Linear scaling DFT accelerated with GPUs and ML

Ole Schütt[a], Alice Shoshana Jakobovits[b], and Joost VandeVondele[b*]

[a] *DMATL, ETH Zurich*
[b] *CSCS, ETH Zurich*
*Joost.VandeVondele@cscs.ch

Linear scaling density functional theory with a demonstrated capability to simulate millions of atoms has been available for a couple of years in CP2K.[1] With increasing computer resources and refinements in algorithms, such simulations have become quite feasible. With tight-binding basis sets on workstations, with DFT quality basis sets on supercomputers. This progress is enabled by software, in particular a library for sparse matrix matrix multiplication. This enabling library, DBCSR, is now stand-alone and freely available.[2] It is highly efficient, MPI parallel, threaded and GPU-accelerated.[3]

Recently, the challenge of generating optimal, domain specific GPU kernels has been tackled with a combination of just-in-time (JIT) compilation and machine learning (ML). We show how ML parameters reach near optimal performance for 10000s of possible kernels.[4] Finally, we discuss how machine learning can be used to construct adaptive basis sets. This technique is based on the intuitive observation that an optimal basis set must adapt to its atomic environment. This optimal adaptation can be inferred from available data. This technique allows for large speedups (up to 50x), essentially allowing DFT calculations at tight-binding cost.[5]

References

1. J. VandeVondele, U. Borstnik, J. Hutter, *J. Chem. Theory Comput.* (2012), 3565.
2. https://www.cp2k.org/dbcsr/
3. O. Schütt, P. Messmer, J. Hutter and J. VandeVondele, *Electronic Structure Calculations on Graphics Processing Units: From Quantum Chemistry to Condensed Matter Physics* (2016), 173.
4. S. Jakobovits and J. VandeVondele *unpublished*.
5. O. Schütt and J. VandeVondele, *J Chem Theory Comput.* **14** (2018), 4168.