

FINAL REPORT for DE-FG05-08OR23335

**Quantum Monte Carlo Endstation for Petascale Computing**

University of Illinois Urbana-Champaign

PI David Ceperley (Physics and NCSA)

Dates 1/1/2008 -12/31/2010 (12 month extension)

The major achievements enabled by QMC Endstation grant include

- Performance improvement on clusters of x86 multi-core systems, especially on Cray XT systems
- New and improved methods for the wavefunction optimizations
- New forms of trial wavefunctions
- Implementation of the full application on NVIDIA GPUs using CUDA

The scaling studies of QMCPACK on large-scale systems show excellent parallel efficiency up to 216K cores on Jaguarpf (Cray XT5). The GPU implementation shows speedups of 10-15x over the CPU implementation on older generation of x86.

We have implemented hybrid OpenMP/MPI scheme in QMC to take advantage of multi-core shared memory processors of petascale systems. Our hybrid scheme has several advantages over the standard MPI-only scheme.

- Memory optimized: large read-only data to store one-body orbitals and other shared properties to represent the trial wave function and many-body Hamiltonian can be shared among threads, which reduces the memory footprint of a large-scale problem.
- Cache optimized: the data associated with an active Walker are incache during the compute-intensive drift-diffusion process and the operations on an Walker are optimized for cache reuse. Thread-local objects are used to ensure the data affinity to a thread.

- Load balanced: Walkers in an ensemble are evenly distributed among threads and MPI tasks. The two-level parallelism reduces the population imbalance among MPI tasks and reduces the number of point-to-point communications of large messages (serialized objects) for the Walker exchange.
- Communication optimized: the communication overhead, especially for the collective operations necessary to determine ET and measure the properties of an ensemble, is significantly lowered by using less MPI tasks.

The multiple forms of parallelism afforded by QMC algorithms make them ideal candidates for acceleration in the many-core paradigm. We presented the results of our effort to port the QCMPACK simulation code to the NVIDIA CUDA GPU platform. We restructured the CPU algorithms to express additional parallelism, minimize GPU-CPU communication, and efficiently utilize the GPU memory hierarchy. Using mixed precision on GT200 GPUs and MPI for intercommunication and load balancing, we observe typical full-application speedups of approximately 10x to 15x relative to quad-core Xeon CPUs alone, while reproducing the double-precision CPU results within statistical error.

We developed an all-electron quantum Monte Carlo (QMC) method for solids that does not rely on pseudopotentials, and used it to construct a primary ultra-high-pressure calibration based on the equation of state of cubic boron nitride. We computed the static contribution to the free energy with the QMC method and obtained the phonon contribution from density functional theory, yielding a high-accuracy calibration up to 900 GPa usable directly in experiment. We computed the anharmonic Raman frequency shift with QMC simulations as a function of pressure and temperature, allowing optical pressure calibration. In contrast to present experimental approaches, small systematic errors in the theoretical EOS do not increase with pressure, and no extrapolation is needed. This all-electron method is applicable to first-row solids, providing a new reference for *ab initio* calculations of solids and benchmarks for pseudopotential accuracy.

We compared experimental and theoretical results on the momentum distribution and the quasiparticle renormalization factor in sodium. From an x-ray Compton-profile measurement of the valence-electron momentum density, we derived its discontinuity at the Fermi wavevector finding an accurate measure of the renormalization factor that we compared with quantum-Monte-Carlo and G0W0 calculations performed both on crystalline sodium and on the homogeneous electron gas. Our calculated results are in good agreement with the experiment.

We have been studying the heat of formation for various Kubas complexes of molecular hydrogen on Ti(1,2)ethylene-nH<sub>2</sub> using Diffusion Monte Carlo. This work has been started and is ongoing. We are studying systems involving 1 and 2 Ti bonding sites with up to 10 hydrogen molecules in numerous configurations. This work will establish a benchmark that will test the accuracy of density functional calculations and establish the feasibility of our methods for similar systems.

## Publications

- K. P. Esler, J. Kim , L. Shulenburger L and D. M. Ceperley, "Fully accelerating quantum Monte Carlo simulations of real materials on GPU clusters", submitted to Computing in Science and Engineering
- K. P. Esler, J. Kim and D. M. Ceperley, "Fully accelerating quantum Monte Carlo simulations of real materials on GPU clusters", to be published in the proceeding of 2010 Symposium on Application Accelerators in High Performance Computing, Knoxville, Tennessee
- J. Kim, K. P. Esler, and D. M. Ceperley, "Quantum Monte Carlo algorithms: making most of large-scale multi/many-core clusters", to be published in the proceeding of SciDac 2010, Chattanooga, Tennessee
- S. Huotari, et. al., J. McMinis, J. Kim, K. Esler, D. M. Ceperley, "Momentum distribution and renormalization factor in sodium and the electron gas," Phys. Rev. Lett. 105, 086403 (2010)
- K. P. Esler, R. E. Cohen, B. Militzer, J. Kim, R. J. Needs and M. D. Towler, "Fundamental High-Pressure Calibration from All-Electron Quantum Monte Carlo Calculations," Phys. Rev. Lett. 104, 185702 (2010).
- K. Esler, J. Kim et al. SciDAC 2008, "Quantum Monte Carlo algorithms for electronic structure at the petascale; the Endstation project," SciDAC 2008, J. of Physics: Conference Series 125, 012057 (2008).
- T. D. Beaudet, M. Casula, J. Kim, S. Sorella, R. M. Martin, "Molecular hydrogen adsorbed on benzene: Insights from a quantum Monte Carlo study," T. D. Beaudet, M. Casula, J. Kim, S. Sorella, R. M. Martin, J. of Chem. Physics, 129, 164711 (2008).

## Invited Talks and posters

- J. Kim, K. P. Esler and D. M. Ceperley,"QMC, harnessing computing powers of today and beyond", CECAM 2010, Lugago, Switzerland

- K. P. Esler, J. Kim and D. M. Ceperley, "Fully Accelerating Quantum Monte Carlo Simulations of Real Materials on GPU Clusters", 2010 Symposium on Application Accelerators in High Performance Computing, Knoxville, Tennessee
- J. Kim, K. P. Esler and D. M. Ceperley, "Quantum Monte Carlo algorithms: making most of large-scale multi/many-core clusters", SciDac 2010, Chattanooga, Tennessee
- "Quantum Monte Carlo Simulation of Materials using GPUs," K. P. Esler, J. Kim and D. M. Ceperley, Poster, SC09, Portland, OR, 2009.
- Talks by D. Ceperley at SCIDAC (Seattle) and Fall Creek Workshop (2009)