Institution: Queen's University Belfast

Unit of Assessment: 9 (Physics)

Title of case study: Impact of the HELIUM code on UK Government procurement and provision of national high-performance computing facility HECToR.

1. Summary of the impact (indicative maximum 100 words)

The High Performance Computing (HPC) application code HELIUM, developed at Queen's University Belfast to assist the development of attosecond technology, has impacted on the provision of public services through guiding procurement and acceptance testing of the high-performance computer facility HECToR. This facility was funded by UK Government with a total expenditure of £113M during 2007 – 2013. The HELIUM code was used for procurement and acceptance testing for the initial HECToR service in 2007 (Phase 1, 11k cores), and its upgrades in 2009 (Phase 2a, 22k cores), 2010 (Phase 2b, 44k cores) and 2011 (Phase 3, 90k cores). The HELIUM code was particularly invaluable in demonstrating that the Phase 2b and Phase 3 systems perform correctly at pre-agreed performance levels, since this code can be adapted to run for several hours over >80k cores.

2. Underpinning research (indicative maximum 500 words)

Development of the HELIUM code started in 1993, when Cray Research Inc. provided Professor K.T. Taylor (Professor of Physics, QUB 1993-2013) with an award for a three-year postdoctoral fellowship to develop computational codes to answer prominent scientific questions. Following the purchase by the Research Councils in 1994 of a 256-processor Cray T3D machine, Prof. Taylor decided to investigate what happens to the dynamics of individual electrons in a strong laser field when their strong repulsion is taken into account.

This prominent question remains an important one to date for atomic physics and for laser physics. Electron repulsion cannot be neglected in dynamics on the sub-femtosecond timescale. Detailed understanding about the coupled electron dynamics requires simultaneous analysis of time-dependent dynamics of individual electrons and detailed treatment of the interactions between them. The scale of this problem exceeds traditional problems by several orders of magnitude: its numerical solution requires massively parallel computational facilities. Professor Taylor and Dr J.S. Parker (postdoctoral research fellow, QUB 1993-present) developed a finite-difference code which could efficiently exploit these facilities, including extensive computation and extensive data transfer: the HELIUM code. In 1996, a paper demonstrated successful initial application of the HELIUM code for a laser wavelength of 248 nm [1].

The application of finite-difference techniques enabled Dr. Parker to construct the code such that it could handle the projected exponential growth in HPC power. The atomic wavefunction is described on a multi-dimensional radial grid. This grid can be distributed efficiently across (tens of) thousands of processors. Dr. Parker and E.S. Smyth (Ph.D. student 1996-1999) further implemented an Arnoldi-Krylov algorithm which provided stable time propagation using only (parallelisable) matrix-vector multiplications [2]. The code therefore scales very well from hundreds to tens of thousands processors. This computational excellence of the HELIUM code was recognised through the award to Prof. Taylor, Dr. Parker and B.J.S. Doherty (Ph.D. student 2004-2007) in 2006 of a UK Research Councils' HEC Strategy Committee High Performance Computing Prize in the category of fastest application improvement, for their work on re-engineering the HELIUM code on HPCx.

The dedication to numerical techniques that allow efficient exploitation of future HPC facilities enabled Prof. Taylor, Dr. Parker and co-workers to explore new science from 1997 to 2006 following subsequent investment into HPC facilities (T3E, HPCx) using the HELIUM code. Between 1999 and 2003 many fundamental scientific discoveries were made, most notably in the general features of double ionization [3], the prediction of double above-threshold ionization [4] and in time delays in double ionization [5]. Subsequently, emphasis on comparison with experiment resulted in a 2006 publication on the energy distribution of ejected electrons, which showed almost unparalleled agreement between theory and experiment [6].

The techniques used within the HELIUM code continue to be exploited to this day. Prof. Taylor and Prof. H.W. van der Hart (Professor in Theoretical Physics, QUB 1999-present) combined their expertise to lead a 5-year EPSRC HPC software-development programme (UK-RAMP, involving University College London, the Open University and STFC Daresbury Laboratory), in which the algorithms underpinning the HELIUM code are combined with R-matrix techniques to explore ultra-fast multi-electron dynamics in atoms and molecules.

3. References to the research (indicative maximum of six references)

The papers which we identify to best illustrate the quality of the research are papers 2, 4 and 6.

[1] J. Parker, K.T. Taylor, C.W. Clark, and S. Blodgett-Ford:
'Intense-field multiphoton ionization of a two-electron atom',
Type: paper in peer-reviewed journal, cited 94 times (Web of Knowledge, 29-8-2013)
J. Phys. B.: At. Mol. Opt. Phys. **29**, L33 (1996).
DOI: 10.1088/0953-4075/29/2/002

[2] E. S. Smyth, J. S. Parker and K. T. Taylor,
'Numerical integration of the time-dependent Schrodinger equation for laser-driven helium',
Type: paper in peer-reviewed journal, cited 90 times (Web of Knowledge, 29-8-2013)
Comput. Phys. Commun., **114** 1 (1998).
DOI: 10.1016/S0010-4655(98)00083-6

[3] D. Dundas, K.T. Taylor, J.S. Parker and E.S. Smyth:
'Double-ionization dynamics of laser-driven helium',
Type: paper in peer-reviewed journal, cited 67 times (Web of Knowledge, 29-8-2013)
J. Phys. B.: At. Mol. Opt. Phys. **32**, L231 (1999).
DOI: 10.1088/0953-4075/32/9/107

[4] J. S. Parker, L. R. Moore, K. J. Meharg, D. Dundas and K. T. Taylor:
'Double-electron above threshold ionization of helium',
Type: paper in peer-reviewed journal, cited 99 times (Web of Knowledge, 29-8-2013)
J. Phys. B.: At. Mol. Opt. Phys., **34** L69 (2001).
DOI: 10.1088/0953-4075/34/3/103

[5] J. S. Parker, B. J. S. Doherty, K. J. Meharg and K. T. Taylor:
'Time delay between singly and doubly ionizing wavepackets in laser-driven helium',
Type: paper in peer-reviewed journal, cited 51 times (Web of Knowledge, 29-8-2013)
J. Phys. B.: At. Mol. Opt. Phys., **36** L393 (2003).
DOI: 10.1088/0953-4075/36/21/L04

[6] J. S. Parker, B. J. S. Doherty, K. T. Taylor, K. D. Schultz, C. I. Blaga and L. F. DiMauro:
'High-energy cutoff in the spectrum of strong-field nonsequential double ionization',
Type: paper in peer-reviewed journal, cited 101 times (Web of Knowledge, 29-8-2013)
Phys. Rev. Lett., **96** 133001 (2006).
DOI: 10.1103/PhysRevLett.96.133001

4. Details of the impact (indicative maximum 750 words)

High-end computation tools and techniques are an essential complement to theory and experiment for many of today's scientific and technological challenges. Challenges investigated include, amongst many others, improving the accuracy of weather forecasting [1] and the relationship between genetic markers and cancer [2]. High-end computational facilities are thus indispensable for scientific and technological progress in many disciplines. This strategic importance was underpinned by the investment of £113M made by the UK Government to install and operate the HECToR high-performance computing (HPC) facility over the period 2007- March 2014 [3,4]. The facility provided 11328 cores in Phase 1 (2007), and expanded to 22656 cores in Phase 2a (2009), to 44544 cores in Phase 2b (2010), and to 90112 cores in Phase 3 (2011) [5].

The HECToR HPC facility has been of strategic importance to the UK, as it enables the UK to be recognised as an international leader for computational science and engineering. The repeated upgrades of the facility are of major importance to ensure that the facility maintains its position as an internationally leading HPC facility. Usage statistics further demonstrate the significance of HECToR to HPC users, including industrial users. Typical usage from September 2012 – August 2013 shows that over 60000 processors are in use on average, with peak demand well over 85000 processors [6].

To guarantee that the investment provided by UK Government was spent judiciously, the procurement of the HECToR facility was overseen by EPSRC on behalf of the UK Research Councils. NAG Ltd. conducted the benchmarking process and operated the various acceptance tests for the different stages of HECToR, to ensure that the pre-agreed performance levels of the facility were met. To assess the vendors, a scoring system was put in place, in which the benchmarking process formed a substantial part of the overall score. (For the next UK high-performance computing facility ARCHER, benchmarking accounted for 35% of the overall score, and vendors were informed that those vendors that put significant effort into the benchmark tests tended to score better overall [7].) Five application codes were selected to participate in the HECToR benchmarking tests, one of which was the HELIUM code [8].

The HELIUM code was selected for inclusion amongst the benchmark codes for the following reasons [8]. With HECToR being the replacement facility for CSAR and HPCx, the most heavily used codes from this service were prime candidates for the procurement benchmarks. These candidate codes were assessed by NAG Ltd. for suitability. HELIUM was assessed to be a cleanly structured code, without complicated software dependencies. This was a major advantage for porting to the wide range of different architectures that were offered by the competing hardware vendors. In addition, good support was available from Prof. Taylor and Dr. Parker at QUB, and one of the members of the NAG benchmarking team was familiar with an earlier version of the code. Finally, the use of finite-difference methods in the HELIUM code complemented the different approaches pursued by other benchmark applications well. This is important in the procurement of a general purpose system, which will be used by a wide variety of users in many application areas. It is also important during acceptance testing: the system must be capable of performing at pre-agreed levels at each Phase of the installation for this wide variety of users.

Another important consideration for inclusion of the HELIUM code amongst the benchmark codes was its scalability [8]. The scale of the HECToR facility increased by a factor 8 during its lifetime. It was thus essential that benchmark codes offered the potential to run efficiently on next-generation facilities. NAG Ltd. assessed that the HELIUM code can be configured with ease to run on very large core counts. Hence, the HELIUM code played a particular important role in the acceptance tests for Phase 2b and Phase 3, where the hardware vendor was required to demonstrate that benchmark codes, such as the HELIUM code, could run for several hours on >90% of the total number of cores available (~40000 for Phase 2b and ~80000 for Phase 3), to give confidence that the system was sufficiently stable and reliable for production use on large scale, capability jobs.

The HELIUM code has helped EPSRC to demonstrate expected value for money for the £113M investment provided by the UK Government in the high-performance computing facility HECToR. Immediately after delivery of each phase, the HECToR facility has consistently ranked as a top-20 facility in the world, and a top-5 facility in the EU, (17th in November 2007 (phase 1), 20th in November 2009 (phase 2a), 16th in June 2010 (phase 2b), and 19th in November 2011) [9]. These rankings demonstrate a leadership position of the UK in the strategically important area of High-Performance Computing.

5. Sources to corroborate the impact (indicative maximum of 10 references)

[1] Supercomputing entry on resources for journalists on the Met Office web-site, update 7 Nov 2011
http://www.metoffice.gov.uk/news/in-depth/supercomputers

[2] Overcoming computational barriers: the search for gene - gene interactions in colorectal cancer, High Performance Computing Case Study by the University of Edinburgh Colon Cancer Genetics Group and Edinburgh Parallel Computing Centre, University of Edinburgh
http://www.hector.ac.uk/casestudies/oncology.php

[3] Press release EPSRC, 14 January 2008,
"Launch of new supercomputer for UK researchers"
http://www.epsrc.ac.uk/newsevents/news/2008/Pages/launchofnewsupercomputer.aspx

[4] News release, HECToR, 29 January 2013,
"HECToR Extension until March 2014"
http://www.hector.ac.uk/news-events/news/2013-01-29.php

[5] Hardware description web pages of HECToR web-site, downloaded 29 August 2013,
http://www.hector.ac.uk/service/hardware/ (Phase 3)
http://www.hector.ac.uk/service/hardware/phase1.php (Phase 1)
http://www.hector.ac.uk/service/hardware/phase2a.php (Phase 2a)
http://www.hector.ac.uk/service/hardware/phase2b.php (Phase 2b)

[6] Usage statistics web-pages on HECToR web-site, downloaded 29 August 2013
http://www.hector.ac.uk/service/status/

[7] Briefing notes for vendors for Archer – National HPC facility, 24 May 2012
www.epsrc.ac.uk/SitecollectionDocuments/other/ARCHERSuppliersBriefingMeeting24May2012.pdf

[8] Statement provided by Senior Technical Consultant, NAG Ltd.

[9] Web-site detailing six-monthly lists of top-500 most powerful supercomputers, maintained by Hans Meuer (University of Mannheim), Erich Strohmaier (NERSC / Lawrence Berkeley National Laboratory LBNL), Jack Dongarra (University of Tennessee) and Horst Simon (NERSC/LBNL).
http://www.top500.org/list/2007/11 (November 2007)
http://www.top500.org/list/2009/11 (November 2009)
http://www.top500.org/list/2010/6 (June 2010)
http://www.top500.org/list/2011/11 (November 2011)