Molecular Dynamics Code (DL_POLY_3) Enhanced by HECToR dCSE Team

Ilian Todorov, Bill Smith, Laurence Ellison, STFC Daresbury Laboratory
HECToR CSE Team, Numerical Algorithms Group Ltd (NAG)

HPC developers from STFC Daresbury, working under NAG's Computational Science and Engineering (CSE) support service for HECToR, the UK's national academic supercomputing facility, have enhanced the DL_POLY molecular dynamics package, increasing the number of processors over which it may be parallelized by two orders of magnitude. Other improvements have reduced the DL_POLY solution time for large biomolecules, enabling scientists to study bigger, more complex systems within their existing compute budgets.

The project was focussed on the incorporation of two enhancements: (i) a load balancing scheme, which would improve the efficiency of the code when modelling systems which exhibited spatially non-uniform atomic density, and (ii) rigid body dynamics, which would be a better treatment of the forces and equations of motion. Commenting on the dCSE project, Dr Ilian Todorov of STFC Daresbury Laboratory said: *"[This] was a useful project in terms of improving our understanding of the particular load balancer and appreciating its potential and power. However, the level of intricacy coupled with the limited effort available make this development bear more academic value than actual CSE value. It is just too expensive to develop and incorporate into already big code without any recurrent funding. The effort this dCSE project awarded to the [rigid body dynamics] algorithms development was the invaluable time for testing them extensively and learn more about coupling topological objects dynamically and handling their various degrees of freedom with the domain decomposition framework. Overall, this project taught us how grossly we had underestimated the effort needed for these developments."*

Significant progress has been made in adding two enhancements to the code, which will improve its efficiency for non-uniform systems and for large biomolecules.

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 XE6 supercomputer, located at the University of Edinburgh, is 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 over 60 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 DL_POLY project reported here adds to these success stories with the development of an improved version of the code.

Project Background

DL_POLY is a general purpose classical molecular dynamics (MD) package which is aimed at the modelling of large systems on multicore machines such as HECToR. DL_POLY_3 uses a static, equi-spatial domain decomposition scheme in its parallelization, and scales up to 512 cores for systems with around 1,000,000 atoms. However, such a scheme can lead to inefficiencies when modelling systems with a high degree of spatial non-uniformity because of the uneven distribution of work between cores. The main aim of this project was to address this problem by adding dynamic load balancing, which periodically checks for imbalances between core workloads and rectifies them by relocating atoms from busy cores those that are idler. The project also aimed to incorporate rigid body dynamics into DL_POLY_3, which treats the atomic forces and equations of motion in a more efficient fashion, thereby reducing the solution time for large biomolecules.

Ilian Todorov of STFC Daresbury Laboratory was the Principal Investigator for the project. He, along with Bill Smith and Lawrence Ellison (both of whom are also from STFC Daresbury) out the 18 person-month project, in close collaboration with the NAG CSE team.

Around 50,000 kAUs (thousands of allocation units) are used per year on HECToR to run calculations using DL_POLY.

Project Results

A prototype load balancing algorithm was devised which attempted to minimise the variations in atom density throughout the computational (not physical) domain. This was implemented as a standalone code, whose performance on HECToR was encouraging for tests up to 2744 cores for systems with 5,000,000 atoms. Incorporating load balancing gives good overall speedup, although it only leads to slightly better scaling with increasing number of cores in this test because of the communications overhead associated with distributing atoms across large numbers of idle cores. Further work is required before the code is incorporated into DL_POLY.

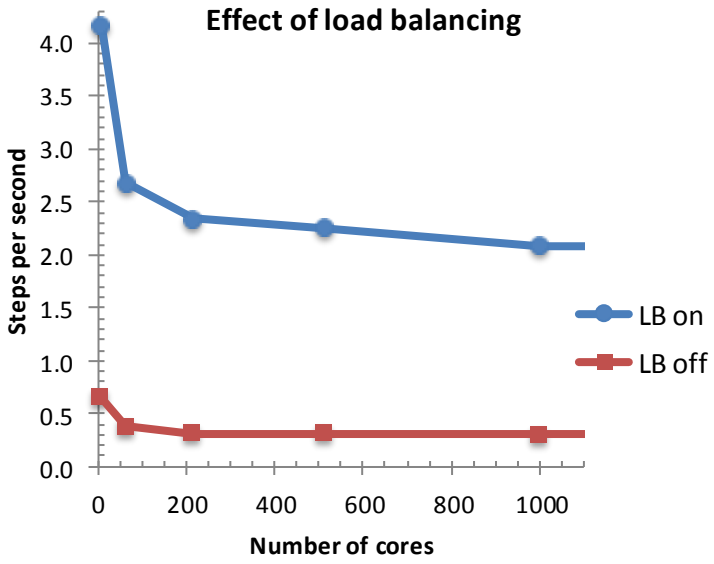

Turning on load balancing increases the speed of the simulation (measured in time steps per second) by a factor of between 6 and 8 (the latter value is the theoretical maximum for this test).

In addition, rigid body dynamics was incorporated into DL_POLY. Insights gained during this part of the project included an understanding of how to handle the molecules' various degrees of freedom within the framework of the domain decomposition.

A full technical report can be found at http://www.hector.ac.uk/cse/distributedcse/reports/
More details about DL_POLY can be found at http://www.stfc.ac.uk/CSE/randd/ccg/software/DL_POLY/25526.aspx

For more information, please 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/