# Scalability and Capability of Materials Science Code (CONQUEST) Enhanced by HECToR dCSE Team

*David Bowler, Lianheng Tong, University College London (UCL)*
*HECToR CSE Team, Numerical Algorithms Group Ltd (NAG)*

**An HPC expert from UCL, working under NAG's Computational Science and Engineering (CSE) support service for HECToR, the UK's national academic supercomputing facility, has successfully improved both scalability and functionality in the CONQUEST density functional theory (DFT) code. CONQUEST may now be used for scalable simulations on bio-molecular systems and magnetic systems, as well as metallic systems.**

CONQUEST is a linear scaling, or O(N), DFT electronic structure code developed jointly by NIMS (National Institute for Materials Science, Japan) and UCL. The code is designed to perform DFT calculations on very large systems (a system containing two millions atoms of silicon has been demonstrated). It can be run at different levels of precision, ranging from ab initio tight binding up to full DFT with plane wave accuracy. CONQUEST is currently used by academics at UCL and Imperial. It is also in use on the Jaguar supercomputer at Oak Ridge National Laboratory (ORNL) and is one of the main codes selected for optimisation on the Japanese 10PF 'K' supercomputer. Prior to this work, CONQUEST use was limited to systems with relatively simple bonding, where spin and van der Waals interactions were not important. As a result of this project, the code can now be used for weak bonding in bio-molecular systems and magnetic systems. Furthermore, as these new application areas require the use of a scalable code and large HPC resource, scalability was also improved. This was achieved by developing hybrid parallelisation within CONQUEST via the implementation of OpenMP with the existing MPI.

Commenting on the dCSE project success, *Dr David Bowler of the Department of Physics and Astronomy at UCL said: 'This dCSE work has served CONQUEST extremely well. Spin polarised calculations are essential for production calculations within density functional theory, particularly for current applications of CONQUEST (dopants in semiconductors). Similarly, recent developments in modelling van der Waals interactions within DFT (as implemented in CONQUEST) are an enormously important new area, particularly for biomolecules. The scaling improvements brought about through the OpenMP/MPI hybrid implementation restore the perfect weak scaling of CONQUEST, enabling efficient use of HECToR. The developments have fitted CONQUEST for the future: both for production calculations and for efficiency.'*

> The van der Waals implementation will enable the accurate modelling of biomolecules, which is a major achievement for the use of CONQUEST.

*'The spin and OpenMP/MPI implementations have enabled calculations on the structure of dopants in bulk silicon, which are important for possible quantum computing implementations. This on-going project has potentially major impact, and would not be possible without the developments.'*

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 70 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 CONQUEST project reported here adds to these success stories with a successful performance improvement.

## Project Background

The first objective of this dCSE project was to extend functionality of CONQUEST to bio-molecular systems and magnetic systems by adding the capability to perform spin polarised calculations and also van der Waals interactions. The second objective was to enable researchers in the electronic structures community to do accurate and efficient ab initio quantum mechanical simulations on modern HPC systems (with multicore processors) to a very large scale by developing CONQUEST to use OpenMP for the main bottlenecks. This objective was to restore the perfect weak scaling found on earlier HPC platforms.

David Bowler of the Department of the Department of Physics and Astronomy at UCL was the Principal Investigator for the project. Lianheng Tong also of UCL carried out the 20 person-month project, in close collaboration with the NAG CSE team.

## Project Results

The new functionality in CONQUEST was tested by performing calculations of spin polarisation on two systems: bulk silicon in a face-centered cubic lattice and bulk Iron, in a body-centered cubic lattice. Results were validated against known experimental values and calculations from the SIESTA code. Subsequently, the implementation has been used to calculate the electronic structure of dopants in large unit cells of bulk silicon (up to 262,144 atoms). Scalability of CONQUEST with the was demonstrated using 1,2 and 4 HECToR Phase 3 nodes – linear weak scalability was maintained for the spin polarisation part of the calculation.

To improve the weak scalability of CONQUEST for general use of the code on many-core HPC resources, hybrid parallelisation within CONQUEST was developed via the implementation of OpenMP. The matrix multiplication subroutines within the code were identified as the main bottleneck and a new MPI approach was developed, to facilitate use of OpenMP. Weak scaling was restored to CONQUEST on HECtoR, a test case of bulk Ice using a fixed work-load with 48 atoms per core demonstrated perfect weak scaling to 1000 cores.

The van der Waals implementation will enable accurate modelling of biomolecules, which is a major step to widen the applications for future CONQUEST modelling. The OpenMP / MPI implementation will benefit all future work with CONQUEST.

---

Full technical reports on this work can be found at http://www.hector.ac.uk/cse/distributedcse/reports/conquest02/ and http://www.hector.ac.uk/cse/distributedcse/reports/conquest03/

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/