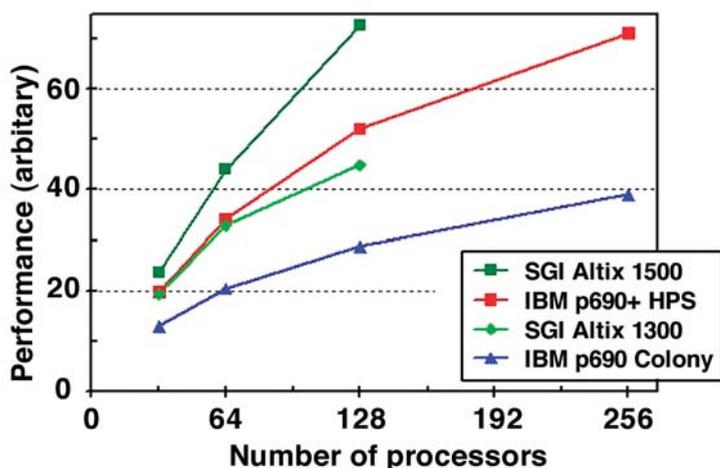


Figure 2: Performance of the DL-POLY3 distributed data code for the 792,960-atom Gramicidin-A simulation on the two IBM clusters and two SGI Altix systems.

The HPCx Phase2 performance is seen to be intermediate between that measured on the two SGI Altix systems, with the Altix 1500 exhibiting far better speedup than the 1300. In addition to slightly faster processors (1.5 GHz vs. 1.3 GHz), the 1500 system has double the Level 3 cache size (6MB vs. 3MB) and runs a single operating system over all 256 processors (the CSAR machine has multiple 64-way operating systems), which all combine to deliver a much improved performance. The overall level of performance represents a significant advance over that exhibited by both DL\_POLY 2 and CHARMM.

## Summary

The key goal behind our work is to enable complex and increasingly coupled and multi-disciplinary scientific applications to harness the potential performance of current and future high-performance parallel systems in order to meet critical scientific objectives. We have identified from the outset three key issues: management of the memory hierarchy, expression and management of concurrency and efficient sequential execution. These are key issues on all current systems.

We find that applications which scale best across all systems are those where the developers have most successfully optimised the global data management for multi-level cache architectures and maintained a high compute to communications ratio. There is no substitute for careful design of data structures, data access (loop indexing) and data distribution. For legacy codes this may require a complete re-design.

Our focus on algorithm development and the drive to remove dependencies on collective, global operations have been successful in several cases (including in this article CASTEP and DL\_POLY) in improving the scalability of these codes. Where all these issues have been addressed we find excellent levels of scalability and performance, with the platform-dependence far less critical.

M. F. Guest, M. Ashworth, I. J. Bush and M. Plummer  
CSED, Daresbury Laboratory.

## Acknowledgements

We are grateful to CSAR for access to SGI Altix 3700/1300 and to Oak Ridge National Laboratory for access to the SGI Altix 3700/1500 system.

# Atomic Physics, Astrophysics and HPCx



The solution of many important problems in astrophysics requires accurate and detailed atomic data. Atomic target ions in the iron-peak region of the periodic table, particularly singly ionised ions of iron (FeII), cobalt (CoII) and nickel (NiII), are of paramount importance in these applications. The FeII electron impact excited lines, for example, are observed in many astronomical objects; transitions are seen in the Sun, gaseous nebulae, quasars, Seyfert galaxies and supernovae. Forbidden Ni II and Fe II emission is observed to be particularly strong in gaseous nebulae such as the Orion nebula.

There are important outstanding questions on the Ni/Fe abundance ratios in gaseous nebulae. The accuracy of the atomic data may be a key factor in understanding this mystery.

Electron collisions with atoms and atomic ions have been studied both experimentally and theoretically since the 1930s. Highly accurate data is available for many hundreds of target species; theoretical calculations are able to predict the outcome of electron collisions with light atomic targets with an accuracy of a few per cent (even more precisely in some cases). It might seem that further investigation is superfluous. In fact, the calculation of the electron correlation effects that occur during the collision remains a computational challenge.

Relativistic corrections are necessary to obtain accurate electron-FeII rate coefficients. So far even non-relativistic LS calculations have failed to obtain converged rates. There is little or no experimental information on the important transitions. The FeII ion has 25 electrons, 7 of which are in open shells. As a consequence, the ion has a dense spectrum of low-lying excited states. Typically hundreds of these states will contribute to the collision when a low-energy electron with a kinetic energy of a few eV impacts a FeII ion.

During the past year the group at Daresbury has collaborated with researchers at Queen's University Belfast to obtain accurate low-temperature (up to 30,000 °K) FeII rate coefficients required in studies of the interstellar medium. Even the lowest transitions, between the  $3d^6 4s a^6 D^e$  ground state and the  $3d^7 a^4 F^e$ ,  $3d^6 4s a^4 D^e$  and  $3d^7 a^4 P^e$  excited states, must include allowance for virtual transitions to large numbers of excited states.

A previous three configuration R-matrix calculation by Ramsbottom *et al* considered all sextet to quartet transitions among the 38 LS states derived from the  $3d^7$ ,  $3d^6 4s$  and  $3d^6 4p$  configurations. The results included the discovery of important resonance effects arising from bound configurations representing the collision complex. The discovery implies that the physics of the collision depends upon an accurate and balanced representation of the correlation effects in both the target and collision systems. The result also calls into question the widely used practice in collision calculations of shifting the energy associated with an approximate target excited state wave function into agreement with experiment. The present work has focused on obtaining close to converged low-energy partial wave collision strengths including configuration interaction (CI) effects. All terms of the configurations  $3d^6 4s$ ,  $3d^7$ ,  $3d^6 4p$ ,  $3d^5 4s^2$  and  $3d^5 4s 4p$  are included giving 113 LS terms and 336 scattering channels.

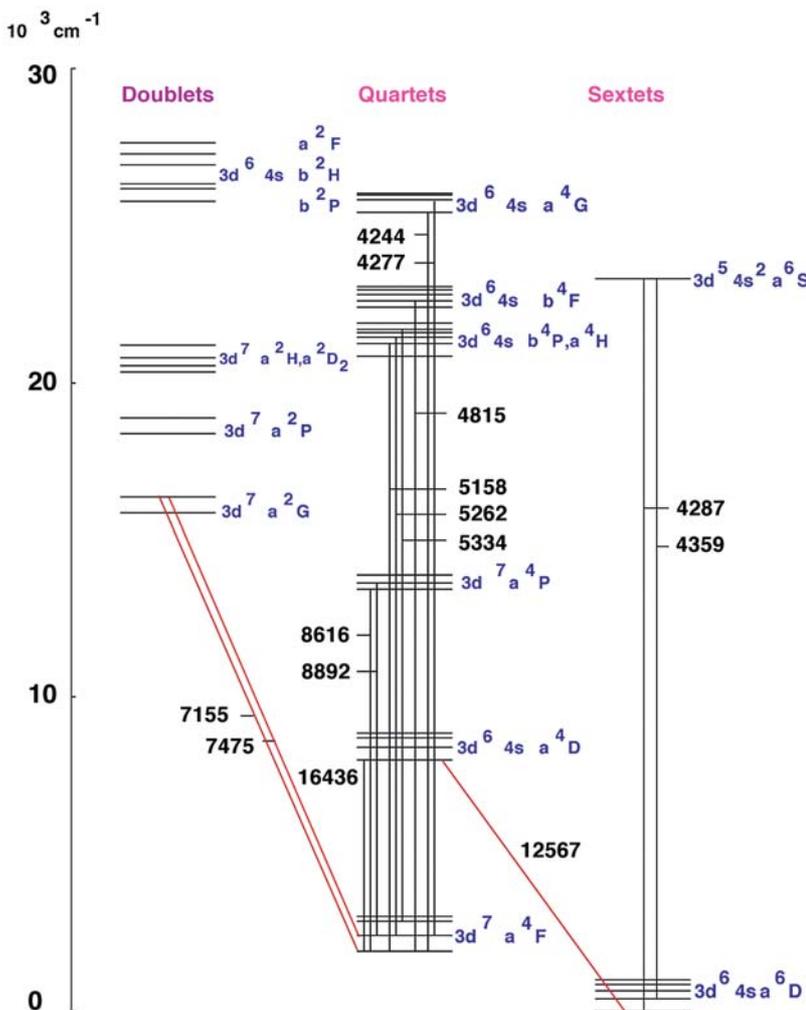


Figure 1: The 16 LS terms of Fe II below 30,000 cm<sup>-1</sup>

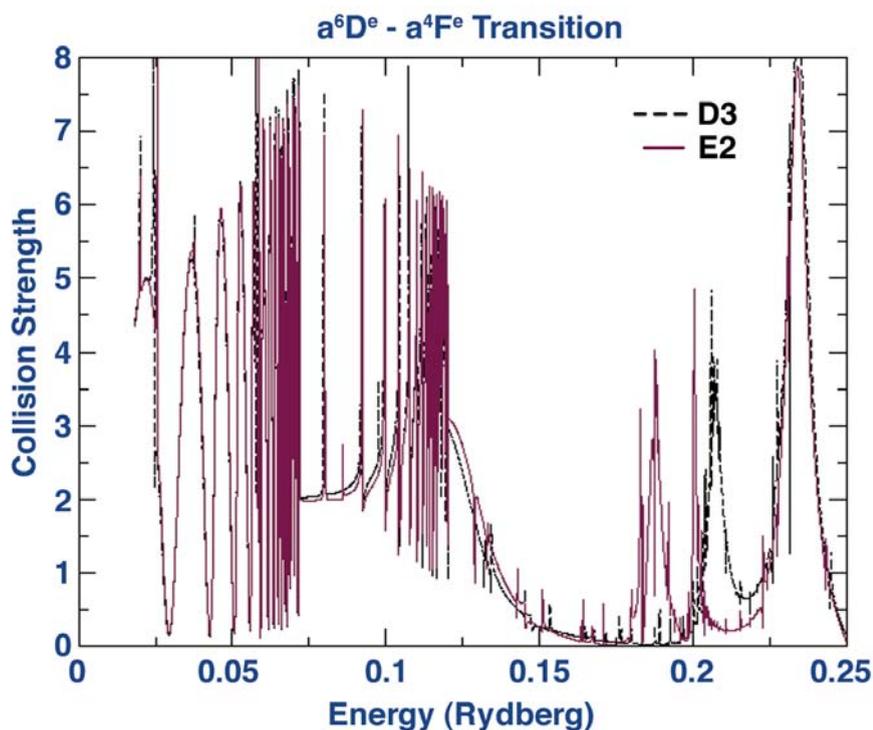


Figure 2: Contribution of  $^5D^e$  symmetry to collision strength.

A further 20 configurations were included in the calculation as CI resulting in over 2000 terms. Such calculations are unprecedented but have become feasible using a new R-matrix program developed at Daresbury. This general program is applicable to electron collisions with both atoms and atomic ions. It has been used successfully in calculations of near threshold electron collisions with neutral oxygen atoms.

The new parallel R-matrix code, is a re-implementation of the program RMATRIX-II developed at Belfast, and is aimed at exploiting the greatly expanded computing power available on the HPCx system. The program is written in an object oriented style using fortran95. A wide range of detailed changes have been made including the introduction of sparse matrix algebra and full memory management. Even serial versions of the new program allow big improvements in the description of CI effects. Current parallel versions of the program employ combined OpenMP and MPI to fully parallelize the angular algebra computation. MPI2.0 one-sided communication functions are used extensively. Good scaling has been observed. Modules have also been written using the BLACS, PBLAS and Scalapack libraries to distribute and diagonalise Hamiltonian matrices. A parallel file-handling library has been written using non-blocking MPI2.0 I/O functions.

Configuration selection procedures have been implemented. Target models of FeII that are significantly more accurate than previous work have been developed by allowing one - and two-electron excitations from the 3s and 3p inner shells to the 3d shell and by taking complete sets of configurations. The inclusion of configuration interaction effects in both the targets and in the collision complex has been systematically studied. The balance of correlation is crucial for determining the energy of the important  $^5D^e$  resonance. This resonance is seen in the figure showing the contribution from the  $^5D^e$  symmetry to the collision strength for the transition between the ground state and the  $a^4F^e$  state.

Completion of this study requires application of Breit-Pauli (BP) relativistic corrections to both the target and collision systems. The algebraic techniques employed in RMATRIX-II facilitate the inclusion of both one- and two-body BP corrections. It is estimated that the final calculation may involve more than 1000 target states and 10000 scattering channels – a challenging calculation even on the HPCx phase2 system.

*C. J. Noble and V. M. Burke  
CSED, Daresbury Laboratory.*

# Building an Infrastructure for Computational Quantum Chemistry

## Introduction

Traditionally quantum chemists, calculating the properties of molecular matter from first principles, have been at the forefront of the exploitation of computer hardware. However, the discipline has not seen the same level of innovation when it comes to the introduction of new informatics approaches as exemplified by, for example, the bioinformatics community.

Despite the ubiquity of *ab-initio* programs such as Gaussian, MOLPRO, and the UK and US versions of GAMESS, there are no agreed data standards for the transfer of many of the input data and results, much less intermediate data. There is also no way to archive the computational parameters in a manner that can be readily (let alone automatically) regenerated for subsequent calculations. Most workers in the field manage input and output using only their computer O/S file system with very little supporting tools.

The availability of data exchange standards would open up the possibility for users to take advantage of particular features of different packages, rather than performing the whole calculation in a single package. The issue of data standards is now being addressed within the UK's CCP1 (the oldest of the Collaborative Computational Projects) by a collaboration with the CCLRC e-Science Centre. In the long run, this has the potential to change the way software is designed and to maximise collaboration within the community, this is one of the primary aims of the CCPs.

## Data Model Development

The first stage in the development of a common format for computational chemistry commands and data is the agreement of the initial data structures that need to be represented. This would include data such as structures, atomic basis sets and molecular orbitals. The format for each type of data is then expressed in a data model, and we are using the Unified Modelling Language (UML) for this specification. Tools are being developed that enable quantum chemistry data documents, formatted in a way expressed by the data model, to be both validated and read into quantum chemistry applications. The UML can be converted to XML schema to enable the document validation and converted to code, in a number of different languages, for use by the quantum chemistry codes in parsing the formatted data. A component oriented approach has been taken for the design of the data model, making it easily extensible and re-usable by different groups. The relationship between the data model and some of the documents and tools is illustrated in Figure 1. Where possible, components are built upon the Chemical Markup Language and its extensions, such as CMLComp (<http://cml.sourceforge.net>). This enables compatibility with tools that already support this format. Since quantum chemistry is a computationally and data intensive discipline, there are a number of technical challenges. One such is the handling of large datasets which will need to be held in binary form; a number of possible technologies exist for this, including NetCDF, HDF, BinX and DFDL. Concurrently with the development of the data model we are also looking into approaches for storing and handling the data. Once we have standard

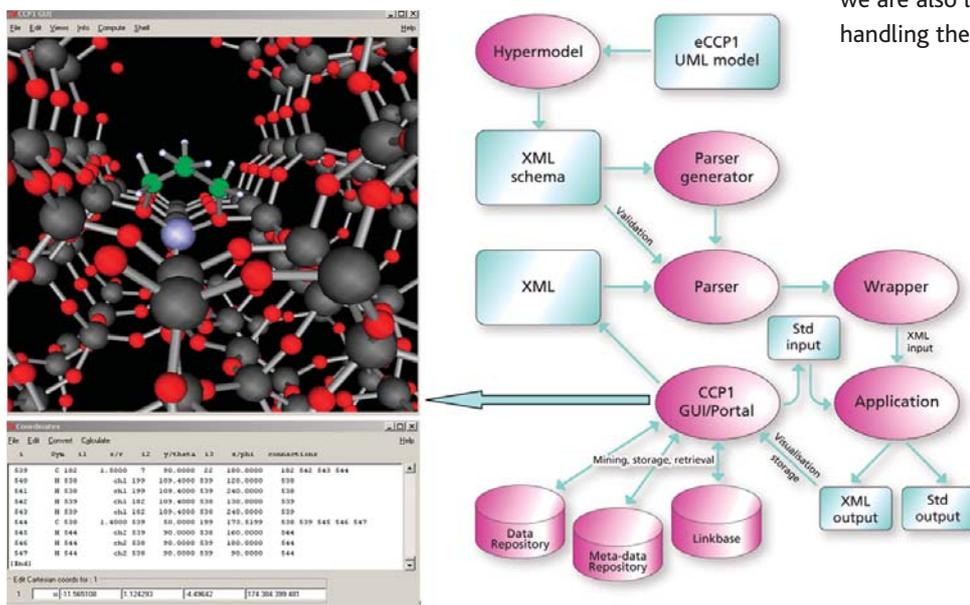


Figure 1: A Data Management Framework for Quantum Chemistry

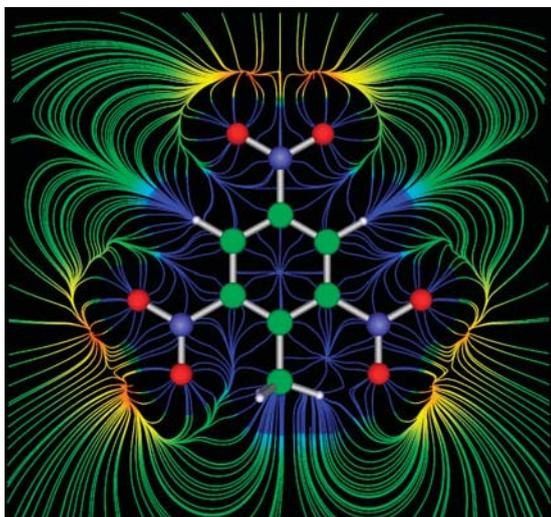
representations and tools it becomes possible to greatly increase data re-use and data sharing among the community, both within closely collaborating communities and globally (as the WWW has done for document sharing).

As consensus is reached on parts of the data model, we will provide a reference implementation which allows existing software to read and write data in the new representation. In view of the current quantum chemistry software base, it is likely that Fortran and Python will be the first languages to be addressed. An important part of that reference implementation will be the visualisation and model-building tools also under development within CCP1, and these are discussed in the next section.

### Visualisation and data preparation tools

A further component of our development programme aims to simplify the preparation of input to the quantum chemistry codes and provide support for graphical analysis of the output. Within the commercial sector, a number of powerful GUIs exist, but these are tied to particular programs and are not designed to facilitate interoperability.

In fact, in the commercial sector the quality of the interface is often the differentiator between different software providers. In contrast the academic community, as exemplified by CCP1, has a focus on scientific software and GUI development has been a low priority. We are now working to provide an interface to the quantum chemistry codes used within CCP1, and the current software has the status of a usable prototype, written in Python.



*Figure 2: The electric field around the TNT molecule. The colour denotes the electrostatic potential.*

The visualisation is provided by the open source VTK library and a number of other Python modules, such as Tkinter. Once these tools are installed the GUI should run in any modern environment.

The GUI development has the following characteristics:

- The nature of the target user base (quantum chemists) has led to an emphasis on internal coordinate representations and model building capabilities as well as conventional Cartesian-coordinate based tools.
- Use of a powerful open visualisation system offers the possibility to extend the visualisation tools as visualisation technology develops (e.g. volumetric analysis, vector fields etc, as shown in Figure 2).
- The object oriented nature of the interface and modular construction allows new code interfaces to be built readily, but also allows components of the GUI to be incorporated into other Python-based projects.
- Use of the scripting framework allows for easy extensibility by the expert user without recompilation.
- The use of an open source model ensures that distribution with any other software (for profit or otherwise) is possible and thereby encourages participation by the whole community.

### Getting Involved

If you are involved in the development of quantum chemistry software or are an interested user, you are encouraged to get involved in the discussions on data model development now. This will maximise the chance that the data models will serve the needs of a wide community and thus encourage uptake and our target of interoperability. The eCCP TWiki (<http://www.grids.ac.uk/eccp>) is a web site which allows users to register and then contribute to the content. We have posted draft data models and the derived schemas and started a mail list (described there). The current version of the CCP1 GUI can be downloaded from the CCP1 web site (<http://www.ccp1.ac.uk>).

*P. Sherwood, P. A. Couch, M. F. Guest, H. J. J. van Dam, J. M. H. Thomas and R. J. Allan, CSED, Daresbury Laboratory; P. J. Knowles, Cardiff University.*

# The Use of Partitioning in Solving Sparse Linear Systems

In many industrial applications the solution of sparse linear systems of equations  $Ax = b$  is the single most computationally expensive step. Thus it is essential that they are solved both accurately and as rapidly as possible. As the requirements of computational scientists for more accurate models increases so too do the sizes of the systems to be solved. In recent years this has led to the use of parallel algorithms and parallel computers. In this article, we consider exploiting coarse-grained parallelism using partitioning schemes.

Direct methods of solution of large, sparse linear systems involve factorizing the system matrix  $A$ . Such methods can exploit parallelism through the parallelization of operations on full submatrices and the simultaneous execution of tasks corresponding to separate elimination operations that are independent because of the sparsity structure of the system. Additionally, partitioning can be used to split the system so that its solution can be formed from the solution of smaller subsystems. Solving the subproblems independently yields the twin benefits of reduced complexity and added parallelism.

Clearly the best partitioning would be a block diagonal form where the subproblems are totally independent and the solution is affected just by solving the independent subproblems corresponding to each block on the diagonal. Close generalizations are the singly bordered block diagonal (SBBD) form

$$\begin{pmatrix} A_{11} & & & C_1 \\ & A_{22} & & C_2 \\ & & \dots & \cdot \\ & & & A_{NN} & C_N \end{pmatrix}$$

and the doubly bordered block diagonal (DBBD) form

$$\begin{pmatrix} A_{11} & & & C_1 \\ & A_{22} & & C_2 \\ & & \dots & \cdot \\ R_1 & R_2 & \dots & A_{NN} & C_N \\ & & & R_N & E \end{pmatrix}$$

that is associated with domain decomposition methods in the solution of partial differential equations and nested dissection methods.

The crucial issue in each case is the size of the border. The border size should be small compared with the overall size of the system so that there is little communication between the blocks on the diagonal and those in the border and so that the cost of solving the interface problem that links the subproblems is significantly less than that of factorizing the subproblems.

A symmetrically structured problem can be ordered to DBBD form using a standard graph partitioning tool such as MeTiS. We can then attempt to apply a linear solver to the blocks  $A_{ll}$ ,  $l = 1, \dots, N$ . But there are potential difficulties in doing this. In particular, the  $A_{ll}$  may be singular even if the problem is nonsingular. It may be possible to use a direct solver that incorporates threshold-based partial pivoting for stability and moves rows and columns into the border as necessary, but the border size may then grow too much. However, it is always possible to stably factorize the rectangular submatrix of an SBBD form into the product of a lower trapezoidal matrix and an upper triangular matrix. We can write this as

$$\begin{pmatrix} A_{ll} & C_l \end{pmatrix} = P_l \begin{pmatrix} L_l & \\ \tilde{L}_l & I \end{pmatrix} \begin{pmatrix} U_l & \tilde{U}_l \\ & S_l \end{pmatrix} \begin{pmatrix} Q_l & \\ & I \end{pmatrix},$$

where  $P_l$  and  $Q_l$  are permutation matrices, and  $L_l$  and  $U_l$  are lower and upper triangular matrices, respectively. Expressing the factorization in this way enables us to identify a DBBD form in which the coupling columns are as in the SBBD form and the coupling rows comprise the rows in the border blocks. This DBBD form is rather special inasmuch as the set of rows in each border block is disjoint from the rows in the other border blocks. The borders of this DBBD form are generally larger than for the DBBD form obtained directly using a graph partitioning tool, but the main point is that it is stabilized in the sense that pivots can be stably chosen from within the blocks on the diagonal and no further alterations to the DBBD form will be necessary.

Identifier	Order	MA48	HSL_MP48 ( $N = 8$ )							
			No. of processors							
			1	2	4	8				
10cols	29496	16.4	2.75	1.60	1.72	0.93	2.96	0.65	4.23	
Matrix35640	45640	NS	50.3	34.8	1.44	25.6	1.96	20.7	2.43	
1hr71c	70304	50.6	71.2	39.8	1.79	22.3	3.19	12.4	5.74	

*Table 1. Timings in seconds on an Origin 2000. NS denotes not solved. The numbers in italics are the speedups for HSL\_MP48 compared with using a single processor.*

The MONET algorithm was developed at Daresbury by Yifan Hu and colleagues for ordering unsymmetric matrices to SBBB form. MONET produces well-balanced SBBB forms with narrow borders, but it can require more CPU time than the subsequent matrix factorization. An alternative approach is a two-phase algorithm in which the matrix is first ordered to DBBD form then, in the second phase, row stretching is used to obtain an SBBB form. Stretching is a sparse matrix preprocessing technique that makes matrices sparser but, by adding additional rows and columns, larger. In practice, we find that with up to 8 blocks, the order of the stretched matrix is less than 5% larger than that of the original system while the border sizes are competitive with (and, for more symmetrically structured problems, better than) those obtained using the MONET algorithm.

We illustrate the use of the SBBB form using the new HSL sparse parallel solver HSL\_MP48. Timings for HSL\_MP48 and the serial code MA48 on an SGI Origin 2000 are given in Table 1. Problem Matrix 35640 was not solved using MA48 because more than one hour CPU time was needed. We observe that, on a single processor, the HSL\_MP48 times can be faster than those for MA48. As the number of processors is increased, HSL\_MP48 achieves good speedups. However, if we increase the number of blocks to more than 16, in most cases we find that the larger interface system causes a bottleneck and thus we conclude that this approach is most suited to SBBB forms with a small number of blocks. For larger numbers of blocks, we are currently looking at combining the methods of this article with iterative methods.

*I. S. Duff and J. Scott,  
CSED, Rutherford Appleton Laboratory*

# Non-intuitive Flow Phenomena in Gas-phase Microflows

Many micro-scale devices involve the manipulation of small volumes of fluid, typically nanolitres, and this has led to the emergence of an exciting new field of research called *microfluidics*. Examples include miniaturised heat-exchanges to cool integrated circuits, micro-reactors, bio-chemical sensors, and hand-held gas chromatography systems for the detection of trace concentrations of air-borne pollutants. A common link between these examples is the requirement to transport fluid through or over the device in a controlled manner. However, one of the emerging issues is the realisation that the fluid mechanics at such small scales is not the same as that experienced in the macroscopic world. This is especially true for gases where, even under atmospheric conditions, the continuum assumption associated with the Navier-Stokes equations breaks down and non-equilibrium or rarefaction effects are present. Understanding and predicting the consequences of these effects is proving to be a formidable challenge and, as will be demonstrated, the resulting flow behaviour can be very non-intuitive.

To illustrate some of the implications of scaling down, we consider the case of Couette flow between two concentric rotating cylinders. This is a classical fluid dynamics problem discussed in many textbooks since it forms the underlying theory for the analysis of rotating bearings. However, over the last few years there has been a renewed interest in the rotating Couette problem for the specific case of air-lubricated micro-bearings. Recent analytical and numerical studies<sup>1,2,3</sup> have demonstrated that under certain conditions of rarefaction, the flow between the cylinders can behave in a completely non-intuitive manner. For example, in the case of a stationary outer cylinder and a rotating inner cylinder, the velocity profile can become 'inverted' with the velocity *increasing* away from the inner rotating wall. This effect is contrary to the normal velocity profile expected within a cylindrical Couette flow.

Our investigation has reassessed the velocity inversion process using Maxwell's slip-boundary condition<sup>4</sup> to define the velocity discontinuity between the gas and the rotating cylinder walls. We consider isothermal flow for the case where the inner and outer cylinders have radii,  $R_1$  and  $R_2$ , and rotate at angular velocities,  $\omega_1$  and  $\omega_2$ , respectively.

In a cylindrical polar co-ordinate  $(r, \theta)$  reference frame, the circumferential momentum expression in the Navier-Stokes equations reduces to

$$\frac{d^2 u_\theta}{dr^2} + \frac{d}{dr} \left( \frac{u_\theta}{r} \right) = 0$$

where  $u_\theta$  is the tangential velocity component and  $r$  is the radius. Under rarefied conditions, Maxwell's slip-velocity boundary condition has to be applied and can be written in cylindrical form as follows<sup>5,6</sup>:

$$u_{\theta (gas)} - u_{\theta (wall)} = \pm \frac{(2-\sigma)}{\sigma} \lambda \left( \frac{du_\theta}{dr} - \frac{u_\theta}{r} \right) \Big|_{wall}$$

where  $\lambda$  is the mean free path of the gas molecules and  $\sigma$  is the tangential momentum accommodation coefficient (TMAC) which accounts for the average streamwise or tangential momentum exchange between the molecules and the solid surface. Theoretically, the TMAC can vary from zero (for specular reflection) up to unity (for complete or diffuse accommodation).

The general solution for the velocity profile between the two cylinders can be written as

$$u_\theta(r) = ar + \frac{b}{r},$$

where

$$a = \frac{A\omega_1 - B\omega_2}{A - B} \quad \text{and} \quad b = \frac{\omega_1 - \omega_2}{B - A}.$$

The parameters  $A$  and  $B$  can be derived as

$$A = \frac{1}{R_2^2} \left( 1 - \frac{(2-\sigma_2) 2\lambda}{\sigma_2 R_2} \right) \quad \text{and} \quad B = \frac{1}{R_1^2} \left( 1 + \frac{(2-\sigma_1) 2\lambda}{\sigma_1 R_1} \right).$$

For a stationary outer cylinder ( $\omega_2 = 0$ ) and a rotating inner cylinder, the velocity profile can be written in non-dimensional form as

$$u_\theta^* = \frac{u_\theta}{\omega_1 R_1} = \frac{1}{(A - B) R_1} \left( Ar - \frac{1}{r} \right).$$

To be compatible with the DSMC results presented by Tibbs *et al.*<sup>2</sup>, the radii of the inner and outer cylinders are chosen to be  $3\lambda$  and  $5\lambda$ , respectively.

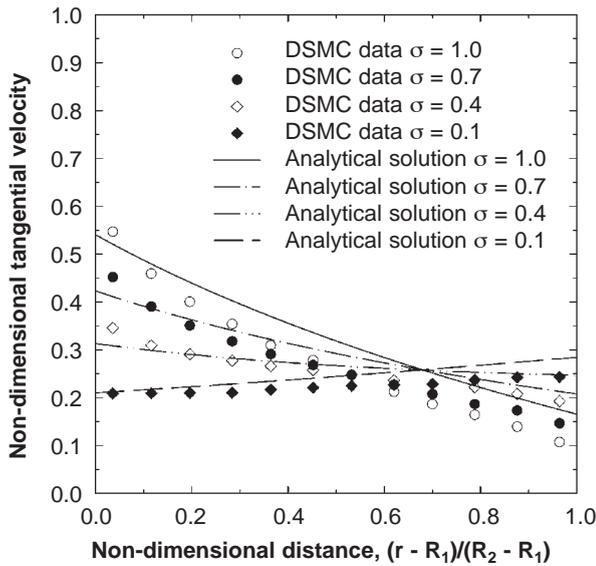


Figure 1: Non-dimensional velocity profiles for a cylindrical Couette flow ( $\sigma_1 = \sigma_2 = \sigma$ )

Figure 1 presents a comparison between the present analytical solution and the DSMC results obtained by Tibbs *et al.* for accommodation coefficients ranging between 0.1 and unity. Exact quantitative agreement between the DSMC data (symbols) and the analytical predictions (lines) is not expected since the separation distance between the cylinder walls is only two mean free paths, implying a Knudsen number (based upon the annular clearance) of 0.5. At such a high Knudsen number, continuum-based (Navier-Stokes) flow models are likely to be at the limit of their applicability. Nevertheless, the results show that the analytical formulation and the DSMC data follow the same basic trends and predict an inverted velocity profile for an accommodation coefficient of  $\sigma = 0.1$ .

For the first time, our analysis has been able to show that the velocity inversion process depends only upon the accommodation coefficient of the stationary outer cylinder<sup>7</sup>. This surprising result has allowed us to define a new limiting criterion that specifies the conditions under which the velocity inversion phenomenon will occur:

$$\sigma_2 < 2 \left( 1 + \frac{(R_1^2 + R_2^2) R_2}{2\lambda R_1^2} \right)^{-1}.$$

It is clear that this analytical result shows how non-intuitive gas-phase microflows can be and independent experimental investigation is now required.

*D. R. Emerson, Y. Sun and R. W. Barber  
CSED, Daresbury Laboratory.*

## References:

- [1] Einzel D, Panzer P, Liu M. *Phys Rev Lett*, 64(19), pp. 2269 - 2272, 1990.
- [2] Tibbs KW, Baras F, Garcia AL. *Phys Rev E*, 56(2), pp. 2282 - 2283, 1997.
- [3] Aoki K, Yoshida H, Nakanishi T, Garcia AL. *Phys Rev E*, 68:016302, 2003.
- [4] Maxwell JC. *Philos Trans Roy Soc Part 1*, 170, pp. 231 - 256, 1879.
- [5] Lockerby DA, Reese JM, Emerson DR, Barber RW. *Phys Rev E*, 70:017303, 2004.
- [6] Barber RW, Sun Y, Gu XJ, Emerson DR. *Vacuum*, 76(1), pp. 73 - 81, 2004.
- [7] Sun Y, Barber RW, Emerson DR. *Proc. 24th Int. Symp. on Rarefied Gas Dynamics*, 2004.

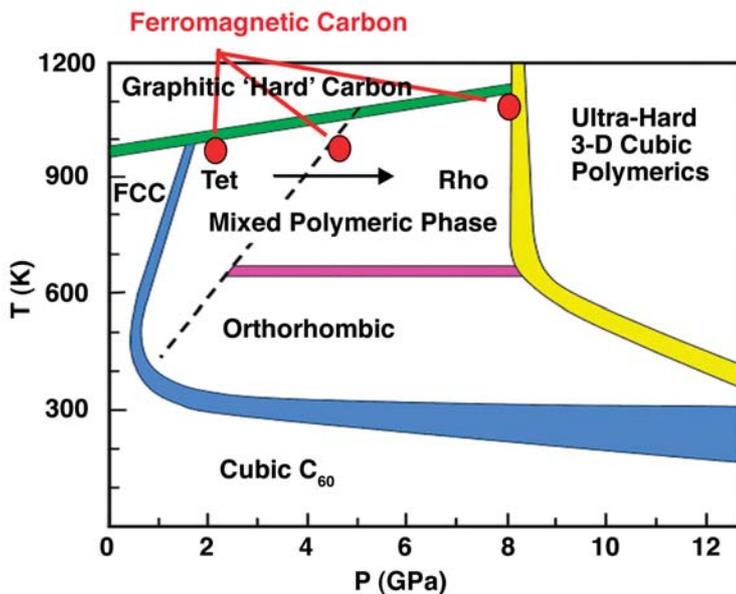
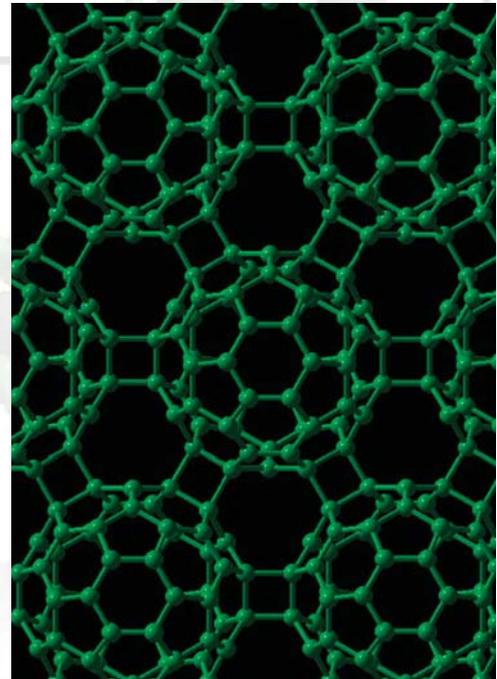
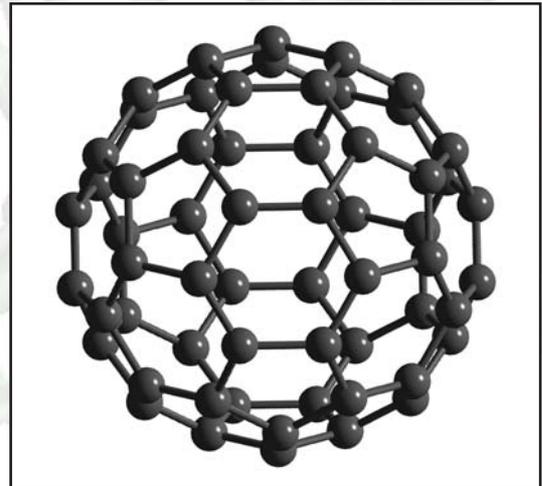
# Magnetic Carbon

High temperature ferromagnetism has recently been observed in carbon materials. This new class of magnetic materials introduces the possibility of highly tunable materials for use in magnetic devices. The materials are based on fullerene cages (bucky balls). The bucky balls form well ordered bulk crystals consisting of close packed balls.

As the pressure and temperature are raised a variety of phases are now known to occur.

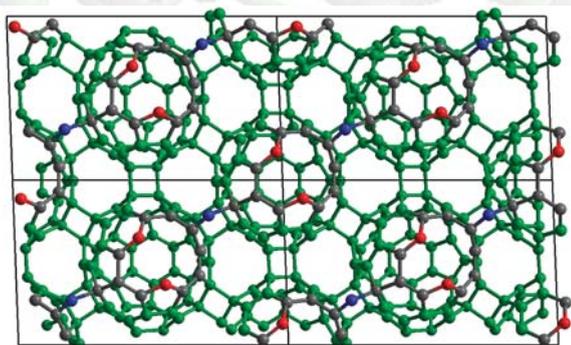
At around 900 K additional bonds form between the Bucky balls and cubic (FCC), tetragonal and rhombohedral arrangements of the balls are observed. At higher temperatures there is a phase boundary (the green line in the phase diagram below) above which the balls break and graphite like phases are formed.

It is close to this boundary that magnetic phases have recently been observed. Apparently just as the balls begin to fracture a new structure forms. The characterisation of these phases has proven to be extremely difficult and their structure and the reason for the remarkable magnetic properties remain a mystery.



Computer simulation can play a major role in predicting possible structures and interpreting their properties. Here molecular dynamics simulations have been used to generate a variety of structures at elevated temperatures and pressures.

One of the predicted structures is displayed below.



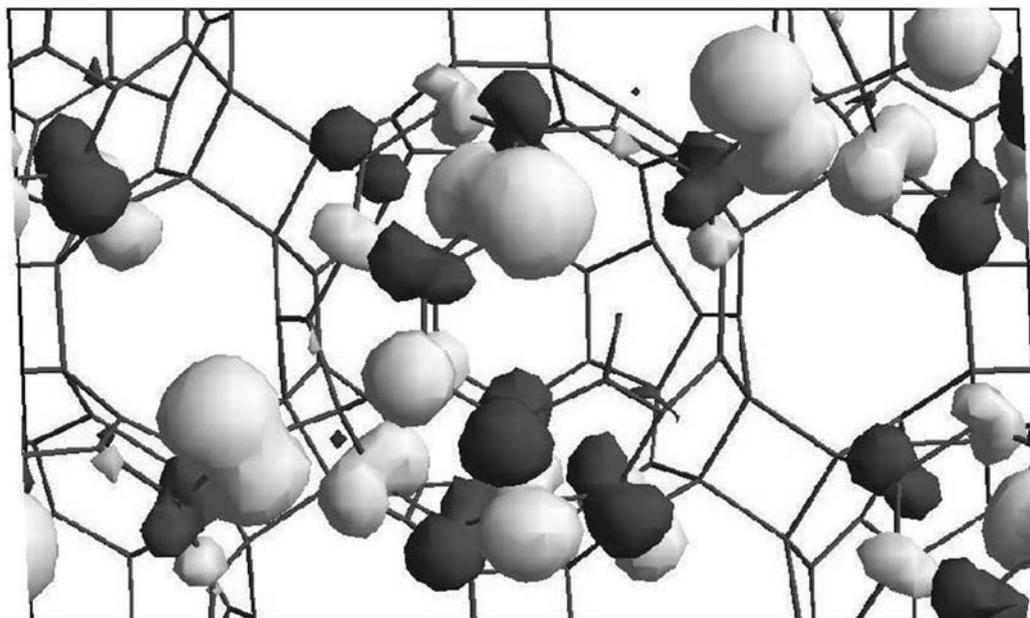
The balls have just started to fracture: in order to relieve the local stress a carbon vacancy occurs in each ball with the liberated atom forming a bridge like structure (the blue atom) to the neighbouring ball.

The magnetic properties of this structure can be computed using very large scale quantum mechanical calculations - the predicted spin density is shown below.

These calculations predict that this structure will indeed be ferromagnetic with the bridges facilitating long range magnetic coupling between the balls.

This is the first realistic model of the magnetic coupling mechanism and opens up the possibility of controlling the magnetism in these systems.

*J. A. Chan, Imperial College London;  
B. Montanari, Imperial College London, and  
CSED, Rutherford Appleton Laboratory;  
J. D. Gale, Curtin University of Technology, Western Australia;  
S. M. Bennington and J. W. Taylor, ISIS, Rutherford Appleton  
Laboratory;  
N. M. Harrison, Imperial College London, and CSED, Daresbury  
Laboratory.*



### References:

*J. A. Chan, B. Montanari, J. D. Gale, S. M. Bennington,  
J. W. Taylor and N. M. Harrison, Phys. Rev. B70 041403 (2004).*

# First Principles Study of the Stability of PuO<sub>2</sub> Under Oxidation

Until recently, PuO<sub>2</sub> has been the preferred choice for long term storage of Pu. In its Pu (IV) oxidation state, this compound forms a stable solution over a wide temperature range with no sign of reaction when exposed to air. In addition it has been relatively safe to handle and has played a major role in recycling Pu arising from disarmament as it can be blended with slightly enriched uranium to form a nuclear reactor fuel called "mixed-oxide" (MOX), a combination of UO<sub>2</sub> and PuO<sub>2</sub>.

In 2000, J.M. Haschke, T.H. Allen and L.A. Morales (Science **287**, p.285) reported the discovery of higher composition binary oxides, PuO<sub>2+x</sub> (x ≤ 0.27). The reaction was found to occur in the presence of water, i.e.

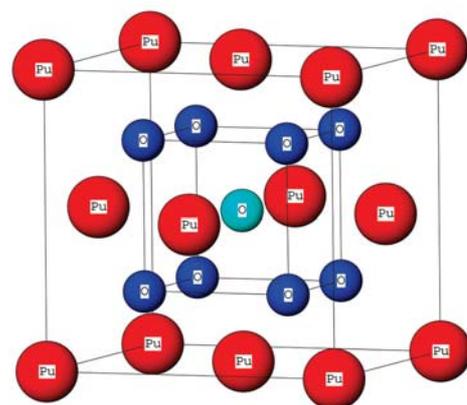


for temperatures in the range of 25 °C to 350 °C. If confirmed, this finding could have far-reaching implications for any future nuclear waste disposal. In a higher oxidation state, Pu oxide could for example be more water soluble, and therefore more mobile in the underground geological environment. This is because Pu (IV) moves 2 to 3 orders of magnitude slower than Pu (V) and Pu (VI). Furthermore the generation of hydrogen gas would constitute a serious risk of rupture for the containment vessels, unless these are properly vented. Despite considerable experimental evidence however, the existence of PuO<sub>2+x</sub> remains controversial mainly because many earlier attempts to prepare these oxides failed.

We have investigated the changes in the Pu oxidation state induced by the oxidation process from PuO<sub>2</sub> to PuO<sub>2+x</sub> based on ab-initio electronic structure calculations. (L. Petit, A. Svane, Z. Szotek and W.M. Temmerman, Science **301**, p.498, 2003) The quantum-mechanical understanding of the physics of actinides and their compounds is a challenge due to the intricate nature of the correlations of the partially filled 5f shell. In particular, the calculations need to determine how many of the 5f electron states are localized and this leads to the evaluation of the Pu valence.

The 5f electrons are less localized than the 4f electrons in the rare earths and their compounds, and at the beginning of the actinide series, they can play a significant role in bonding. SIC-LSD calculations were performed which were quite successful in applications to the study of the valence in f-electron systems. (P. Strange, A. Svane, W. M. Temmerman, Z. Szotek, H. Winter, Nature **399**, p.756 (1999))

Depending on the chemical surrounding, Pu f-states exhibit a more or less localized character. The SIC-LSD approach is in particular successful in determining the valence of rare earths and actinide compounds. From the energetics of the bonding to the O ligands, we determine the ground state configuration of the Pu f-electron manifold, and thereby the preferred oxidation state of Pu. Our calculations find stoichiometric PuO<sub>2</sub> to be an insulator, with the Pu atoms being in the 5f<sup>4</sup> ground state configuration, in agreement with experiment.



*Figure 1: The conventional cubic unit cell of plutonium dioxide in the fluorite structure. The cube contains four PuO<sub>2</sub> units (red Pu and dark blue O atoms), while the central O atom (light blue) marks the octahedral interstitial position, where excess O in PuO<sub>2+x</sub> is accommodated. The modelling of PuO<sub>2.25</sub> is facilitated by occupying the central position in all cubic units.*

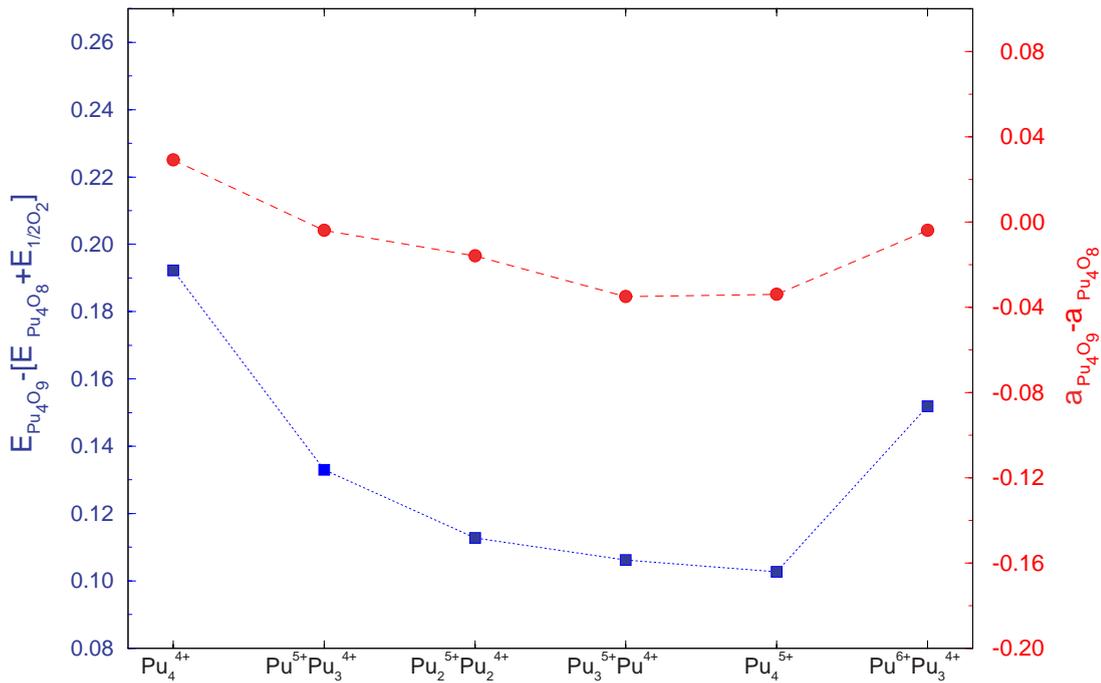
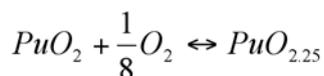


Figure 2: Total energy minima (blue and left-hand side axis) and corresponding theoretical equilibrium lattice constants (red and right-hand side axis) for different Pu f-configurations of Pu<sub>4</sub>O<sub>9</sub>. The energy (in Ry units) and lattice parameter data (in Å units) are given with respect to the corresponding values for Pu<sub>4</sub>O<sub>8</sub> + [1/2] O<sub>2</sub>

To study oxidation, one additional O is placed on an octahedral interstitial site of a Pu<sub>4</sub>O<sub>8</sub> super cell. (Figure 1) This has a profound effect on the electronic structure, as the added O impurity adsorbs the electrons released from neighbouring Pu(IV) ions in the process of f-electron delocalization, leading to the formation of Pu(V) ions. The combined effects of lattice expansion through O inclusion, and lattice contraction due to the increased number of bonding electrons results in a lattice parameter that is almost unchanged. (Figure 2) These calculated properties of PuO<sub>2.25</sub> are in good agreement with the experimental findings by J. M. Haschke, T. H. Allen and L. A. Morales.

Whether the oxidation process actually takes place depends on the total energy balance



Our calculations find the reduction (left side) to be more favourable by 25 mRy. Thus we conclude that at T = 0K the dioxide is stable with respect to interactions with free O<sub>2</sub> molecules, consistent with the results from decades of research on this subject. However, given the small energy barrier, it is conceivable that the oxidation may become favourable under specific experimental or even environmental conditions (the experimental process uses H<sub>2</sub>O as oxidant, in which case oxidation is accompanied by production of H<sub>2</sub> which is favoured by high entropy).

L. Petit, ORNL, USA;  
 A. Svane, University of Aarhus, Denmark;  
 Z. Szotek and W. M. Temmerman, CSED, Daresbury Laboratory.

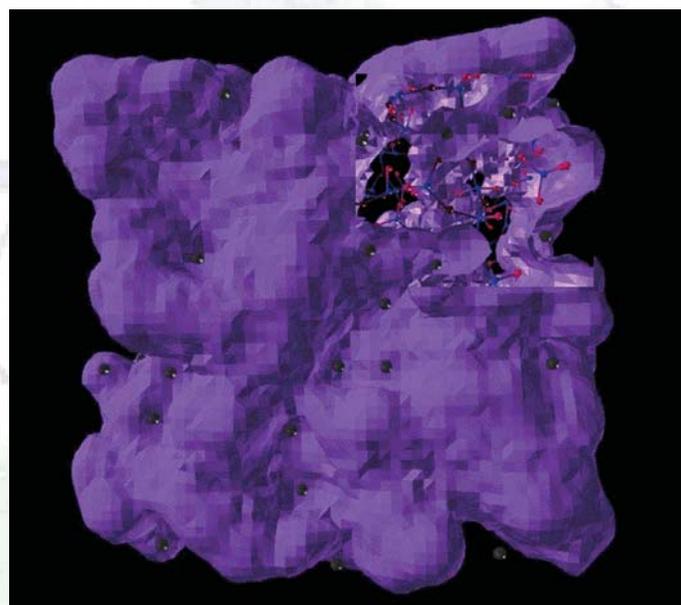
# Virtual Reality Molecular Dynamics of Ion Diffusion in Glass

Until relatively recently, our understanding of the structure and properties of glass has come from experimental investigation.

The Laboratories of Daresbury and Rutherford and Appleton, which are homes to synchrotron radiation and neutron scattering research respectively, have been in the forefront of the scientific investigation. By such methods a great deal has been learned about the atomic structure of glass and this provides a basis for predicting how glasses will behave in the bulk. However it is fair to say that most of this knowledge is based on a static picture of glass structure, which cannot represent the whole truth. Thus there are phenomena occurring in glasses that are not easily interpreted from this picture because they are manifestations of dynamical process at the atomic level. A good example of this is the so-called 'mixed alkali' effect, in which the electrical conductivity of alkali glasses falls precipitously if different alkali ions are present in the same glass.

An interesting way of investigating the dynamical properties of glasses is through computer simulation. The method known as molecular dynamics solves Newton's equations of motion at the atomic level and thus provides a wealth of information about the dynamical processes of atoms. This approach offers profound insights into complex atomic systems in general and has opened up many new avenues of study, but there are disadvantages to it. The main disadvantage is that, of necessity, many of the properties calculated are statistical in nature – the method generates so much data that statistical averages offer the best means of handling the results.

Today however a new approach is possible – Virtual Reality (VR)[1]. This is a computational technique which renders the wealth of data generated by computer simulations into a visual representation. In such a visualisation, the scientist can literally be placed into the heart of the atomic system and observe its behaviour as if it was happening before his eyes. Thus the tools to analyse the phenomena are not statistical but visual and they capitalise on the processing ability of the human eye and brain, which



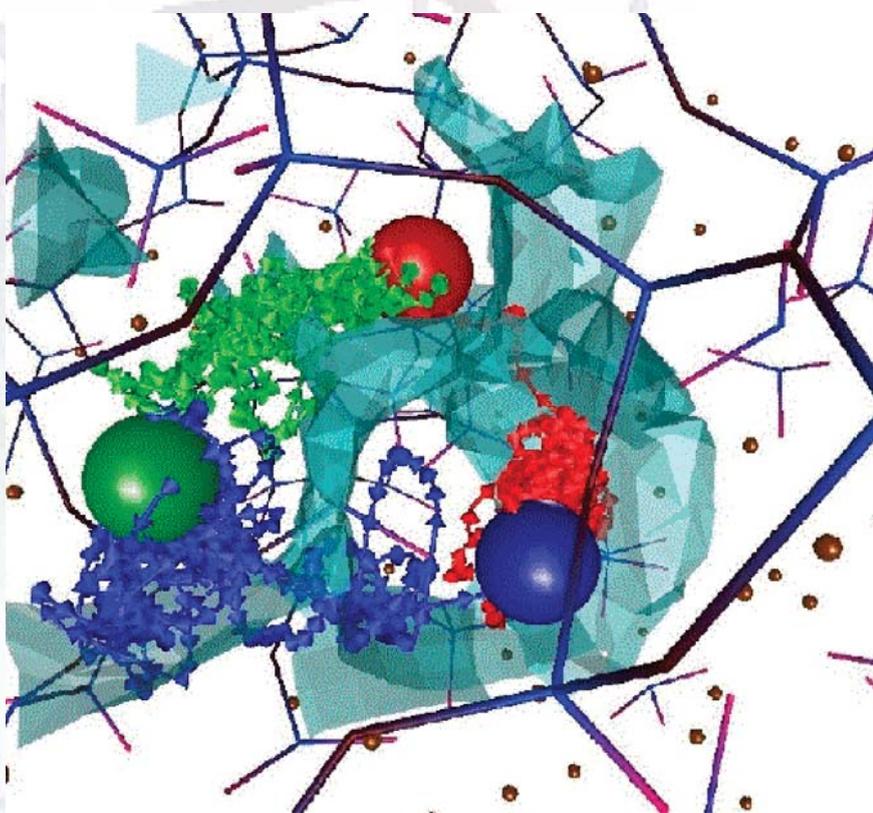
*Figure 1: Isosurface defining the overall network morphology in sodium disilicate glass (see text).*

are able to apprehend what is happening in a complex environment, albeit one created artificially.

At Aberystwyth a sophisticated VR system has been built and harnessed to a molecular dynamics program called DL\_POLY[2], which was developed at Daresbury. DL\_POLY provides a numerical description of the motion of the atoms in a glass and the VR system renders this visually in three dimensions on a human scale. The result needs to be seen to appreciate its astonishing clarity. Atoms can be seen to vibrate, collide and diffuse through the complex glass structure. Co-operative motions can be readily observed, in which atoms seemingly work together to accomplish the process of diffusion.

Special surface rendering techniques can be used to reveal the otherwise hidden paths through the bulk of the material and which permit the alkali ions to diffuse at all. On top of this, quantitative measurements of available volume and its associated fluctuation can be performed to gain support for or against various theories of the diffusion mechanisms. Thus by linking the free volume available in the structure to observed co-operative events it has been possible to show how the available free volume in single-alkali and mixed-alkali glasses differ in a way that accounts for the observed mixed alkali effect.

G. N. Greaves, A. J. Shaw, P. Summers, B. Bevan, and O. Majerus  
University of Wales, Aberystwyth;  
W. Smith, CSED, Daresbury Laboratory;  
R. J. McGreevy, ISIS, Rutherford Appleton Laboratory.



*Figure 2: The detail revealed by VR is apparent in this static picture. The framework of a silicate glass (represented by connected rods in the above picture) contains mobile alkali ions (represented by coloured spheres). Also shown is part of the isosurface that delineates the volume surrounding the framework that the ions are forbidden to enter. The tracks of the ions as they move through the available volume are traced out in arrowed paths of a matching colour. When viewed in VR mode, the co-operation between the motions of the ions can be clearly seen.*

## References:

- [1] G. N. Greaves, A. J. Shaw, P. Summers, B. Bevan, O. Majerus, R. J. McGreevy and W. Smith (Daresbury) (2004) submitted.
- [2] Website: [http://www.cse.clrc.ac.uk/msi/software/DL\\_POLY/index.shtml](http://www.cse.clrc.ac.uk/msi/software/DL_POLY/index.shtml)

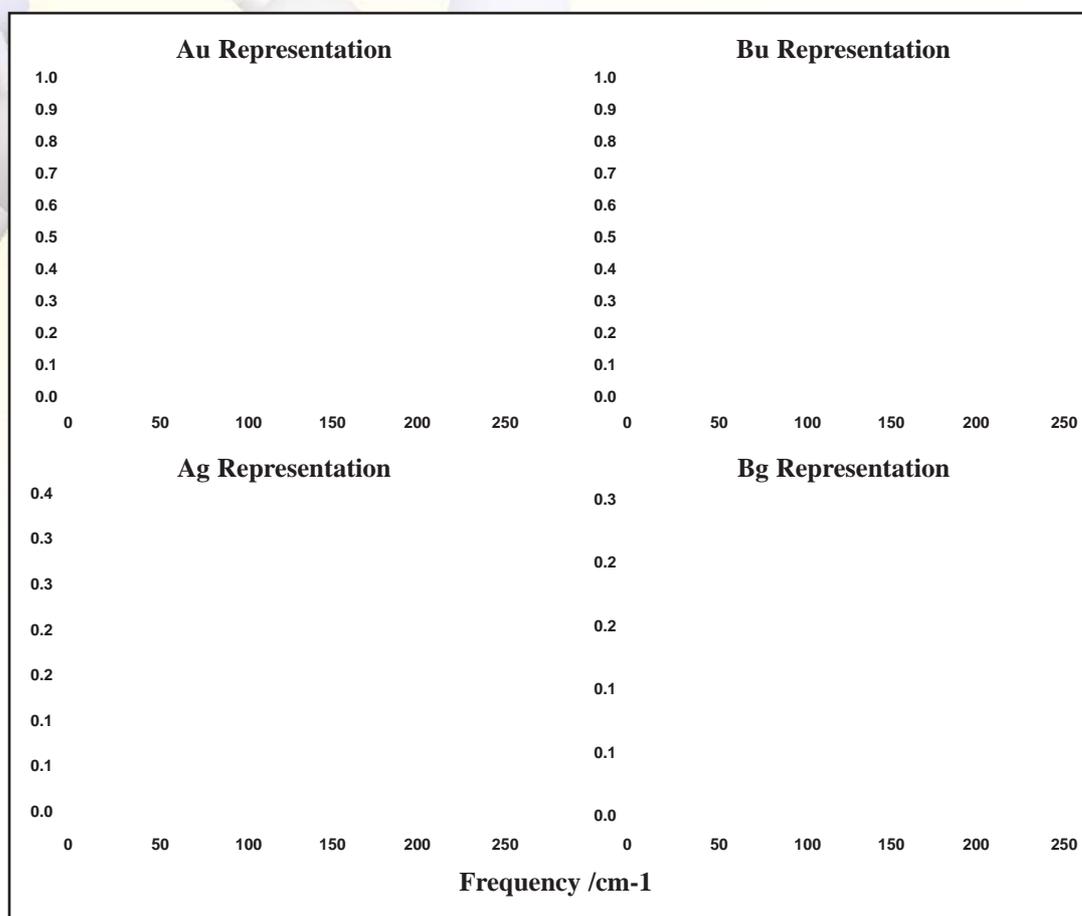
# Dynamics in Crystals of Rigid Organic Molecules

The phenomenon of polymorphism (the existence of two or more different crystalline forms of the same substance) is of great importance for the pharmaceutical industry and other applications where crystalline small organic molecules are used.

The discovery of a new polymorph can have regulatory and processing consequences, as happened in the case of the pharmaceutical ritonavir, where the more stable form was first discovered during manufacture leading to a very costly redesign of the manufacturing process.

The extension of computer modelling techniques to the organic solid state and the prediction of polymorphism has lagged behind inorganic materials. Polymorph stability will be determined by the free energy, however kinetic factors may be important in some cases in determining which phase is obtained. Marginally metastable structures that have a sufficient kinetic advantage in nucleation, or growth,

or a barrier to transformation to more stable structures, can be observed before the thermodynamically most stable structure is discovered. Molecular mechanics techniques have been used to obtain energy minimised structures from a very large number (up to several thousand) of trial structures in different crystal symmetries. The model used describes the electrostatic contribution to the lattice energy using distributed multipoles. However, these calculations only model internal energy and do not model the thermal motion of the molecules. Thermal properties of crystals can be modelled using two methods. Harmonic lattice dynamics can be used to calculate the rigid-body  $k = 0$  (long wave length) phonon frequencies. We have recently calculated these phonon frequencies for a range of small organic molecules where the frequencies were known experimentally. One of these molecules is imidazole, which is widely used as an intermediate in the synthesis of pharmaceuticals and other organic compounds. The predicted frequencies



*The power spectra of the velocity autocorrelation functions of the four representations of imidazole. Also highlighted as vertical lines are the calculated harmonic frequencies.*

were sufficiently good to suggest they would provide worthwhile estimates of the relative free energies of different polymorphs by using a standard thermodynamic expression for the entropy. This method is computationally very quick but has limitations, as there may be significant anharmonicity in the crystal.

The harmonic lattice dynamics calculation would also not detect any transformation of the crystal to a new structure which was not among the trial set, nor any rapid switching between two possible structures.

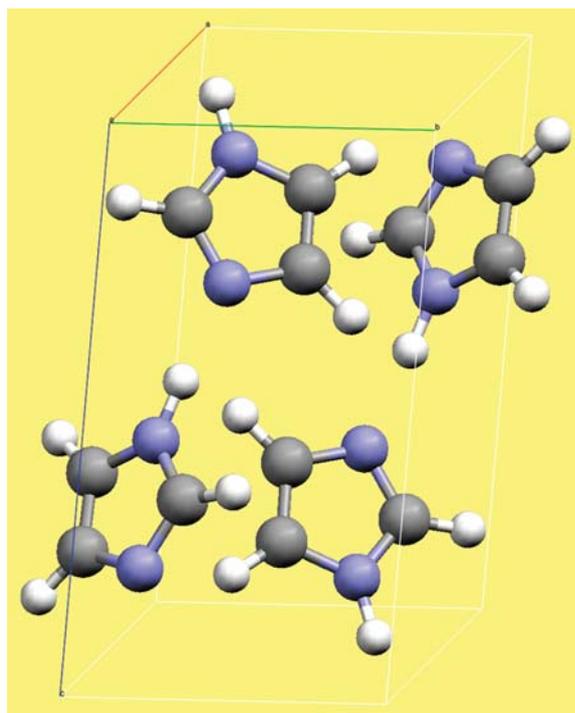
For this reason a new Molecular Dynamics package DL\_MULTI has recently been developed as an extension of DL\_POLY using the same distributed multipole models for the electrostatics. This technique allows us to validate the lattice dynamics method and overcomes some of the limitations of the lattice dynamics, but is computationally much more expensive.

Molecular Dynamics simulations have been successfully performed to model crystalline imidazole at 100 K. The simulation was compared with a previous potential that had been specifically derived using the lattice dynamics of imidazole. The distributed multipole electrostatic model gives a far more successful simulation demonstrating the necessity of using of them. Symmetry adapted velocity autocorrelation functions were calculated from the MD simulation. The Fourier transform then gives the power spectrum and hence characterises the phonon motions and frequencies, for each symmetry representation of the crystal. The frequencies and characteristics obtained are in agreement with those found from harmonic methods. (See figure opposite)

In the case of imidazole, no indication of any phase transformation was apparent in the Molecular Dynamics run itself, or in large shifts of the phonon frequencies when including anharmonic effects.

The results of this simulation shows that harmonic Lattice Dynamics calculations give sufficiently reasonable estimates of frequencies, to compare the relative free energies of hypothetical crystal structures. The ability to predict the many situations when hypothetical low energy structures may

transform readily to slightly more stable structures, and thus avoid being observed, will be a major advance in polymorph prediction. Thus, understanding the varied dynamical motions in organic crystals will require Molecular Dynamics studies of different organic crystals, to develop a practical theory of phase transformation behaviour for use in polymorph prediction.



*The crystal structure of imidazole  
(C - black, N - blue, H - white)*

*E. Gray, University College London;  
G. M. Day, University College London and  
University of Cambridge;  
M. Leslie, CSED, Daresbury Laboratory;  
S. L. Price, University College London.*

# High Performance Computing Molecular Dynamics Simulations

Until the emergence of parallel computing, large scale Molecular Dynamics (MD) simulations of non-equilibrium systems were limited in realism and accuracy due primarily to the large memory requirement and excessive computing time demanded. The parallel MD program DL\_POLY\_3 [1] now offers an exciting way forward based on a blend of concepts such as:

1. Domain decomposition (DD) – a parallel technique that shares the memory requirement and work load of a simulation across computer nodes.
2. Link cells (LC) – a method for fast search and location of atoms interacting within a specified short-range cutoff.
3. The Daresbury advanced Fourier transform (DAFT) routine – a discrete fast Fourier transform (DFFT) routine particularly adapted to the DD strategy and used to calculate Coulombic forces.
4. Symplectic integration algorithms – based on the Velocity Verlet scheme, which ensures time-reversibility and excellent numerical stability.

The program has enabled scientists to extend research to more realistic large scale systems with greater complexity and to carry out MD simulations that were unthinkable just a few years ago.

Our investigation of DL\_POLY\_3 performance [1] on HPCx [2] indicates that the performance of the program is extremely good. The load balancing of the simulation across processors achieved by the link cells ensures efficient use of both memory and computing power and the integration algorithms exhibit high stability. Simulations of a model system containing approximately **three million atoms** showed close to linear scaling of performance with numbers of processors ranging from 4 to 1024. This clearly indicates that MD simulations in excess of multi-million atom systems are possible using DL\_POLY\_3 on HPCx.

Most scientists that investigate large systems on a molecular level, refrain from using systems larger than a million atoms since the output requires storage volumes as well as graphical analysis to render the results intelligibly. However, research on systems of a few hundred thousand atoms is carried out routinely [3], especially in such areas as radiation damage (RD). For such systems DL\_POLY\_3 is the primary means available for MD simulations, in which a highly energetic particle is driven into the system, giving rise to cascades of displaced atoms. Such studies are motivated by the growing need for a method to immobilize radioactive nuclear waste safely, by embedding it in a material form that can be an effective barrier to its escape into the environment. Materials to be immobilized come primarily from nuclear power stations, with the remaining part coming from decommissioned nuclear weapons. Safe immobilization of nuclear waste is crucial to the future of nuclear power industry, but regardless of this, the amount of currently stored non-immobilized waste is already large enough to pose serious concerns.

In our studies of RD effects [3,4] we considered materials related to those proposed for encapsulating highly radioactive nuclear waste (waste forms, glasses) as well as new ceramic materials that promise considerably higher durability. The encapsulating material inevitably gets damaged by irradiation from radioactive decay and it is important to assess what consequences this irradiation may have on the performance of the waste form over time, which can vary from about one hundred to tens of thousands of years for different isotopes.

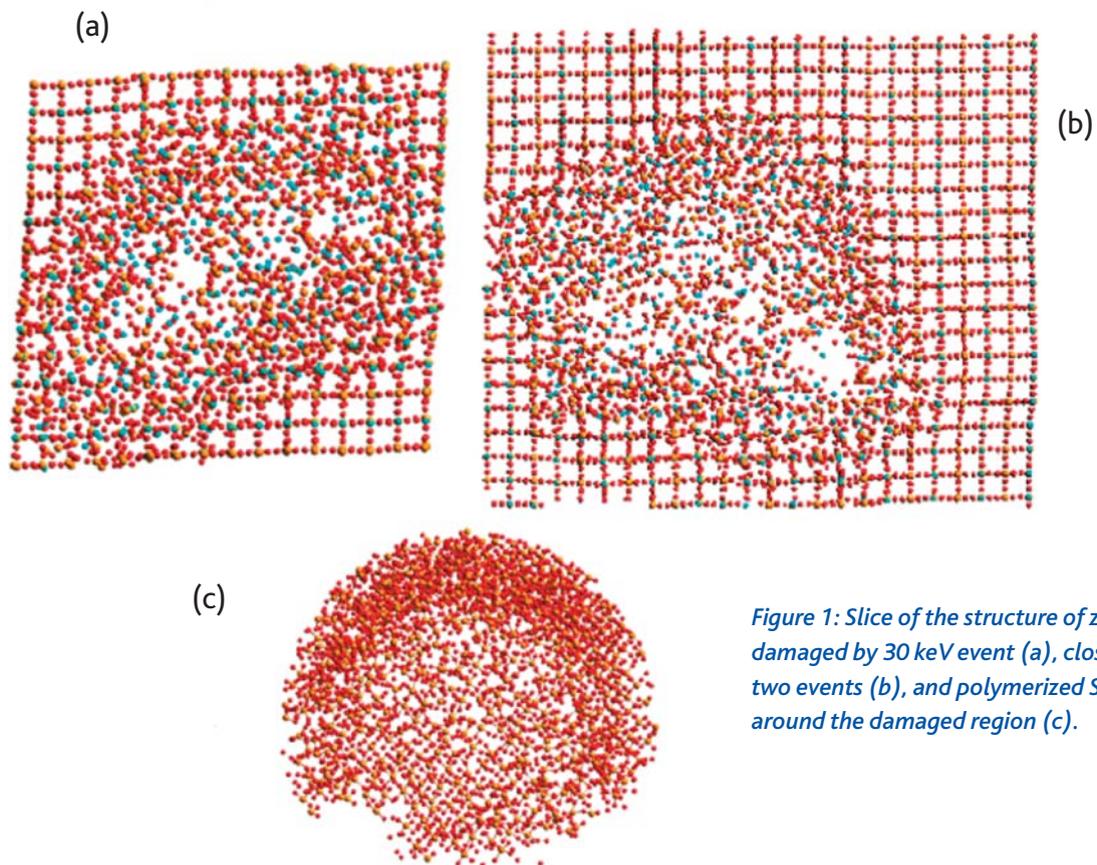
In order to study how properties of waste forms can change over time, we perform massive parallel MD simulations of high-energy recoils in materials of interest.

These recoils simulate an alpha-decay event, during which a heavy recoil causes considerable damage in the structure, resulting in several thousand permanently displaced atoms.

An example of the damaged structure of zircon ( $\text{ZrSiO}_4$ ), a potential waste form, is shown in Figure 1. A highly disordered region, consisting of about five thousand permanently displaced atoms is seen, which is surrounded by the largely undamaged crystalline lattice.

The animation of the MD simulation showing a propagation of a high-energy recoil can be downloaded from [5]. It shows that only a portion of the initially displaced atoms are able to regain crystalline positions, while the rest form stable alternative arrangements, constituting the final damage in the structure.

*I. T. Todorov, and W. Smith, CSED, Daresbury Laboratory;  
K. Trachenko, and M. T. Dove, University of Cambridge.*



*Figure 1: Slice of the structure of zircon, damaged by 30 keV event (a), close overlap of two events (b), and polymerized Si-O phase around the damaged region (c).*

### References:

- [1] I. T. Todorov and W. Smith, *Phil Trans R Soc: Lond. A* **362**, 1835 (2004);
- [2] IBM SP4 cluster - <http://www.hpcx.ac.uk>;
- [3] K. Trachenko, M. T. Dove, T. Geisler, I. T. Todorov and W. Smith, *J Phys: Cond Matt*, **16** (27), 2623 (2004) (in press);
- [4] T. W. Geisler, K. Trachenko, S. Rios, M. T. Dove and E. K. H. Salje, *J Phys: Cond Matt*, **15**, L597 (2003);
- [5] <http://www.esc.cam.ac.uk/movies>;

# Studies on Manganese Oxide Perovskites in Conjunction with the CDS

The Chemical Database Service (CDS) gives UK academics ready access to a wide range of high quality datasets, and the major crystallographic data collections have always been a key feature. Locally written retrieval codes have been developed over the years to access the Cambridge Structural Database of organics and organometallics (CSD), the Inorganic Crystal Structure Database (ICSD), the Metals Data File (CRYSTMET) and the NIST Crystal Data Identification File (CDIF).

CDS has developed its own specifically tailored retrieval and analysis software package, the CrystalWeb system. It uses Web Browser technology, performs searches over all datasets simultaneously, and has a common graphics display option giving good quality display facilities for all structure types. In addition there is direct "single click" access to locate sources of the primary electronic literature, where these are available.

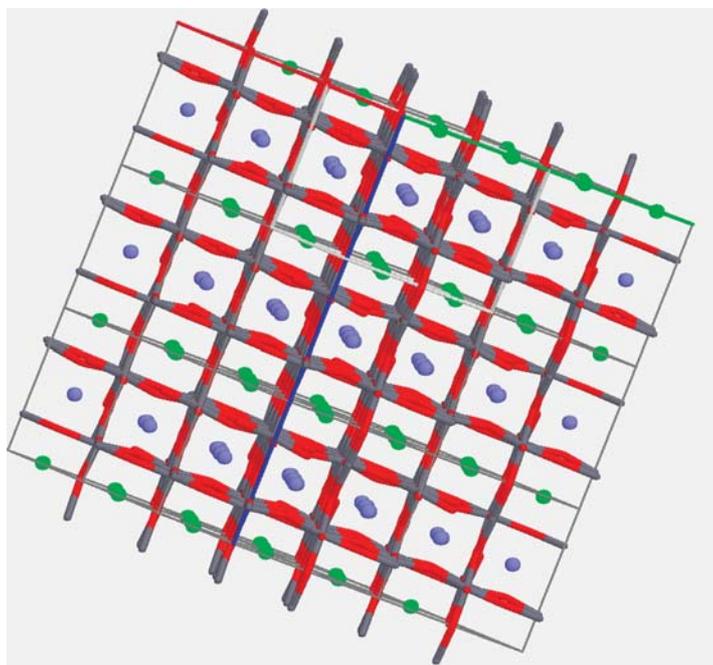
Members of the Attfield group, now in Edinburgh, routinely use the ICSD for identification of secondary phases in the synthesis of novel metal oxides. Their access to the ICSD is via the Chemical Database Service where they use both the ICSD-WWW and CrystalWeb interfaces.

Since the discovery of colossal magnetoresistances in manganese oxide perovskites,  $\text{Ln}_{1-x}\text{M}_x\text{MnO}_3$  ( $\text{Ln}$  = trivalent lanthanide,  $\text{M}$  = Ca, Sr, Ba), they have been the focus of much research over the last decade, with a view to better understanding and improving their properties. They exhibit a very rich phase diagram, with a wide variety of electronic and magnetic states which can be interconverted by changes in chemical composition at the A ( $\text{Ln}$  or  $\text{M}$ ) site. Starting originally in Cambridge, the Attfield group have worked extensively on chemical tuning of the manganites, in which three parameters have been identified as important – doping concentration (Mn oxidation state), tolerance factor and the ( $\text{Ln}$ ,  $\text{M}$ ) cation size disorder.

More recently, their research has focused on engineering an ordering of cations on the perovskite A site, and thereby potentially obtaining novel physical states.

The key tool in the structural analysis and characterisation of their manganese oxide perovskites is powder neutron diffraction, and they rely principally on Rietveld refinement of the data. The compression, however, of three dimensions of reciprocal space onto one dimension of a powder diffraction profile can lead to some very complex and subtle crystallographic models being refined on the basis of a limited amount of information. Here, the ICSD has proved invaluable in allowing them to analyse other analogous known crystal structures, and more specifically the bond lengths and angles around the cations in these materials.

This enables them subsequently to constrain their refinements to maintain reasonable metal-oxygen distances, and also, crucially, to screen their final models to ensure that all internal bond lengths and angles are realistic.



*Figure 1: The framework structure for  $\text{TbBaMn}_2\text{O}_6$  illustrating its basic Perovskite structure. The manganese (grey) centres are represented as 6-coordinate bridged by oxygens (red). A site barium (green) and terbium (blue) cations are shown as small spheres for clarity. This orientation illustrates the layered ordering of the  $\text{Tb}^{3+}$  and  $\text{Ba}^{2+}$  cations.*

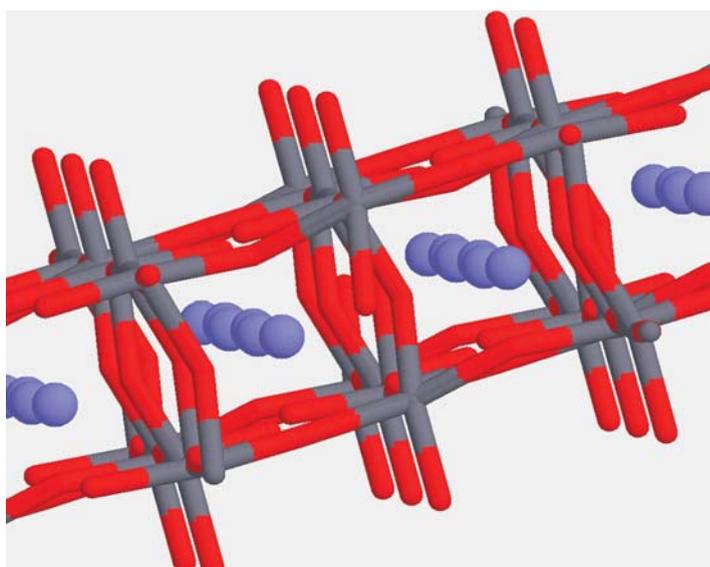
One specific example of this is recent work on the A cation ordered perovskite  $\text{TbBaMn}_2\text{O}_6$ , a highly distorted material which displays both orbital ordering and a novel charge ordering polymorph in its room temperature crystal structure. The superstructure reflections in the neutron powder pattern evidencing many of these distortions were extremely weak, rendering it impossible to freely refine all the atomic parameters in the structure. Comparison with typical Tb-O and Mn-O bond lengths in similar transition metal oxides through the ICSD allowed them to construct a superstructure model to fit the profile, whilst maintaining coordination of their metal ions consistent with that previously observed. Some details of the refined structure are shown in the *Figures*.

The CrystalWeb system continues to be developed. Data from its component databases can already be exported in a wide range of file formats. Initial support is available for XML/CML formats in line with developments and emerging standards for the Semantic Web or GRID (see the Murray-Rust Group site – (<http://wwmm.ch.cam.ac.uk>)). Display of structure using the Java-based Jmol system (<http://jmol.sourceforge.net>) in addition to Chime is also available. The underlying machinery is in place which will allow routine import of structures using a variety of formats for display via Chime/Jmol, and if required the user can then export them with a new choice of format.

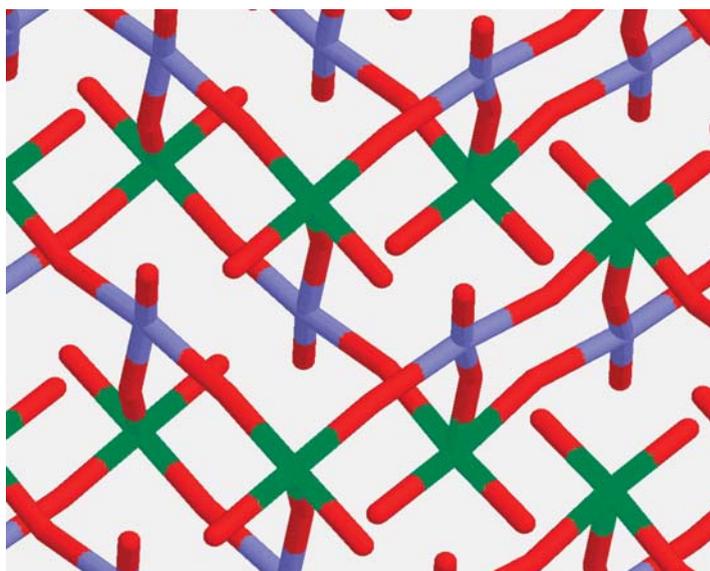
*Further information on all aspects of the CDS can be obtained by visiting the website <http://cds.dl.ac.uk>.*

*Further Research Highlights produced by users of the CDS and a fuller version of the current one with interactive links are available at [http://cds.dl.ac.uk/report/res\\_high.html](http://cds.dl.ac.uk/report/res_high.html).*

*R. F. McMeeking, and D. A. Fletcher,  
CSED, Daresbury Laboratory;  
J. P. Attfield and A. J. Williams,  
University of Edinburgh.*



*Figure 2: If the criterion for Mn coordination is tightened the  $\text{TbBaMn}_2\text{O}_6$  structure can be considered to be composed of discrete Mn/Tb/O layers separated by the bulkier  $\text{Ba}^{2+}$  cations. This figures show a fragment of such a layer with the  $\text{Ba}^{2+}$  cations omitted for clarity.*

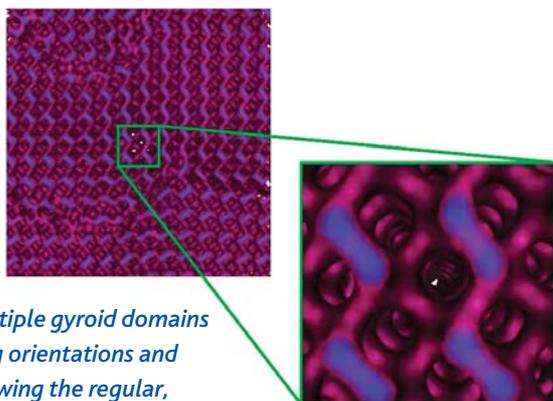


*Figure 3: The Mn/Tb/O layer in Figure 2. is shown here in a slightly different orientation and with the  $\text{Tb}^{3+}$  cations removed for clarity. Also the  $\text{Mn}^{3+}$  (blue) and  $\text{Mn}^{4+}$  (green) centres within the layer are designated by different colours. With the contact criterion selected here the  $\text{Mn}^{3+}$  sites are shown as 4-coordinate with the  $\text{Mn}^{4+}$  5-coordinate.*

# The TeraGyroid Experiment - Supercomputing 2003

The TeraGyroid project was an ambitious experiment to investigate the new opportunities for computational science created by the Grid - for example, demonstrating new scientific capabilities and international collaborative working - at the same time as establishing technical collaborations to support the development of national and international Grids. The federated resources of the UK High-end Computing facilities and the US TeraGrid were harnessed in an accelerated programme of computational materials science that peaked during the Supercomputing 2003 conference (SC'03). This project was the result of an international collaboration jointly funded by EPSRC and NSF. Trans-Atlantic optical bandwidth was supported by British Telecommunications plc. On the UK side the project team was led by:

Richard Blake - Project Coordinator,  
CSED Daresbury Laboratory;  
Peter Coveney - Principal Investigator RealityGrid,  
Dept of Chemistry, University College London;  
Peter Clarke, Networking, Dept of Physics and  
Astronomy, University College London;  
Stephen Pickles, RealityGrid Technical Director,  
University of Manchester;  
Peter Beckman (Argonne National Laboratory)  
coordinating the US involvement.



*Figure 1: Multiple gyroid domains with differing orientations and close-up showing the regular, crystalline, gyroid structure within a domain.*

Amphiphiles are chemicals with hydrophobic (water-avoiding) tails and hydrophilic (water-attracting) heads. When dispersed in solvents or oil/water mixtures, they self assemble into complex shapes; some (gyroids) are of particular interest in biology.

The gyroid conformation depends on parameters such as the abundance and initial distribution of each

chemical species and the strength of the surfactant-surfactant coupling. Desired structures can sometimes only be seen in very large systems. Smaller regions form gyroids in different directions (as shown in Figure 1) and how they then interact is of major significance in determining the constitutive properties of the mixture.

The LB3D lattice-Boltzmann model provides a hydrodynamically correct mesoscale fluid simulation method which describes the equilibrium, kinetic, and flow properties of complex surfactant-containing fluids with applications to bulk and confined geometries. A distinguishing feature of the model used in the RealityGrid work is that it takes a "bottom-up" approach, simulating the behaviour of fluids by specifying mesoscale interactions between fluids, rather than imposing postulated continuum behaviour.

The compute resources that contributed to the TeraGyroid testbed are listed in Table 1. The calculations exploited a total compute capability with some 6 TBytes of memory and some 5K processors in an integrated resource. Many  $128^3$  grid point simulations were undertaken with  $512^3$  simulations run on HPCx, and the world's largest 3 dimensional time dependent simulation with some  $1024^3$ -grid points run at PSC. The experiment generated some 2Tbytes of useful data in a very compressed period of time. The Project also used a number of visualization systems including SGI Onyx systems at Manchester, UCL, NCSA and Phoenix (on loan from SGI and commissioned on site), and the TeraGrid visualization cluster at ANL. The testbed and networks are depicted schematically (in much simplified form) in Figure 2.

The project was a major success as it resulted in the first transatlantic federation of major high performance computing facilities through the use of Grid technology. The demonstration of the experiment at SC'03 resulted in the project winning the HPC Challenge competition in the category of "Most Innovative Data Intensive Application". The TeraGyroid experiment represents the first use of collaborative, steerable, spawned and migrated processes based on capability computing accelerating scientific discovery. The TeraGyroid Project was a significant achievement, which helped to confirm that the UK does have a world-class presence in the area of High-end Computing and e-Science.

Scientific scenarios investigated within the experiment included:

- exploration of the multi-dimensional fluid coupling parameter space with  $64^3$  simulations accelerated through steering.
- study of finite size periodic boundary condition effects, exploring the stability of the density of defects in the  $64^3$  simulations as they are scaled up.  $128^3$  and  $256^3$  simulations were found to be clear of finite size effects.
- long time scale stability with simulations from 100K to 1,000K time steps. Perfect crystals were found not to form in  $128^3$  systems even after 600K time steps.
- exploring the stability of the crystalline phases to perturbations and variations in effective surfactant temperature. The statistics of the number of defects, velocity and lifetimes requires large systems as these have sufficient defects, data analysis is ongoing.

In terms of Grid technology the project demonstrated the migration of simulations (using Globus middleware) to and fro across the Atlantic exploiting the availability of resources. Integration of the systems considerably accelerated the time to insight. Calculations exploring the parameter space of the

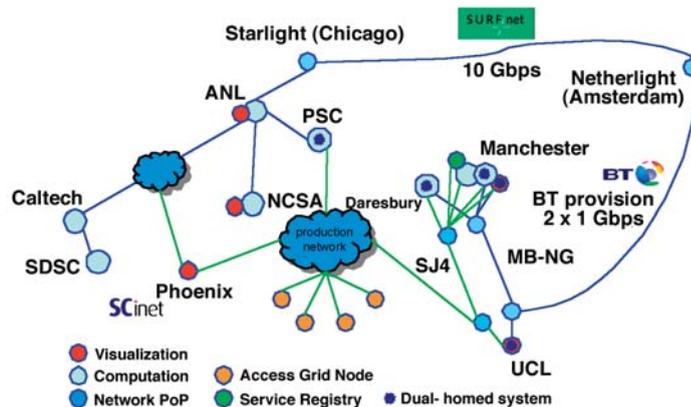


Figure 2: Schematic of TeraGrid testbed for SC'03 demonstration.

interactions within the complex fluid were steered from University College London and Boston, informed by scientific discourse over the Access Grid. Steering requires reliable near-real time data transport across the Grid to visualization engines. The output datasets were visualized at a number of sites including Manchester University, University College London and Argonne National Laboratory using both commodity clusters (Chromium at ANL) and SGI Onyx systems. The experiment was also demonstrated live before an Access Grid audience during the SC Global showcase "Application Steering in a Collaborative Environment". Figure 3 shows a snapshot of the Access Grid screen as seen in Phoenix during Super Computing Global.

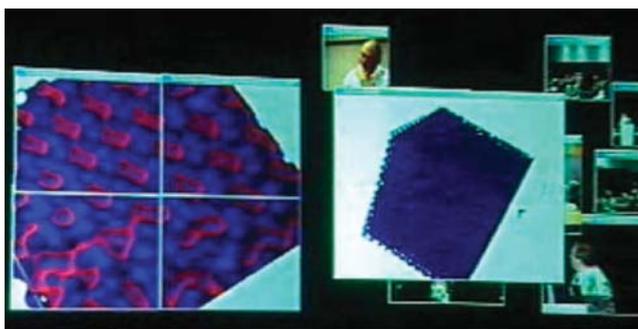


Figure 3: Access Grid screen as seen in Phoenix during the Super Computing Global session on application steering.

*R. J. Blake, CSED, Daresbury Laboratory;  
P. Coveney and P. Clarke, University College London;  
S. Pickles, University of Manchester.*

Site	Architecture	Processors	Peak (TF)	Memory (TB)
HPCx (Daresbury)	IBM Power 4 Regatta	1024	6.6	1.024
CSAR (Manchester)	SGI Origin 3800	512	0.8	0.512 (shared)
CSAR (Manchester)	SGI Altix	256	1.3	0.384
PSC	HP-Compaq	3000	6	3.0
NCSA	Itanium 2	256	1.3	0.512
SDSC	Itanium 2	256	1.3	0.512
Total		5304	17.3	5.944

Table 1: Principal computational resources in the TeraGrid testbed.