Quantum Chemistry on Supercomputers

Dmytro Bykov, Ashleigh Barnes and Dmitry Lyakh

National Center for Computational Sciences, Oak Ridge National Laboratory (ORNL), 1 Bethel Valley Rd, Oak Ridge, TN 37831, USA

During the last decades, computers have taken an increasingly prominent role in science and technology, reaching the stage where results from computer simulations are often more reliable and faster to achieve than experimental results, especially when the molecules and properties of interest are difficult to study experimentally. Nevertheless, harnessing this modern computing power may take significant effort due to the fact that modern computers have extremely complex architecture. For instance, Summit,[1] currently the biggest supercomputer in the world hosted at Oak Ridge Leadership Computing Facility at ORNL, has 4608 compute nodes. Each node contains 2xCPU (22 cores each) and 6xGPU devices (NVIDIA Volta). Summit has complex memory hierarchy on a node, which includes RAM, HBM2, solid state memory. The 250 PB parallel file system is mounted. Obviously, implementing computational chemistry algorithms on Summit takes careful designing efforts. One effective method to successfully utilize Summit supercomputer for computational chemistry applications is the Divide-Expand-Consolidate (DEC) scheme.[2] The scheme can be applied to evaluate molecular energy and properties in a linear-scaling and embarrassingly parallel manner using a set of local Hartree-Fock molecular orbitals. The essence of the method lies in the fact that all manipulations with the intermediate four-dimensional quantities are carried out independently within small local orbital fragment spaces. The sizes of the orbital fragment spaces are determined in a black-box manner to ensure that the error in the DEC implementation is proportional to a single input threshold, denoted as the fragment optimization threshold (FOT). GPU-acceleration of the DEC scheme is achieved using both OpenACC pragmas and the TAL-SH library,[3] which allows for offloading of tensor contraction operations to GPUs. The scaling behavior, performance and benchmark studies as well as series of showcase calculations using the LS-Dalton program[4] proved the DEC method to be a highly effective tool for approaching large molecular systems.

ACKNOWLEDGMENTS: This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

[1] https://www.olcf.ornl.gov/summit/
[2] T. Kjærgaard et al., *WIREs Comput. Mol. Sci.*, 7, e1319. doi:10.1002/wcms.1319, (2017)
[3] https://github.com/DmitryLyakh/TAL_SH
[4] K. Aidas et al., *WIREs Comput. Mol. Sci*., 4, 269 (2014)