REDUCING THE COMPUTATIONAL COST OF COUPLED CLUSTERY THEORY

Sara Kokkila Schumacher, Stanford University
2015-2016 Graduate Fellow

RESEARCH SUMMARY

Electronic structure methods are capable of describing the electron arrangement in molecules. In order to accurately describe the absorption spectra of molecules, it is necessary to use electronic structure methods. Unfortunately, few electronic structure methods are capable of efficiently calculating the excited states of molecules. The equations-of-motion second-order approximate coupled-cluster singles and doubles (EOM-CC2) method is an accurate electronic structure method for studying the absorption of light by molecules. [1,2] However, this approach costs $O(N^5)$ where N is the number of basis functions used to represent a molecule. In order to improve the efficiency of this method, the tensor hypercontraction (THC) approximation is introduced to reduce the computational cost to $O(N^4)$. [3, 4, 5]

Graphical processing units (GPUs) have proven to be useful for accelerating quantum chemistry codes. [6] In order to further improve the efficiency of THC-EOM-CC2, we developed a GPU accelerated version of the grid based THC-EOM-CC2. Unlike other GPU implementations of coupled-cluster theory, which often focus on using GPUs only to accelerate the most computationally demanding step, the algebraic flexibility of THC can be used to form contributions to THC-EOM-CC2 "on-the-fly" on the GPU. Because the THC formation allows us to represent high order tensors as a product of lower order tensors, there are many alternative ways to implement THC-CC2 and THC-EOM-CC2. In order to parallelize these methods, we developed THC-CC2 and THC-EOM-CC2 algorithms that focus on blocking over a grid point index for a series of intermediates. This allows for parallelization at multiple levels. Figure 1 demonstrates the general parallelization scheme for THC-CC2 where Parts 1-5 each represent a series of intermediates formed through Compute Unified Device Architecture (CUDA) kernels and NVIDIA CUDA Basic Linear Algebra Subroutines (cuBLAS). The THC-EOM-CC2 method can be parallelized in a similar manner.

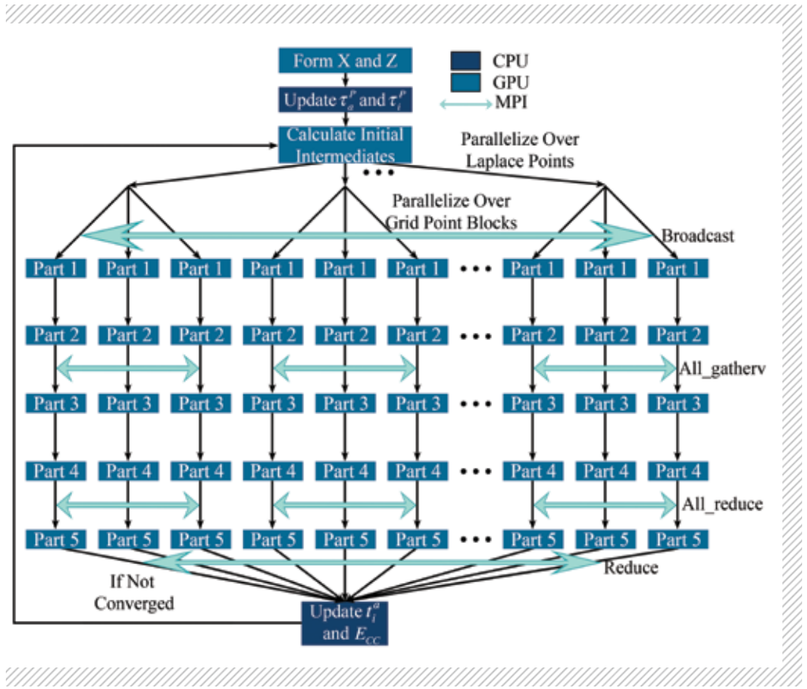

FIGURE 1: Parallelization outline for the GPU-based THC-CC2 algorithm. The five main parts represent a series of linear algebra and tensor algebra routines performed on GPUs.

Timings for the GPU accelerated THC-CC2 and THC-EOM-CC2 approaches are shown in Figure 1. This shows that the recently developed GPU accelerated algorithms do achieve $O(N^4)$ cost for both the ground state calculation (THC-CC2) and the excited state calculations (THC-EOM-CC2). These timings were computed on a single GPU on a single node of Blue Waters.

We have already shown that the THC approach can be used to reduce the computational cost of different electronic structure methods. This work is the first to demonstrate how one can take advantage of the algebraic flexibility offered by THC to develop massively parallel GPU accelerated THC based electronic structure methods. This approach can be used to extend the applicability of electronic structure methods to larger chemical systems and to rapidly compute energies of sampled geometries from dynamics simulations. We expect this to lead to the development of other THC based electronic structure methods in different quantum chemistry software packages.

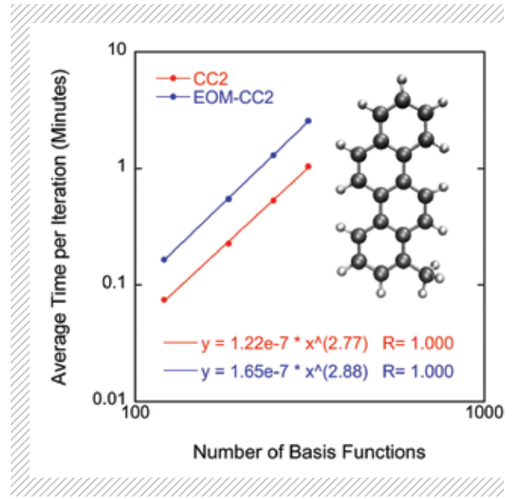

FIGURE 2: Timings of THC-CC2 and THC-EOM-CC2 iterations for a series of methylated polycyclic aromatic hydrocarbons. These timings were performed on a single node of Blue Waters. Each computation was performed in triplicate and the averaged timings are shown here.

WHY BLUE WATERS

Access to Blue Waters allowed for the development of the MPI enabled GPU accelerated THC-EOM-CC2 approach. With Blue Waters, we were able to rapidly develop and test different THC-EOM-CC2 implementations. As our first foray into combining MPI and GPU parallelization approaches, the Blue Waters project staff offered helpful advice during the design process. Additionally, we are using the THC-EOM-CC2 code to study the absorption of light by different protein chromophores. This requires sampling many different geometries and evaluating the THC-EOM-CC2 energies. The use of Blue Waters is necessary to compute enough geometries, ground state energies, and excited state energies to accurately describe the absorption spectra of protein chromophores.

PUBLICATIONS AND DATA SETS

Kokkila Schumacher, S.I.L, E. G. Hohenstein, R. M. Parrish, L.-P. Wang, and T. J. Martínez, Tensor Hpercontraction Second-Order Møller-Plesset Perturbation Theory: Grid Optimization and Reaction Energies, *J. Chem. Theory Comput.*, 11, pp. 3042-3052 (2015).

Sara Kokkila Schumacher graduated fall 2016 from Stanford with a Ph.D. in Chemistry, where she worked with advisor Todd Martinez. She is currently a postdoctoral researcher in high performance computing at IBM. She ultimately aims to become a professor where she can develop a course that introduces high-performance computing to science majors.