

Science & Technology Facilities Council

# Contacts and Further Information

Dr Richard Blake	Acting 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	r Damian Jones Administrative Assistant		
Miss Nina Woodall	Administrative Assistant	n.m.woodall@dl.ac.uk	

### Group leaders

Dr Steve Andrews	Systems and Operations	s.j.andrews@dl.ac.uk	
Dr Mike Ashworth	Advanced Research Computing	m.ashworth@dl.acuk	
Dr Miles Deegan	Distributed Computing	m.j.deegan@dl.ac.uk	
Dr Jennifer Scott	Numerical Analysis	j.a.scott@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	
Dr Bob McMeeking	Chemical Database Service	r.f.mcmeeking@dl.ac.uk	
Dr Chris Morris	PIMS	c.morris@dl.ac.uk	
Dr Martin Plummer	Atomic and Molecular Physics	m.plummer@dl.ac.uk	
Dr Paul Sherwood	Computational Chemistry	p.sherwood@dl.ac.uk	
Dr Martin Lüders	Band Theory	m.lueders@dl.ac.uk	
Dr Martyn Winn	CCP4	m.d.winn@dl.ac.uk	

### Websites

Computational Science and Engineering Department	www.cse.scitech.ac.uk	
HPCx Service	www.hpcx.ac.uk	
EPSRC HPC Programme	www.epsrc.ac.uk/hpc	
Science and Technology Facilities Council	www.scitech.ac.uk	

# Addresses

### Science and Technology Facilities Council

### **Rutherford Appleton Laboratory**

Harwell Science and Innovation Campus Didcot OX11 0QX Tel: 01235 445000 Fax: 01235 445808

### Science and Technology Facilities Council

### Daresbury Laboratory

Daresbury Science and Innovation Campus Daresbury Warrington WA4 4AD Tel: 01925 603000 Fax: 01925 603100

# Contents

Welcome1
Computational Science and Engineering Department CSED Software to Leap Ahead
Advanced Research Computing Group Towards the Petascale
Atomic and Molecular Physics Group Calculating Antimatter Physics and Chemistry: Collisions and Annihilation at Low Energies 6-7
Band Theory Group
E electron Valence of Carium Embedded in a Erea Electron Gas
-Science Infrastructure for the Complete Study of Phase Diagrams in High T. Superconductors 10,11
Current Induced Switching, Spin Hall Effects and Spin Orbit Coupling
Lapthapide Contraction and Magneticm in the Heavy Pare Farth Elements
CCPA Group
CCP4: Automation Developments and Improvements to the User Interface 16-17
Computational Chemistry Group
The Ab Initio based Docking Algorithm (ABIDA): From High Performance
Computing to In Silico Screening
Finding Minima, Transition States, and Conical Intersections: the Development of a Geometry Optimiser for Atmomistic Simulation Codes
Computational Engineering Crown
Computational Engineering Group
Computational Materials Science Group
Possible Mechanisms for the Catalytic Activity of Aluminium Chloro-Fluorides:
High Efficient Intermediate Band Solar Cells Based on Quantum Dot Arrays 26-27
Electronic and Magnetic Properties of Graphitic Ribbons 28-29
Numerical Analysis Group
Numerical Analysis Group
The Challenge of Larger and Larger Linear Systems
Software Engineering Group
SUBACE Engineering Group
EUKACE - European Agent-Based Computational Economics Using FLAME

# Welcome to the 2007 issue of CSE Frontiers, an Annual Report of STFC's (CCLRC's) Computational Science and Engineering Department (CSED)

Plus ça change (plus c'est la même chose) .. well things are changing and things will certainly not be the same but more of this in future reports ...

In April 2007, the Council for the Central Laboratory of the Research Councils (CCLRC) merged with the Particle Physics and Astronomy Research Council (PPARC) to form the Science and Technology Facilities Council (STFC) with responsibility for funding research in astronomy, nuclear and particle physics and enabling access to the best facilities in the world supporting the lifesciences, engineering and physical sciences and environmental sciences in particular. In Frontiers 2007 we report on recent scientific and technical highlights from work within the Department (within CCLRC) and look forward to the future.

The Department strives to ensure that UK researchers benefit from the best computational methods, supporting them through research and collaboration, theory and software development, facilities and training and support of international and industrial outreach. A highlight over the past year was the very successful International Review of our core activities in support of Engineering and Physical Science Research Council's communities, a significant fraction of our funding, which reported very high levels of Customer Satisfaction and recommended continuation of the activities at the current levels through to 2013.

The articles in this report present an insight into recent scientific and technical highlights but I suggest that you look back at Frontiers from previous years as well to gain an overview of the full scale and scope of our activities. The major themes highlighted within the current report include:

- how do we develop codes that can model large systems for long periods of time?
- what algorithms do we need to develop to efficiently and effectively exploit current and next generation HPC systems?
- how can we exploit Grid systems to compute trends in properties and phase diagrams?
- what are the best algorithms for calculating the optimum geometries for atomistic codes?
- how do we automate and accelerate the determination of protein crystal structures to inform structural genomics?
- what are the best methods for screening target molecules in rational drug design?
- simulations of antimatter physics and chemistry;
- insights into strongly interacting electron systems such as f-electron quantum impurities, lanthanide contraction and magnetism in the heavy rareearth elements

- theoretical developments that will allow us to investigate current induced switching and spin-Hall effects and possible mechanisms for the catalytic activity of aluminium chloro-fluorides;
- the magnetic properties of graphite ribbons and simulations of the electronic structure of solar cells based on quantum dot arrays;
- droplet transport in digital microfluidic systems involving electro-wetting-on-dielectric systems;
- agent based modelling applied to economics but of increasing importance in simulating coupled Lagrange-Eulerian systems; and;
- improving software quality through metrics a vital activity for the whole computational science and engineering community.

**Plus ça change**, what is going to change? Well, STFC will develop an integrated strategy on High Performance Computing covering the strategic complementary science themes of the former CCLRC and the former PPARC. These cover the gamut of computational science activities in particular multiple length and time-scales and extreme high performance computing performance. Within the new Council, we can study quarks and gluons, electrons and nuclei, nanoparticles and mesoscale flows, cells, tissues, organs and organisms; laboratory, global and astrophysical systems. Plans are being developed for the Hartree Centre – a home for UK CSE community as part of the Daresbury Science and Innovation Campus. Watch this space!

Well I hope that there is plenty to interest you in this issue of Frontiers. If it does then I hope you will look at our web site: http://www.cse.scitech.ac.uk/ and come and see us at many of the events, workshops and meetings that we hold at Daresbury, Rutherford and elsewhere around the UK and in Europe, we are constantly looking for new collaborations.



1 Rick

Richard Blake, Acting Director of CSED

# **CSED Software to Leap Ahead**

The availability of large computers, with 10,000+ processors, such as HECTOR and Blue Gene, offers challenges to computational science. The chief challenge is that the exploitation of such machines requires software to have high parallel efficiency. Early in 2007 the EPSRC put out a call for grants entitled 'HPC Software Development'. CSED was extremely successful, obtaining funding for five projects detailed here.

# Enhancing HSL for HPC Architectures (Ref: EP/F006535/1)

This project focuses on the mathematical library HSL (www.cse.scitech.ac.uk/nag/hsl) and the need to enhance its performance for HPC architectures. The emphasis is on the solution of large sparse linear systems of equations. We will investigate the benefits for these solvers of using mixed precision. HSL offers both single (SP) and double precision (DP) versions of subroutines. DP is preferred because of its increased accuracy. However, the performance of SP can be significantly better than DP on systems such as the IBM Cell BE processor. The Cell, designed for the Sony Playstation, will be produced in large volumes and may be an affordable option for future HPC systems. The Cell is optimized towards SP computation (peak performance 204.8 GFlop/s at 3.2GHz). The SPEs are capable of performing DP calculations with an order of magnitude performance penalty (14.6 GFlop/s).

The challenge is to exploit the SP performance while achieving DP accuracy. Recent studies suggest that it is possible to perform computationally expensive tasks using



Figure 1: Left: Snapshot from the Pacific of HiGEM showing Tropical Instability Waves (TIWs). Colours are the windstress divergence, indicating that the atmosphere is coupling with

SP and switch to using DP at critical stages. This project will investigate the feasibility of mixed precision solvers, leading to the development of new HSL solvers that use a combination of low and high precision arithmetic in conjunction with a combination of direct and iterative methods.

# DL\_POLY\_4: Large Scale Molecular Dynamics for Inhomogeneous Systems (Ref: EP/F010877/1)

For more than a decade the DL\_POLY package has been used in UK universities to perform molecular dynamics simulations on parallel computers. Its usefulness is due to its continuing development to exploit high performance computers, such as HPCx, where the DL\_POLY\_3 variant has proved successful. However, to progress to machines such as HECToR a new approach is required. Many modern atomistic simulations are both large and inhomogeneous; atoms in the system are often localized in clusters of differing kinds. Such systems are hard to model on parallel computers because load balancing the processors is difficult. To accommodate such systems The DL\_POLY\_4 project will develop a new molecular dynamics simulation code with a decomposition strategy that will also incorporate dynamic load balancing. With this capability, DL\_POLY will be able to scale with the numbers of processors and achieve fast, efficient simulations of large systems.



the oceanic TIWs. Right: lower resolution HadGEM model, showing that TIWs are poorly resolved. (Figures courtesy L.Shaffrey, NCAS)

# Towards Generic Scalability of the Unified Model (Ref: EP/F010885/1)

The Unified Model (UM) is used by the UK academic community for a range of research. Projects undertaken using the UM vary widely, from low resolution models to very high resolution models. The value of high resolution studies from the HiGEM project is shown in Figure 1. Such projects push the performance and scalability of the UM system on the machines available to the community. Coupled atmosphere-ocean simulations are increasingly important both for forecasting and for climate research. In future the UM will run with the NEMO (Nucleus for European Modelling of the Ocean) model.

This project will provide improved scaling of the UM and NEMO, running separately and as a coupled system, on high-performance architectures, bringing quantitative improvements to the capabilities of these codes. In particular we will demonstrate the performance benefits of i) a more loosely coupled strategy for the models, ii) a more flexible approach for input/output, and iii) a less synchronous and more latency tolerant approach. For details see the web site at http://gsum.nerc.ac.uk/.

### The Accurate Computation of Thin Film Flows and the Motion of Droplets on Real Functional Surfaces (Ref: EP/F010915/1)

Fluid flow over "real" surfaces containing micro-scale heterogeneities is ubiquitous throughout nature, engineering and the precision manufacturing industries. This project will extend the investigators' world leading software algorithms for analysing such problems to exploit massively parallel computing architectures. This will enable the computation of flows over realistic surfaces to be carried out for the first time, allowing the fundamental limits of predictive capabilities, such as those on the time-scales of droplet spreading and coalescence, to be investigated. The insight gained into this important class of flow problems will be of benefit both to academic researchers, working in the fields of coating process fundamentals and droplet motion, to engineers and scientists working in fields such as in the production of micro-electronic components and microfluidic devices for medical diagnostics, and to those interested in developing more sustainable alternatives to chemical pesticides.

# Development of Parallel Search Algorithms (Ref: EP/F01130X/1)

Enormous effort has been devoted, in the field of materials chemistry, to the efficient parallelization of the calculation of the energy and its gradients for a given geometry. Most studies also involve geometry optimization; either minimization of the energy with respect to the atomic positions, or the location of saddle points on the energy surface. One way to increase the level of parallelism is to implement parallel geometry optimization algorithms, including genetic algorithms and stochastic searches. The numerical evaluation of second derivatives, while expensive, may also be possible with large processor counts. The project will develop code to be used within a variety of CSED's packages, including CRYSTAL, DL\_POLY, GAMESS-UK and ChemShell. This should result in accelerated time to solution of many problems on largescale facilities.

> A. Wander, P. Sherwood, W. Smith, J. Kästner, C.L. Bailey, D. R. Emerson, X. J. Gu, I.T. Todorov, I.J. Bush, M. Ashworth, CSED, Daresbury Laboratory; C. Cartis, C. Greenough, J.A. Scott, I.S. Duff, J.K. Reid, CSED Rutherford Appleton Laboratory; L. Steenman-Clark, Reading. J. Gurd, C. Pilov, P. Ford, Manchester, P. H. Caskell and

G. Riley, R. Ford, Manchester. P. H. Gaskell and H. Thompson, Leeds.

### **Towards the Petascale**

High Performance Computing (HPC) has been highly successful in raising the credibility of computer simulations to such an extent that computational science now functions on equal terms with experiment and theory. As raw computer power continues to grow, one of the key challenges for HPC is to deliver on the capability aspirations of the community across a broad spectrum of scientific and engineering disciplines. New scales of modelling and new generations of simulation for scientific applications remain critically dependent on advances in high-end computing. The international standing of many UK research groups and communities relies on their access to and successful exploitation of the latest in HPC technology.

World-leading facilities are now offering around 50-100 Teraflop/s and there are several initiatives looking to Petascale computing (10<sup>15</sup> flop/s). The UK's HPC Strategy set out in "A Strategic Framework for High-End Computing" states:

"... the UK should aim to achieve sustained Petascale performance as early as possible across a broad field of scientific applications, permitting the UK to remain internationally competitive in an increasingly broad set of high-end computing grand challenge problems."

Computing at the Petascale raises a number of significant challenges, not least in the development of computational software to exploit systems of around 100,000 processors. CSED has recently been successful in securing funding from STFC's internal Strategic Initiative Programme to begin to address these challenges. We present here the first results from codes that run on >10,000 processors.

### DL\_POLY\_3 - Molecular Dynamics Studies of Radiation Hard Materials

DL\_POLY\_3 is a classical molecular dynamics package developed at STFC Daresbury Laboratory [1], and is used to study the atomistic behaviour of materials. Though explicitly designed for HPC machines the code had never been run before on machines with appreciably more than 1000 processors. However, recently an opportunity arose to test it on the Blue Gene/L machine at the Forschungszentrum Jülich, which has 16384 processors. The physical system chosen for this test was a model of radiation damage in a pyrochlore. This is of interest as these materials have been proposed as suitable media for long term storage of radioactive waste. To model this system one must study very large supercells, and the total system size was approximately 14.6 million particles.

For DL\_POLY\_3 the scaling of the time to solution depends on the complexity of the force field employed, but it is always approximately O(N). To achieve both the time and memory scaling a link-cell algorithm [2] is used. The force field can be conceptually divided in two terms:

- Short range repulsion and van der Waals (VDW) attraction
- 2. Coulomb forces

The first is short ranged, and scales very well with processor count. On the other hand Coulomb forces are long range terms, and have to be handled differently. As is standard, the Ewald sum technique is used in DL\_POLY\_3. This splits the evaluation into two terms – one short ranged, one long. The former can be handled in a similar manner to the VDW terms, and so scales well. The long range term, however, has to be handled differently. DL\_POLY\_3 uses the Smooth Particle Mesh Ewald (SPME) algorithm [3], the key feature of which is a Fast Fourier Transform (FFT). For this DaFT is used, a package written at Daresbury [4]. The use of an FFT implies a considerable amount of communication, so the scaling is ultimately controlled by this portion of the code.



Figure 1: Scaling of DL\_POLY\_3 on BG/L

### Results

It can be seen that the scaling for the various elements of the force field is good. As expected the long range Ewald terms, i.e. those terms that require an FFT, scale the least well. However, given the comparatively small size of the FFT grid, 512\*512\*512, the scaling is still good. On 16384 processors it takes just under 0.5 seconds per time step. This is sufficiently small to allow full simulations to be performed in a realistic amount of time, or to put it another way, the code runs fast enough to allow science to be done. This must be the ultimate criterion of performance!

PDNS3D - a Direct Numerical Simulation of Turbulence Fluid flows encountered in real applications are invariably turbulent. There is, therefore, an ever-increasing need to understand turbulence and, more importantly, to be able to model turbulent flows with improved predictive capabilities. As computing technology continues to improve, it is becoming more feasible to solve the governing equations of motion - the Navier-Stokes equations - from first principles. The direct solution of the equations of motion for a fluid, however, remains a formidable task and simulations are only possible for flows with small to modest Reynolds numbers. Within the UK, the Turbulence Consortium (UKTC) has been at the forefront of simulating turbulent flows by direct numerical simulation (DNS). UKTC has developed a parallel code to solve problems associated with shock/boundary-layer interaction [5].

The PDNS3D code is a sophisticated DNS code that incorporates a number of advanced features: namely highorder central differencing; a shock-preserving advection scheme from the total variation diminishing (TVD) family; entropy splitting of the Euler terms and the stable boundary scheme. The code has been written to be efficient, scalable and portable across a wide range of high-performance platforms.

The benchmark is a simple turbulent channel flow benchmark using a grid size of 360x360x360 run for 100 iterations. The most important communications structure is a halo-exchange between adjacent computational sub-



Figure 2: Performance of the PDNS3D code on two large-scale systems: the Cray XT4 (jaguar) and the IBM p5-575 (hpcx).

domains. Providing the problem size is large enough to give a small surface area to volume ratio for each sub-domain, the communications costs are small relative to computation and do not constitute a bottleneck and we see almost linear scaling from all systems. Performance on HPCx (the UK's IBM POWER5 national service) has been known to be linear out to 1024 processors. Scaling tests on the Cray XT4 at ORNL (a system very similar to the UK's nascent HECToR system) shows performance just starting to flatten off at around 12288 processors. This is for a fixed-size benchmark; larger domains should scale to even higher processor counts.

#### **References:**

[1] I.T. Todorov and W. Smith, 2004, Phil Trans R Soc Lond, A 362, 1835.

[2] M.R.S. Pinches, D. Tildesley, W. Smith, 1991, Mol Simulation, 6, 51.

[3] U. Essmann, L. Perera, M.L. Berkowtz, T. Darden, H. Lee and L.G. Pedersen, 1995, J Chem Phys, 103, 8577.

[4] I.J. Bush, I.T.Todorov, W. Smith , Comp. Phys. Commun., 175, 323.

[5] Direct Numerical Simulation of Shock/Boundary Layer Interaction, N.D. Sandham, M. Ashworth and D.R. Emerson, http://www.cse.clrc.ac.uk/ceg/sbli.shtml

### Acknowledgements:

We are grateful to the Research Centre, Jülich, and Oak Ridge national Laboratory for access to machines.

I. J. Bush, I. T. Todorov, W. Smith, M. Ashworth, CSED, Daresbury Laboratory.

### Calculating Antimatter Physics and Chemistry: Collisions and Annihilation at Low Energies.

There has been a noticeable increase of interest in the study of antimatter in recent years, spurred on by increasingly sophisticated experiments investigating both the production of anti-atoms and the mechanisms by which positrons interact with 'normal' matter leading to annihilation. In the case of anti-atoms, the ATHENA and ATRAP projects at CERN [1] are managing to cool and trap antiprotons and positrons created in high energy collisions sufficiently to form antihydrogen atoms which may then be cooled further and studied. Precision measurements of the properties of anti-atoms require very cold (of the order of a few Kelvin) anti-atoms to be available. Similarly, positrons passing through gaseous matter can be attracted close enough to the electrons of an atom or molecule to annihilate, form positronium or even form complexes with individual molecules if their kinetic energy is low or can in some way be reduced. The study of antimatter interactions thus takes place in the context of low-energy atomic and molecular physics and chemistry as well as in the high energy collision physics domain.



Figure 1: antihydrogen (A) and helium (B)

Trapped antihydrogen provides the opportunity to study fundamental principles of physics: tests of CPT (chargeconjugation, parity, time-reversal) and Lorentz invariance of relativistic quantum field theory, and also tests of various aspects of general relativity. Hydrogen and antihydrogen should have identical spectral properties. The above experiments aim to measure, for example, the antihydrogen 1s-2s transition frequency to the same accuracy as has been performed for hydrogen. Knowledge of interactions of antihydrogen with other molecules or atoms which may be particularly suitable for cooling is thus important both for the experiments and as interesting new matter-antimatter chemistry. Similarly, helium may be present as an impurity, leading to scattering and rearrangement.

The Nottingham group lead by Professor E A G Armour has been performing variational scattering calculations of lowenergy antihydrogen helium interactions: on figure 1,  $\alpha$  is the helium nucleus or alpha-particle,  $\overline{p}$  is an antiproton and *R* is the internuclear distance. As well as scattering, rearrangement can occur (figure 2).

Figure 2: rearrangement.  $\rho$  is the Jacobi-type vector of the centre of mass of positronium with respect to the centre of mass of  $\alpha$   $\overline{p}$ 



He + anti-H  $\rightarrow$  He<sup>+</sup> $\overline{p}$  + Ps He + anti-H  $\rightarrow$  He  $\overline{p}$  + e<sup>+</sup> He + anti-H  $\rightarrow \alpha \overline{p}$  + Ps<sup>-</sup>

The rearrangement scattering matrix elements at individual values of R involve integrals over products of the separated final states and the full (variational) leptonic scattering wavefunction for the incoming arrangement. These integrals also include interaction Hamiltonians or in certain cases can be recast as overlap integrals [2].

The variational calculations rely on fully-correlated treatment of the lepton interactions. The two-centre coordinate systems may be conveniently expressed in terms of prolate spheroidal coordinates [3] while Jacobitype coordinates may be used for the positronium atom and ion. In order for the wavefunctions to be accurate, functions of the attractive positron-electron distance (Hylleraas-type functions) must be included explicitly: Hylleraas-type functions of the repulsive electron-electron distance are also included explicitly for high accuracy [3]. Inter-lepton Coulomb potentials and odd powers of these distances are expanded in prolate spheroidal coordinates in non-separable infinite series: the Neumann expansion expresses distance  $r_{ii}^{-1}$  as a sum over products of associated Legendre functions (equivalent to the expansion over products of spherical harmonics in spherical coordinates) [3]. Following the variational calculation the scattering matrix elements are formed using up to 8dimensional numerical integration (some azimuthal integrations are carried out analytically).

The collaboration with CSE (M Plummer) first involved a root and branch optimization of a standalone code to form 3-particle integrals involving Hylleraas functions. The code

was analyzed both by eye and with HPCx performance analysis tools, split into component parts, modularized and rewritten to ensure no unnecessary repetitions of calculations took place. Options were put in to calculate user-specified groups of variational trial function elements simultaneously and efficiently. Memory allocation was rationalized and an automated procedure was built in for calculating component functions of the integrand using the most accurate of various alternative mathematical expressions for the particular values of the integration coordinates. The program was rewritten with a straightforward but complete user interface to be available as a fully-portable module in a variational controlling code (A Todd, Nottingham) designed to calculate and use all the necessary integrals (the CSE-optimized code forms part of a library of two- and three-particle leptonic integral modules available for general use in similar calculations [4]).



Figure 3: variation of Z<sub>eff</sub> with trial input parameters.

This controlling code was designed to experiment 'intelligently' and continuously with sets of variational trial functions, retaining a core of functions that are most important for the scattering calculation [4]. With the CSEoptimized code running at least a hundred times faster than the original code, the intelligent code was able to examine many more basis functions with the complicated but essential Hylleraas-type character. Accurate and costefficient scattering cross sections for the Ps rearrangement channel are now available [4]. Following this successful optimization, M Plummer is now collaborating with Professor Armour to optimize the

(8-dimensional numerical integration) code that forms rearrangement scattering matrix elements for the Ps<sup>-</sup> channel and which involves long-range Coulomb interactions, necessitating a particular combination of prolate spheroidal and Jacobi-type coordinate systems to minimize computational effort.

The leptonic codes used for particle antiparticle scattering at particular values of *R* are fairly straightforwardly

adapted for the case of positron hydrogen molecule scattering. An ongoing CCP2 flagship project based jointly at Nottingham and University College London is studying mechanisms that produce unexpectedly high annihilation rates for positrons in the vicinity of organic polymer molecules [5]. The effective number of electrons available for the positron to annihilate with (called  $Z_{eff}$ ) is much higher than average density calculations would suggest. It is thought that the positron is able to form an excited combined state with a molecule with its kinetic energy absorbed into the vibrational motion of molecule. Molecular electrons are then drawn towards the positron to give the high Z<sub>eff</sub>. Studies of such ionization rates and positron-molecule complexes are important to understand both new chemistry possibilities on earth (with applications, for example, to medical imaging) but also provide clues to astrophysical observations of antimatter near the centre of the Milky Way. If annihilation signals are due to positron interactions with cold molecular clouds (thought to be mainly hydrogen but also larger molecules including hydrocarbons) then correct interpretation of the signals provides important information concerning the galactic centre and antimatter chemistry.

The CSE-adapted Hylleraas code is currently being used in calculations of scattering and Z<sub>eff</sub> by J Cooper at Nottingham. Results will be of interest in themselves and will also be used as benchmarks for a new general positron molecule R-matrix code currently being developed by J Franz at UCL [3,6]. Figure 3 shows how Z<sub>eff</sub> changes for a range of input parameters to the variational calculation, allowing optimal values for these parameters to be set.

*M. Plummer, Daresbury Laboratory; E. A. G. Armour, A. Todd and J. Cooper, University of Nottingham.* 

#### References:

[1] M. Amoretti et al, Nature 419 (2002) 456-9; Gabrielse et al, Phys. Rev. Lett. 89 (2002) 213401 & 233401.
[2] E. A. G. Armour and S. Jonsell, J. Phys. A 40 (2007)

701-710. [3] 'Explicitly Correlated Wavefunctions', eds. E. A. G. Armour, J. Franz and J. Tennyson, CCP2 Workshop Proceedings (2006), CCP6 (on behalf of CCP2).

[4] A. Todd, Thesis (2007), University of Nottingham, UK

[5] 'When antimatter attacks', E. S. Reich, New Scientist, 24 April 2004, 34-37.

[6] 'Mathematical and Computational Methods in R-matrix Theory', eds. M. Plummer, J. Gorfinkiel and J. Tennyson, CCP2 Workshop Proceedings, to be published 2007, CCP2.

# F-electron Valence of Cerium Embedded in a Free Electron Gas

The electronic structure of a single ion embedded in a metal has been one of central problems in many-body physics since the 1960's. It is now 40 years since the "Kondo effect" became one of the key discoveries in condensed matter physics, when it was found that a single quantum impurity atom in a metallic host could not be described within the framework of conventional many-body physics.



Figure 1, the atom-induced electron density calculated within the local spin density functional approximation (LSDA) for a Ce atom embedded in a jellium host. This is defined as the difference between the total electron density and the electron density of the jellium host without the atom.

To study such quantum impurity atoms in a metallic host we have developed a series of computational tools using fully first principles quantum mechanics. Unlike earlier simplified approaches, such as the famous Anderson impurity model, we solve the full quantum mechanical system and there are no adjustable parameters which can be tuned to fit experiments.

We call our approach the "atom in jellium model"[1]. The atom is represented by a point charge nucleus, for example Z=58 for Ce. This nucleus is placed in a "jellium", which is an infinite region with constant positive charge density,  $n_0$ (in units where the electron charge, e=1). This can be thought of as the nuclei of the metallic host, which have been averaged (or smeared out) into a constant uniform positive charge density. Finally electrons are added until the whole system is charge neutral.

This atom in jellium system can be solved with a variety of computational methods[1]. The simplest is to use the well-known local spin density approximation (LSDA) of density functional theory (DFT). The result of such a calculation is shown in figure 1, which shows the atom induced electron

density of a Ce atom embedded in a jellium host. One can see the region of high electron density near the atom, arising from bound core states of the atom, as well as from scattering states of s, p, d, and f character. Outside of the immediate vicinity of the atom the jellium electron density is modulated by Friedel oscillations induced by the perturbation of the atom, visible in the outer region in figure 1.

The LSDA approximation is widely used in computational studies of materials, but it is known to perform poorly in systems with highly localized quantum states. This is the case for Ce and the other lanthanide and actinide rare earth elements and their compounds. For such "strongly correlated" systems the LSDA approximation makes all of the f-electrons band like and places all of them very close to the Fermi level. This leads to an unphysically high density of states at the Fermi level, compared to experimental results from photoemission spectroscopy and other probes. A more realistic approach can be achieved within DFT by replacing the LSDA exchange correlation functional with one which more accurately reflects the energetics of localized quantum states. The Self-Interaction-Correction (SIC) is designed to subtract out the "self-interaction" errors which occur in the LSDA for localized states. It has previously been shown to provide a very successful description of localization of f-electron and d-electron systems in a wide range of studies by the Daresbury Band Theory group of W. M. Temmerman, Z. Szotek, M. Lüders and their collaborators[2].

We have adapted the recent development of a new implementation of the SIC methodology, called local-SIC, by Lüders et al<sup>[2]</sup>, to the case of an atom in jellium. Specifically we apply the SIC correction to electrons within a very narrow resonance, as well as for true bound states. The resulting scattering phase shifts  $\delta^{\sigma}_{lm}(\varepsilon)$  as a function of energy,  $\varepsilon$ , are shown in figure 2. Here *l* is the angular momentum quantum number, m is the magnetic quantum number, and  $\sigma$  is the spin (up or down) of the scattering channel. The SIC correction is applied, in this case, to one of the f scattering channels, resulting in a sharp resonance (the step-like increase in the curve) shifted to well below the Fermi energy. The remaining f scattering channels are assumed band-like and so are not SIC corrected, corresponding to the two resonances pinned at the Fermi energy for spin up and down electrons. These remain resonant at the Fermi level, but with less occupation than would have been the case in a pure LSDA calculation.

Similar results are obtained by applying SIC corrections to two or three scattering channels.

Figure 3 shows the results for the total energy of a bulk Ce metal using the "effective medium theory" (EMT) from our L-SIC atom in jellium results. Our results correspond to no SIC correction, i.e.standard LSDA (tetravalent Ce), one f state SIC corrected (trivalent Ce), two f states corrected (divalent Ce) and three states corrected (monovalent Ce). The energies of each state are found as a function of the effective Wigner Seitz atomic radius, obtained within the EMT theory. The fcc- $\gamma$  phase of Ce is believed to be trivalent[2], corresponding to one localized f state.



Figure 2, scattering phase shifts for a Ce atom in a jellium host, in the case where local-SIC corrections are applied to a single f electron state. The remaining band-like f states have phase shifts with resonances near the Fermi energy, while the d states have a much broader resonance above the Fermi energy, and the s and p states are non-resonant.

Our calculations predict the minimum energy Wigner Seitz radius for trivalent Ce to be 1.99 Å, in very good agreement with the Wigner-Seitz radius of 2.02 Å fcc- $\gamma$  phase Ce.

We also find that as the Wigner Seitz radius is decreased (corresponding to increasing pressure) the tetravalent valence state becomes more favourable than trivalent, consistent with the observed high pressure transition to the  $\alpha$ -phase of Ce, believed to be tetravalent[2]. Work is continuing to explore the consequences of these f-electron localization effects across the full lanthanide rare earth series of elements.



Figure 3, total energy of a Ce "effective medium" plotted as a function of the neutral sphere radius of the atom in jellium, in units of the Bohr radius. This is the analogue of the Wigner-Seitz radius in a periodic crystal, and is proportional to the crystal lattice constant. The predicted zero pressure equilibrium lattice constant corresponds to the position of the minimum in each curve.

Acknowledgements: We are grateful for funding from STFC during this project, and have benefited from many helpful discussions with W. Temmerman, Z. Szotek, M. Lüders and B. Gyorffy.

J.F. Annett, Dept of Physics, University of Bristol, and CSED, Daresbury Laboratory; A.I. Duff, CSED, Daresbury Laboratory.

### **References:**

[1] A. I. Duff and J. F. Annett, "Variational QMC study of a Hydrogen atom in jellium with comparison to LSDA and LSDA-SIC solutions," Phys. Rev. B 76, 115113 (2007).

[2] M. Lüders, A. Ernst, M. Däne, Z. Szotek, A. Svane, D. Ködderitzsch, W. Hergert, B. L. Györffy, and W. M. Temmerman, "Self-interaction correction in multiple scattering theory", Phys. Rev. B 71, 205109 (2005).

# eScience Infrastructure for the Complete Study of Phase Diagrams in High T<sub>C</sub> Superconductors

The high  $T_c$  materials have a complex phase diagram as a function of doping, temperature and pressure. Computational studies of this multi-dimensional parameter space involve many computational runs, which are impractical to be done individually. For example, a complete study of the system can require a very large number of separate calculations, including:

- scans over different concentrations of dopants,
- scans over different configurations of localized and delocalized states,
- scans over the lattice constant,
- scans over different magnetic structures,
- scans over the temperature.

All results must then be analysed in order to extract the relevant information.

This class of problems is ideally suited to the use of Grid technologies, which can automate many of the steps involved in running parametric studies. The Grid can provide a number of highly advantageous features such as:

- · (Semi-) automatic generation of input data.
- Metascheduling across suitable compute resources in order to minimise queuing time and hence maximise throughput of the calculations.
- Simulations can be run across multiple heterogeneous compute resources, using the abstraction provided by the grid middleware.
- Both the input and output files are automatically archived within the data grid, creating a complete audit trail for the calculation.
- Key data values from the calculation are harvested from the output files and stored within a metadata database, which significantly facilitates the processing of data from a large number of related calculations.

The key point is that the system provides integrated computation, data management and metadata management, where most of the 'bookkeeping' tasks associated with the calculations are handled automatically by the Grid technology. This automation means that large scale studies can be undertaken that would otherwise be intractable simply due to the overhead of tracking the large number of calculations and managing their outputs. In order to implement this methodology the following eScience infrastructure was used:

- For the actual calculation: Fortran90 code (SIC-LSD)
- Generation of XML output: FoX-Library (F95)
- Automatic metadata insertion: RCommands
- Storage of data files: SCommands (SRB client tools)
- Data extraction: AgentX: the AgentX framework aims to facilitate the inter-operability of applications that form part of complex workflows and to promote the development of common tools that can be used to analyse and post process scientific data.
- Job submission: Remote My\_Condor\_Submit (RMCS)
- Data Analysis: Rgem



The plot shows the energy difference between the  $d^9$  and the  $d^8$  configurations: positive values indicate the metallic phase, while negative values correspond to the magnetic, insulating phase. It can be seen that the critical hole concentration depends on the lattice constant, and can be lowered by compressing the system.

As a demonstrator calculation we took  $CaCuO_2$ as a prototype of the high temperature superconductors. In its pure form, it is an anti-ferromagnetic insulator. Upon doping with Sr, which introduces holes in the  $CuO_2$  layers, it undergoes a Mott metal-insulator transition to a nonmagnetic metallic phase, which occurs at a hole concentration of about 0.15. This cannot be explained with the local spin-density (LSD) approximation, which would predict a nonmagnetic metal also for the undoped system. However the SIC-LSD method is a very powerful tool to investigate the ground state configuration of strongly correlated materials, and which correctly describes the undoped system.

10

More specifically, the self-interaction corrected LSD (SIC-LSD) allows for:

- the correct description of the magnetic phase through localization of all Cu *d*-electrons (d<sup>9</sup> configuration).
- the description of the transition to the metallic state through the delocalization of one *d*-electron (d<sup>8</sup> configuration).

Computationally, the transition can be investigated by studying the energy difference between the d<sup>8</sup> and the d<sup>9</sup> configurations. Here this energy difference is plotted as a function of the hole doping (through a rigid band description to mimic Sr doping) for different lattice constants.

The XML output from the simulation codes simplifies the automatic extraction of key information at the postprocessing stage. This is then stored within the metadata database, providing a highly efficient method of summarising the results from a large number of individual calculations, e.g. during a study to explore a phase space. In addition, XML output can be transformed to XHTML, using XSLT tools such as ccViz. This rendered output facilitates information delivery as, for example, it can include embedded SVG graphs showing the variation of key properties during the self-consistency iterations. After this successful feasibility study, the Grid-technology will be applied to a series of high-profile applications:

- ground state configuration of U compounds:
   spin-orbit coupling reduces the symmetry of the system: more configurations
- finite temperature phase diagram of Pu:
  - at T=0K: a wealth of configurations which are close in energy
  - at T>OK: include thermal fluctuations between these configurations
- study of the stripe phases of LaCu<sub>2</sub>O<sub>4</sub>:
  - striped phases require a large supercell: large number of d<sup>8</sup>/d<sup>9</sup> configurations
- metal insulator transition in LaCu<sub>2</sub>O<sub>4</sub>:
  - explore phase space spanned by hole concentration and lattice constant
  - use CPA to model the hole doping and a mixed valence of Cu

Also, we aim at data interoperability in collaboration with Oak Ridge, involving four different LMTO codes. The combined efforts include:

- ORNL have developed an XML schema
- ORNL uses Xerces library for XML parsing and validation
- CCLRC has developed the AgentX library to map from an appropriate ontology to the specific XML data format.
- CCLRC and ORNL have organised a common workshop and have regular AccessGrid meetings

Recent progress was reviewed at a Computational Science Conference (CompSci07) in London, June 25-26, 2007.



Figure showing raw text output, XML output and transformed XHTML output from the Daresbury SIC-LMTO code. The XML facilitates extraction of key data from the output, which can then be stored as metadata. In addition, the XML output can be readily transformed to XHTML, including embedded SVG graphs that can simplify the analysis of the output data. P. Couch, R. Tyer, eScience, Daresbury Laboratory; M. Lüders, Z. Szotek, W. M. Temmerman, CSED, Daresbury Laboratory; J. F. Annett, Bristol University and CSED, Daresbury Laboratory; L. Petit, T. C. Schulthess, G. M. Stocks, M. Summers, Oak Ridge National Laboratory.

### Current Induced Switching, Spin-Hall Effects and Spin-Orbit Coupling

Motivation: Conventional electronics is based on manipulating the flow of electrons with electric fields. Of course, we could also influence the motion of electrons by coupling to their spins with exchange or magnetic fields. Indeed, this is what the nascent technology of spin-based electronics, often referred to as Spintronics, aims to do. A spectacular example of this strategy is the Magnetic Random Access Memory (MRAM) devices based on the Giant Magneto Resistance (GMR) phenomena. While these are already a successful industry, the scientific significance and technological potentials of its inverse, 'Current Induced Switching', is only now beginning to be explored. The basic, surprising, physics here is that a spin-polarized current through a magnetic heterostructure can exert a torque which changes the relative orientation of its parts. A summary of the experimental situation is depicted in Figure 1.



Figure 1: Spin transfer by current flow.

(M.D. Stiles and J. Milat in 'Spin Dynamics in Confined Magnetic Structures III', Eds. B. Hillebrands and A. Thiaville, Springer-Verlag, Berlin 2006)

Evidently, in a device this could mean 'writing ' information into its magnetic state. Normally, such 'writing' is done by turning a magnetic bit over using a magnetic field. What makes the prospects of a new technology based on 'Current Induced Switching' attractive is that here there is no need for an external magnetic field, a current pulse will suffice, and this facilitates the miniaturazion to nanoscale.

Other unexpected current induced magnetic phenomena are the **Spin-Hall Effects**. In the simplest of these a current through a metallic strip will result in a spin polarization at the edges with opposite polarization on the opposite sides as illustrated in Figure 2.



*Figure 2: Spin accumulation at lateral boundaries.* (*H.-A. Engel, E.I. Rashba and B.I. Halperin arXiv: Condmat/0603306v1* 10 Mar 2006)

The Conceptual Difficulty Arising From Spin-Orbit Coupling: Classically speaking, an electron moving with a velocity  $\vec{v}$  carries a spin  $\vec{s}$  as well as a charge e and thus it corresponds to a spin current  $\vec{J}_s$  and a charge current  $\vec{J} = e\vec{v}$ . Indeed the spin current tensor,

$$\overrightarrow{J_s} = \overrightarrow{s} \otimes \overrightarrow{v}, \tag{1}$$

where  $\otimes$  denotes a tensor product, and its quantum mechanical generalization is the central concept for understanding all the above spitronics phenomena. Thus it is a general and fundamental difficulty in the field that this latter quantity has no well defined meaning in the presence of spin orbit coupling. Evidently, an adequate theory of spintronics must deal with the spin and orbital degrees of freedoms of the current carrying electrons on equal footing and this in relativistic quantum mechanics implies a spinorbit coupling. The source of the problem is that, unlike classically as above or in non-relativistic Schrödinger Quantum Mechanics, in Dirac's theory one can not specify the spin of an electron without reference to its state of motion. In short the spin operator  $\overline{\Sigma}$ , relativistic version of the Pauli spin operator  $\vec{\sigma}$ , does not commute with relativistic velocity operator  $c \vec{\alpha}$ . Therefore, the spin polarization is not a conserved quantity and hence there is no Noether Theorem and no spin-current defined by a continuity equation.

The resolution of the dilemma: Vernes, Gyorffy and Weinberger [1] observed that the dilemma disappears if one describes the polarization of a moving electron using the Wigner-Bargmann covariant polarization operator well known in accelerator physics (for review see Fradkin and Good [2]),

12

$$\vec{T} = \beta \tilde{\Sigma} + \frac{\gamma_{5}}{m_{e}c} \vec{P}, \qquad (2)$$

where  $\beta$  and  $\gamma_s$  are Dirac matrices and  $\vec{P}$  is the momentum operator  $-i\hbar\vec{\nabla}$ . The point they make is that unlike the spin operator in the rest frame  $\vec{\Sigma}$ , the polarization operator  $\vec{T}$  commutes with the free electron Dirac Hamiltonian and hence the polarization density  $\vec{\tau}(\vec{r},t)$ , defined as the expectation value of  $\vec{T}$  with respect to a four component solution of the Dirac equation  $\psi(\vec{r},t)$ , satisfies a continuity equation

$$\partial_t \vec{\tau}(\vec{r},t) + \vec{\nabla} \cdot \vec{J}_s(\vec{r},t) = 0 , \qquad (3)$$

which defines a spin current tensor  $\vec{J}_{,i}(\vec{r},t)$ . Using the time dependent Dirac equation to find  $\partial_t \vec{\tau}(\vec{r},t)$  leads to

$$\vec{J}_{s}(\vec{r},t) = c \,\psi^{*}(\vec{r},t)\vec{T} \otimes \vec{\alpha} \,\psi(\vec{r},t)\,, \tag{4}$$

which is a very satisfactory generalization of the simple classical formula in equation (1).

To appreciate the physical significance of  $\vec{J}_s(\vec{r},t)$  note that a simple general consequence of equation (3) is that the total polarization within a volume  $\Omega$  enclosed by a surface *A* experiences a torque given by the formula

$$torque = -\oint_{A} \overrightarrow{J_s} \cdot \overrightarrow{dA}, \qquad (5)$$

where  $\overline{dA}$  is the differential surface element normal to the surface at each point on the surface. Namely, if there is net polarization current flowing out of a volume  $\Omega$  the total polarization inside will rotate even if its magnitude does not change in time.

This is what happens in current induced switching: a current carried by spin polarized electrons pumps 'Spin Current' into a magnetic target whose magnetization will experience a torque. The above conceptual framework describes this process even in situations when spin-orbit coupling is important. Moreover, in the presence of an electric filed in the direction and to lowest order in  $1/c^2$  the above arguments lead to a correction to the spin-current tensor

$$\delta J_{s,y}^{x} = \frac{e\hbar}{2mc^{2}} n_{e} E_{z} , \qquad (6)$$

where  $n_e$  is the electron density. Evidently, this contribution has the Spin Hall symmetry indicated in Figure 2 implying that Spin Hall effects are generic features of Relativistic Quantum Mechanics.

**References:** 

[1] A. Vernes, B. L. Gyorffy and P. Weinberger, Phys. Rev. B76, 012408, (2007).

[2] D. M. Fradkin and R. H. Good, Rev. Mod. Phys.33, 343 (1961).

B. L. Gyorffy, Bristol University and CSED Daresbury Laboratory.

# Lanthanide Contraction and Magnetism in the Heavy Rare Earth Elements

The late rare earths (REs), namely gadolinium (Gd) to Thullium (Tm) all crystallize in the same hexagonal close packed (hcp) structure. They share the same outer shell electron configuration and differ only in the nuclear charge and the number of 4f electrons, which are chemically inert. However, despite their similarities, they show a rather complex magnetic behaviour, ranging from a ferromagnetic ground state in Gd to complicated, incommensurate spin spiral structures from Terbium (Tb) to Thullium. Although experiments indicate that it is the dependence on the c/a ratio that leads to this diversity in magnetic structures, the nature of this complexity is far from being understood. Here we use an approach that combines the disordered local moment (DLM) theory and the self-interaction corrected local spin density (SIC-LSD) approximation to density functional theory (DFT) to investigate the finite temperature magnetism of these materials, where electron correlations are treated better than in LSD. The physical picture that emerges from this study is that the localised 4f electrons constitute local magnetic moments, which polarise the conduction electrons.



Figure 1: Magnetic structure of Gd as function of the c/a ratio and Wigner-Seitz redius. Other late rare earths are also marked, indicating their magnetic structures, which coincide with the structures predicted for Gd at different lattice parameters. Red marks incommensurate order, blue refers to a ferromagnetic state and the green/yellow defines the region when the ferromagnetic and incommensurate phases compete against each other.

The latter mediate the exchange interaction between the local magnetic moments. Our study confirms the paramount importance of the c/a ratio in determining magnetic structures of the heavy REs, but in addition we discover that the unit cell volume plays a separate, distinct

role. Treating Gd as a prototype for all the late REs, we have been able to construct a unified phase diagram (see Figure 1) which has true predictive power and correctly describes magnetic structures not only for the late RE elements, but also their alloys, based on their lattice parameters and volumes.

It is a non-trivial task to determine complex magnetic structures in a conventional way, using LSD, as one has to explore a variety of possible magnetic structures to find the one with the lowest energy. For incommensurate structures, this is a formidable task. In our method, the ground state magnetic structure is approached starting from the high temperature paramagnetic phase, where the local magnetic moments are randomly oriented. The central quantity of this DLM method, which is rigorously formulated within DFT, is the paramagnetic spin susceptibility. It describes the response of the system to a magnetic perturbation, and hence the spin fluctuations present in the paramagnetic phase. A magnetic instability of the system is indicated by a divergence of the susceptibility as a function of temperature and wave vector **q** of the perturbation, which then determine respectively the Curie (or Néel) temperature and the magnetic order, assumed by the system just below the transition.

We have applied this methodology to Gd in its equilibrium structure (both experimental and theoretical). In order to properly describe the electronic structure and the ferromagnetic ground state of Gd, it has been necessary to take into account the strong correlations of the localized 4f electrons by means of SIC-LSD. The standard LSD calculations fail to describe the correct ferromagnetic ground state of Gd, predicting incorrectly a commensurate antiferromagnetic order. The SIC-LSD method has been implemented within a KKR-CPA multiple scattering code, which utilizes the coherent potential approximation (CPA) to perform the ensemble averages over the random spin and valence configurations. With this method, the calculations carried out at the experimental lattice parameters have given a Curie temperature of 324K, while for the theoretically optimized lattice parameters we have obtained the value of 280K, in fair comparison with the experimental value of 293K. The overestimation is due to the mean-field treatment of the spin fluctuations.



Figure 2: Spin susceptibility of Gd for different c/a ratios at a fixed Wigner-Seitz radius. One can see that for c/a<1.597, the shoulder at q wins over the ferromagnetic peak.

The calculated susceptibiliy for different c/a ratios is presented in Figure 2. Note that with decreasing c/a ratio a shoulder develops at an incommensurate q vector, which eventually wins over the ferromagnetic peak at q=0. This shows that changing the c/a ratio leads to a change of magnetic structure from ferromagnetic to some, in general, incommensurate spin modulation. An important finding of this study is that a similar dependence is also found when changing the volume of the unit cell. The shoulder and the dependence of its height on the lattice parameters can be understood in terms of topological features of the Fermi surface (see Figure 3).

One important contribution to the spin susceptibility comes from the conduction electrons and is equal to:

$$\chi_0^{conduct}(\mathbf{q}) = \int d\varepsilon \int d\varepsilon' \int_{BZ} d^3k \frac{f(\varepsilon) - f(\varepsilon')}{\varepsilon - \varepsilon'} A_B(\mathbf{k}, \varepsilon) A_B(\mathbf{k} + \mathbf{q}, \varepsilon')$$

Since the integrals are dominated by the  $\varepsilon_F$  contributions from the Fermi energy, the above expression can be approximately written as a convolution of Bloch spectral functions at the Fermi energy,  $A_B(\mathbf{k},\varepsilon_F)$  which is the generalization of the Fermi surface to disordered systems:

$$\chi_0^{conduct}(\mathbf{q}) \approx \int_{BZ} d^3 k A_{B}(\mathbf{k}, \varepsilon_{F}) A_{B}(\mathbf{k} + \mathbf{q}, \varepsilon_{F}).$$

From this expression one can see that large values of  $\chi_0^{conduct}(\mathbf{q})$  can arise for those  $\mathbf{q}$  vectors for which large parts of the (generalized) Fermi surface coincide, when translated by  $\mathbf{q}$ . This is called a nesting of the Fermi surface.

When the *c*/*a* ratio is varied, the topology of the Fermi surface changes and the nesting feature eventually disappears, as seen in Figure 3. It corresponds to the supression of the shoulder in the susceptibility. As already mentioned, this dependence of the spin susceptibility of Gd on the *c*/*a* ratio and the volume has given rise to the unified phase diagram of Figure 1, which allows one to predict magnetic structures of the late REs and their alloys. By compressing the volume and reducing the *c*/*a* ratio one can realize a transition from the ferromagnetic order to some incommensurate spin phase. What is reassuring here is that for all the cases marked in Figure 1, the predicted magnetic structures agree well with the experimentally determined magnetic order. It is also important to stress here that with this phase diagram we have established a clear link between the gradual decrease of the lattice parameters, known as the lanthanide contraction, and the arrival of incommensurate structures at the end of the REs series.



Figure 3: The Bloch spectral functions of Gd at the Fermi energy for three different c/a ratios.

### **References:**

I.D. Hughes, M. Däne, A. Ernst, W. Hergert, M. Lüders, J, Poulter, J.B. Staunton, A. Svane, Z. Szotek and W.M. Temmerman, Nature 446, 650 (2007). I. D. Hughes, J. B. Staunton, University of Warwick; M. Däne, University of Halle and MPI Halle; A. Ernst, MPI Halle; W. Hergert, University of Halle; M. Lüders, Z. Szotek, W. M. Temmerman, CSED, Daresbury Laboratory; A. Svane, Aarhus University; J. Poulter, Bangkok University.

# **CCP4: Automation Developments and** Improvements to the User Interface

CCP4 [1] is a highly successful collaborative computational project and involves the development, maintenance and distribution of a suite of software for solving protein structures by processing X-ray diffraction image data. The software is used extensively by protein crystallographers and other researchers world-wide. The process of going from the X-ray images to the final 3-dimensional structure of the target protein involves several steps with several possible routes through the software depending on how the X-ray data collection experiment was carried out. Each of the steps is catered for by several different programs and the use of a library of helper applications for manipulating the various data formats used.

As the methods have become more robust, Protein X-ray

Crystallography (PX) has become a more and more important scientific tool to an ever increasing variety of scientists and commercial users. Many of today's researchers require PX as a tool in a larger project and have less time to gain a deep understanding of the PX structure solution methods. In addition, with the realisation of the importance of X-ray synchrotron radiation for

Appleton Laboratory. In Phase 1 of its commissioning 3 of the first 7 beamlines are dedicated to Protein X-ray Crystallography.

scientific research, there has been a rapid expansion in the number of synchrotron radiation sources being built around the world. A large percentage of the capacity of these sources is usually devoted to PX. For example, three out of the first seven beamlines to be commissioned at the UK Diamond synchrotron (figure 1) have been devoted to PX [2]. The availability of these facilities and the revolution taking place in Biology due to the success of the various gene sequencing projects, has engendered the development of high-throughput projects such as the JCSG Structural Genomics project [3] and the eHTPX e-Science project for high throughput protein crystallography [4].

These developments have put increasing pressure on the software used for solving protein structures to become both more automated and easier to use. CCP4 has sought to address these requirements by initiating several automation projects for automating different parts of the structure solution pipeline. To facilitate these projects and to allow for the accurate recording and tracking of the large amount of data that these automated pipelines will produce, a separate project, the dbCCP4i [5] project tracking system, has been developed. This project has three main aims, to provide a database for the storage of the results and parameters from the various CCP4 programs used in each pipeline, to provide a database handler allowing for efficient interactions between the database and the pipelines and to provide an interactive data visualiser allowing users to easily view their projects and help them to infer the best result possible from the data processing.

One example of an automation project to adopt the dbCCP4i developments is the MrBUMP project [6] (Molecular Replacement with Bulk Model Preparation) for

> automating the process known as Molecular Replacement (MR). In MR, pre-solved protein structures that are related to the target protein through a common evolutionary ancestry can be used to aid the determination of the target structure. It is assumed that the phase information from the related or homologous structure will be similar to that of the target due to

their structural similarities. Combined with the X-ray diffraction data, the related protein is used to determine the orientation of the target protein structure in the unit cell of the crystal lattice and hence the phase information relating to the target structure. This then gives an interpretable 3-D electron density map for the target from which the shape of its structure can be determined.

The success of the Molecular Replacement process depends on several factors such as the 'closeness' of the template and target amino-acid sequences and the resolution to which the X-ray diffraction data has been collected, typically of the order of Angstroms (10<sup>-10</sup> metres). MrBUMP automates the process of seeking out the best template structures from the database of previously solved structures, the Protein Data Bank (PDB) [7], and preparing them for use in MR against the target data. It can prepare many possible templates and will try all of them in MR until a solution is found. For routine Molecular Replacement problems, the pipeline automates what a



crystallographer might do and its value is simply one of convenience. For more difficult cases, the pipeline aims to discover the particular template structure and model edits required to produce a viable search model and may succeed in finding an efficacious combination that would be missed otherwise. To speed up the process, the program can take advantage of compute clusters to farm out the MR and structure refinement processing. XIA2 is designed to work from the raw images measured at the beamline and a little metadata explaining how the measurements were made. This is used to determine the protocols for the data reduction, and the form of the final results. For straightforward cases it is possible to go from the raw images to interpretable electron density *via* experimental phasing or molecular replacement with very little user input – the final goal of automation.



Figure 2. The dbCCP4i dbviewer in use in MrBUMP. Users can easily access any of the input, output or log files from any part of the pipeline allowing them to quickly decide on the best result and proceed with completing the 3-dimensional structure of their target protein.

A typical run of MrBUMP can produce a large amount of data. Guidance is given as to which solutions are good, marginal or poor but it is left to the user to decide which of the possible solutions is the best one to proceed to further processing with. The dbCCP4i system has been adopted in MrBUMP to help keep track of all this data and to present it to the user in an easy to interpret and informative way via the dbCCP4i dbviewer application (figure 2).

XIA2 [8] is another example of automation for protein crystallography. This is an expert system for performing the early stages of data reduction - from the images measured at the beamline to the reduced data required for structure solution with e.g. MrBUMP - making use of existing data reduction packages including CCP4. In the past, crystallographers performing diffraction experiments would have a clear understanding of the data reduction process and the various steps needed to get the best from the data. Now, however, the typical user understands more biology than crystallography and is primarily interested in the results, and may not wish to spend time learning the ins and outs of data reduction. Equally, with the advent of high throughput crystallography, it is becoming more common to measure dozens of complete data sets on a single visit to a synchrotron. To properly reduce this data would take many days of careful processing by hand.



Figure 3. An example of a search model for use in molecular replacement. This is an ensemble of the top candidate search model structures superposed upon each other. Using this type of model in MR can sometimes improve the success rate of the search for a solution to the target protein structure. MrBUMP generates an ensemble as one of it's search models.

R. Keegan, G. Winter, W.Yang, P. Briggs, CSED Daresbury Laboratory.

### **References:**

[1] CCP4 Collaborative Computational Project, Number 4. (1994), 'The CCP4 Suite: Programs for Protein Crystallography'. Acta Cryst. D50, 760-763. [2] Phase 1 of the Diamond Light Source,

http://www.diamond.ac.uk/Beamlines [3] Joint Centre for Structural Genomics,

http://www.jcsg.org

[4] e-Science Resource for High Throughput Protein Crystallography, http://www.e-htpx.ac.uk

[5] Project Tracking System for Automated Structure Solution Software Pipelines, W. Yang, R.Keegan, P. Briggs, CCP4 Newsletter, Summer 2007.

[6] Automated search-model discovery and preparation for structure solution by molecular replacement, R. Keegan, M. Winn, Acta Cryst (2007). D63, 447-457. [7] Protein Data Bank, Nature Biotechnology (2007), 25, 845-846.

[8] xia2 Automated Data Reduction, http://www.ccp4.ac.uk/xia

### The Ab Initio Based Docking Algorithm (ABIDA): From High Performance Computing to In Silico Screening

Molecular docking is one of the best established computational methods to predict the affinity of small molecules towards a target enzyme and is widely used for rational drug design. The basic tenet of the method is that for a biologically relevant effect the molecule has to find an energetically favourable binding configuration to the target. Typically a molecular mechanics based classical model is used for the interaction energy and a quality of fit measure is derived from known protein-substrate complexes. Although docking methods have seen many improvements recently, still even the most elaborate energy functions determine electrostatic potentials from a point-charge model of the enzyme. With current computers we think this particular simplification is no longer warranted, instead we propose to use ab initio computed electrostatic potentials. Below we will explain our approach as presented recently at the CompLife'07 conference [1].

In practical docking problems the interaction energy between a few thousand protein atoms of the target and up to a hundred atoms of the ligand needs to be evaluated many times. Starting from a simple Coulomb plus Lennard-Jones potential the total energy is approximated so that this can be expressed in terms of protein potentials that can be pre-calculated on a grid:

$$E = \sum_{j=1}^{N_{ligand}} E_j \quad E_j \approx V_{klm}^{Cou} q_j - V_{klm}^{disp} \sqrt{b_j} + V_{klm}^{rep} \sqrt{a_j}$$

Here  $E_j$  is the interaction energy of a ligand atom at grid point k,l,m with all the target atoms. The Van der Waals parameters  $A_{ij}$  and  $B_{ij}$  have been approximated as

$$A_{ij} \approx \sqrt{a_i} \sqrt{a_j} \qquad B_{ij} \approx \sqrt{b_i} \sqrt{b_j}$$

The potentials are

$$V_{klm}^{Cou} = C \sum_{i=1}^{N_{\text{target}}} \frac{q_i}{r_{i,klm}} \quad V_{klm}^{disp} \approx \sum_{i=1}^{N_{\text{target}}} \frac{\sqrt{b_i}}{r_{i,klm}^6} \quad V_{klm}^{rep} \approx \sum_{i=1}^{N_{\text{target}}} \frac{\sqrt{a_i}}{r_{i,klm}^{12}}$$

Traditionally the Coulomb potential V <sup>Cou</sup> is computed using point charges from molecular mechanics calculations but it could equally well be calculated from an ab initio charge density as

$$V_{klm}^{Cou} = C\left(\sum_{i=1}^{N_{\text{target}}} \frac{Z_i}{r_{i,klm}} - \int \frac{\rho(r)}{|r - r_{klm}|} dr\right)$$

provided the calculation of the electron density  $\boldsymbol{\rho}$  is possible.

As a test system we'll consider the binding site of mycobacterium tuberculosis Isocitrate Lyase (ICL, PDB entries 1F61, 1F8I, 1F8M). The whole enzyme contains about 18000 atoms, but we know that a sub-region of about 16Å around the  $Mg^{2+}$ -ion (i.e. about 2000 atoms) is sufficient for docking studies. The size is reduced further by assigning different basis sets to concentric shells from TZVP in the central region, through 6-31G between 6Å and 12Å, to STO3G beyond 12Å, as shown in Figure 1. The resulting calculation involves some 10000 basis functions. The main challenge is to do ab initio calculations from this scale up at all.



Figure 1: Whole ICL enzyme in light grey, and the 16Å cluster with its different basis sets colour code. The pink atom is the  $Mg^{2+}$ -ion in the active site.

In this study we have computed the Coulomb potential with GAMESS-UK [2]. The most time consuming part is the calculation of the density using a direct DFT implementation. This is parallelized in a replicated data fashion: The same code runs on every processor of the system, each maintaining an identical copy of the data structures. The work is divided across the processors, and communication phases are introduced to share the results so that all processors have the same data at the start of the next phase. Also various linear algebra operations such as matrix transformation and diagonalisation have been parallelised using the Global Arrays [3]. Although a large number of processors can now be used efficiently the fact that a complete copy of all data structures is kept on every processor presents a memory bottleneck. The recently written Fortran module FIPC [4] may provide a solution to this problem by enabling the use of shared memory segments in Fortran programs. We are currently investigating this as a way forward.

Having calculated the electron density and the potentials with the above code we performed docking studies of succinate for which correct binding geometries are known from crystal structures. In Figure 2 we compare the results obtained with the point charge derived potential (a) against those obtained with the ab initio derived potential (b).

The optimum binding geometry of the point charge based docking study differs from that of the ab initio based study in three critical points: A hydrogen bond with the coordinated water molecule next to the pink  $Mg^{2+}$ -ion (arrow b-1) is lost (arrow a-1), in favour of a 'promiscuous' Coulomb interaction within the not very discriminating wide positive potential region (arrow a-2). In the ab initio based field the docked ligand retains the interactions that

are expected from the crystal structure (a hydrogen bond at b-1 and with an arginine residue at b-3). Furthermore the point charge based docking study yields three more incorrect local minima, as well as the correct one but at 2 kCal/mol above the lowest energy one. This means that this study predicts that less than 3.5% of the ligands would occupy the correct position. By contrast, the ab initio based study yields only one, and the correct one, optimal binding mode. This is an important improvement as docking studies often find more optimal structures than the crystal structures indicate. Therefore we conclude that the ab initio derived potentials offer a feasible route to significantly improve the reliability of docking studies and that they are to be preferred over the traditional point charge derived potentials.



Figure 2: Docking result of succinate (green carbon atoms) in a point charge potential field (a) and in the ab initio potential field (b). See text for the arrows 1, 2, and 3.

H. J. J. van Dam, I. J. Bush, P.Sherwood CSED, Daresbury Laboratory;

M. R. de Jonge, MolMo Services BVBA, Turnhout, Belgium;

J. H. van Lenthe, Utrecht University, Utrecht, The Netherlands;

H. Maarten Vinkers, MolMo Services BVBA, Turnhout, Belgium;

F. Daeyaert, MolMo Services BVBA, Turnhout, Belgium;

M. F. Guest, ARC, Cardiff University, Cardiff.

### References:

[1] Ab Initio potential grid based docking: From High Performance Computing to In Silico Screening M. R. de Jonge, H. Maarten Vinkers. F. Daeyaert, J. H. van Lenthe, I.J. Bush, H. J. J. van Dam, P. Sherwood and M. Guest. Submitted to CompLife07, proceedings in Lecture Notes in Computer Science. http://www.complife.org/

[2] GAMESS-UK is a package of ab initio programs. See: "http://www.cfs.dl.ac.uk/comment

"http://www.cfs.dl.ac.uk/gamessuk/index.shtml", M.F. Guest, I. J. Bush, H.J.J. van Dam, P. Sherwood, J.M.H. Thomas, J.H. van Lenthe, R.W.A Havenith, J. Kendrick, "The GAMESS-UK electronic structure package:

J.H. van Lenthe, R.W.A Havenith, J. Kendrick, "The GAMESS-UK electronic structure package: algorithms, developments and applications", Molecular Physics, Vol. 103, No. 6-8, 20 March-20 April 2005, 719-747.

[3] http://www.emsl.pnl.gov/docs/global/[4] http://www.hpcx.ac.uk/research/hpc/technic al\_reports/HPCxTR0701.pdf

# Finding Minima, Transition States, and Conical Intersections: the Development of a Geometry Optimiser for Atomistic Simulation Codes

Geometry optimisation is a major task in atomistic simulations of materials, chemical systems, and biological matter. Several types of optimisation are possible depending on the feature of the energy surface that needs to be located. Minima of the energy surface refer to structures the system is most likely to adopt. A transition state (first-order saddle point on the energy surface) is the average crossing point in a transition from one minimum to another, which occurs for example during a chemical reaction. A conical intersection is a geometrical region where two electronic states are degenerate (that is, they have the same energy). The lowest energy point of conical intersection is the most likely geometry for a radiationless transition to occur between the two states.



Figure 1: Converged nudged-elastic band path on an example surface (Müller-Brown potential). The green spheres indicate minima, the blue sphere the climbing image which converged to the transition state.

We have designed and implemented a novel geometry optimiser, DL-FIND, which provides several algorithms for finding minima and conical intersections. Its main strength, however, are the transition state optimisation routines. Besides the standard P-RFO method, the nudged-elastic band and the dimer method are implemented. The modular design of DL-FIND separates the handling of the coordinate system and the optimisation algorithm. This facilitates new combinations of optimisation techniques, such as using the dimer method with second-order optimisers (see below).

The code is included in the CSED software packages ChemShell and GAMESS-UK. It is designed to present a well-defined minimal interface to the calling code, which should facilitate its inclusion into other codes as well.

### **Coordinate Systems**

DL-FIND geometry optimisations can be carried out in Cartesian coordinates, mass-weighted Cartesian

coordinates, redundant internal coordinates, and hybrid delocalised internal coordinates (HDLC). HDLC coordinates are less coupled than Cartesians but still scale linearly with system size. Internal constraints (bond lengths, angles, and dihedrals) can also be applied in internal coordinates.

The coordinate transformation can be extended to a combination of multiple images of the system, which is required for transition state optimisations using the nudged-elastic band and dimer methods.

### Nudged-Elastic Band Method

NEB connects two energy minima with a chain of images of the system connected by springs. The spring forces are modified so that when converged, the images are aligned along the minimum-energy path between the minima. One image can be defined as the climbing image: it does not experience spring forces, but its energy is maximised along the NEB path while minimised in all other directions. Thus, it converges to a transition state. NEB with a climbing image is implemented in DL-FIND. It can be combined with any of the optimisation algorithms mentioned below. Particularly promising is the combination with the L-BFGS algorithm, which leads to super-linear convergence.



Figure 2: Energy and geometries of a nudged-elastic band path of a simple chemical system.

### Dimer Method

The dimer method can be used to find transition states without calculating the Hessian (the second derivatives of the energy). It considers two images of the system separated by a constant distance in configuration space (i.e. by a constant root mean square distance). The problem of finding a transition state is converted into two minimisation problems. First, the dimer rotation is optimised: the dimer is rotated around its midpoint until (1) the sum of the energies of the dimer endpoints is minimised, or (2) the curvature of the energy surface along the dimer axis is minimised, or (3) the dimer is aligned along the eigenvector of the Hessian with the lowest eigenvalue (the softest vibrational mode), or (4) the rotational force – the projection of the forces on the dimer endpoints onto the dimer normal – vanishes. All four criteria are equivalent. In practice, the rotational force is minimised.

With the dimer aligned along the softest mode, the transition state can be found by minimising the energy perpendicular to the dimer axis and maximising the energy along the dimer axis, similar to the climbing image in NEB. Dimer rotation and translation are optimised iteratively. However, in practice only a few rotational iterations are necessary after the initial optimisation of the dimer direction.



Figure 3: Trajectories of two transition-state searches using the dimer method. The dimer midpoint converges to the transition state (blue sphere).

### **Optimisation Algorithms**

Once the coordinate transformation is defined (including the definition of the NEB and dimer systems if applicable), the optimiser can be used with any type of system. Atoms, internal degrees of freedom, or simply 2-dimensional functions (as in Figures 1 and 3) can all be optimised. The standard minimisation algorithms implemented in DL-FIND include steepest descent, conjugate gradient, Newton-Raphson, Broyden-Fletcher-Goldfarb-Shanno (BFGS), and the limited-memory version of BFGS (L-BFGS). The L-BFGS method is particularly useful for minimisation problems, as it is a second-order optimiser but scales linearly in terms of CPU time and memory requirement. To search for a transition state by traditional algorithms (i.e. other than the NEB or dimer methods), we have implemented the partitioned-rational function optimisation (P-RFO) algorithm. It turned out to be superior to the dimer method for the most part only in cases where an initial analytic Hessian is available.

Conical intersections may be located and minimised by a penalty function method, a gradient-projection method, and a Lagrange-Newton method.

DL-FIND is currently included in GAMESS-UK and ChemShell. To encourage external contributions, the source code is available under an open source license at http://ccpforge.cse.rl.ac.uk/projects/dl-find.

### Outlook

DL-FIND was designed and written with the philosophy that efficient optimisations can best be achieved by a flexible code that makes the implementation of new algorithms easy. A highly specialised code may run faster, but at the cost of making it more difficult to include new functionality. DL-FIND will enable us to locate minima, transition states and conical intersections more efficiently and thus in larger systems than previously envisaged.



Figure 4: The dimer method applied to a biological system (the enzyme PHBH) in HDLC coordinates. The resulting transition mode is indicated by the transparent spheres in the insert.

J. Kästner P. Sherwood, CSED, Daresbury Laboratory; T. Keal, M.P.I. Mülheim, Germany;

# Digital Microfluidics Using Electrowetting-on-Dielectric Actuation

The phenomenon of electrocapillarity was first described by Gabriel Lippmann [1] in 1875 when he showed that the capillary depression of mercury in contact with an electrolyte solution could be controlled by the application of a voltage. As well as formulating the basic theory for electrocapillarity, he also proposed several practical applications that made use of the phenomenon including a very sensitive electrometer. Unfortunately, the application of the technique to aqueous electrolytes is severely restricted due to electrolysis.



Figure 1: Sequence of images showing EWOD-induced droplet splitting. The initial droplet has an approximate volume of 0.4  $\mu$ l.

In the early 1990s, however, Berge [2] realised that the problem of electrolysis could be overcome by the introduction of a thin insulating layer to separate the conductive liquid from the metal electrode. This concept has subsequently become known as electrowetting on dielectric (EWOD) and involves applying a voltage to modify the wetting behaviour of a liquid in contact with a hydrophobic, insulated electrode [3].

One particularly promising application area for electrowetting is the manipulation of individual droplets in digital microfluidic systems [4, 5]. The droplets are usually sandwiched between two parallel plates with the bottom plate housing an array of individually addressable buried electrodes and the upper plate fabricated from conductive glass. By modifying the voltages applied to the electrodes, it is possible to manipulate the droplets by setting up surface tension gradients along the plate. The advantage of this approach is that each droplet can be individually transported, split, merged and stored simply by the application of a suitable sequence of voltages to the control electrodes.

The change in contact angle over the buried electrodes can be evaluated using the Lippmann-Young equation:

$$\cos\theta = \cos\theta_{\rm Y} + \frac{\varepsilon_0\varepsilon_{\rm d}}{2{\rm d}\sigma_{\rm ly}}V^2$$

where  $\theta$  is the contact angle over the active electrode,  $\theta_{Y}$ is the equilibrium (or Young's) contact angle in the absence of an applied voltage,  $\varepsilon_{0}$  is the permittivity of free space,  $\varepsilon_{d}$ is the dielectric constant of the insulating layer, d is the thickness of the insulating layer,  $\sigma_{IV}$  is the surface tension between the liquid and the vapour, and V is the voltage applied to the electrode. The parabolic variation of  $\cos\theta$  in the Lippmann-Young equation is only applicable when the applied voltage lies below a threshold value. If the voltage is increased above this threshold, then *contact angle saturation starts* to occur and  $\cos\theta$  eventually becomes independent of the applied voltage [3, 4].



Figure 2: Simulation of droplet creation from a large 6 mm diameter reservoir droplet. The liquid "bridge" just prior to the detachment of the droplet has been confirmed experimentally.



*Figure 3: Simulation of droplet transport between two adjacent electrodes.* 

A collaboration between the Testable Design and Testing of Integrated Systems Group at the University of Twente and the Computational Engineering Group is currently investigating the factors affecting EWOD droplet manipulation. The study is being supported through an EU Framework 6 Network of Excellence entitled "DfMM PATENT: Design for Micro & Nano-Manufacture" and the ultimate aim of the project is to develop a working prototype for a microfluidic peptide synthesiser based on EWOD droplet manipulation.

Understanding the time-dependent flow of the droplets under the action of the electrowetting force is a significant modelling challenge since it is necessary to account for the complex interfacial forces and the deformable boundary of the droplet. In the current study, the motion of the interface has been modelled using the volume-of-fluid method in conjunction with a finite-volume Navier-Stokes solver. The numerical simulations were performed using the commercial computational fluid dynamics code, CFD-ACE+ [6]. The four fundamental unit operations, namely, droplet creation, transportation, splitting and merging have each been modelled.

The initial results show that the volume-of-fluid method can accurately capture the droplet handling operations typically used in digital microfluidic systems. The ability to provide reliable numerical simulations of EWOD-actuated droplet motion will clearly benefit the development of the next generation of digital microfluidic devices.

> D. R. Emerson, R. W. Barber, CSED, Daresbury Laboratory; H. G. Kerkhoff, X. Zhang, University of Twente, The Netherlands.



Figure 4: Sequence of images showing a droplet splitting simulation.



Figure 5: Computed velocity field during the merging of two droplets.

#### References:

[1] G. Lippmann, "Relations entre les phénomènes électriques et capillaires", Ann. Chim. Phys. 5, 494-549, 1875.

[2] B. Berge, "Electrocapillarity and wetting of insulator films by water", C. R. Acad. Sci. Paris, Sèrie II, 317, 157-163, 1993.

[3] F. Mugele and J.-C. Baret, "Electrowetting: from basics to applications", J. Phys.: Condensed Matter 17, R705-R774, 2005.

[4] S. K. Cho, H. Moon, and C.-J. Kim, "Creating, transporting, cutting, and merging liquid droplets by electrowetting-based actuation for digital microfluidic circuits", JMEMS 12, 70-80, 2003.

[5] R. B. Fair, "Digital microfluidics: is a true lab-on-a-chip possible?", Microfluidics and Nanofluidics, 3, 245-281, 2007.

[6] CFD-ACE+ User Manual: Version 2007 (ESI-CFD, Huntsville, USA).

## Possible Mechanisms for the Catalytic Activity of Aluminium Chloro-Fluorides: Investigation of HCI and HF interactions with β Aluminium Fluoride

The catalysis of chlorine and fluorine exchange reactions involving chlorofluorcarbon (CFC) and hydrofluorocarbon (HFC) compounds is an area of scientific importance. CFCs cause depletion of the ozone layer and they also have a high global warming potential. This has lead to the phasing out of these compounds, enforced by the Montreal protocol. HFCs are used as replacements for CFCs. These compounds do not cause depletion of the ozone layer and



Figure 1: Full monolayer coverage of HF on the  $\beta$ - AlF3 (100) surface.

also have a much lower global warming potential. HFCs are synthesised via halogen exchange reactions of CFCs and HF. Recently there has been considerable experimental interest in these reactions. Materials displaying strong Lewis acidity, that is materials that have a high electron affinity, are needed as catalysts. One of the commonly used Lewis acidic materials is aluminium chloride (AlCl<sub>3</sub>). This is used in the chemical industry as a homogenous catalyst for Friedal-Crafts reactions and polymerisation and isomerisation reactions of hydrocarbons. However, its performance as a heterogeneous catalyst for CFC halogen exchange reactions is poor. The catalyst requires a long exposure to CFCs before it becomes activated. The active species has been shown to be a solid phase based on AlCl<sub>3-x</sub>F<sub>x</sub><sup>1</sup>.

This led Dupont to investigate and patent the solid catalyst aluminium chlorofluoride (ACF). It is formed by fluorination of  $AlCl_3$ . ACF shows extraordinarily strong Lewis acidic properties and reaches the reactivity of the

homogeneous catalyst SbF<sub>5</sub>, one of the strongest known Lewis acids. The use of ACF in fluoro-organic chemistry has been extensively tested. One of the major failings of this catalyst is that it must be kept in a dry environment. Contact with water results in an irreversible phase transition to a catalytically inert substance. Recently very high surface (HS) area aluminium fluoride has been synthesised that shows Lewis acidity and reactivity comparable to ACF. This material also becomes inactive upon exposure to water. However, unlike ACF the interaction is reversible. This makes it a promising alternative to ACF in the catalysis of many Cl/F exchange reactions. It is thought that the structure of HS-AlF<sub>3</sub> is based on the  $\beta$  phase of crystalline AlF<sub>3</sub>. This phase, although less catalytically active than the amorphous HS-AlF<sub>3</sub>, still catalyses many Cl/F exchange reactions.

Despite extensive research into the development of these new catalytic materials little is known about their atomic scale chemical structure or the mechanisms by which they catalyse reactions which is essential if we are to tune their properties. In an attempt to achieve this we have previously studied the structure of the  $\beta$ -AlF<sub>3</sub> (100) surface<sup>2</sup>. In the current study we investigate the interaction of the surface with molecules. HF and HCl are often present in many catalytic reactions. Furthermore, it is likely that larger more complex CFCs will interact with the surface in ways analogous to HF and HCl. We have used state of the art hybrid density functional theory to investigate the interaction of HF and HCl with  $\beta$ -AlF<sub>3</sub>. The CRYSTAL code, developed by our group and with collaborators at the University of Turin, was used to perform these calculations.

We have previously shown that there are two low energy (100) surfaces that are likely to exist on crystalline  $\beta$ -AlF<sub>3</sub>. We have calculated the adsorption of HF and HCl at half and full monolayer coverage to both of these surfaces, although we will only show results from one surface here. A full description of this study can be found elsewhere<sup>3</sup>. Figure 1 shows the adsorption of HF to the T1 surface at full monolayer coverage. HF molecules adsorb above under coordinated Al ions via their F atom, and bind via their H atom with a surface F<sup>-</sup> ion to form FHF<sup>-</sup>. At this coverage the Fs from the HF and surface are indistinguishable. At half monolayer coverage HF does not chemically bind to nearby surface F ions but instead forms

a hydrogen bond. The binding energy of the HF to the surface is 1.1 eV irrespective of coverage.

The adsorption of HCl results in many possible adsorption geometries, all of which have very similar binding energies of  $\approx 0.65$ eV. Hence it is likely that they will all co-exist on the surface of real samples. At low coverage the HCl can adsorb to an under coordinated Al ion and form a hydrogen bond to a neighbouring surface F ion or it can dissociate with the H ion being transferred to a nearby F ion forming an HF molecule and leaving a Cl<sup>-</sup> as shown in figure 2. At full monolayer coverage these same processes occur. This results in the formation of FHF<sup>-</sup>, FHCl<sup>-</sup> and ClHCl<sup>-</sup> species as shown in figure 3.



Figure 2: Half monolayer coverage of HCl on the  $\beta$ - AlF3 (100) surface.

### Acknowledgements:

We thank the EU for support of this work through the 6th Framework Programme. (FUNFLUOS, Contract No. NMP3-CT-2004-5005575). Calculations were performed on STFC's SCARF and NW-Grid systems.

C.L. Bailey, A. Wander, B.G. Searle, N.M. Harrison, CSED, Daresbury Laboratory; S. Mukhopadhyay, Imperial College, London. It is predicted that this rich variety of structures for HCl adsorption may be correlated to the high catalytic activity of AlF<sub>3</sub>. CFC molecules could adsorb to these surfaces in analogous ways to the adsorption of HCl enabling chlorine and fluorine exchange reactions. In particular the formation of exposed Cl<sup>-</sup> ions at the surface may result in a similar surface structure to aluminium chlorofluoride (ACF). This type of process may activate the surface, allowing the catalysis of many desirable reactions. In the future we plan to study reaction mechanisms and calculate the pathways and kinetic barriers of a Cl/F exchange reaction on AlF<sub>3</sub>. Longer term it is hoped that understanding the relationship between the structure of the catalyst and its reactivity, will enable the design of efficient, selective and robust catalysts.



Figure 3: Full monolayer coverage of HCl on the  $\beta$ -AlF3 (100) surface.

### **References:**

[1]. J.K. Murphy, U. Gross, S. Rudiger, V.V. Rao, V.V. Kumar, A.Wander, C.L. Bailey, N.M. Harrison and E. Kemnitz, J Phys Chem. B. 110 (2006) 8314

[2]. A. Wander, B.G. Searle, C.L. Bailey and N.M. Harrison, J. Phys. Chem. B. 109 (2005) 22935

[3]. C.L. Bailey, S. Mukhopadhyay, A. Wander and N.M. Harrison, submitted to J. Phys. Chem. B

# High Efficient Intermediate Band Solar Cells Based on Quantum Dot Arrays

The efficiency of a single junction solar cell can be exceeded by splitting the solar spectrum in such a way that each *pn* junction only converts a narrow spectral region. Theoretically, tandem or multiple junction solar cells with an infinite number of *pn* junctions can reach the thermodynamic efficiency limits of solar energy conversion. While tandem solar cells can theoretically exceed 50% efficiency, tandems with large numbers of junctions face increasing complexity and materials issues (for instance the accumulated strain between different layers with different energy gaps), in conjunction with diminishing efficiency improvements. Therefore, significant attention has been paid to developing new approaches in which a single solar cell exceeds the efficiency of a conventional *pn* junction. One of the most promising and



Figure 1: Schematic of the QD array.

challenging methods for achieving high power efficiency (>60%) solar cells is to utilize periodic 3D arrays of the semiconductor quantum dots (QD), figure 1. This is the practical implementation of the so called "Intermediate Band Solar Cell" (IBSC) scheme proposed by Luque and Marti [1]. The IBSC has a predicted power efficiency of 63% under ideal conditions, exceeding the limiting single gap solar cell efficiency of ~41%, overcoming the problem of increasing the solar cell photocurrent without degrading its voltage. The IBSC principle of operation relies on the existence of a material characterized by an electronic band (intermediate band, IB) located within the semiconductor bandgap, figure 2. The introduction of the IB can be provided by periodicity of the QD's nanostructures that produces a band structure in the Brillouin zone of the QD super-lattice (SL). The IB allows two photons with energies below the barrier material band gap to be absorbed, to produce one electron-hole pair and ideally the IB also prevents non-radiative recombination mechanisms.

The theoretical model of the QD array electronic structure is based on the 8-band **k**·**p** method including band mixing between states in the conduction band (CB) and the heavy

hole, light hole and spin-orbit states in the valence band (VB), strain and piezoelectric field. The Hamiltonian is expanded in a plane-wave (PW) basis set [2] and the calculation is parallelized by using ScaLAPACK and MPI libraries. The PW based kp method with periodic boundary conditions is particularly suited for analysis of the QD array structures. In this approach the wave functions are expressed as a linear combination of the coefficients in the PW expansion. To calculate the electronic structure of an QD array the only modification to the basis set is to replace the PW-vectors in the wave function expansion with those arising due to periodicity of the QD-SL, i.e.,  $k_v \rightarrow k_v + K_v^{SL}$ . The  $K_v^{SL}$  is expressed in terms of the period,  $(\pi / L_v^{SL})$ , of the SL in particular direction . This allows the sampling along the K points of a QD-SL to be done at several points at the cost of the single QD calculation at each point.



Figure 2. Schematic of the IBSC band diagram.

In Figure 3. the variation of the first three mini-bands with the period  $L_z$ , is shown. The lower and upper boundary corresponds to  $K_z = 0$  and  $K_z = \pi / L_z$  respectively. The width of the e0 mini-band varies from ~86 meV at  $L_z = 10$ nm and almost vanishes by  $L_z = 10$  nm. The trend that e0, e1, and e2 are still rising in energy even when the quantum coupling between the dots vanishes, is attributed to the slow decay of the strain influence from QD in surrounding layers [3]. There are two contradictory requirements that the mini-band should fulfil for successful operation of an highly efficient IBSC: (a) the IB should exhibit finite energy width so that it can be partially occupied and facilitate simultaneous excitaion from IB to CB and VB to IB, and (b) the IB should be as narrow as possible to reduce carrier transport through the mini-band. To understand the optical characteristics of the IB better, and to assess the extent to which these requirements are met we focus our attention on the variation of the optical dipole matrix element inside the IB.

In Figure 4 the variation of the optical dipole matrix element between the e0 state and five topmost states in the VB is shown for  $L_z = 4$  nm. At that distance the miniband energy width is 14 meV. Each state in the VB is labelled according to its dominant character (heavy or light hole). The first two states in the VB are of clear HH character followed by a LH state. It is possible to identify clear band anticrossing between HH1 and HH2 at  $K_z \approx 0.75\pi / L_z$ . Another remarkable similarity with compressively strained quantum well (QW) structures is that the the e0-LH1 matrix element decrease as the  $K_z$ 



Figure 3. Variation of the first three mini-bands width in the CB with the vertical spacing of the QDs. The QD dimension in array is given in the main text.

vector increases. Based on reasonable uniformity of dipole matrix element for e0-h0 along  $\hat{\mathbf{e}}_{X}$  polarization, it seems that the radiative lifetime should also vary slowly across the mini-band. This statement needs further confirmation by detailed calculations of the dipole matrix elements along other light polarizations  $\hat{\mathbf{e}}_{y}$  and  $\hat{\mathbf{e}}_{z}$  as well as taking into account the variation of the transition energies inside the IB.

Our detailed calculations reveal the dispersion of the states inside the IB and the complex structure of the allowed and forbidden transitions between mini-bands and their absorption strengths. The model presented here will serve to provide accurate information on dipole matrix elements, energy transitions, and radiative life times for the description of the carrier dynamic processes in realistic IB solar cell structures.



*Figure 4. Variation of the different dipole matrix element inside IB mini-band.* 

S. Tomic and N. M. Harrison, CSED, Daresbury Laboratory.

References:

[1] A. Luque and A. Marti, Phys. Rev. Lett. 78, 5014 (1997).

[2] S. Tomic, A.G. Sunderland, and I.J. Bush, J. Mat. Chem. 16, 1963 (2006).

[3] S. Tomic, P. Howe, N.M. Harrison, and T.S. Jones, J. Appl. Phys. 99, 093522 (2006).

### Electronic and Magnetic Properties of Graphitic Ribbons

Recently, room temperature ferromagnetism has been reported for samples of highly oriented pyrolytic graphite irradiated with high energy protons. These observations have fuelled interest in the magnetic properties of carbononly materials, which are of both great technological and fundamental importance. The experimental evidence gathered thus far indicates that the magnetism does not originate from d or f electrons provided by chemical impurities but states of s-p symmetry unpaired at structural defects. The interest in s-p magnetism predates these recent discoveries, and so does the interest in the technological exploitation of nanosized particles of graphite, or nanographites. Their tunable and controllable properties open the way to implementing electromagnetic devices that are simultaneously high-tech, lowcost and easy-to-process.



Figure 1. A Zig-zag graphene ribbon.

Graphitic ribbons, created by cutting a graphene sheet along two parallel lines, are particularly interesting as the presence of zig-zag edges disrupts the diamagnetic delocalisation of the  $\pi$ -electron system, thus creating an instability that manifests itself as a sharp peak in the electronic density of states at the Fermi level. The work summarised here addresses three possible stabilisation mechanisms: geometric distortion, electronic spin polarisation and electronic charge polarisation (see Ref.1 for the discussion of the latter). The CRYSTAL code, developed by our group and with collaborators at the University of Turin, was used to perform these first principles calculations.

A graphene ribbon of the type shown in Figure 1 is obtained by cutting a graphene sheet along two parallel zig-zag lines. The ribbon is periodic in the x direction only, and the unit cell is delimited by dashed lines. The width of the ribbon along the non periodic dimension y, is defined by the number, N, of trans-polyacetylene-like rows of carbon atoms that run along x. We have computed the structural, electronic and magnetic properties of this system as a function of the ribbon's width, N.

The geometric distortion is well known to be the stabilisation mechanism for the N=1 ribbons, i.e., trans-polyacetylene. In this system a Peierls' distortion occurs and the chain of C atoms displays alternating long and short C-C bond distances. This mechanism is hindered in ribbons with N > 1 because the C atoms nearest neighbours of each edge C atom are involved in very strong and stable bonds with the inner C atoms. In order for the edge states to dimerise, the bulk bonding must be disturbed causing an energy loss which competes with the energy gained from dimerisation. The energy cost of dimerisation increases with ribbon thickness and therefore one expects dimerisation to be suppressed at some critical ribbon width. Our fully relaxed, first principles calculations, predict that the bond alternation length along the polyacetylene-like chain at the edge is 0.027 Å for N=2 and smaller than 0.001 Å for N=3. The sharp decrease of the magnitude of the dimerisation with N is a measure of the rigidity of the underlying  $\Sigma \pi$  bonding structure.



Figure 2. Isovalue surfaces of the spin density for the antiferromagnetic (a) and the ferromagnetic (b) case. The blue and red surfaces represent spin up and spin down density, respectively.

By allowing the system to be spin-polarised, the eventuality of stabilisation through spin polarisation is investigated. Stable magnetic states are found, whose spin densities are shown in Figure 2(a) and Figure 2(b). Here, the state depicted in Figure 2(a) is referred to as antiferromagnetic (AF) because the spin moments on the C atoms on one edge are anti-aligned to the spin moments on the opposite edge. Figure 2(b) shows the ferromagnetic (FM) configuration, where the spin moments on both edges point in the same direction. Both the AF and FM configurations are found to have a total energy lower than the non-magnetic state for all ribbon widths, indicating that spin polarisation is a possible stabilisation mechanism. The ground state of the system is found to be the AF case, as predicted by the spin alternation rule.

As the AF and FM states differ by flipping the spin moments of one edge, the range of the magnetic interaction in the ribbons can be measured by computing how the energy difference between these two states varies as a function of the ribbon's width. A width of N~8, corresponding to approximately 14-16 Å, still gives an energy difference of ~25 meV and thus room temperature magnetic ordering. This explicit demonstration of the long ranged nature of the magnetic interaction mediated by the delocalised  $\pi$  orbitals in graphitic systems is consistent with recent observations of magnetism induced by low defect concentrations in a number of materials.

The analysis of the electronic structure and energetics of these two stable magnetic states as a function of the ribbon's width documents the range of the magnetic interaction between the edges and the nature of the stabilisation mechanism via spin polarisation. In Figure 3 and Figure 4 the band structures of the AF (upper panel) and FM (lower panel) states are presented, respectively. The instability of the non magnetic ribbons is resolved into two completely different but well-known scenarios: the AF solution gives a Slater insulator, whilst the FM solution gives a Stoner metal.

From a nanotechnological point of view, these results have direct implications for the control of the spin dependent conductance in graphitic nano-ribbons. In fact, the prediction that at the nanometer level an AF allignment of the edges' spins produces a strongly insulating state and that a FM alignment is conducting, implies that one is able to control the conducting properties of a nano-sample by applying appropriately modulated magnetic fields.



Figure 3. Band structure of N=10 ribbon in the antiferromagnetic configuration (ground state)



*Figure 4. Band structure of N=10 ribbon in the ferromagnetic configuration.* 

### **References:**

1. L. Pisani, J. A. Chan, B. Montanari, and N. M. Harrison, Phys. Rev. B 75, 064418 (2007).

B. Montanari, N.M. Harrison, CSED, Daresbury Laboratory; L. Pisani, J. A. Chan, Imperial College, London.

### Acknowledgements:

We thank the EU for support of this work through the 6th Framework Programme. (FERROCARBON, CEC 012881). Calculations were performed on STFC's SCARF system and HPCX.

# The Challenge of Larger and Larger Linear Systems

The solution of large sparse linear systems of equations is a cornerstone of computational science and engineering. In many industrial applications, such as the analysis of very large structures, industrial processing of complex non-Newtonian fluids, and the simulation of car bodies, the solution of large, sparse linear systems of equations is the single most computationally expensive step. As timedependent three-dimensional simulations are now commonplace, there is a need to develop algorithms and software that can be used to efficiently solve linear systems with several millions of unknowns. These systems take the form AX = B, where the right-hand sides B are known and X must be computed.



Model for problem bone010

Since the 1960s, significant effort has been expended on the development of linear solvers; the Numerical Analysis Group at the Rutherford Appleton Laboratory has been at the forefront of this work. In broad terms, linear solvers can be divided into two main classes: iterative solvers and direct solvers. Iterative methods attempt to solve the linear system by finding successive approximations to the solution starting from an initial guess, while a direct method factorises the matrix into a product LU of a lower triangular matrix and an upper triangular matrix, and then completes the solution process by solving two simple triangular systems (LY = B followed by UX = Y).

Iterative solvers are attractive for large-scale problems because they require little memory (typically, a small number of vectors of size equal to the number of unknowns). However, a preconditioner is generally required to speed up the convergence. For many of the tough systems that arise from practical applications, the difficulties involved in finding and computing a good preconditioner can make iterative methods infeasible. By contrast, direct solvers are much more robust for a wide range of problems but even when well implemented, the memory they require generally increases much more rapidly than the problem size so that memory can quickly become exhausted and the computation consequently fails. One possible way to circumvent this limitation is to develop direct solvers that are able to hold the data structures on disk, that is, out-of-core solvers. This is the approach we have recently adopted in the development of direct solvers for the HSL Library.

We have designed and developed two new out-of-core solvers: HSL\_MA77 and HSL\_MA78. HSL\_MA77 is for symmetric systems; it allows the system matrix A to be supplied either row-by-row or as a sum of element matrices (the latter is convenient for finite element applications since assembly and storage of the whole system matrix A is avoided). HSL\_MA78 is for unsymmetric element problems. Both codes are written in Fortran 95 and implement a multifrontal algorithm; they have many features in common and employ a similar reverse communication user interface, which we describe fully in the accompanying user documentation. The codes allow the matrix A and its factors to be held in files on disk, together with some of the intermediate workspace. Since disk access is much slower than access to main memory, for good performance it is essential that the writing to and reading from these files is done as efficiently as possible. A key design decision for the new solvers was to isolate all the input and output operations in a separate set of Fortran 95 subroutines. This set of subroutines which is available within HSL as the package HSL\_OP01, manages a virtual memory system that ensures actual input/output occurs only when really necessary.

Facilities are provided for reading from and writing to direct-access files using a user-supplied buffer (workarray). The buffer is divided into pages of length equal to the size of the file records. All actual input/output is performed by transfers of whole pages between the buffer and records of the file. The data in a file is addressed as a virtual array of rank one. Because it may be very large, long integers (64 bits) are used to address it. Any contiguous section of the virtual array may be read or written without regard to page boundaries. HSL\_OF01 does this by first looking for parts of the section that are in the buffer and performing a direct transfer for these. For any remaining parts, there has to be actual input and/or output of pages of the buffer. If room for a new page is needed in the buffer, the page that was least recently accessed is written to its file (if necessary) and is overwritten by the new page.

In the table we report on some numerical experiments which were run on a single 3.6 GHz Intel Xeon processor of a Dell Precision 670 with 4 Gbytes of RAM. We report the times (which are elapsed times in seconds) for the different phases of HSL\_MA77 for four problems taken from practical applications. During the initialisation phase, the user must input the problem data and compute an elimination order. The analyse phase analyses the sparsity of the matrix and sets up the data structures that are needed by the factorisation phase. Once the factors are computed, any number of calls may be made to the solution phase, which performs the triangular solves. In the table, MA77\_solve(k) indicates we solved for k right-hand sides simultaneously. Since the factor data must be read from disk once for the lower triangular solve and once for the upper triangular solve independently of the number of right-hand sides, it is clearly much more efficient to solve for multiple right-hand sides rather than repeatedly solving for a single right-hand side.

Phase	inline_1 (n=503,712)	bones10 (n=914,898)	nd24k (n=72,000)	bone010 (n=986,703)
Initialisation	19.1	29.0	19.3	42.7
MA77_analyse	4.20	6.70	22.1	26.7
MA77_factor	90.6	174.6	1284	1491
MA77_solve(1)	5.30	36.0	10.4	311
MA77_solve(8)	10.6	41.3	20.7	331
MA77_solve(64)	60.5	141.0	90.2	499

*Timings for the different phases of the HSL\_MA77 out-of-core solver (n is the number of unknowns)* 

### HSL

HSL (formerly the Harwell Subroutine Library) is a collection of fully supported, tested, and documented ISO Fortran codes for large scale scientific computation, written by members of the Numerical Analysis Group and other experts. A new release of the Library is made available every 2-3 years. Our new out-of-core solvers together with HSL\_OF01 are part of the latest release HSL 2007. All HSL packages are available to UK academics, without charge, for academic purposes. Further details of the contents of HSL and how to obtain HSL packages are available at www.cse.scitech.ac.uk/nag/hsl/hsl.shtml

References:

Further details of the new solvers and the virtual memory management system together with more results may be found in following technical reports, which are available at www.numerical.rl.ac.uk/reports.

[1] An out-of-core sparse Cholesky solver. RAL-TR-2006-013.

[2] An efficient out-of-core multifrontal solver for large-scale element problems. RAL-TR-2007-014.

[3] HSL\_OF01, a virtual memory system in Fortran. RAL-TR-2006-026.

J.K. Reid and J.A. Scott, CSED, Rutherford Appleton Laboratory.

# EURACE - European Agent-Based Computational Economics Using FLAME

EURACE is a three-year EU Project funded under the Framework 6 Complex Systems Programme. EURACE is seeking to develop an economic simulator based on co-operating intelligent agents.

Macroeconomic policy design plays a fundamental role in social welfare and requires a coordinated application of economic policy measures: e.g. fiscal and monetary strategies, knowledge exchange, R&D incentives etc. Generally speaking, the interplay of such different measures is not completely understood and macroeconomic design follows a classical approach. Conversely, there is considerable interest in the development of an alternative paradigm to the rational representative agent model.

The EURACE project tackles this complex problem and proposes an innovative approach to macroeconomic modelling and economic policy design within the agentbased computational economics framework. The project objectives are characterized by scientific, technological and societal scopes. From the scientific point of view, the main effort regards the study and the development of multi-agent models that reproduce, at the aggregate economic level, the emergence of global features as a self-organized process from the complex pattern of interactions among heterogeneous individuals.

From the technological point of view, the project will develop, with advanced software engineering techniques, a software platform in order to realize a powerful environment for large-scale agent-based economic simulations.

Key issues will be the definition of formal languages for modelling and for optimizing code generation, the development of scalable computational simulation tools and the standardisation of data with easy to use humanmachine interfaces.

FLAME (Flexible Agent Modelling Environment) is the basis of the X-Machine simulator. The FLAME framework has been developed at Sheffield University and was originally



Graphical representation of the C@S model showing the agent-to-agent communications (left) and the average wage/price trend graph (right) after 100 iterations of the model



The agent state dependency graph for the C@S model

C. Greenough, D.J. Worth and L.S. Chin, CSED, Rutherford Appleton Laboratory; M. Holcombe, S. Coakley, Computer Science, Sheffield University. targeted toward biological systems and in particular cellular biology. It comprises three main components: the X-Machine parser, a set of template files and the agent models described in XMML.

XMML(X-Machine Markup Language) is an X-Machine dialect of XML. As with XML it has predefined tags to define elements of the model such as an agent's structure – its internal memory, its states - and the messages that can be posted to and read from message boards.

The C@S model was the first economic model to be simulated using FLAME. In this model there are three agent types representing firms (*Firm*), people (*Person*) and shopping malls (*Mall*). The model simulates the processes and exchanges involved in employment and goods purchasing by the population and the firms producing the goods. The XMML defines the actions and interactions between these and FLAME uses a dependency graph generated from the XMML to sequence the agent state changes.

Because of the possible large size of the agent populations the applications generated by FLAME have been designed to run on parallel computing systems using essentially domain decomposition.

SEG has used its experience in developing large-scale parallel programs to improve the original parallel implementation of FLAME. Using this new implementation parallel simulations involving 5 million agents distributed over 2000 processors are being successfully performed on a variety of parallel systems such as HPCx and Blue Gene. Although the communications bandwidth between each agent group is high, by a careful distribution of the agents over the processors and the use of dynamic load balancing it will be possible to run simulations with a reasonable efficiency.

### **References:**

T. Balanescu, A.J. Cowling, M. Georgescu, M. Holcombe and C. Vertan: Communicating stream X-machines are no more than X-Machines, Journal of Universal Computer Science, 5(9), 494-507, 1999.

S. Coakley, A Formal Software Architecture for Agent-Based Modelling in Biology, PhD Thesis, Department of Computer Science, University of Sheffield, 2007.

M. Catalano, F. Clementi, D. Delli Gatti, C. Di Guilmi, et al The C@S Project, EURACE Working Paper, Sept 22, 2006.

# Using Software Metrics to Improve Software Quality

A prerequisite for improving the quality of software is a quantitative measure of the quality of the source code. This is usually achieved through the definition of metrics, values for which can be calculated by analysing the source code. There are many simple metrics, for example:

- Number of statement lines in the source.
- Ratio of comment lines to statement lines.
- Number of statement lines in each procedure.

These can give a rough idea of how complex the source code might be but say nothing about the complexity of the code design and the way a particular function has been implemented. More complex metrics have been defined including:

- McCabe (Cyclometric) Complexity Calculated as the number of linearly independent paths through the code.
- Number of knots A measure of complexity since a knot is defined as a crossing in a path through the code.
- Halstead Metrics Metrics defined in terms of number of operators, number of operands, number of unique operators and number of unique operands. These include 'volume' of a program, effort required to produce the program, estimated number of bugs in the program and a complexity measure for the program.
- Cohesion Used in object oriented programming to measure how related and focused the responsibilities of a class are. High cohesion is good as it makes code easier to understand and increases the chance of code re-use.
- Coupling Low coupling between entities in a code is desirable as the impacts of any change can be minimised and the entities can be understood and tested in isolation.

Software engineers have defined acceptable values for many metrics where values that are 'out of range' indicate code that may be difficult to understand, test or maintain (and frequently all three). We have analysed the source of a Fortran based finite element library that has undergone three revisions from its first release and has used successive Fortran standards during development. We hope to see improvements as the library has been developed. The metrics were calculated by the nag\_metrics tool – part of the NAGWare suite available from NAG Ltd. Values of the following metrics were calculated for the 'worst' procedures over successive releases:

- McCabe complexity Measure of logical complexity for a procedure.
- Absolute complexity Measure of logical complexity relative to the number of lines of code in a procedure.
- Estimate of difficulty Measure of how difficult it will be to develop the procedure further.

In Figure 1 we see that the metric values are below the threshold value of 50 for McCabe complexity. We also see that they change little through the developments since the algorithms used in the procedures have changed in only small ways.

Figure 2 shows that absolute complexity has increased slightly over time. This may be a cause for concern if we had not also looked at the difficulty estimate.

Figure 3 shows the difficulty estimate decreasing as newer features of the language and additional code have simplified the procedures. This has caused the slight rise in absolute complexity but we can live with that if the code is easier to maintain.

The metrics tool from NAG Ltd works for Fortran 77 code, however new Fortran standards have arisen bringing with them modules and increasingly object oriented features. To address this, the Software Engineering Group is collaborating with NAG Ltd and the University of Kent to develop a new metrics tool that will be able to analyse Fortran 95 and 2003 code.

The analysis will be based on the parse tree and symbol table provided by the front end of the NAG compiler. Following advice from Kent we will choose a set of metrics suitable for Fortran 95 and 2003 code and implement the code that calculates the metric values on top of the analysis.

34



Figure 1: McCabe Cyclometric complexity for selected routines



Figure 2: Absolute complexity for selected routines



Figure 3: Measure of difficulty for selected routines

C. Greenough, D.J. Worth and L.S. Chin, CSED, Rutherford Appleton Laboratory;

T. Hopkins, Kent University;

*I. Hounam, Numerical Algorithm Group Ltd.*