

CP2K: Highly scalable atomistic simulation for all


Iain Bethune^{1*}, Gordon Gibb², Lev Kantorovich³, Matt Watkins⁴, Sergey Chulkov⁴, Ben Slater⁵

*iain.bethune@stfc.ac.uk ¹STFC Hartree Centre ²EPCC ³King's College London ⁴University of Lincoln ⁵University College London

Capability

CP2K (<http://www.cp2k.org>) is an popular open-source atomistic simulation program featuring a number of state-of-the-art methods designed for use on modern supercomputers. Building on the QUICKSTEP[1], an efficient approach to Density Functional Theory using a dual Gaussian and Plane Wave basis approach, CP2K has a wide range of functionality:

- Fast Hybrid Functionals based on ADMM [2]
- Post Hartree-Fock methods (MP2, RPA, GW)
- Linear-scaling DFT[3] for 10,000 atom+ systems
- Time-dependent DFT
- Molecular Dynamics, Geometry Optimisation, Vibrational Analysis, ...



With support from an ARCHER eCSE grant, we have added the Non-Equilibrium Green's Function (NEGF) method, enabling quantum transport simulations, such as electronic nano-device modelling and Scanning Tunneling Microscopy.

We have carried out tests on up to 1536 cores on ARCHER. The code is available in the CP2K trunk.


[1] VandeVondele *et al*, Comp Phys Comm 167(2), 2005

[2] Guidon *et al*, JCTC 6(8), 2010

[3] VandeVondele *et al*, JCTC 8(10), 2012

Scalability

CP2K has support for multi- and many-core CPUs via hybrid MPI/OpenMP parallelism as well as NVIDIA GPUs with CUDA. The Distributed Block Compressed Sparse Row (DBC SR) matrix library enables scalability to 1,000s of nodes [4].



Performance on Piz Daint, up to 1024 nodes [5]

We have recently adapted DBCSR for efficient use on Intel Xeon Phi Knight's Landing, and find that the KNL outperforms a Haswell CPU node by up to 24%, but is still

10% slower than an NVIDIA P100 node [6].

Work is currently underway to develop a dynamic parallelism scheme for sparse-matrix multiplication based on OpenMP tasks. Early results show a potential performance improvement of 10% due to reduced thread load imbalance.


[4] Borstnik *et al*, Parallel Computing 40(5-6), 2014

[5] Bethune *et al*, Cray User Group 2014

[6] Bethune *et al*, ParCO 2017

Usability

The CP2K-UK network provides support and training to the CP2K user community in the UK. We offer bespoke hands-on training, classroom tutorial and schools, and tools for use with CP2K.



We implemented a pair of plugins - TETR and LEV00 - for UCSF Chimera to allow easy system setup and visual analysis of results. Both the plugins are available under a BSD license on Github:

https://github.com/gpsgibb/tetr_lev00_Chimera_plugin

Acknowledgements

This work is funded by EPSRC grant EP/K038583/2 and by the embedded CSE programme of the ARCHER UK National Supercomputing Service (<http://www.archer.ac.uk>).

CP2K Summer School 2018

19-22 June, Hartree Centre, Sci-Tech Daresbury
More info at <https://www.cp2k.org/news>



Hartree Centre
Science & Technology Facilities Council