Simulations of Catalytic Chemistry with ChemShell 8x Faster after Extra Parallelism Implemented by HECToR dCSE Team

Richard Catlow, University College London (UCL)
Thomas Keal, STFC Daresbury Laboratory
HECToR CSE Team, Numerical Algorithms Group Ltd (NAG)

HPC experts from STFC, working under NAG's Computational Science and Engineering (CSE) support service for HECToR, the UK's national academic supercomputing facility, have added an extra layer of parallelism to the ChemShell package, enabling simulations to scale to thousands of cores and complete up to eight times faster than previously.

Describing science impact, the Principal Investigator Professor Catlow (UCL) said, *"In a number of projects, particularly within the themes of chemical reactivity, and nanomaterials and defects, the work of the Materials Chemistry Consortium depends crucially on the advanced computational chemistry tools provided by the code ChemShell. One important functionality of this software is the development and utilisation of complex computational models including hybrid quantum-mechanical / molecular-mechanical embedding. Members of the Consortium apply these models to study the structure and reactivity of heterogeneous catalysts, physical and chemical processes on the defects and active sites in the bulk and at the surface and interfaces of advanced functional materials. One of the ubiquitous and non-trivial tasks in this work is the CPU-intensive investigation of the potential energy hypersurfaces, or landscapes. The latest software development within the dCSE project allows us to run efficiently across in excess of 1,000 cores and there is also a reduction of this most time-consuming part of approximately 50 to 90% compared to the older implementations."*

HECToR

HECToR is managed by EPSRC on behalf of the participating Research Councils with a mission to support capability science and engineering in UK academia. The Cray XT supercomputers, located at the University of Edinburgh, are managed by UoE HPCx Ltd. The CSE Support Service is provided by NAG Ltd and ensures users have access to appropriate HPC expertise to effectively exploit advanced supercomputers for their science. A critical feature of the CSE Support Service is the distributed CSE (dCSE) programme which, through lightweight peer review, delivers dedicated performance and scalability projects on specific codes in response to proposals from users. The dCSE programme now consists of nearly 40 focused projects complementing the traditional HPC user applications support and training also provided by NAG.

The dCSE projects completed so far have delivered outstanding examples of the cost savings and new science that can be enabled through dedicated CSE effort. The ChemShell project reported here adds to these success stories with a successful performance improvement.

Project Background

The goal of this project was to enable ChemShell to scale to over 1000 cores for more efficient study of heterogeneous catalysis on HECToR. ChemShell integrates quantum mechanical (QM) and molecular mechanical (MM) software packages to perform combined QM/MM calculations. External programs are used for energy and gradient calculations while ChemShell performs higher level tasks such as geometry optimisation. Professor Richard Catlow of UCL was the PI. Thomas Keal of STFC undertook the 9 month project in collaboration with the NAG CSE team and users.

ChemShell

ChemShell provides a means of integrating quantum mechanical (QM) and molecular mechanical (MM) software packages to perform combined QM/MM calculations. External programs are used for energy and gradient calculations while ChemShell performs higher level tasks such as geometry optimisation. http://www.chemshell.org.

Project Results

The existing implementation of ChemShell exploits the parallelism of the external programs but does not run in parallel itself. An extra layer of parallelism was implemented at the ChemShell level using a task-farming approach, enabling multiple gradient evaluations to be carried out in parallel. This resulted in significantly better scaling than relying on the external programs alone. In a 1024-core test case, speed-up factors of up to 7x were observed for a finite-difference Hessian calculation using the modified code compared to the original (non-task farmed) implementation. The parallel implementation of DL-FIND geometry optimisation library was interfaced to the task-farmed version of ChemShell. In addition, the nudged elastic band (NEB) method for transition-state determination in DL-FIND was parallelised. A speed-up factor of over 8x was found for a QM/MM surface chemistry NEB test case and up to 6x for a stochastic search benchmark.

Commenting on the successful project, Professor Catlow said "*Some 6 million AUs (allocation units) of HECToR supercomputer resources are allocated to projects utilising ChemShell in the current six-month period. The technological development undertaken in this project will tremendously improve the efficiency of our current work using ChemShell and will allow us to expand the studies of complex systems and processes both on the current HECToR supercomputers, and on emerging HPC platforms. Thus we expect the time allocation to such projects to increase over the lifetime of HECToR as more ambitious projects are also now possible in the field of energy materials.*"

A full technical report can be found at http://www.hector.ac.uk/cse/distributedcse/reports/

For more information contact: HECToR CSE Team
The Numerical Algorithms Group Ltd, Wilkinson House, Jordan Hill Road, Oxford, OX2 8DR, United Kingdom
Telephone: 01865 511 245 Email: hector-cse@nag.co.uk Web: http://www.hector.ac.uk/cse/