# High Performance Computing with CP2K

Iain Bethune - ibethune@epcc.ed.ac.uk

## Background

CP2K [1] is a freely available atomistic simulation program, able to study a wide range of molecular and bulk materials with methods from classical potentials through Density Function Theory (DFT) and QM/MM to Hartree-Fock and Møller-Plesset 2nd order perturbation theory (MP2). In particular, CP2K contains a highly efficient linear-scaling DFT implementation based on the Gaussian and Plane Waves (GPW) method [2] known as QUICKSTEP. Built on top of these force evaluation methods are many different tools such as molecular dynamics, Monte Carlo, geometry optimisation, Nudged Elastic Band (NEB), path integrals and metadynamics.

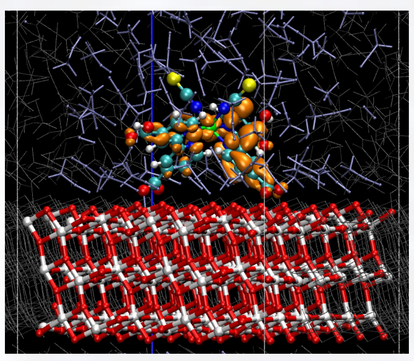

Figure 1: Dye-sensitised Solar Cell treated using CP2K. Image from J. VandeVondele.

Developed for over 10 years by a distributed team led by Prof. Jürg Hutter (Uni. Zurich) and Prof. Joost VandeVondele (ETH Zurich), CP2K contains over 800,000 lines of Fortran 95 code, including a sophisticated mixed-mode parallelisation strategy, scalable sparse matrix multiplication library and GPGPU support. It achieves excellent performance on HPC architectures from linux clusters through to the Cray XE/XK and IBM BlueGene.

CP2K is the third most heavily-used code on the UK national service HECToR and dominates usage on the Cray supercomputers at CSCS, the Swiss national HPC centre. It was also found to be popular among users of the PRACE Research Infrastructure.

[1] CP2K Website, www.cp2k.org
[2] Quickstep: fast and accurate density functional calculations using a mixed Gaussian and plane waves approach, VandeVondele *et al*, Comp. Phys. Comm. 167, 103 (2005)

## Parallelisation and Optimisation

Since 2008, Iain Bethune has led work at EPCC to optimise and extend the parallelisation strategy in CP2K to support leading-edge scientific applications on modern HPC systems. We have minimised communication in key kernels of the QUICKSTEP algorithm – transforming from a Gaussian to a Plane Wave basis, and improved the efficiency of Fast Fourier Transformations (FFT) and real-space grid operations. In addition, we have implemented a new distributed load balancing algorithm, taking advantage of the multi-grid hierarchy to rebalance work by re-ordering MPI ranks at each grid level. For non-homogenous problems this can provide significant speedups – with a non-periodic MD calculation of a 648-atom water cluster in vacuo, a 341% speedup was obtained on 1024 cores of a Cray XT4 [3].

From 2010 onwards, we have concentrated on the implementation of a mixed-mode MPI and OpenMP parallelisation strategy, which is now complete through all performance-critical areas of the code. For current multi-core CPU architectures, using OpenMP within a compute node can often provide better performance than pure MPI by reducing the amount of communication required, limiting the impact of poorly-scaling algorithms, and providing finer-grained load balance. Particularly on large core counts, using mixed-mode can result in significant performance benefits (Fig 2).

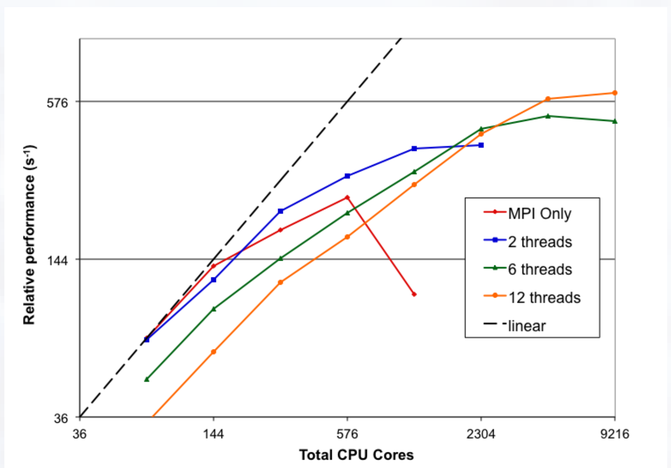


Figure 2: Performance of CP2K on a Cray XT5 (12 cores per node AMD Opteron), 216 $H_2O$ DFT MD calculation, using MPI and OpenMP

[3] Improving the Performance of CP2K on the Cray XT, I. Bethune, Proceedings of the Cray User Group (2010)

## Materials Science at the Petascale

While the performance of CP2K using QUICKSTEP DFT has increased steadily, effort is now being concentrated on extending CP2K's capabilities with more accurate, and more computationally expensive, methods to take advantage of the HPC hardware capable of PetaFLOP/s performance available today in Europe. In 2008 hybrid functionals (Hartree-Fock Exchange) was added to CP2K [4], and by using mixed-mode parallelisation, excellent performance and scalability was obtained (Fig 3).

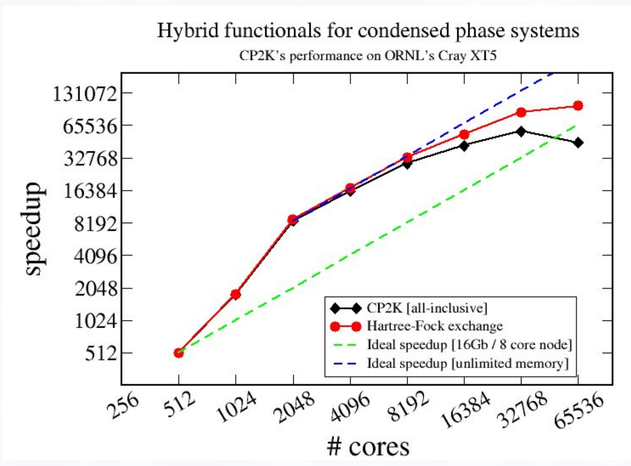


Figure 3: Scalability of HFX calculation of LiH crystal (216 atom cell). Image courtesy J. VandeVondele.

In 2010, a new sparse matrix multiplication library DBCSR (Distributed Block Compressed Sparse Row) was developed at Zurich and optimised with the assistance of EPCC [5]. Again distributed memory parallelisation using MPI was augmented by OpenMP threading, and the addition of an optional CUDA GPGPU implementation continued to provide good performance and scalability up to 10,000s of cores. DBCSR provides the basis for a new linear-scaling DFT implementation based solely on sparse matrix multiplication [6], successfully scaling to 2 million atoms on 46,656 cores – a truly groundbreaking capability.

Most recently, the implementation of Møller-Plesset 2nd order perturbation theory (MP2) based on the GPW method has been added to CP2K [7]. EPCC is leading work under the PRACE project to optimise the key computational kernels for this method. We have developed an auto-tuning framework which allows the code to be adapted efficiently to specific CPU architectures, allowing portable performance across a range of HPC platforms.

Also under the PRACE project, we have deployed and tested CP2K on the CURIE (Bull/Intel Cluster) and Jugene (BlueGene/P) systems, in addition to the Cray and Linux systems used for code development and testing. We believe that CP2K will continue to be a key code for the European materials science and computational chemistry communities and we will continue to improve and extend the code as High Performance Computing moves from Petascale to Exascale over the next 10 years.

[4] Robust Periodic Hartee-Fock Exchange for Large-Scale Simulations Using Gaussian Basis Sets, Guidon *et al*, J. Chem. Theory Comput. (2009)
[5] CP2K – Sparse Linear Algebra on 1000s of cores, I. Bethune, 2012
[6] Linear Scaling Self-Consistent Field Calculations with Millions of Atoms in the Condensed Phase, VandeVondele *et al*, J. Chem. Theory Comput. (2012)
[7] Second Order Møller-Plesset Pertubation Theory in the Condensed Phase: An Efficient and Massively Parallel Gaussian and Plane Waves Approach, Del Ben *et al*, In Preparation

## Acknowledgements

We have been financially supported by EPSRC via the HECToR distributed Computational Science and Engineering (dCSE) support programme, and the PRACE project funded in part by the EU's 7th Framework Programme.

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.

The support of Prof. Jürg Hutter and Prof. Joost VandeVondele by providing access to HPC systems at CSCS is also gratefully acknowledged.

epcc www.epcc.ed.ac.uk THE UNIVERSITY OF EDINBURGH