Annual Report for Blue Waters Allocation

- **Project Information**
  - Blue Water Professorship
  - So Hirata, University of Illinois at Urbana-Champaign
  - sohirata@illinois.edu

- **Executive summary (150 words)**

We have fully developed two highly scalable and systematically accurate electronic structure methods. One is the embedded-fragment method for various structural, thermodynamic, and spectroscopic properties of molecular crystals, amorphous solids, and liquids. The other is the Monte Carlo many-body perturbation and Green's function methods, running efficiently on thousands of CPUs or hundreds of GPUs or more. With Blue Waters, we aim to perform grand-challenge large-scale applications of both methods.

- **Description of research activities and results**

  - *Key Challenges:* Our group made two breakthroughs in computational chemistry for large molecules, solids, and liquids. (1) One is the method that allows such high-level calculations to be applied to an infinitely extended molecular crystals, amorphous solids, and liquids, including nature's most important solids and liquids such as ice and liquid water. It divides them into fragments embedded in their electrostatic field and treat these fragments by well-developed molecular theories and software in a highly parallel algorithm. (2) The other weds systematic (i.e., *ab initio*) many-body electronic structure methods with quantum Monte Carlo, enabling massively parallel, systematically accurate electronic structure calculations for large molecules and solids. With Blue Waters, we aim to perform grand-challenge applications to demonstrate these methods' full capacity. The applications may include (1) computationally explaining the anomalous thermodynamic properties of ice Ih such as thermal contraction, anomalous volume isotope effect, and pressure-induced glass transition to high-density amorphous phase as well as (2) near-exact calculations of ionization energies and electron affinities of large conjugated molecules, which may be used as advanced organic materials for optoelectronic devices.

  - *Why it Matters:* Project (1) advances geochemistry, astrophysics, and planetary science where probing high-pressure phases of ices of atmospheric species on Earth or other planets are important but experimentally difficult and expensive. The systems studied by project (2) include solids that serve as bases of advanced materials such as bulk heterojunction organic solar cells, batteries, sensors, smart windows, field-effect transistors, and light emitting diodes. The optoelectronic parameters are the quantities of prime importance determining the solids performance and functions, but the usual density-functional approximations are

known to be poor for these properties. Here, our new method is uniquely useful and accurate. The impact of this research is broadly in the area of energy science and technology.

- *Why Blue Waters:* Today's workhorse computational methods for solids (density-functional methods) and liquids (classical molecular dynamics) are routine on a small computer cluster, but with limited accuracy. Therefore, the most meaningful use of Blue Waters in this area is to enhance the accuracy rather than enlarge the system size (which is already formally infinite). In electronic structure theory, this means switching from density-functional methods and empirical force fields to *ab initio* theories, which solve the fundamental equation of motion of chemistry rigorously and in systematic approximations with controlled errors. However, conventional matrix-algebra algorithms of *ab initio* theories are fundamentally non-scalable and extremely expensive. The aforementioned combination of *ab initio* theories with quantum Monte Carlo and the embedded-fragment methods are among the few that may be realistically and usefully deployable on the large number of processors available on Blue Waters, enabling formerly unthinkable predictively accurate calculations for large systems.

- *Accomplishments:*

(1) **Embedded-fragment MP2.**

***Ab initio* MP2 study of the anomalous thermodynamic properties of ice.** Ice Ih displays unusual thermodynamic properties such as negative thermal expansion at low temperatures and the anomalous volume isotope effect, which renders the volume of D_2O ice greater than that of H_2O ice. The former is related to the well-known negative slope of water's melting curve and the pressure-induced collapse to an amorphous solid at low temperatures. We use the embedded-fragment second-order many-body perturbation (MP2) method and the quasiharmonic approximation to determine the crystal structures that minimize the Gibbs energy of proton-disordered ice Ih at the temperatures and pressures spanning the entire domain of its stability. The calculations correctly predict the thermal contraction at low temperatures followed by expansion at higher temperatures, confirming this behavior originates from the negative Grüneisen parameters of acoustic phonons. They also reproduce the anomalous

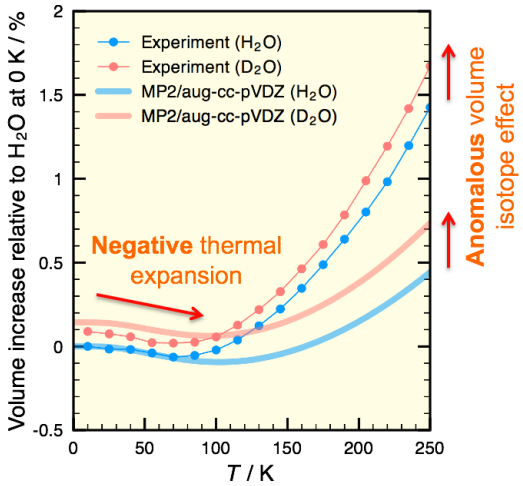

volume isotope effect, but only when atomic partial charges (but not molecular dipole moments) are used in the embedding field. They confirm that it is caused by the large volume-contracting effect of the OH/OD stretching phonons, but reveal that its sign and magnitude are the result of subtle cancellation among closely competing effects of all phonons. The equally unusual temperature dependence of the volume isotope effect, on the other hand, is reliably reproduced by theory and is shown to originate from the librational phonons. We furthermore observe discontinuous changes in the volume and the ratio of lattice constants as well as softening of acoustic phonons at ca. 1.5 GPa, which points to the pressure-induced amorphization. We also perform *ab initio* MP2 molecular dynamics to computationally and directly detect the mechanical instability leading to this amorphization. This work has been published in *J. Chem. Phys.* in 2016 and chosen as a JCP Editors' Pick.

(2) **Monte Carlo MP2.**

Monte Carlo explicitly correlated MP2. *Ab initio* methods for electron correlation, such as MP2, are well known to be slowly convergent with respect to the size of one-electron basis set. The slow convergence is due to the inability of such a basis set to describe the cusped form of the wave function at electron-electron coalescence. The cure is to include two-electron basis functions, particularly, a function that explicitly depends on interelectronic distance (R12), which can describe the cusps exactly. The problem of the methods that include such basis functions, known as explicitly correlated or R12 methods, are that the formalism now involves high-dimensional integrals, that, if such high-dimensional integrals are avoided by the resolution-of-the-identity (RI) approximation, a large auxiliary basis set becomes necessary, that one has to derive and implement analytical Gaussian integrals of R12 basis function, which is difficult and possible only for the simplest such functions, and that their matrix-algebra algorithm is not scalable. The Monte Carlo MP algorithm solves all of these problems. A Metropolis Monte Carlo integration performs particularly well (relative to cubature) as the dimension is higher and hence RI with an auxiliary basis is unnecessary; it works for any functional form of R12 basis; it is fundamentally more scalable than matrix-algebra algorithms with respect to both computer size and system size. We have implemented

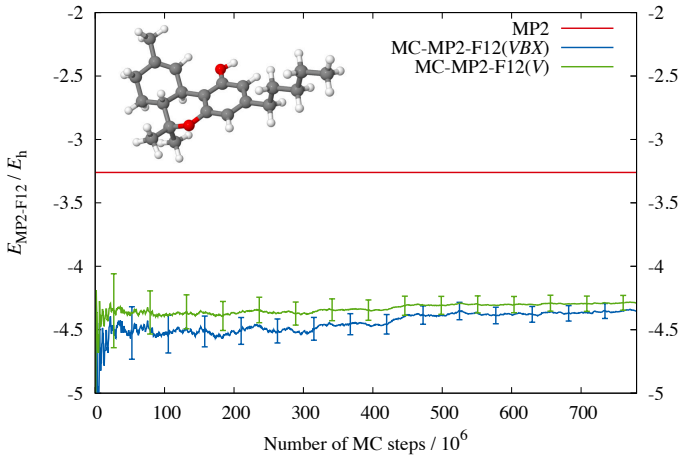

explicitly correlated Monte Carlo MP2, which has at least one-rank lower cost scaling with system size ($O(n^4)$ as opposed to usual $O(n^5)$ where n is the number of orbitals) and naturally parallel. The method gives the MP2 reaction enthalpies that are within a few kJ/mol of the complete basis set limits with just the aug-cc-pVDZ basis set. This work has been published in *J. Chem. Phys.* in 2016.

Using this method, we analyze the performance of 17 different correlation factors in explicitly correlated Monte Carlo second-order many-body perturbation calculations for energies. We find that highly performing correlation factors have near-universal shape and size in the short range of electron-electron distance ($0 < r_{12} < 1.5$ a.u.). The long-range behavior ($r_{12} > 1.5$ a.u.) is insignificant insofar as the factor becomes near constant, leaving an orbital expansion to describe decoupled electrons. An analysis based on a low-rank Taylor expansion of the correlation factor seems limited, except that a negative second derivative with the value of around –1.3 a.u. correlates with high performance. This work will be submitted to *Chem. Phys. Lett.*

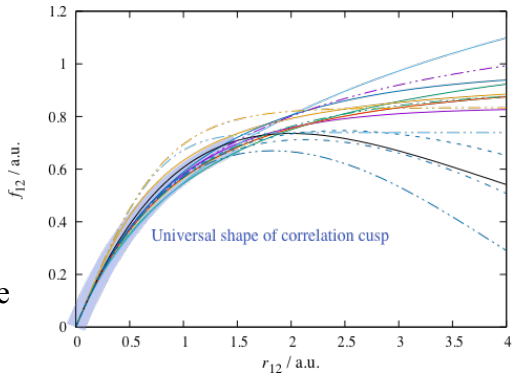

Monte Carlo MP2 on many GPUs. In the Monte Carlo MP2 method, the long sum-of-product matrix expression of the MP2 energy, whose literal evaluation may be poorly scalable, is recast into a single high-dimensional integral of functions of electron pair coordinates, which is evaluated by the scalable method of Monte Carlo integration. The sampling efficiency is further accelerated by the redundant-walker algorithm, which allows a maximal reuse of electron pairs. Here, a multitude of GPUs offers a uniquely ideal platform to expose multilevel parallelism: fine-grain data-parallelism for the redundant-walker algorithm in which millions of threads compute and share orbital amplitudes on each GPU; coarse-grain instruction-parallelism for near-independent Monte Carlo integrations on many GPUs with few and infrequent interprocessor communications. While the efficiency boost by the redundant-walker algorithm on CPUs grows linearly with the number of electron pairs and tends to saturate when the latter exceeds the number of orbitals, on a GPU it grows quadratically before it increases

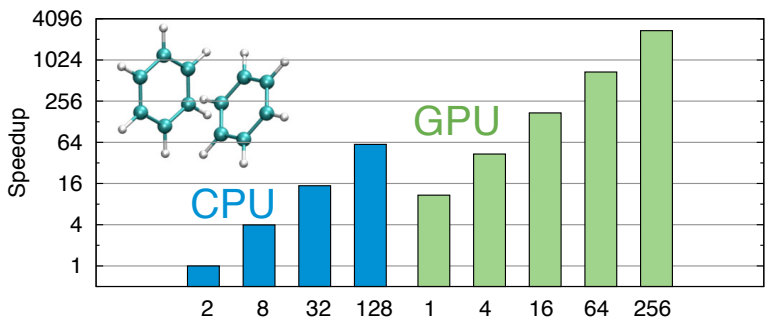

linearly and then eventually saturate at a much larger number of pairs. This is because the orbital constructions are nearly perfectly parallelized on a GPU and thus completed in a near-constant time regardless of the number of pairs. In consequence, a Monte Carlo MP2/cc-pVDZ calculation of a benzene dimer is 2,700 times faster on 256 GPUs (256 XK nodes) (using 2,048 electron pairs) than on 16 CPU cores (1 XE node). We also numerically determine that the cost to achieve a given relative statistical uncertainty in a Monte Carlo MP2 energy increases as $O(n^3)$ or better with system size n, which may be compared with the $O(n^5)$ scaling of the conventional implementation of deterministic MP2. We thus establish Monte Carlo MP2's scalability with both system and computer sizes. This work has been published in *J. Chem. Theory Comput.* in 2016.

- **List of publications and presentations associated with this work**

S. Y. Willow, M. R. Hermes, K. S. Kim, and S. Hirata, *Journal of Chemical Theory and Computation* **9**, 4396-4402 (2013), "Convergence acceleration of parallel Monte Carlo second-order many-body perturbation calculations using redundant walkers."

S. Y. Willlow, M. A. Salim, K. S. Kim, and S. Hirata, *Scientific Reports* **5**, 14358 (2015) (14 pages), "*Ab initio* molecular dynamics of liquid water using embedded-fragment second-order many-body perturbation theory towards its accurate property prediction."

M. A. Salim, S. Y. Willow, and S. Hirata, *The Journal of Chemical Physics* **144**, 204503 (2016), "Ice Ih anomalies: Thermal contraction, anomalous volume isotope effect, and pressure-induced amorphization."

S. Hirata, K. Gilliard, X. He, M. Keçeli, J. Li, M. A. Salim, O. Sode, and K. Yagi, an invited book chapter in *Fragmentation: Toward Accurate Calculations on Complex Molecular Systems* edited by M. S. Gordon (in press, 2017), "*Ab initio* ice, dry ice, and liquid water."

A. E. Doran and S. Hirata, *Journal of Chemical Theory and Computation* **12**, 4821-4932 (2016), "Monte Carlo MP2 on many graphical processing units."

C. M. Johnson, A. E. Doran, J. Zhang, E. F. Valeev, and S. Hirata, *The Journal of Chemical Physics* **145**, 154115 (2016), "Monte Carlo explicitly correlated second-order many-body perturbation theory."

C. M. Johnson, S. Hirata, and S. Ten-no, *Chemical Physics Letters* [A. H. Zewail Commemorative Issue] (to be submitted, 2017), "Correlation factors in explicitly correlated methods."

Work benefitted from Blue Waters has been highlighted in invited talks by the PI in the following conferences and university seminars: American Conference on

Theoretical Chemistry (2014); Low-Scaling and Unconventional Electronic Structure Techniques Conference (2014); International Symposium on "Quantum Chemistry for Extended Systems" in honor of Professor Jean-Marie André (2014); Seminar, Department of Chemistry and Biochemistry, University of Arkansas (2014); Seminar, Department of Chemistry, Virginia Tech (2014); Blue Waters Symposium (2014); Sanibel Symposium (2015); CECAM meeting (2015); Blue Waters Symposium (2015); RAEST (2015); Pacifichem 2015 (2015); Low-Scaling and Unconventional Electronic Structure Techniques Conference (2016); Molecular Quantum Mechanics 2016; International Symposium of Theoretical Chemical Physics (2016); Molecular Electronic Structure in Buenos Aires (2016); Southeast Regional ACS Meeting (2016); Seminar, Institute for Molecular Science (2016); Seminar, Kobe University (2016); Seminar, Waseda University (2016); Seminar, Keio University (2016).

- **Plan for next year**

We request 200,000 node-hours (100,000 node-hours on XE and 100,000 node-hours on XK).

Grand-challenge applications of Monte Carlo MP and Green's function (GF) calculations are well underway for large molecules which conventional matrix-based algorithms cannot handle. We plan to use the entirety of the 100,000 XK node-hours for this purpose using the GPU-based algorithm with demonstrated speedup in excess of 2,700-fold on 256 XK nodes relative to 1 XE node. We will also use 50,000 XE node-hours for calculations of Monte Carlo MP/GF methods that are yet to be ported to GPUs.

We will use the remaining 50,000 XE node-hours on embedded-fragment calculations on liquid water and other molecular liquids and for the development of multi-time-step approach as well as the extension of the method to molecular clusters and crystals with charge transfer. The formulation and pilot implementation of the latter are complete and parallelization is expected to be straightforward.

The usage schedule is, Q1: 25%, Q2: 25%, Q3: 25%, Q4: 25%.