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 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.

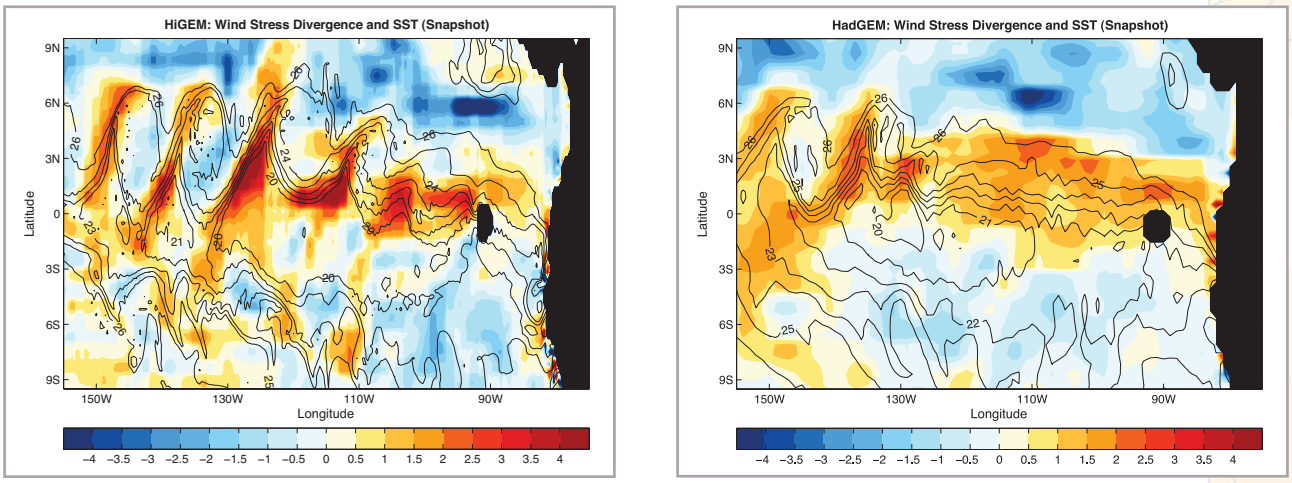

Figure 1: Left: Snapshot from the Pacific of HiGEM showing Tropical Instability Waves (TIWs). Colours are the windstress divergence, indicating that the atmosphere is coupling with the oceanic TIWs. Right: lower resolution HadGEM model, showing that TIWs are poorly resolved. (Figures courtesy L.Shaffrey, 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 large-scale 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, G. Riley, R. Ford, Manchester. P. H. Gaskell and H. Thompson, Leeds.