CP2K: Capability and Scalability

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 this are a wide range of tools such as molecular dynamics, Monte Carlo, geometry optimisation, Nudged Elastic Band, path integrals and metadynamics.

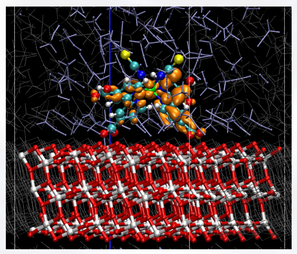

Figure 1: Dye-sensitised Solar Cell treated using CP2K. Image from J. VandeVondele.

Developed for over 10 years by a distributed team led by Prof. Jürg Hutter (Uni. Zürich) and Prof. Joost VandeVondele (ETH Zürich), 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 second most heavily-used code on the UK national service HECToR, dominates usage on the Cray supercomputers at CSCS, the Swiss national HPC centre, and is also widely used across the various systems 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. This is invaluable for non-homogenous problems – 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 mixed-mode MPI and OpenMP, which is now complete through all performance-critical areas of the code. For current multi-core CPU architectures, using OpenMP within a compute node will 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 at large scale.. OpenMP can also help reduce memory usage for use on systems like BlueGene.

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

Materials Science at Petascale and Beyond

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 2).

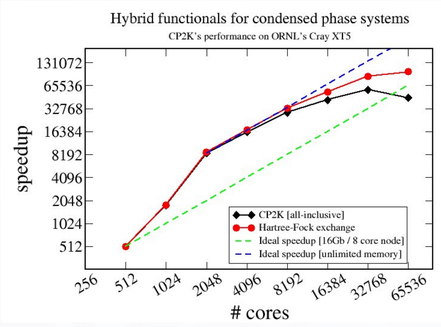

Figure 2: Scalability of HFX calculation of LiH crystal (216 atom cell). Image courtesy J. VandeVondele.

In 2010, a new sparse matrix multiplication library DBCSR (Distributed Block Compressed Sparse Row) was developed at Zürich 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 the implementation of linear-scaling DFT based solely on sparse matrix operations [6], 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]. In MP2 calculations, computation is dominated by the core GPW operation of integrating Gaussian basis functions on real-space grids. Communication is very limited and the code scales linearly to 12,000 cores. We have developed an auto-tuning framework for the local computational kernels which allows the code to be adapted efficiently to specific CPU architectures, enabling portable performance across a range of HPC platforms.

CP2K has been successfully adapted to the major technological trends in HPC over the last decade – multi-core and GPGPU – and with a view towards Exascale we are currently evaluating the use of many-core Intel MIC accelerators (in collaboration with Intel, Cray and CSCS), and are investigating the use of new programming models such as SMPSs, a dynamic task-based parallelisation scheme from BSC planned to be used by the DEEP and Mont-Blanc Exascale projects. While this work is still at an early stage we believe that CP2K will continue to be a key code for the European materials science and computational chemistry communities and so we will continue to improve and extend the code to harness the power of Exascale computing which will be delivered 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*, J. Chem. Theory Comput. (2012)

CP2K-UK: User Support and Training

CP2K-UK is an EPSRC-funded project to foster the growth of an active and vibrant community of CP2K users and developers in the UK. Led by Iain Bethune (EPCC), Prof. Lev Kantorovitch (KCL), Dr. Ben Slater and Dr. Matt Watkins (UCL), our vision is to enable increased uptake and more effective usage of CP2K within the UK atomistic simulation community.

Running from 2013 to 2018, the project will improve the usability of CP2K by the provision of improved documentation and case studies, making it much easier for new users to quickly achieve successful results with their calculations. We will also implement a number of new algorithms in CP2K with the aim of enabling the simulation of larger and more complex systems required to tackle key UK research challenges, for example: understanding the function of large biomolecules, porous systems such as zeolites and Metal-Organic Frameworks (MOFs), design of new materials for fuel cells, nanotechnology including scanning probe microscopies, and next-generation catalysts to create fuels from solar radiation. Finally, we will provide support, training and networking activities for CP2K users, share best practice through a series of annual user group meetings and training courses, and as a result enable the community to not only make effective use CP2K but also to develop and extend the code to support their individual research goals.

The CP2K-UK network already has members from eight research groups around the UK, and we aim to grow! If you use (or would like to use) CP2K in your research please get in touch with us – ibethune@epcc.ed.ac.uk .

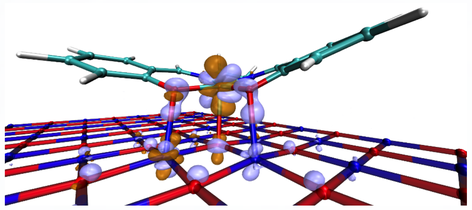

Figure 3: Charge density difference induced by Co-Salen molecule interacting with NiO(100) surface. From A. Schwarz *et al*, J. Phys. Chem. C, 17 1105 (2013). Image courtesy M. Watkins.

Acknowledgements

We have been financially supported by the Engineering and Physical Sciences Research Council 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. CP2K-UK is supported by the EPSRC [grant number EP/K038583/1].

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, Prof. Joost VandeVondele and staff at CSCS for providing access to HPC systems (Cray XE6/XK7, Intel Xeon Phi) is gratefully acknowledged.