CP2K:
HIGH PERFORMANCE ATOMISTIC SIMULATION

Iain Bethune
ibethune@epcc.ed.ac.uk

Outline

- CP2K Overview
- HECToR dCSE Timeline
- Key Achievements of dCSE Projects
  - Improved MPI scaling & load balancing
  - Hybrid MPI/OpenMP
  - DBCSR
- Other Ongoing Work

CP2K Overview

"CP2K is a program to perform atomistic and molecular simulations of solid state, liquid, molecular, and biological systems. It provides a general framework for different methods such as e.g., density functional theory (DFT) using a mixed Gaussian and plane waves approach (GPW) and classical pair and many-body potentials."

From www.cp2k.org (2004!)

CP2K Overview

- Many force models:
  - Classical
  - DFT (GPW)
  - Hybrid Hartree-Fock
  - LS-DFT
  - post-HF (MP2, RPA)
  - Combinations (QM/MM, mixed)
- Simulation tools
  - MD (various ensembles)
  - Monte Carlo
  - Minimisation (GEO/CELL_OPT)
  - Properties (Spectra, excitations …)
- Open Source
  - GPL, www.cp2k.org
  - 1m loc, ~2 commits per day
  - ~10 core developers

epcc | THE UNIVERSITY OF EDINBURGH

CP2K Overview

- Many force models:
  - Classical
  - DFT (GPW)
  - Hybrid Hartree-Fock
  - LS-DFT
  - post-HF (MP2, RPA)
  - Combinations (QM/MM, mixed)
- Simulation tools
  - MD (various ensembles)
  - Monte Carlo
  - Minimisation (GEO/CELL_OPT)
  - Properties (Spectra, excitations …)
- Open Source
  - GPL, www.cp2k.org
  - 1m loc, ~2 commits per day
  - ~10 core developers

CP2K Overview

- HECToR Phase 3 code usage (Nov 2011-Mar 2014)

Rank	Code	Node hours	Fraction of total	Method
1	VASP	5,822,878	19.34%	DFT
2	CP2K	2,222,059	7.38%	DFT
3	GROMACS	1,594,218	5.29%	Classical
4	DL_POLY	1,359,751	4.52%	Classical
5	CASTEP	1,351,163	4.49%	DFT

- CP2K usage £1.6m notional cost
  - (+ £2.4m on Phase 2)

CP2K Overview

- QUICKSTEP DFT: Gaussian and Plane Waves Method (VandeVondele *et al*, Comp. Phys. Comm., 2005)
  - Advantages of atom-centred basis (primary)
    - Density, KS matrices are sparse
  - Advantages of plane-wave basis (auxiliary)
    - Efficient computation of Hartree potential
  - Efficient mapping between basis sets
    - -> Computation of the KS Matrix is O(nlogn)

- Orbital Transformation Method (VandeVondele & Hutter, J. Chem. Phys., 2003)
  - Replacement for traditional diagonalisation to orthogonalise wave functions
  - Cubic scaling but ~10% cost

CP2K Overview

- (A,G) – distributed matrices
- (B,F) – realspace multigrids
- (C,E) – realspace data on planewave multigrids
- (D) – planewave grids

- (I,VI) – integration/ collocation of gaussian products
- (II,V) – realspace-to-planewave transfer
- (III,IV) – FFTs (planewave transfer)

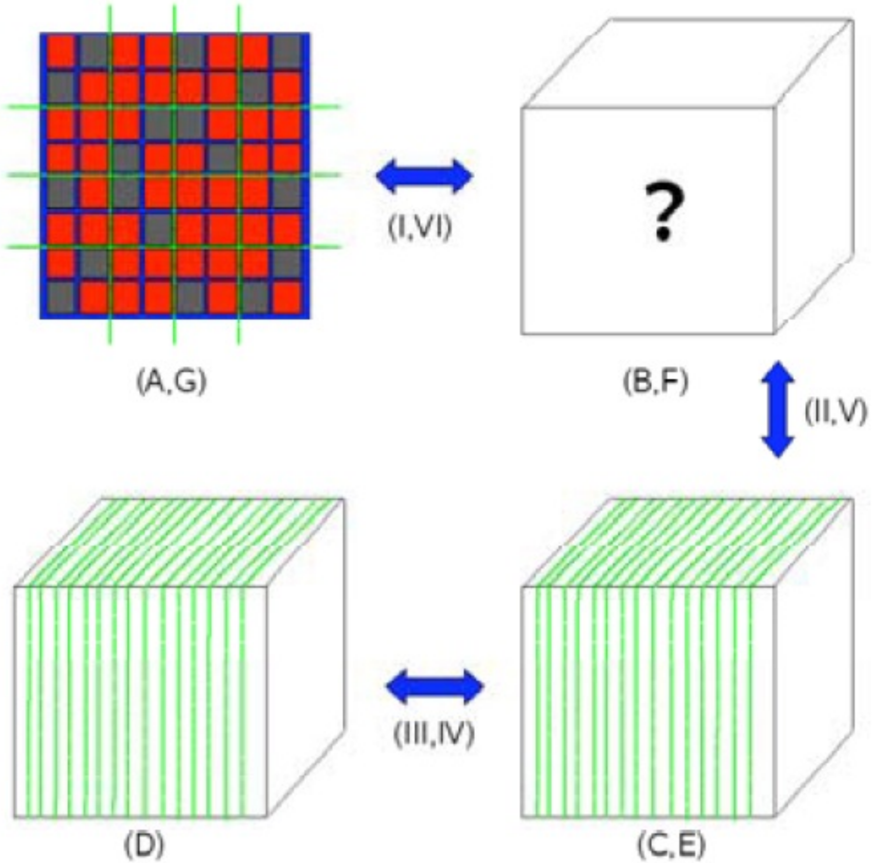

HECToR dCSE Timeline

- 3 dCSE projects during the HECToR service:

- Improving the performance of CP2K on HECToR
  - Jul 08 – Jul 09, 6 PMs
- Improving the scalability of CP2K on multi-core systems
  - Sep 09 – Sep 10, 6 PMs
- CP2K – Sparse Linear Algebra on 1000s of cores
  - Oct 10 – Dec 11, 6 PMs

- Joint proposals with Dr. Ben Slater (UCL), Prof. Joost VandeVondele (U. Zurich, ETHZ)

Key Achievements of dCSE Projects

- Optimised communication scheme for ‘halo swap’ in realspace-to-planewave transfer routine
  - Reduction in halo data volume
  - Non-blocking communication
  - -> 14% speedup on 256 cores (HECToR Phase 1)

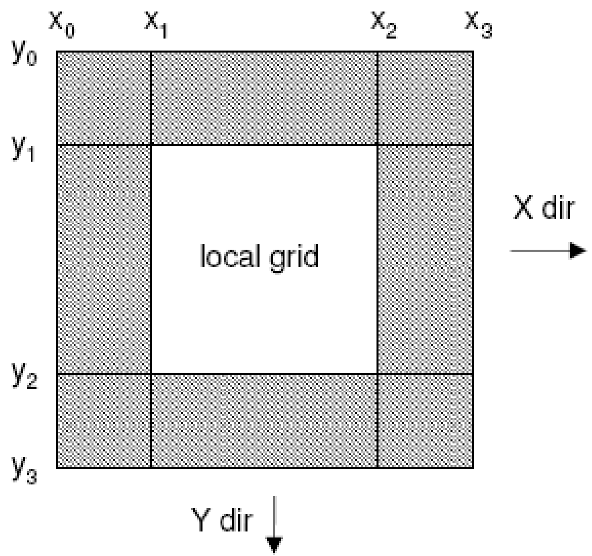

Key Achievements of dCSE Projects

- 3D FFT can be expensive in many calculations
  - 2D or 1D decomposition depending on process count
  - CrayPAT analysis showed repeated `mpi_cart_sub` calls
  - Introduced caching of re-used variables in FFT scratch space
  - Also FFT plan (allowing `FFTW_MEASURE`, `FFTW_PATIENT` etc.)
  - OpenMP added to parallelise on-node 1D FFT operation

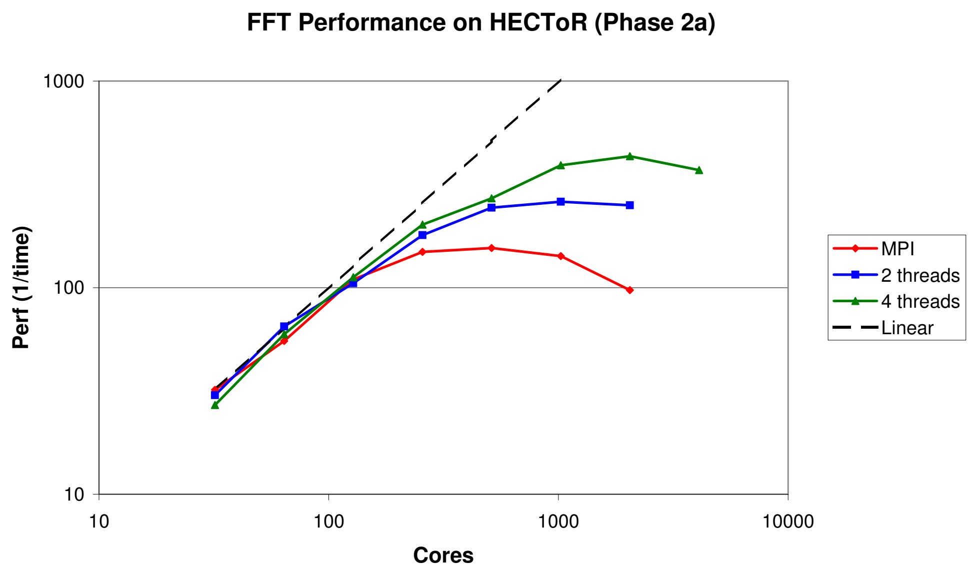

125^3 FFT performance on HECToR Phase 2a

1x Dual-core CPU per node

Key Achievements of dCSE Projects

- Load balancing by re-ordering multi-grids

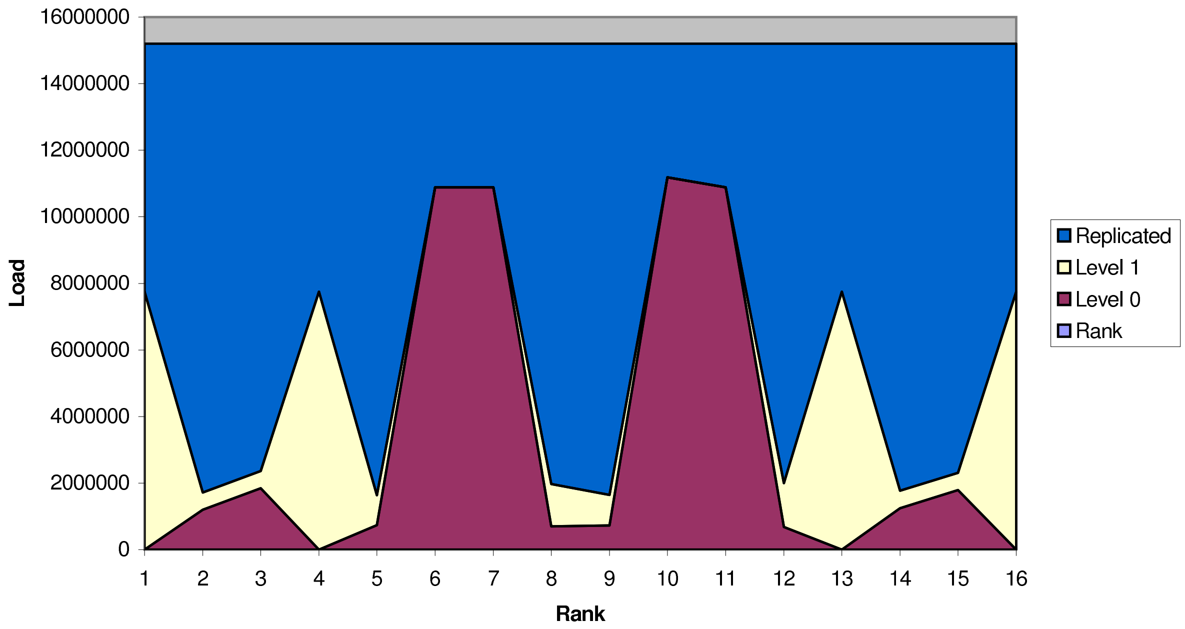

Figure 5: W216 load balance on 16 cores - perfect load balance achieved

Key Achievements of dCSE Projects

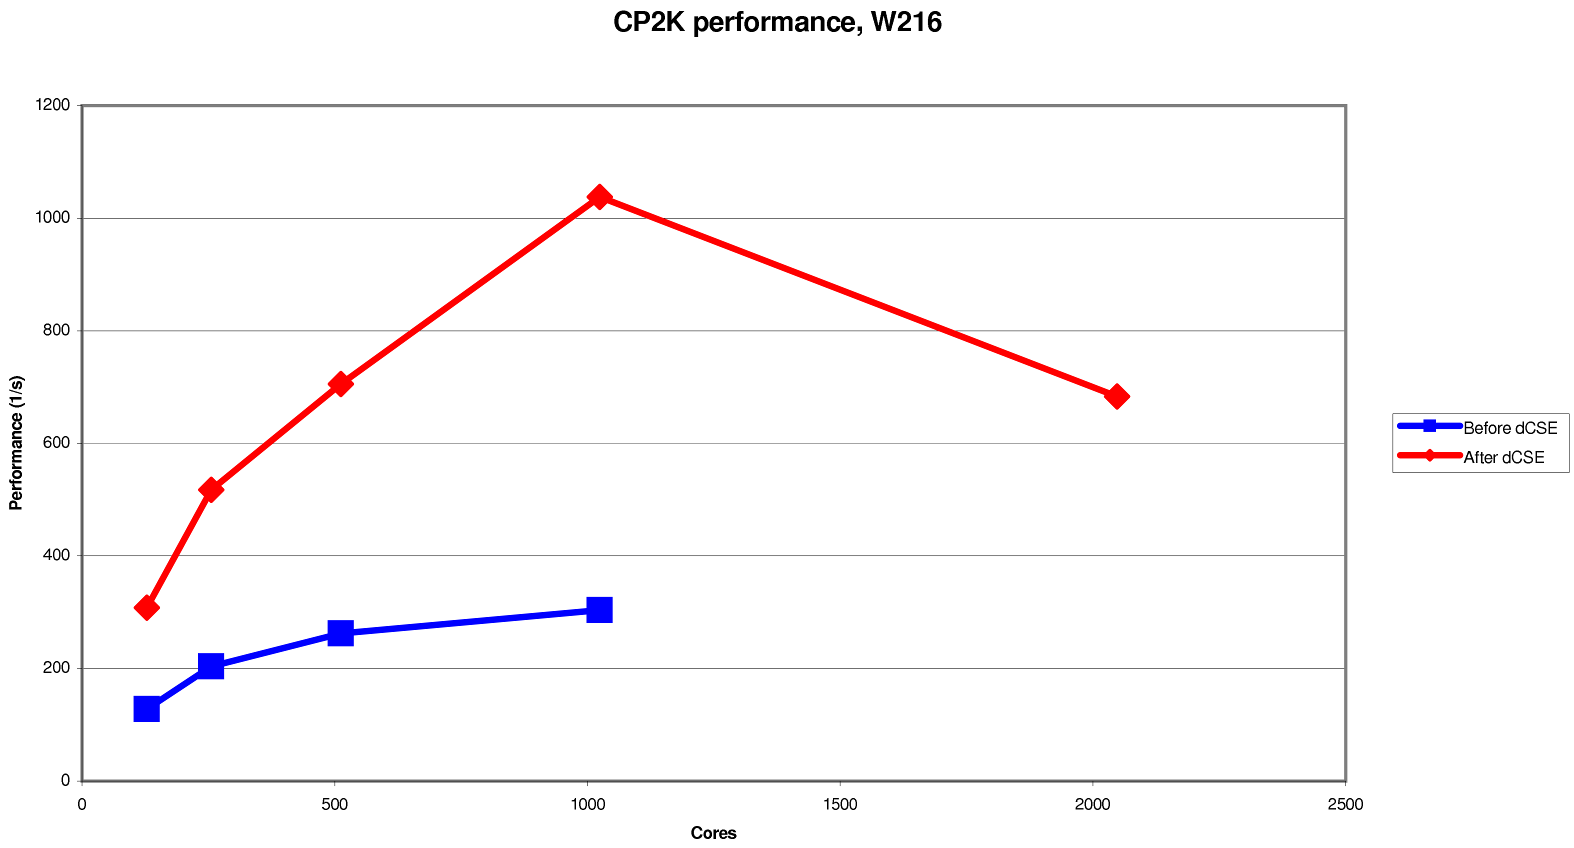

Key Achievements of dCSE Projects

- Hybrid MPI/OpenMP parallelisation
  - Loop-level OpenMP added to performance-critical regions
    - FFT (planewave grids)
    - Realspace grid <-> Sparse Matrix mapping
    - Realspace <-> planewave transfer
    - XC Functional Evaluation (PBE)
  - Efficient OpenMP implementation (GNU >= 4.5, Cray, Intel)
    - Allows micro-OpenMP regions
    - Minimal code restructuring
  - Result – Better scalability on medium/large core counts

Key Achievements of dCSE Projects

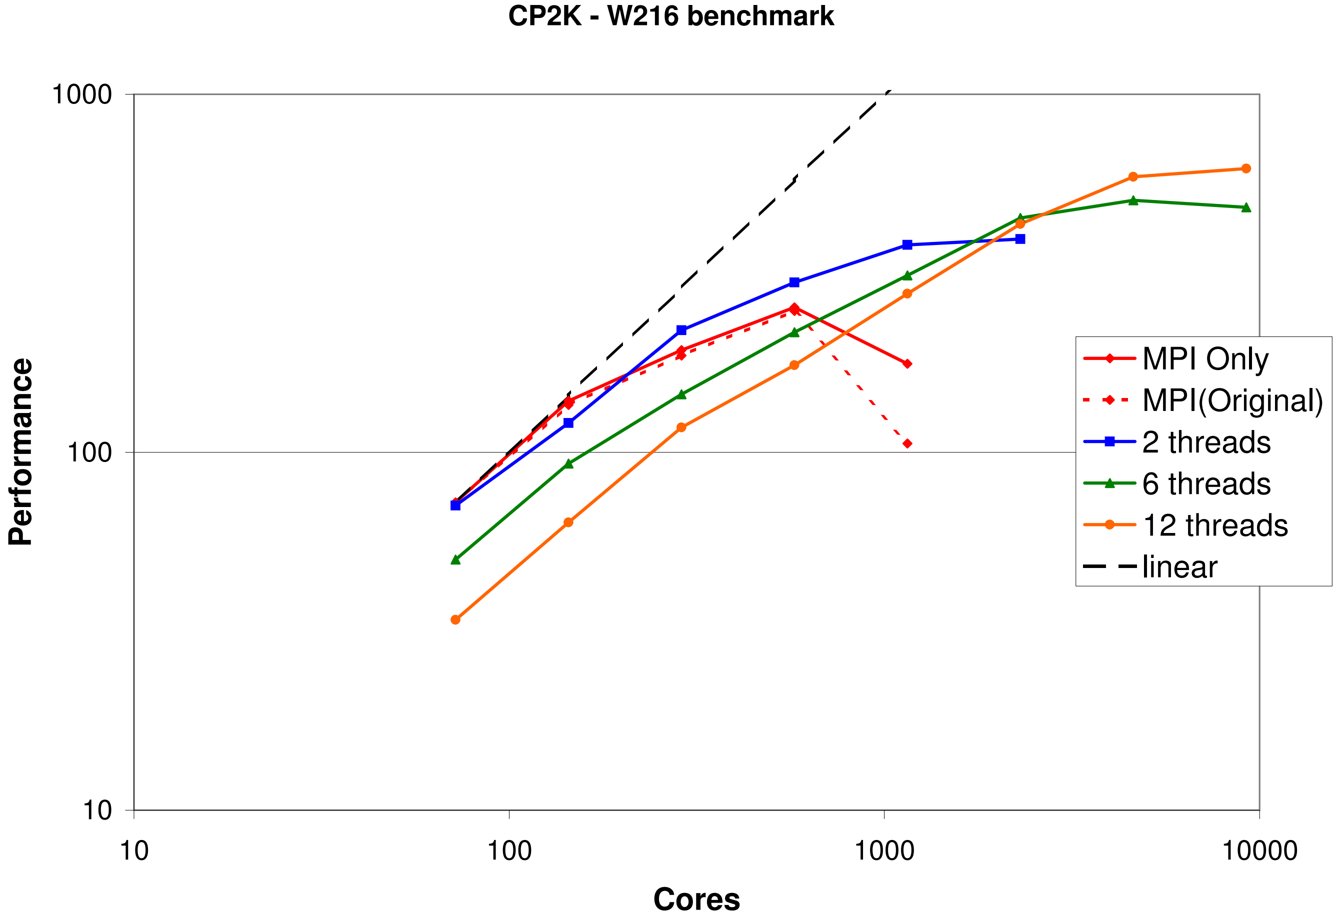

Figure 11: Performance of W216 on Rosa

Benchmark on CSCS Monte Rosa XT5

2x Hexa-core CPU per node

Similar to HECToR Phase 2b

Key Achievements of dCSE Projects

- DBCSR – Distributed Block Compressed Sparse Row
  - Developed in collaboration with Uni. Zurich

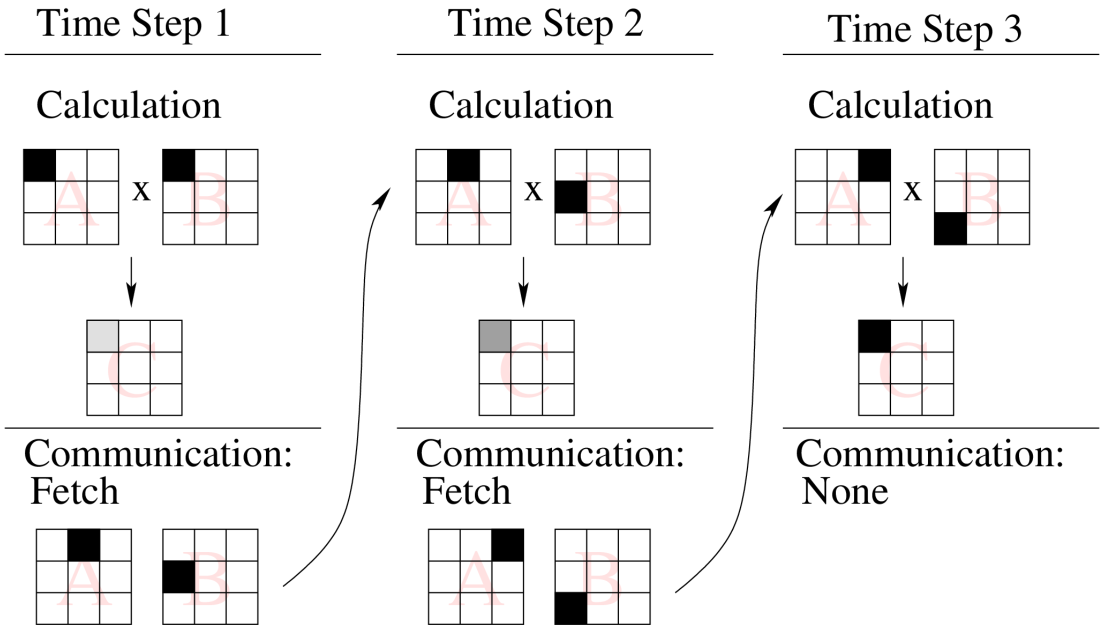

Figure 3: Example of Cannon's algorithm in DBCSR for 9 processes (Image courtesy Urban Borstnik)

Key Achievements of dCSE Projects

- Communication
  - Double-buffering comm. and calc. buffers
  - Non-blocking MPI communication
  - Optionally underload thread 0 to account for MPI overhead

- Computation
  - Local multiplication via cache-oblivious recursive multiplication
  - Auto-tuned matrix multiplication kernels for small block DGEMMs

    e.g. 4x4, 4x7, 4x13, 7x13, 13x13 …
  - Libsmm (idea now implemented in Cray libsci)

Key Achievements of dCSE Projects

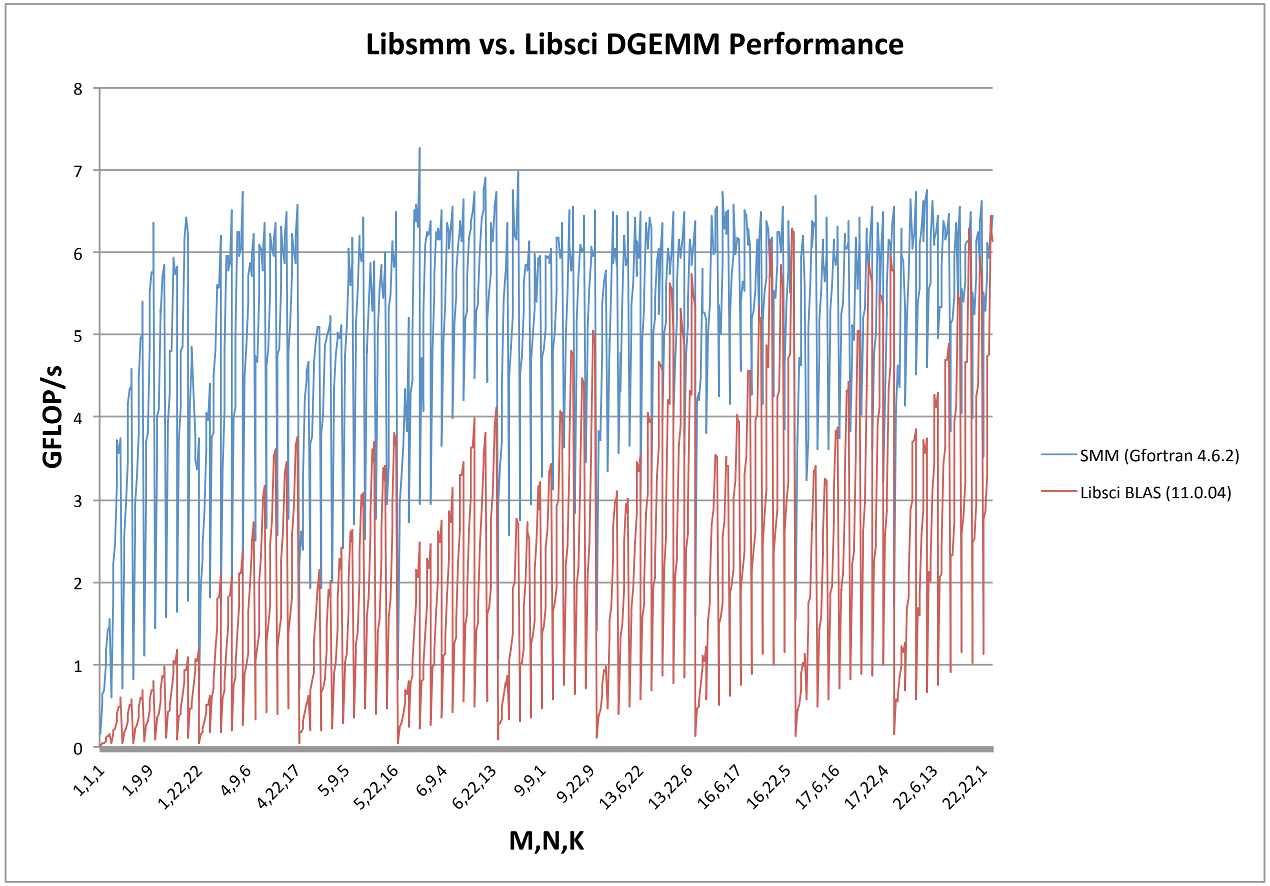

Figure 5: Comparing performance of SMM and Libsci BLAS for block sizes up to 22,22,22

Key Achievements of dCSE Projects

- Research Outputs:
  - 3 HECToR dCSE reports
    - Including user guidelines for achieving good scalability
    - Centrally-installed code with all optimisations
  - dCSE Workshop report (2009)
  - CUG Papers (2010, 2014)
  - Posters at CAMD, CCP5 Summer Schools (2012), Exascale Applications and Software Conference (2013)
- User feedback:
  - *"Using the optimized version of CP2K we were able to routinely model cells containing over 1000 atoms. The CSE investment into CP2K was crucial to making the study tractable.", Dr. Ben Slater*

Other Ongoing Work

- PRACE Application Enabling Projects
  - ~3 person-years effort from PRACE 1-IP, 2-IP, 3-IP
  - Porting and testing on
    - IBM BlueGene/P
    - Intel Clusters
    - Intel Xeon Phi (native mode)
  - Further OpenMP:
    - 24 functionals fully parallelised
    - Threaded 3D FFT for MP2 calculations
    - QS neighbour list operations
  - First-touch memory optimisation for realspace grids
  - Auto-tuned grid and matrix-multiplication kernels

Other Ongoing Work

- CP2K-UK: EPSRC Software for the Future
  - £500,000, 2013-2018
  - EPCC, UCL, KCL + 7 supporting groups

EPSRC
Pioneering research and skills

- Aims
  - Grow and develop existing CP2K community in UK
  - Lower barriers to *usage* and *development* of CP2K
  - Long-term sustainability of CP2K
  - Extend ability of CP2K to tackle challenging systems
- First user meeting in Jan 2014
  - 55 attendees, good feedback!

epcc | THE UNIVERSITY OF EDINBURGH

Other Ongoing Work

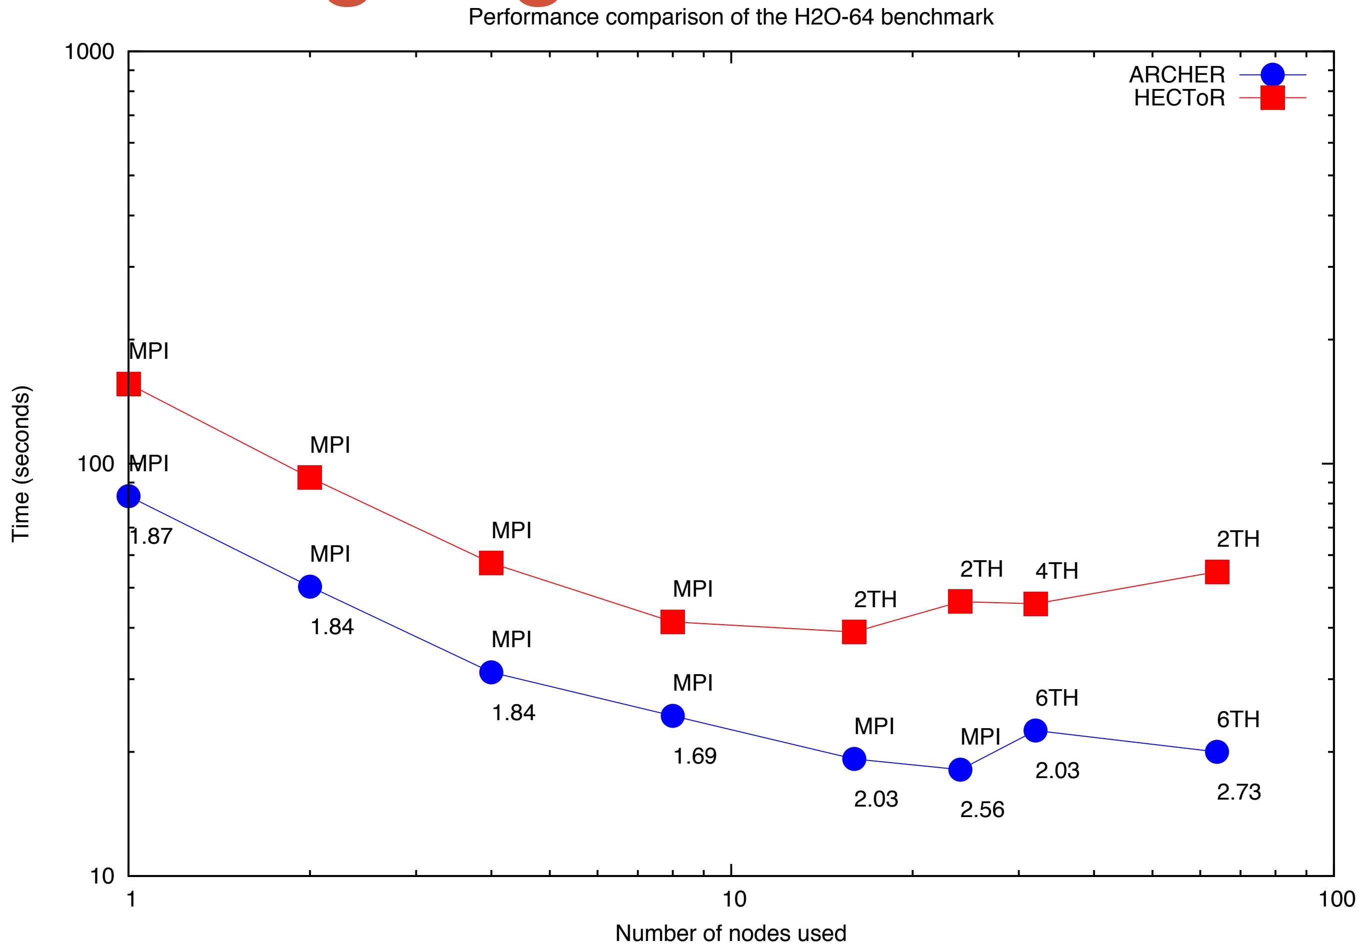

Other Ongoing Work

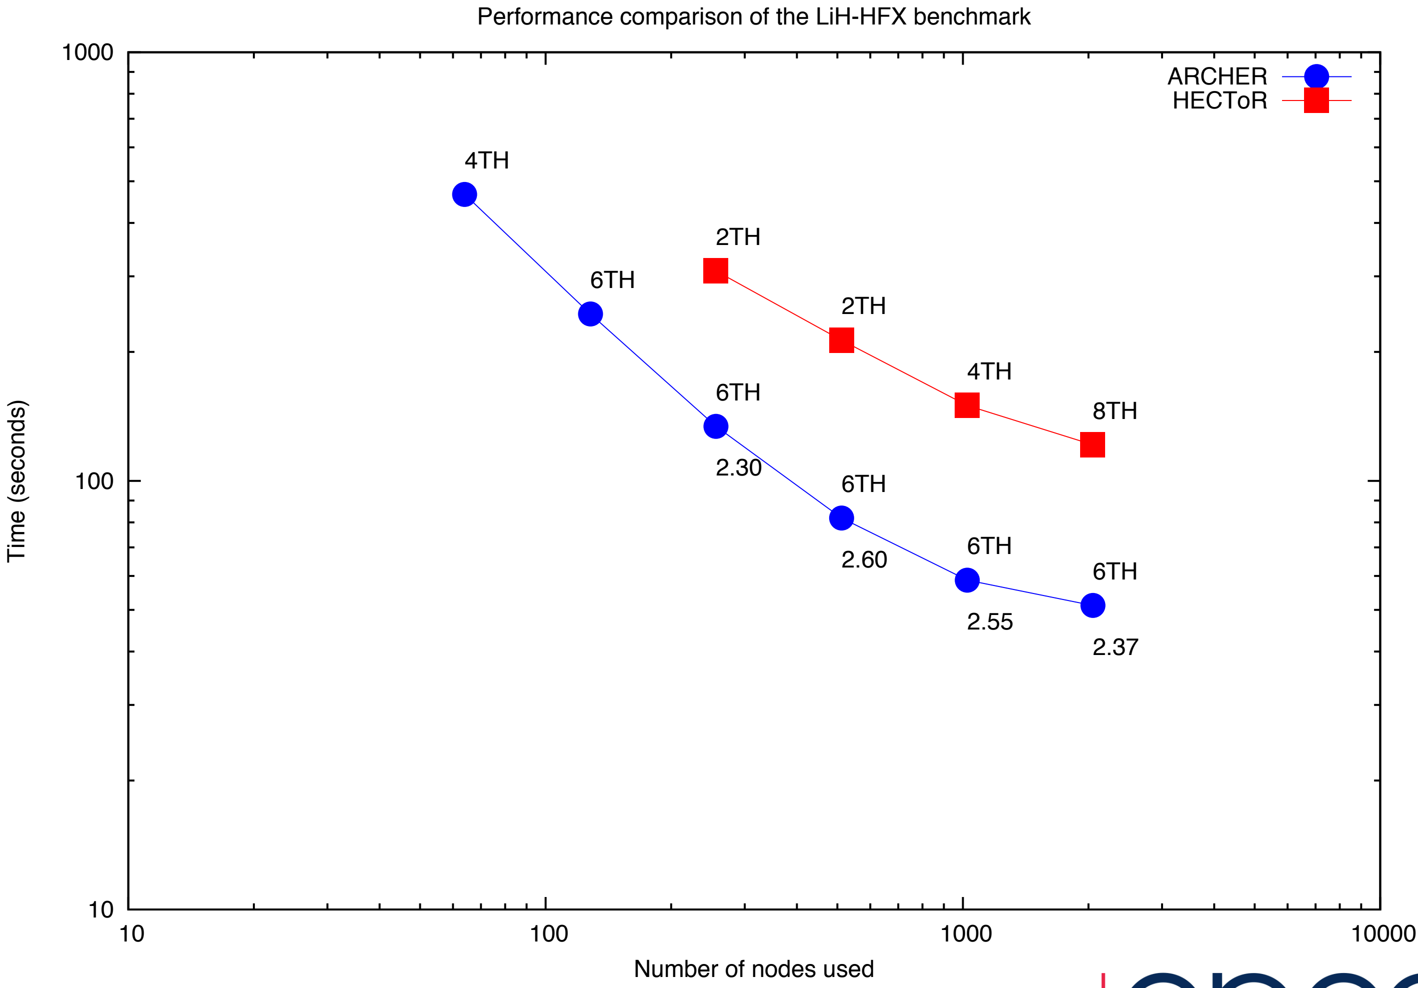

Other Ongoing Work

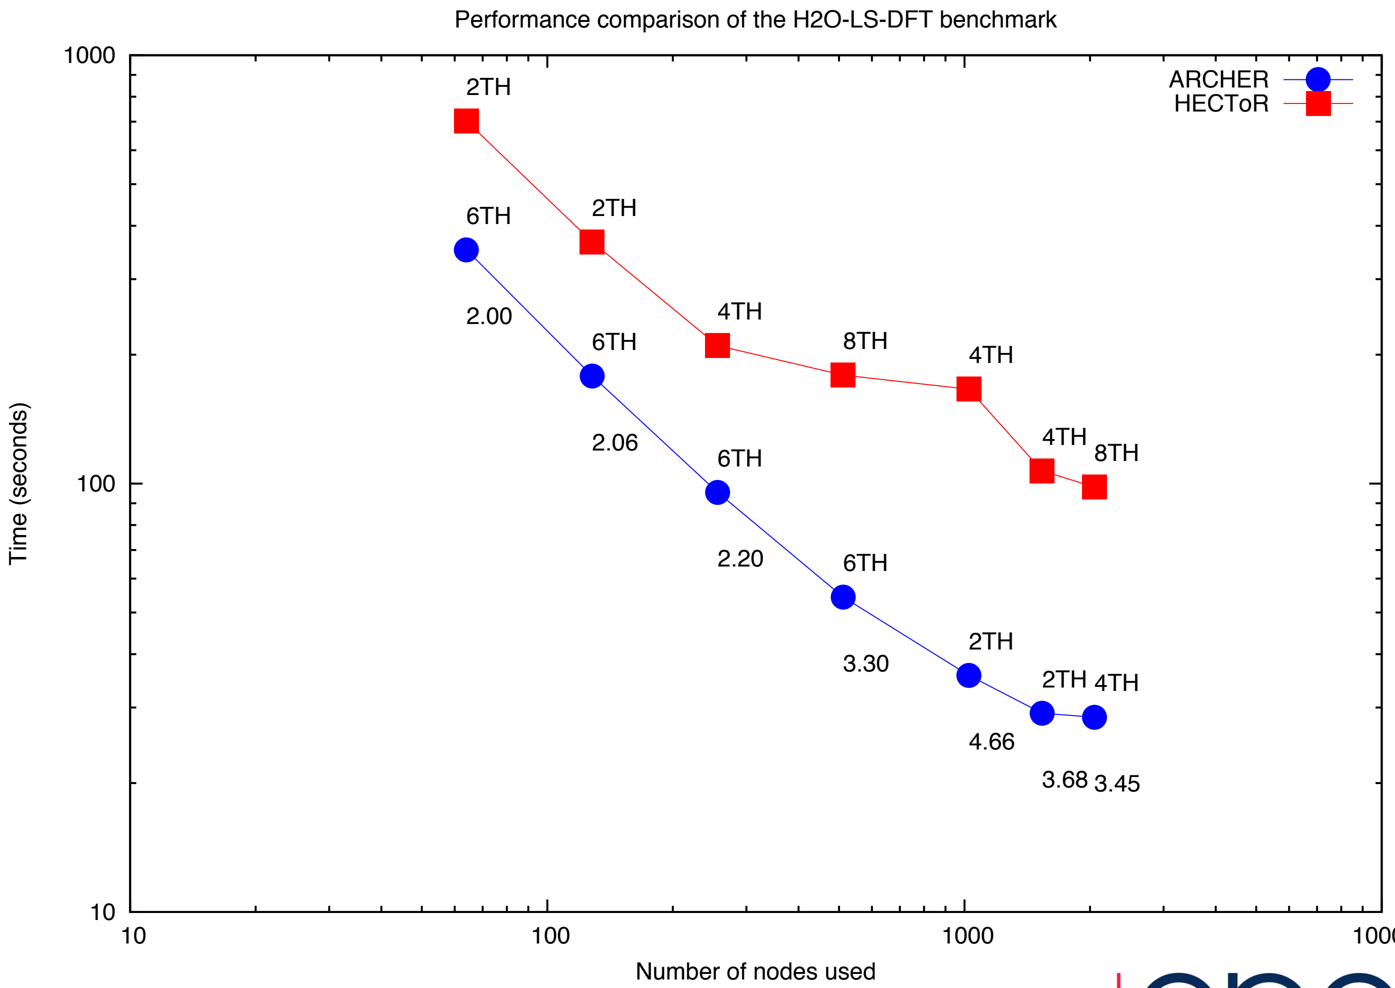

Some Final Thoughts

- CP2K has made huge strides in capability and performance
  - Key performance improvements under HECToR dCSE
  - dCSE projects were a *key enabler* for further funding
- Usability remains an issue
  - Although steadily improving (tutorials, training, support…)
- Long-term funding for user support now in place
  - Looking to spin-off new focused development projects under ARCHER eCSE
- Expect growing usage on ARCHER

Acknowledgement

This work made use of the facilities of HECToR, the UK's national high-performance computing service, which is provided by UoE HPCx Ltd at the University of Edinburgh, Cray Inc and NAG Ltd, and funded by the Office of Science and Technology through EPSRC's High End Computing Programme.

Thanks for your attention, and…

…any questions?