GA

REDUCING THE COMPUTATIONAL COST OF COUPLED CLUSTERY THEORY

Sara Kokkila Schumacher, Stanford University
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

The retinal chromophore, found in the photoactive retinylidene proteins, is a challenging system for many of the excited state methods that are currently available. This chromophore plays an important role in vision in humans and other species. In order to accurately and efficiently simulate the absorption spectra of retinal models, our project aims to develop a parallel implementation of the tensor hypercontraction equation-of-motion second-order approximate coupled cluster singles and doubles (THC-EOM-CC2) method. We implemented an MPI (message passing interface) version of THC-EOM-CC2 that is accelerated by the graphics processing units (GPUs) available on Blue Waters. We successfully applied our new approach to several retinal models. Our initial findings indicate that THC-EOM-CC2 performs better than other available excited state methods that are capable of treating this system size.

RESEARCH CHALLENGE

Retinylidene proteins are photosensitive proteins that serve a variety of biological functions and have been found in all domains of life. [1] For example, the rhodopsin protein and the visual pigments found in human eyes enable vision. [2] Other retinylidene proteins can induce movement in certain types of bacteria or act as light-induced pumps. Retinylidene proteins consist of an opsin protein and the retinal chromophore. Through the use of Blue Waters, we aim to simulate the absorption spectra of different retinal chromophore models at an unprecedented level of accuracy.

Simulating the excited states of the retinal chromophore necessitates the use of electronic structure methods since simulation requires an accurate description of the electronic wave function. Unfortunately, the accuracy and computational demands of an electronic structure method are often at odds with each other. The second-order approximate coupled cluster singles and doubles (CC2) method can be extended to excited states through the equation-of-motion formalism (EOM). [3,4] The EOM-CC2 method is robust against many of the issues found in inexpensive excited state alternatives such as configuration interaction singles or time-dependent density functional theory. [5] However, studying the retinal chromophore with canonical EOM-CC2 is infeasible with typical computational resources. This is because the runtime of canonical EOM-CC2 increases formally on the order of $O(N^5)$ where N is related to the size of the chemical system.

METHODS & CODES

In an effort to reduce the computational cost of EOM-CC2 and other quantum chemistry methods, we developed the tensor hypercontraction (THC) approximation. [6] We applied the THC approximation to EOM-CC2 (THC-EOM-CC2) and showed that this lowered the cost of the calculation from an order of $O(N^5)$ to $O(N^4)$. [7,8] To further improve the efficiency of THC-EOM-CC2, we redesigned the algorithms to take advantage of parallelism. This included developing a code that was parallelized at a higher level to take advantage of multiple compute nodes and at a finer level to take advantage of acceleration with GPUs. The parallel THC-EOM-CC2 approach was developed in TeraChem. [9]

We are working with the grid-based variant of the THC approximation, which allows the new algorithm to block over grid point indices. Additionally, we use a Laplace transformation to express certain terms by numerical quadrature in this method. In the parallel implementation, each MPI task is assigned a quadrature point and set of grid points. The THC-EOM-CC2 approach is memory-intensive, and by blocking over grid points on an MPI task, we can exploit this level of parallelism while ensuring that the computations fit within the limited memory available on the GPU. This new approach allows us to extend the THC-EOM-CC2 method to system sizes that are challenging for canonical EOM-CC2.

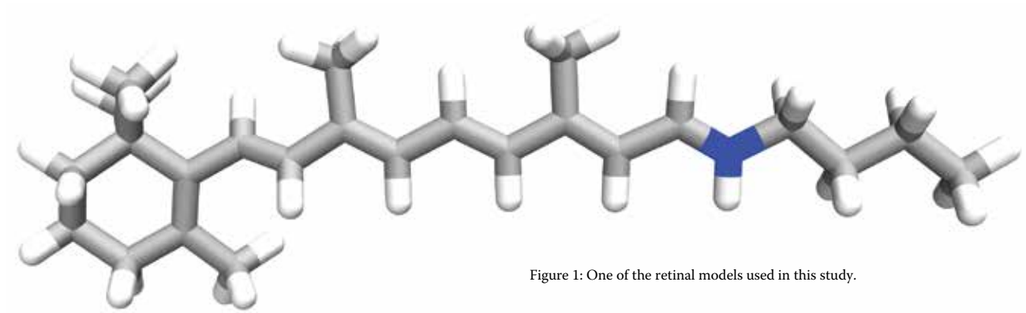

Figure 1: One of the retinal models used in this study.

RESULTS & IMPACT

Our approximation, combined with the parallelization approach developed in this work, offers a new way to parallelize and improve the efficiency of quantum chemistry methods. In addition to its use in CC2 and EOM-CC2, the THC approximation has been shown to reduce the computational complexity of other electronic structure methods. [6] While this work was focused on parallelizing THC-CC2 and THC-EOM-CC2, the general parallelization design outlined can be applied to other quantum chemistry methods using the THC approximation. We expect this new approach will be used to improve the efficiency of many other electronic structure methods in different quantum chemistry software.

The development of the MPI-enabled and GPU-accelerated THC-EOM-CC2 code and the use of Blue Waters allow us to study the absorption spectra of several model retinal chromophores. We find that for a fixed problem size, the parallel implementation of THC-EOM-CC2 scales superlinearly in the regime where we parallelize only Laplace quadrature points across compute nodes. In the regime where we distribute both quadrature points and grid point blocks across compute nodes, we see a computational speedup that levels off as the number of compute nodes increases. We are currently comparing the results of THC-EOM-CC2 to other excited state methods. Our initial findings indicate that the THC-EOM-CC2 method more accurately describes the excited states of the retinal models compared to the results from time-dependent density functional theory. The THC-EOM-CC2 method will be a useful tool for studying the excited states of chemical systems that are otherwise challenging for the available excited state methods.

WHY BLUE WATERS

Access to Blue Waters allowed us to develop the MPI-enabled and GPU-accelerated THC-EOM-CC2 method. Blue Water's code development environment and tools helped us test and debug different implementations more rapidly. This work represents the first time we have tried to combine MPI and GPU acceleration for a method within TeraChem. The Blue Waters project staff offered valuable insight into code development for high-performance computing systems. Additionally, studying the absorption spectra of the retinal models requires sampling different conformations of each retinal model. Blue Waters enabled us to calculate the ground state and multiple excited states of many different configurations.

PUBLICATIONS AND DATA SETS

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

Sara Kokkila Schumacher received a Ph.D. in chemistry from Stanford University in January 2017. She currently is a postdoctoral researcher in high-performance computing at IBM.