CP2K PERFORMANCE FROM CRAY XT3 TO XC30

Iain Bethune (ibethune@epcc.ed.ac.uk)
Fiona Reid
Alfio Lazzaro

Outline

- CP2K Overview
  - Features
  - Parallel Algorithms
- Cray HPC Systems
  - Trends
- Water Benchmarks
  - 2005 – 2013
- Comprehensive Benchmarking
  - XE6 vs XC30
- CP2K with Accelerators

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!)

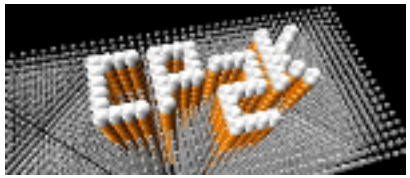

epcc | THE UNIVERSITY OF EDINBURGH

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

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

epcc | THE UNIVERSITY OF EDINBURGH

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)

epcc | THE UNIVERSITY OF EDINBURGH

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

epcc | THE UNIVERSITY OF EDINBURGH

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)

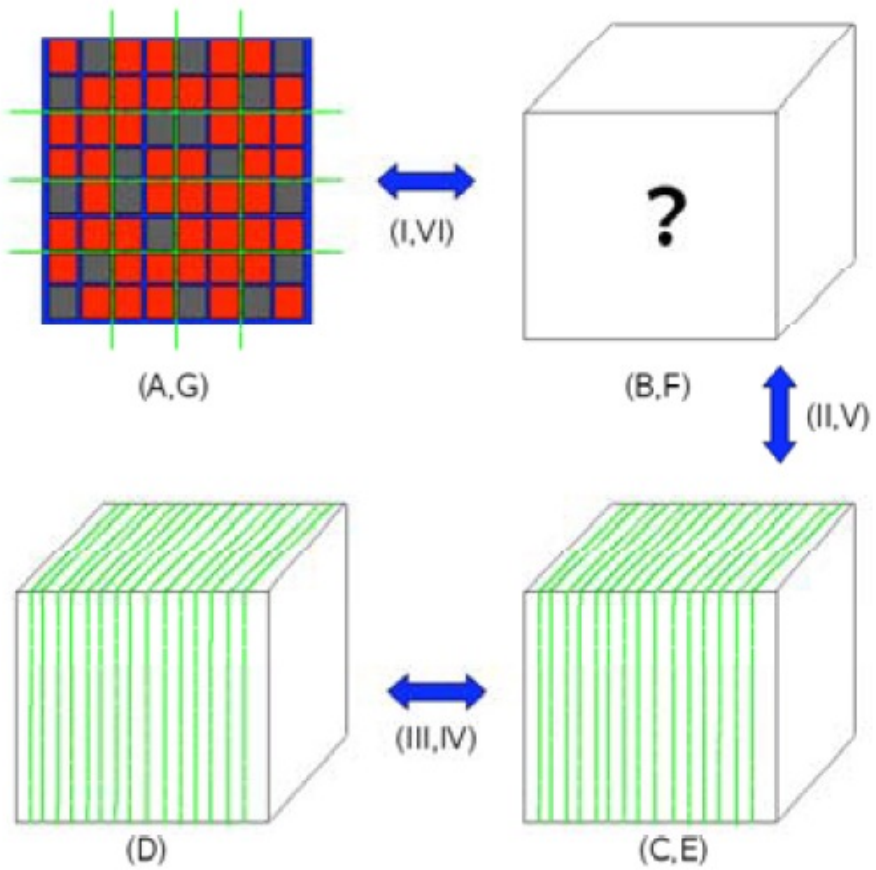

epcc

THE UNIVERSITY OF EDINBURGH

CP2K Overview

- Distributed realspace grids
  - Overcome memory bottleneck
  - Reduce communication costs
  - Parallel load balancing
    - On a single grid level
    - Re-ordering multiple grid levels
    - Finely balance with replicated tasks

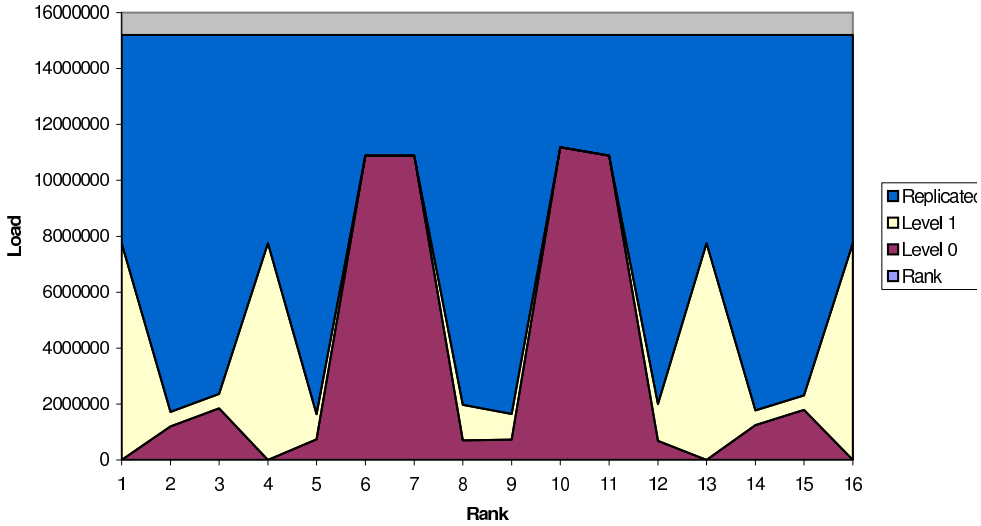

Level 1, fine grid, distributed

1	2	3
4	5	6
7	8	9

Level 2, medium grid, dist

5	6	8
3	1	7
9	4	2

Level 3, coarse grid, replicated

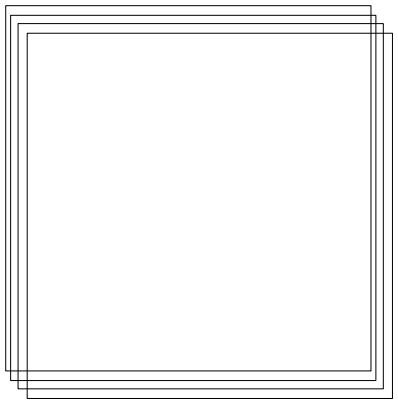

epcc

THE UNIVERSITY OF EDINBURGH

CP2K Overview

- Fast Fourier Transforms
  - 1D or 2D decomposition
  - FFTW3 and CuFFT library interface
  - Cache and re-use data
    - FFTW plans, Cartesian communicators
- DBCSR
  - Distributed Sparse MM based on Cannon's Algorithm
  - Local multiplication recursive, cache oblivious
  - `libsmm` for small block multiplications

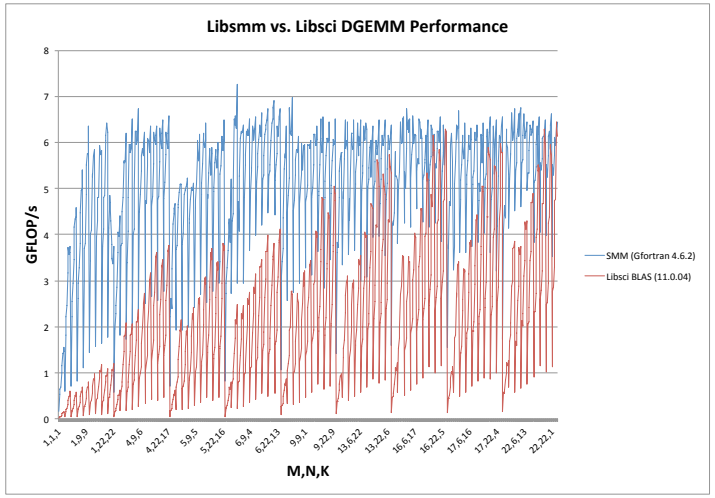

Figure 5: Comparing performance of SMM and Libsci BLAS for block sizes up to 22,22,22

epcc

THE UNIVERSITY OF EDINBURGH

CP2K Overview

- OpenMP
  - Now in all key areas of CP2K
  - FFT, DBCSR, Collocate/Integrate, Buffer Packing
  - Incremental addition over time

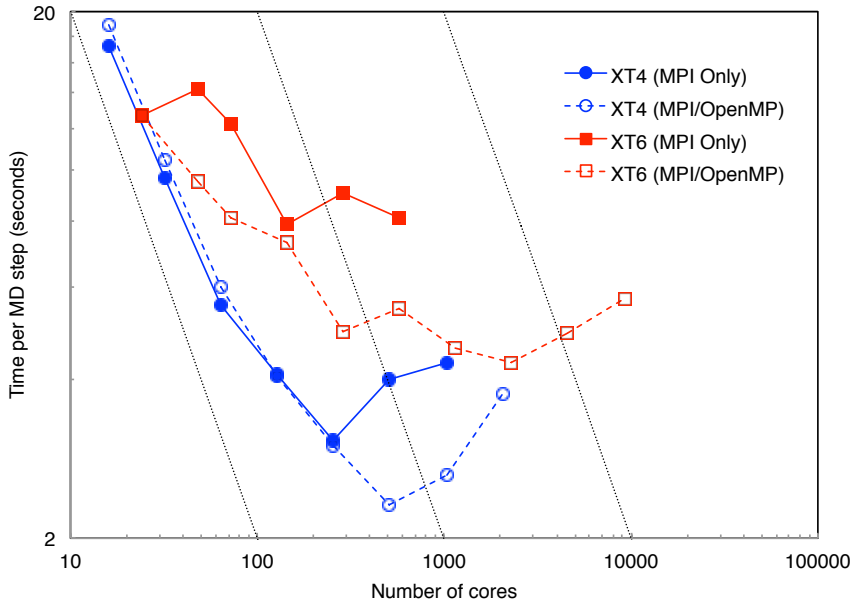

epcc

THE UNIVERSITY OF EDINBURGH

Cray HPC Systems

Name	Arch.	Processor	Clock (GHz)	Nodes	Cores/ Node	Peak TFlop/s	GFlop/s/ Node	Year
XT3 Stage 0	XT3	AMD Opteron 146	2.0	84	1	0.336	4.0	2005
XT3 Stage 1	XT3	AMD Opteron 152	2.6	1100	1	5.72	5.2	2006
Piz Palü	XT3	AMD Opteron 185 Dual Core	2.6	1664	2	17.31	10.4	2007
HECToR Phase 1	XT4	AMD Opteron 1220 “Santa Ana” Dual Core	2.8	5664	2	63.44	11.2	2007
HECToR Phase 2a	XT4	AMD Opteron 2356 “Barcelona” 4-Core	2.3	5664	4	104.22	18.4	2009
Monte Rosa	XT5	AMD Opteron 2431 “Istanbul” 6-Core	2.4	1844	12	212.43	115.2	2009
HECToR Phase 2b	XT6	AMD Opteron 6172 “Magny-Cours” 12-Core	2.1	1856	24	374.17	201.6	2010
Piz Palü[1]	XE6	AMD Opteron 6272 “Interlagos” 16-Core	2.1	1496	32	402.12	268.8	2011
HECToR Phase 3	XE6	AMD Opteron 6276 “Interlagos” 16-Core	2.3	2816	32	829.03	294.4	2011
Tödi	XK7	AMD Opteron 6272 “Interlagos” 16-Core + NVIDIA Tesla K20X	2.1	272	16 (+14)	392.90	1444.5	2012
Piz Daint	XC30	Intel Xeon E5-2670 “Sandy-Bridge” 8-Core + NVIDIA Tesla K20X	2.6	5272	8 (+14)	7788.90	1477.4	2013
ARCHER	XC30	Intel Xeon E5-2697 v2 “Ivy-Bridge” 12-core	2.7	3008	24	1559.35	518.4	2013

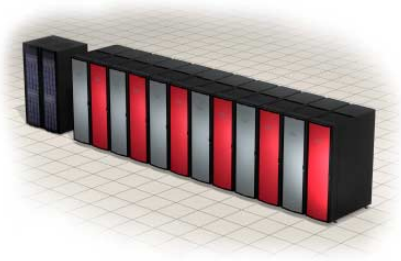

epcc

THE UNIVERSITY OF EDINBURGH

Water benchmarks

- Born-Oppenheimer MD using Quickstep DFT
  - TZV2P basis set
  - 280 Ry planewave cut-off
  - LDA exchange-correlation functional

= typical production settings

  - 32 up to 2048 water molecules

  - `H2O-32` – 96 atoms, 256 electrons, 9.9 Å^3
    - Typical problem size in ~2005

  - `H2O-2048` – 6144 atoms, 49152 electrons, 39.5 Å^3
    - Large, even for 2014!

epcc | THE UNIVERSITY OF EDINBURGH

Water benchmarks

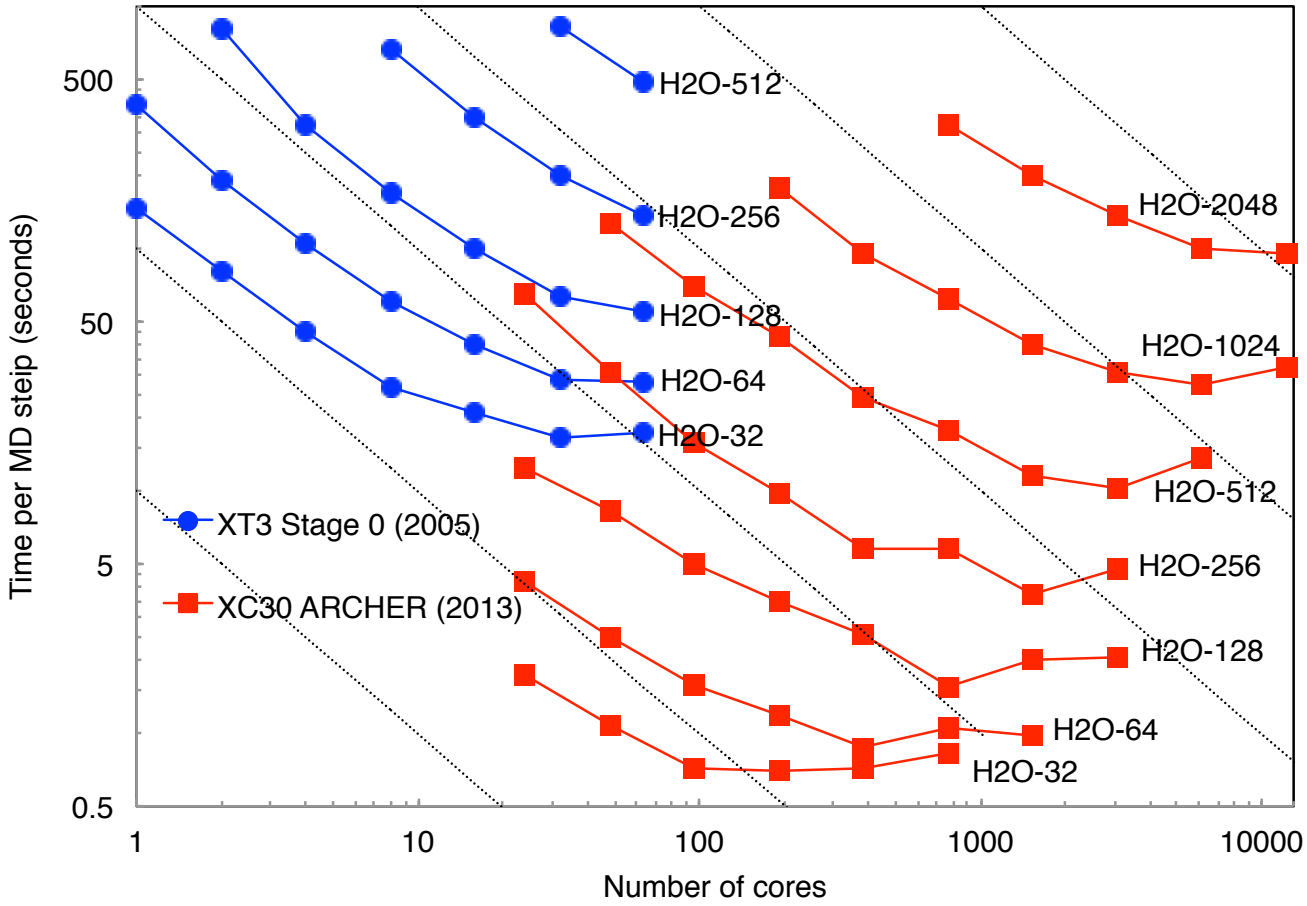

epcc

THE UNIVERSITY OF EDINBURGH

Water benchmarks

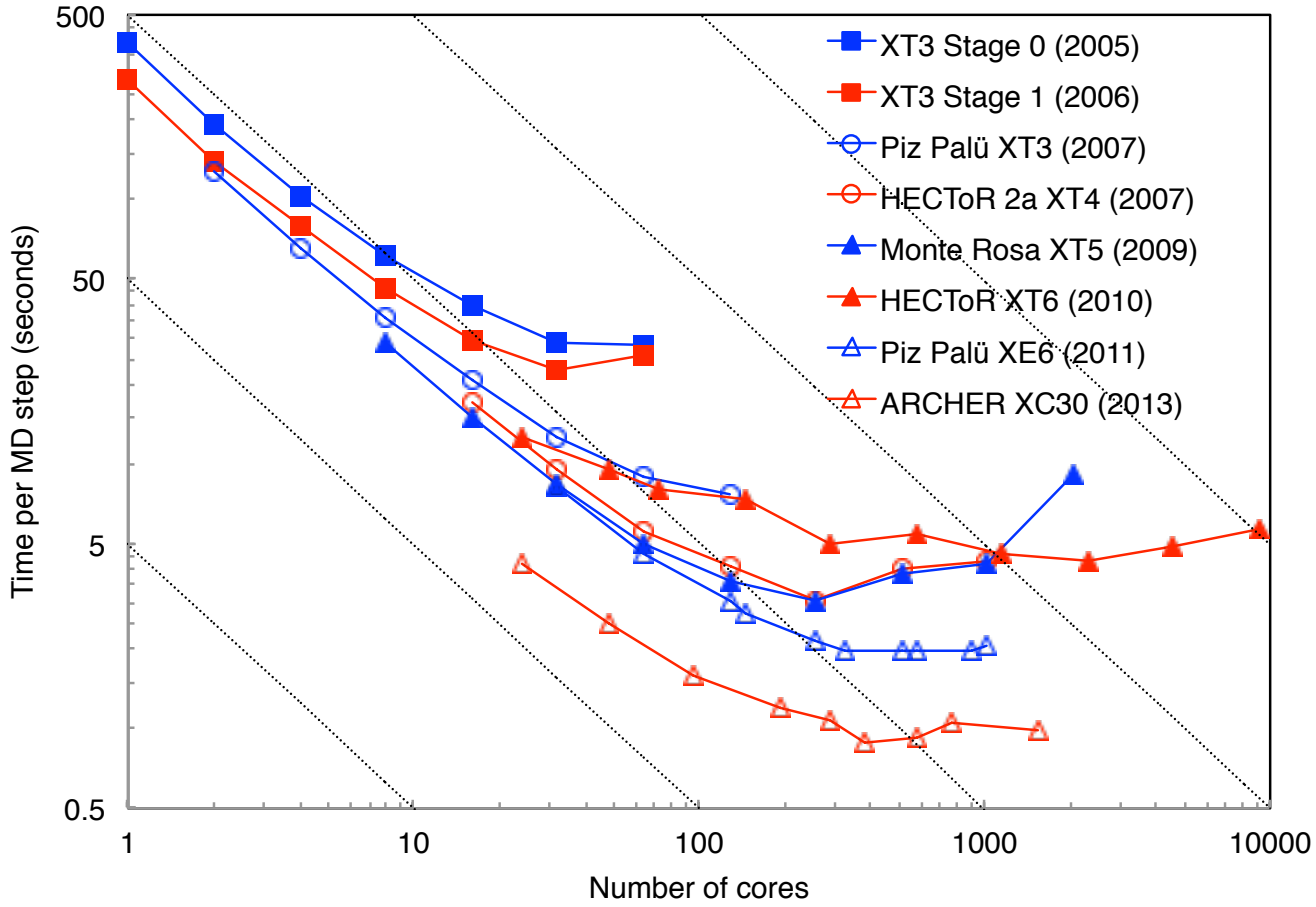

epcc

THE UNIVERSITY OF EDINBURGH

Comprehensive Benchmarking

- `H2O-*` benchmarks do not address the range of features now available in CP2K
  - Classical Force Fields
  - Linear-scaling DFT
  - Hybrid DFT (Hartree-Fock Exchange)
  - Many-body correlation (MP2, RPA)

- Aimed at users
  - Performance expectations HECToR Phase 3 -> ARCHER
  - Presented at 1st Annual CP2K Users Meeting (Jan 2014)

Comprehensive Benchmarking

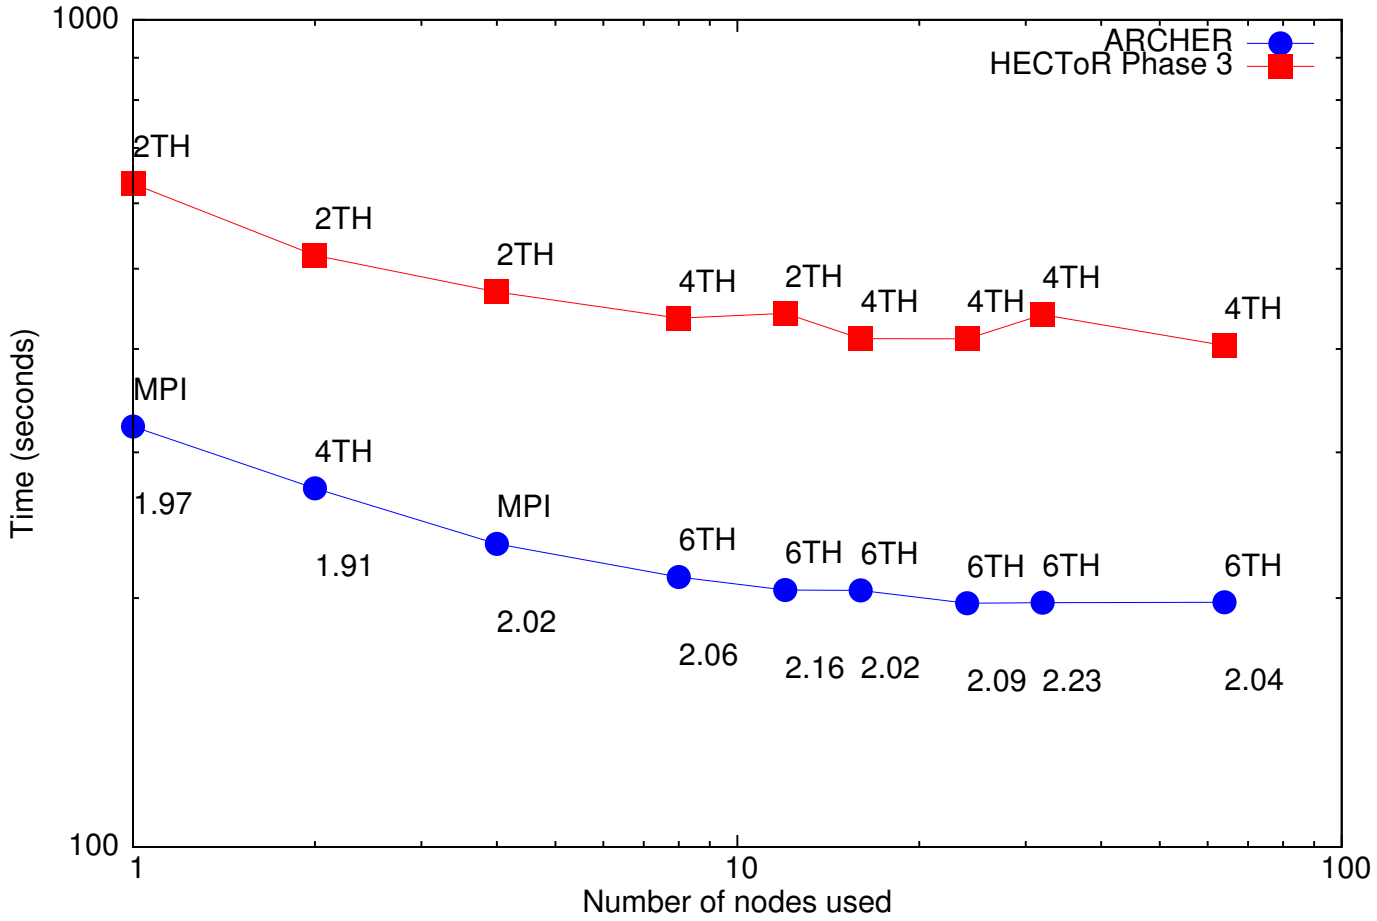

epcc

THE UNIVERSITY OF EDINBURGH

Comprehensive Benchmarking

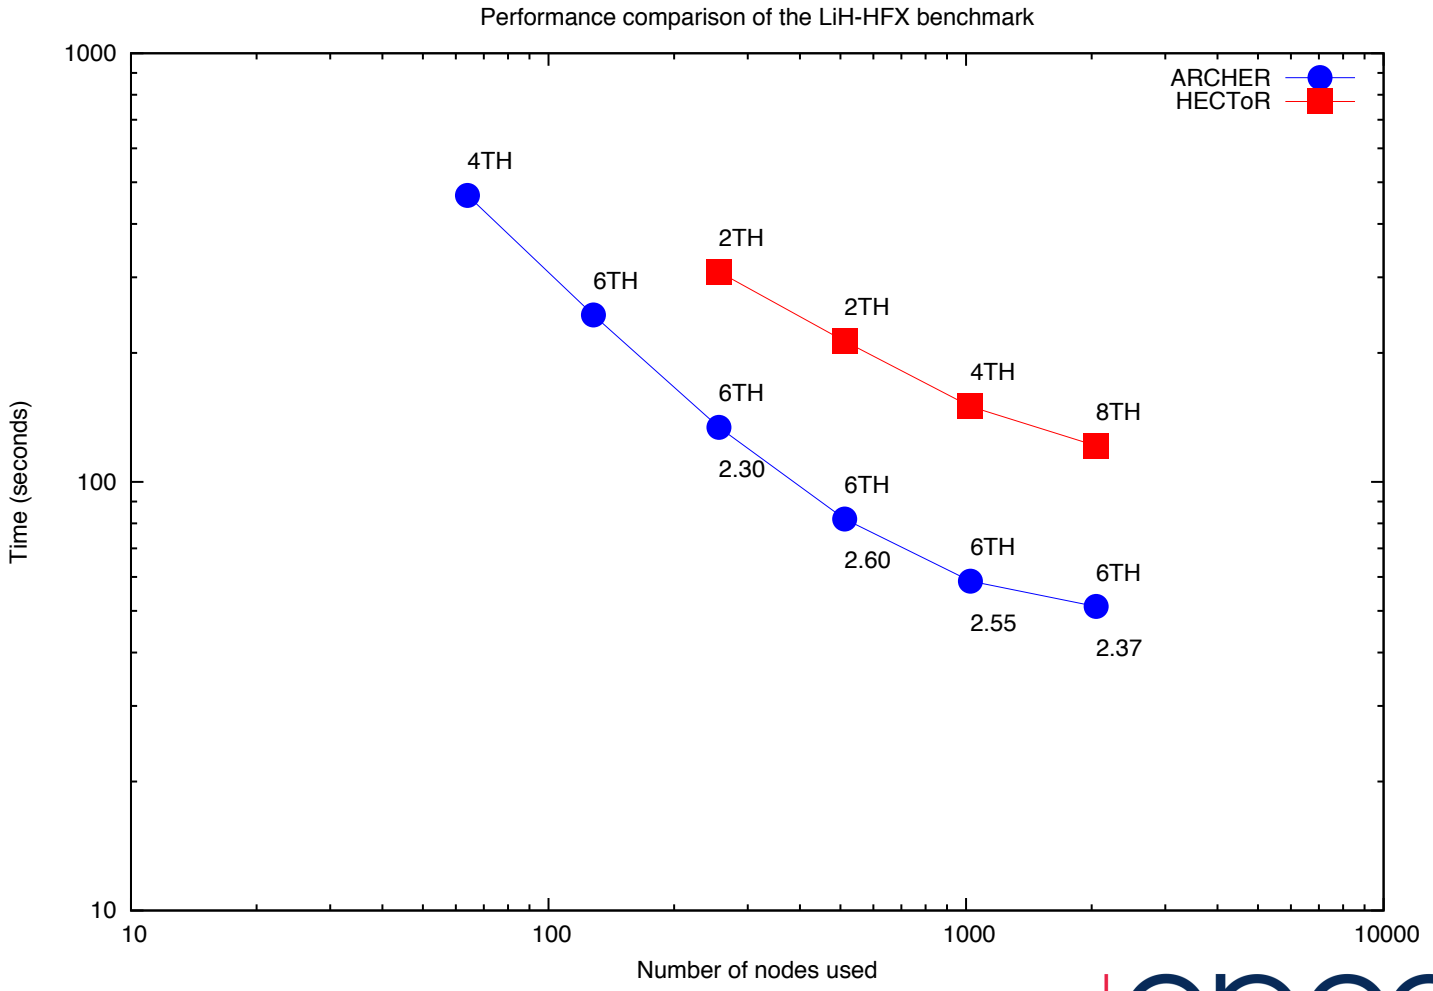

epcc

THE UNIVERSITY OF EDINBURGH

Comprehensive Benchmarking

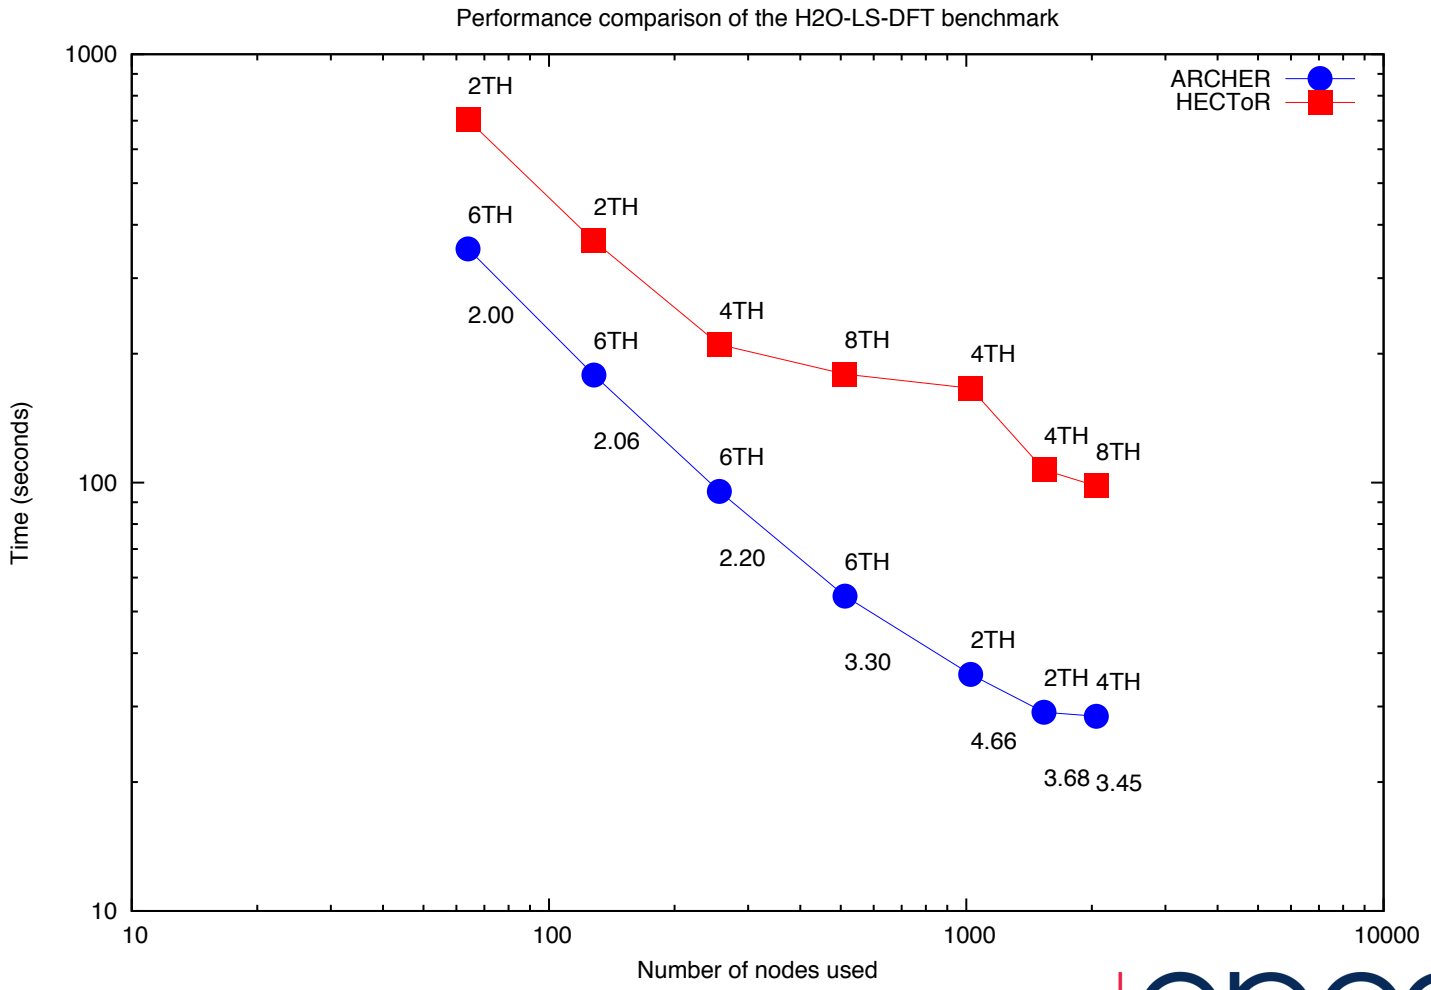

epcc

THE UNIVERSITY OF EDINBURGH

Comprehensive Benchmarking

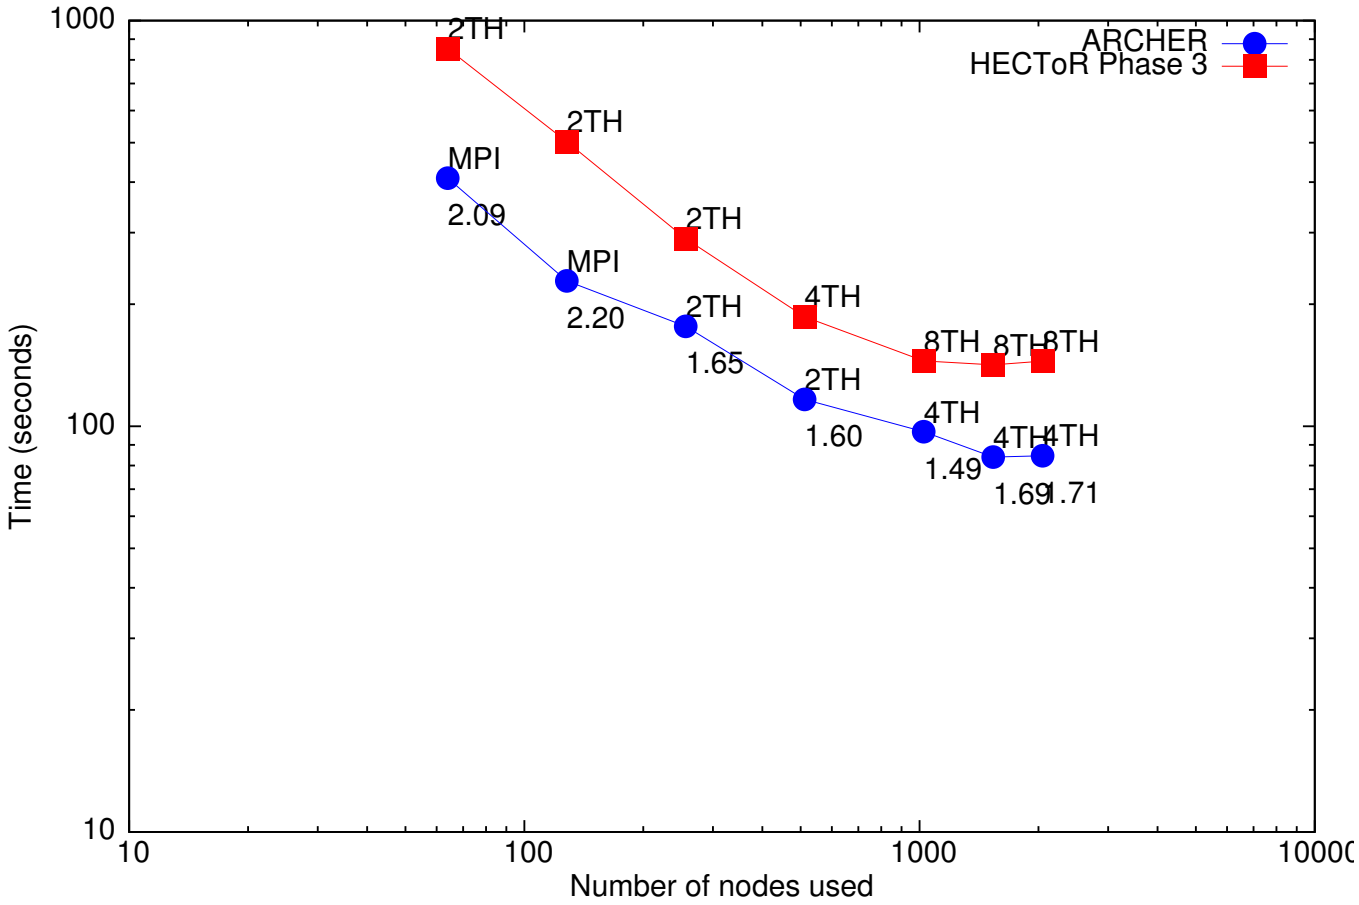

CP2K with Accelerators

- Heterogeneous systems well established
  - #1,2,6,7 in TOP 500 use Intel Xeon Phi or NVIDIA K20x GPU
  - XC30 & XK7 dual socket = 2 x CPU or CPU + GPU

- CP2K used during initial validation tests of Piz Daint
  - CUDA GPU support for DBCSR
  - Best performance obtained for LS-DFT calculations
  - Work by Zurich, Cray, NVIDIA & CSCS

CP2K with Accelerators

- Implementation details:
  - `libcusmm` for block-level of multiplication (4x better than cuBLAS)
  - CPU fills stacks of smm
  - One GPU per MPI process, utilise cores with OpenMP
  - Asynchronous offload to GPU via CUDA streams
  - Asynchronous communication between nodes

- Benchmarks
  - H2O-DFT-LS (6144 atoms, large blocks)
  - TiO2 (9786 atoms, mixed block sizes)
  - AMORPH (13846 atoms, small blocks)

CP2K with Accelerators

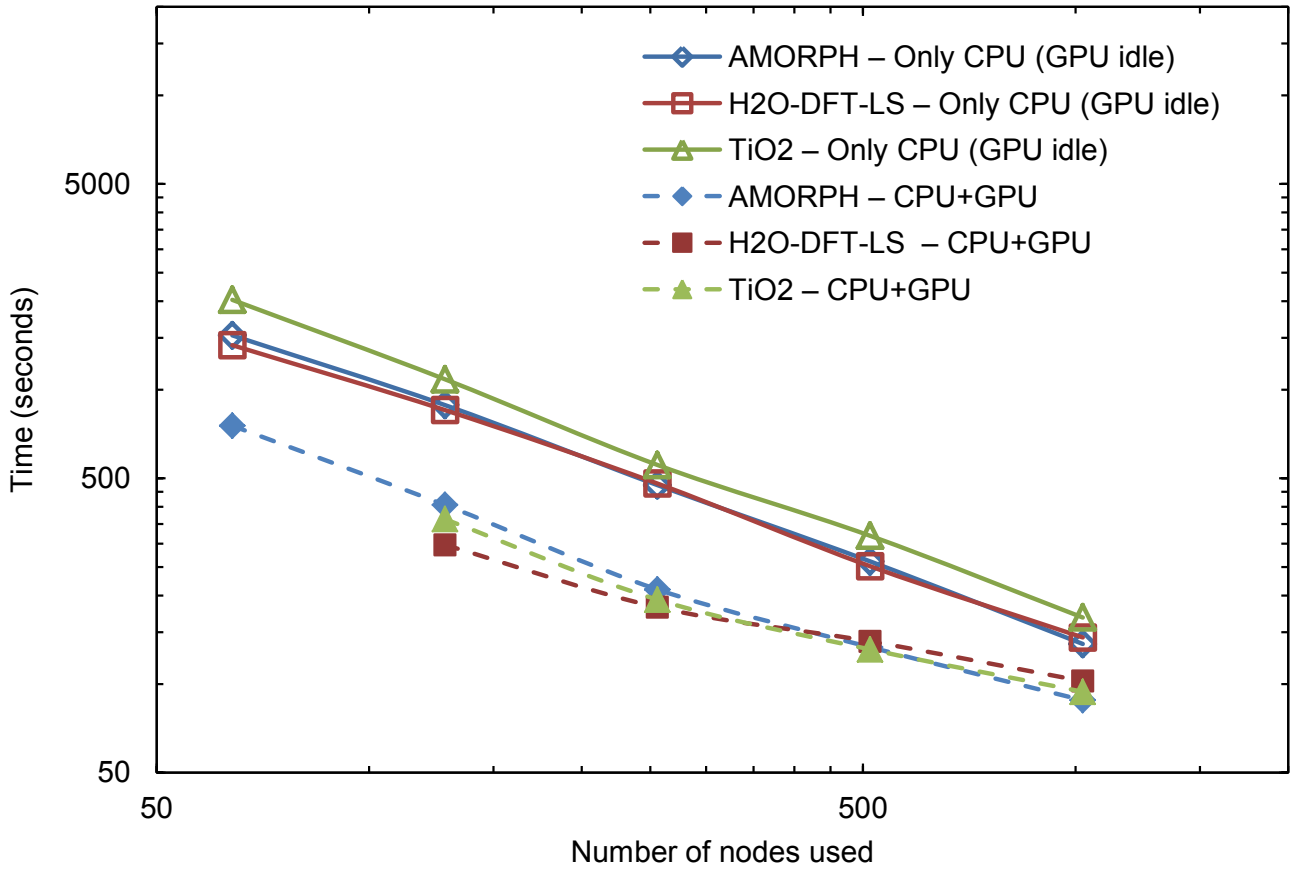

epcc

THE UNIVERSITY OF EDINBURGH

CP2K with Accelerators

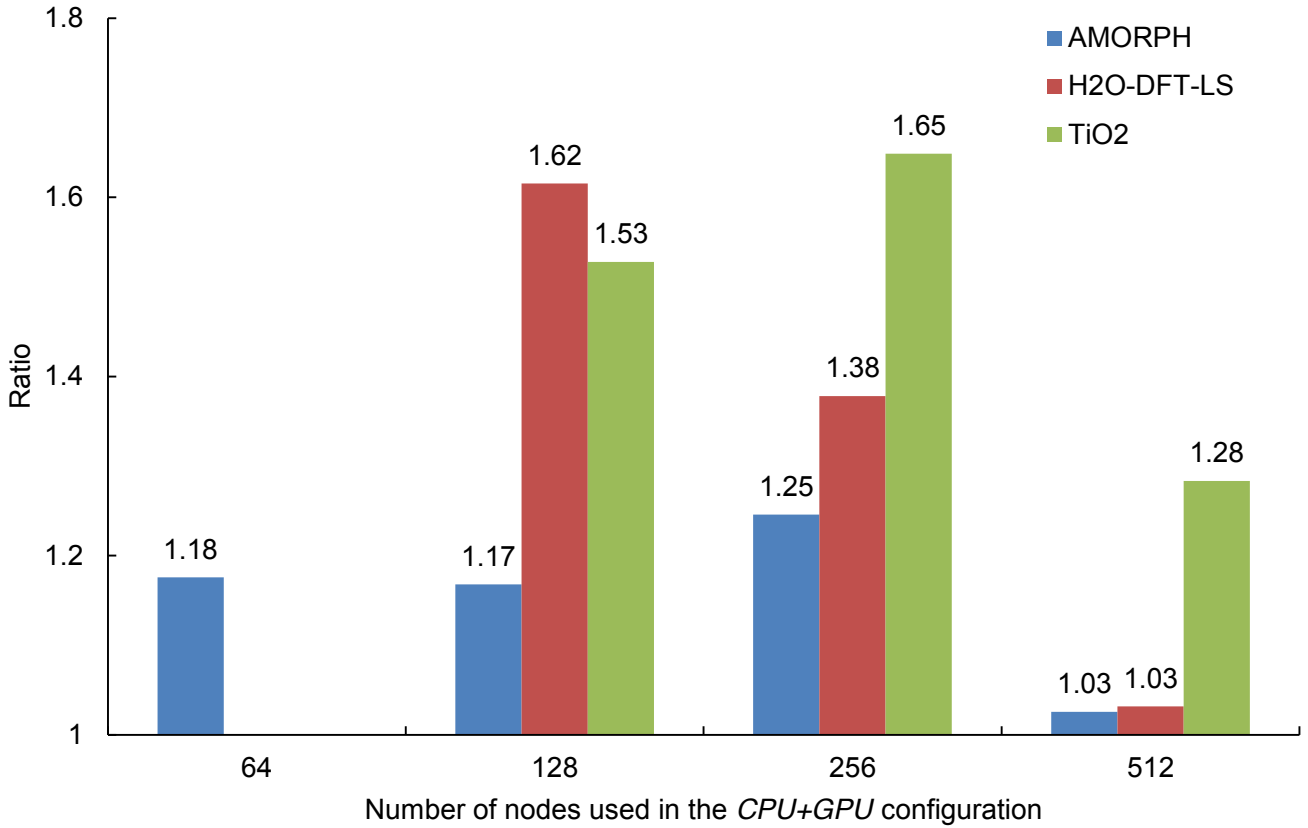

epcc

THE UNIVERSITY OF EDINBURGH

CP2K with Accelerators

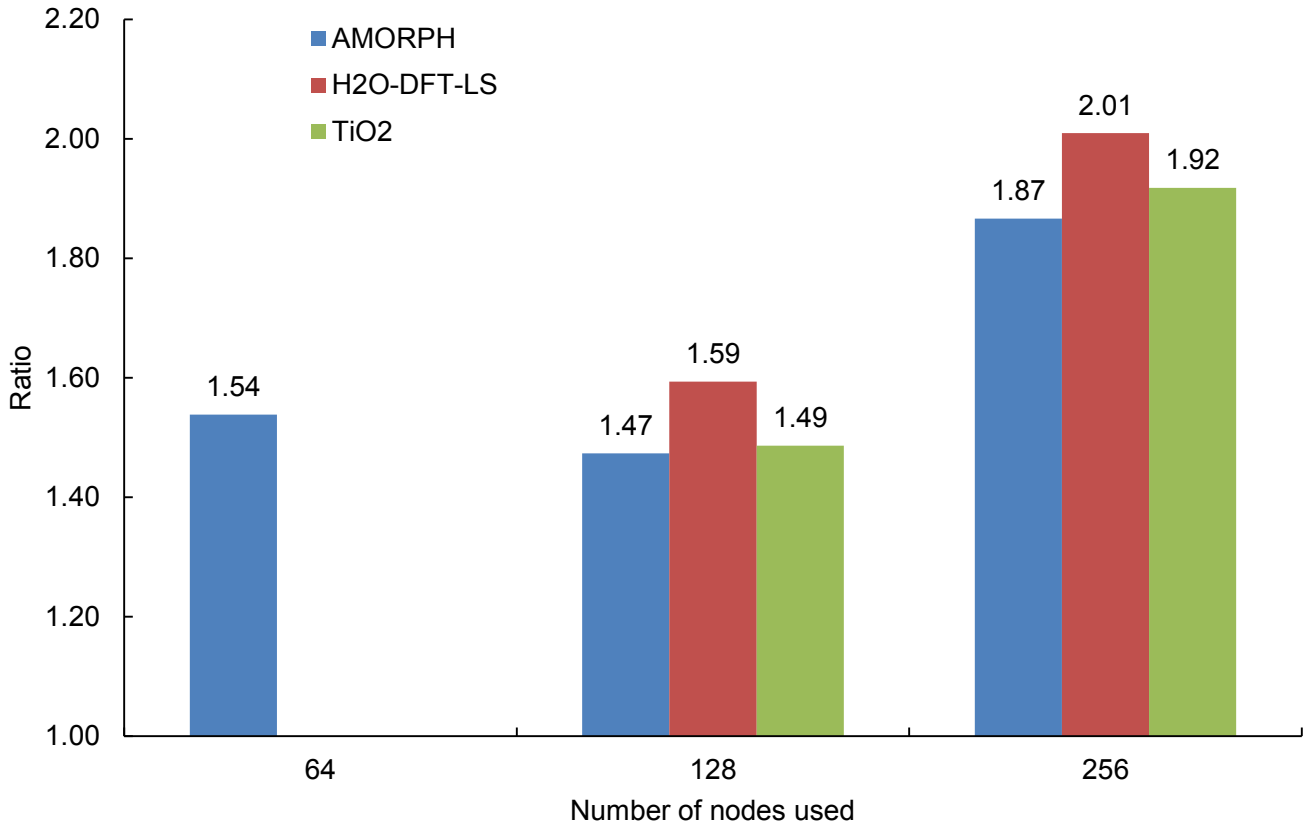

epcc

THE UNIVERSITY OF EDINBURGH

Summary & Outlook

- CP2K performance has increased steadily year by year
  - Hardware, software and algorithms all important
- Development has followed architectural trends
  - Multi-core -> OpenMP
  - Heterogeneous nodes -> CUDA
  - Work on Xeon Phi port ongoing
- Collaborative development (co-design?) model
  - End-users, code authors, HPC centres, vendors
  - Funding from PASC, IPCC, ARCHER eCSE

Acknowledgements

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.

This work used the ARCHER UK National Supercomputing Service (http://www.archer.ac.uk)

epcc | THE UNIVERSITY OF EDINBURGH

Acknowledgement

We are grateful to CSCS for giving us access to and supporting our use of a wide range of HPC systems.

The first two authors are supported by the Engineering and Physical Sciences Research Council 'CP2K-UK' project (grant number EP/K038583/1)

epcc | THE UNIVERSITY OF EDINBURGH

Acknowledgements

Special thanks to Prof. Jurg Hutter and Prof. Joost VandeVondele for historical benchmark data and access to compute time for benchmarking and code development.

Thanks for your attention, and…

…any questions?

epcc | THE UNIVERSITY OF EDINBURGH


epcc
THE UNIVERSITY OF EDINBURGH