Journal of Physics: Conference Series **125** (2008) 012057 doi:10.1088/1742-6596/125/1/012057

Quantum Monte Carlo algorithms for electronic structure at the petascale; the endstation project

Kenneth P. Esler[1], **Jeongnim Kim**[1,2], **David M. Ceperley**[1,2,3]
National Center for Supercomputing Applications[1]
Materials Computation Center[2];Department of Physics[3]
University of Illinois at Urbana-Champaign, Urbana, IL, USA

Wirawan Purwanto, Eric J. Walter, Henry Krakauer, Shiwei Zhang
Department of Physics, College of William and Mary, Williamsburg, VA, USA

Paul R. C. Kent
Center for Nanophase Materials Sciences, Oak Ridge National Laboratory,
Oak Ridge, Tennessee, USA

Richard G. Hennig[4], **Cyrus Umrigar**[5]
Department of Material Science[4]; Department of Physics[5]
Cornell University, Ithaca, NY, USA

Michal Bajdich, Jindřich Kolorenč, Lubos Mitas
Department of Physics, North Carolina State University, Raleigh, NC, USA

Ashok Srinivasan
Department of Computer Science, Florida State University, Tallahassee, FL, USA

Abstract. Over the past two decades, continuum quantum Monte Carlo (QMC) has proved to be an invaluable tool for predicting of the properties of matter from fundamental principles. By solving the Schrödinger equation through a stochastic projection, it achieves the greatest accuracy and reliability of methods available for physical systems containing more than a few quantum particles. QMC enjoys scaling favorable to quantum chemical methods, with a computational effort which grows with the second or third power of system size. This accuracy and scalability has enabled scientific discovery across a broad spectrum of disciplines. The current methods perform very efficiently at the terascale. The quantum Monte Carlo Endstation project is a collaborative effort among researchers in the field to develop a new generation of algorithms, and their efficient implementations, which will take advantage of the upcoming petaflop architectures. Some aspects of these developments are discussed here. These tools will expand the accuracy, efficiency and range of QMC applicability and enable us to tackle challenges which are currently out of reach. The methods will be applied to several important problems including electronic and structural properties of water, transition metal oxides, nanosystems and ultracold atoms.

1. Computational quantum mechanics and quantum Monte Carlo

> "**The underlying physical laws necessary for a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact applications of these laws lead to equations much too complicated to be soluble.**" – Paul Dirac, 1929

This quotation, colloquially dubbed *Dirac's Challenge*, reflects the understanding that, at least in principle, the exact solution of the Dirac equation (or its non-relativistic counterpart, the Schrödinger equation) would yield all the information necessary to predict the behavior of matter at normal energy scales, but finding the solution by traditional pencil and paper methods is very difficult.

For a system containing N electrons, the Schrödinger equation is a partial differential equation in $3N$ dimensions. Because of the high dimensionality of the problem, exact solutions to these equations have thus far been found only simple cases such as one dimensional systems, so that Dirac's Challenge is still unmet. Nonetheless, much of theoretical physics and chemistry has been devoted to finding ever more accurate *approximate* solutions to these equations governing the behavior of matter at the quantum level. Through the decades, the coupling of more advanced methods with the exponential increase in computing power has enabled us to progress from qualitatively correct models based on empirically determined parameters to truly *ab initio* calculations, i.e. those starting *from the beginning*. These calculations, taking the Challenge head on, have as their input only the atomic numbers of system's constituent atoms.

Of these *ab initio* methods, the mostly widely known are based on *density functional theory* (DFT). The basis of these methods is an approximate mapping of the $3N$-dimensional equation onto a coupled set of N three-dimensional equations. In this mapping, the details of the correlated motion of individual electrons are approximated in an average "mean-field" sense. DFT has enjoyed a tremendous degree of success, and continues to do so, with new insights being gained every day.

Nevertheless, there are many outstanding problems for which the level of accuracy provided by DFT has proved insufficient. As an example, DFT predicts a number of known insulators, including important transition metal oxides, to be metallic. In these cases, methods which treat the many-body electron more directly are required. The most accurate method available, known as *full configuration interaction*, aims at a near-exact solution to the Schrödinger equation. Unfortunately, the required computation scales exponentially with N, the number of electrons in the system, limiting its application to atoms and small molecules. Other methods from quantum chemistry address larger systems by truncating the many-particle basis intelligently using chemical insight. These provide highly accurate approximate solutions, but are still very limited in system size since the cost scales with a high power of N.

Quantum Monte Carlo (QMC) methods achieve very high accuracy by treating the Schrödinger equation in the original $3N$-dimensional space using a stochastic sampling of the many-body wave function. Since the cost of Monte Carlo sampling is relatively insensitive to dimensionality, large systems can be simulated at reasonable computational expense. In current implementations, the cost scales as N^2 to N^3, depending on the quantities of interest, although order-N methods are currently under development. Since electron correlation is treated explicitly, QMC can achieve accuracy similar to the best quantum chemistry methods, but can be applied to much larger systems, including nanoclusters, macromolecules, and crystals. For these reasons, it is generally accorded the distinction of being the most accurate method available for systems of more than a few atoms. Furthermore, the methods are extremely general and have been successfully applied to molecules, fluids, perfect crystals, and material defects. Finally, the methods afford several avenues for parallelization which, when combined, will allow very efficient operation at the petascale.

Though still a rapidly evolving method, QMC has already contributed significant advancements across a range of disciplines. The Ceperley-Alder calculation of the correlation energy of the homogeneous electron gas [1] has become the basis of all modern functionals for DFT methods. Calculations of the phase diagram and properties of hydrogen under high temperature and compression have provided invaluable data necessary for models of Jovian planets [2, 3]. QMC was able to predict the correct lowest-energy structures for silicon nanoclusters, where DFT had failed [4]. QMC calculations of the pressure-volume curve and frequencies of lattice vibrations for materials under ultra-high compression have been shown to produce more accurate results than DFT [5].

Looking over the last 50 years, one of the remarkable developments has been the growth of the molecular dynamics (MD) and Monte Carlo simulation of atomic and molecular systems. This started in the 50's with the pioneering applications of Metropolis, Rosenbluth, Teller [6], Wood and Alder to the simplest possible hard sphere model of atoms. In the 60's, Verlet, Rahman, Berne and others developed the methodology to treat the more realistic Lennard-Jones systems, with further developments to treat more complicated systems using semi-empirical forces. A major leap forward occurred in 1985 with the discovery by Car and Parrinello [7], that a CRAY-1 class computer could use DFT calculated forces to perform a MD simulation of bulk liquid silicon. Clearly, improved computer capability can be used either to increase system size and duration *or* accuracy of the interatomic forces. The processing power is now sufficient to replace the DFT evaluation of energy and forces with more accurate QMC calculations. This will allow for simulation of systems undergoing chemical changes, such as during combustion, with a much higher degree of reliability, and validate chemical models that can be used for larger scale calculations.

The significant science advances through the use of QMC methods would not be possible without the phenomenal advance in computer power that has occurred over the past decades. The breakthrough QMC simulations have been enabled by the state-of-art, high-performance computers of every generation: hard-core bosons on a CDC 6600 [8]; ground state of solid hydrogen at high pressures on CRAY XMP and CYBER 205 [2]; electronic and structure properties of carbon/silicon clusters, Grossman-Mitas on HP 9000/715 cluster and Cray Y-MP; MLW, Williamson-Hood-Grossman [9] and Mitas-Grossman continuous DMC-MD [10] on the LLNL teraflop cluster; Coupled Electron-Ion Monte Carlo simulations of dense hydrogen, Ceperley et al. on Linux clusters and most recently on Cray XT3/4 at ORNL. Until recently, QMC methods have been able to exploit the steady increase in clock rates and enhancement of performance enabled by advances in CMOS VLSI technology for faster and bigger simulations without any significant modification in the algorithms and their implementations. However, over the last few years, clock rates have been relatively flat due to problems with heat dissipation and the structure of the computational components has undergone significant changes, including multi-core processors and powerful accelerators such as GPUs (Graphics Processing Units) for general-purpose scientific computing. These hardware trends have significant ramifications on an application performance.

The Endstation Project for quantum Monte Carlo aims to enable ground-breaking science with QMC methods at the petascale systems that have become available at the DOE leadership computing facilities. We are developing new algorithms and software solutions that can take advantage of the current hardware trends based multi-core processors. In this paper, we introduce the various methods grouped under the umbrella term of quantum Monte Carlo, and describe the QMC algorithms that have been effectively employed on current large-scale HPC systems. Then, we address the challenges posed by the recent hardware trends and propose solutions for the QMC to run efficiently at the petascale and beyond.

2. Quantum Monte Carlo (QMC) Methods

In quantum mechanics, all physically observable quantities for a system containing N electrons can be computed from the $3N$-dimensional *wave function*, $\psi(\mathbf{r}_1, \ldots, \mathbf{r}_N)$. The exact wave function for the system satisfies the time-independent Schrödinger equation,

$$\hat{\mathcal{H}}\psi = E_0\psi, \tag{1}$$

where E_0 is the ground-state energy for the system and $\hat{\mathcal{H}}$ is the Hamiltonian, or total energy operator, given by

$$\hat{\mathcal{H}} = \sum_{i=1}^{N}\left[-\frac{1}{2}\nabla_i^2 + V_{\text{ext}}(\mathbf{r}_i)\right] + \sum_{i<j}\frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}, \tag{2}$$

Where $V_{\text{ext}}(\mathbf{r}_i)$ is the *external* potential generated by the nuclei, applied electric fields, etc. For any approximate, or *trial* wave function, $\psi_T(\mathbf{r}_i)$, we compute an energy as the expected value of $\hat{\mathcal{H}}$,

$$E_T = \frac{\int d^{3N}\mathbf{R}\ \psi_T^*(\mathbf{R})\hat{\mathcal{H}}\psi_T(\mathbf{R})}{\int d^{3N}\mathbf{R}\ |\psi_T(\mathbf{R})|^2}, \tag{3}$$

where $\mathbf{R}$ is a $3N$-dimensional vector representing the positions of the N electrons. It can be shown that E_T satisfies a *variational principle* which guarantees that $E_T \geq E_0$ for all ψ_T. This implies that the lower the energy a given ψ_T has, the closer it is to the exact solution of the Schrödinger equation. Several many-body methods, including QMC and quantum chemistry methods, seek to find these low energy solutions, but differ fundamentally in how they evaluate (3).

2.1. Variational Monte Carlo

In variational Monte Carlo (VMC), the integrand in (3) is rewritten as a product of a probability density, $\pi(\mathbf{R}) \equiv |\psi_T(\mathbf{R})|^2$, and the *local energy*, $E_L(\mathbf{R}) \equiv [\hat{\mathcal{H}}\psi_T(\mathbf{R})]/\psi_T(\mathbf{R})$. Metropolis Monte Carlo is then used to generate a large number,N_{MC}, of sample points, $\{\mathbf{R_i}\}$ distributed according to $\pi(\mathbf{R})$. The trial energy, E_T, may then be estimated by

$$E_T \approx \frac{\sum_i^{N_{\text{MC}}} E_L(\mathbf{R}_i)}{N_{\text{MC}}}. \tag{4}$$

This estimate becomes exact in the limit of $N_{\text{MC}} \to \infty$, with a statistical error proportional to $N_{\text{MC}}^{-\frac{1}{2}}$.

In VMC, the trial function is written in an analytic form with a number of optimizable parameters, $\{\alpha_n\}$. These parameters can then be adjusted to minimize E_T, yielding the best wave function possible given the analytic form. The direct evaluation of many-dimensional integrals by stochastic sampling enables us to employ highly accurate variational/trial wave functions which can capture crucial many-body effects in an efficient manner. In contrast, quantum chemistry methods represent ψ_T as a linear combination of many determinants, for which the integral in (3) can be evaluated analytically. This choice leads to the difficulties in scaling those methods to large systems mentioned above.

The trial wave functions most commonly used in quantum Monte Carlo applications to problems in electronic structure are written in the Slater-Jastrow form, $\Psi_T(\mathbf{R}) = \Psi_A(\mathbf{R})\exp[J(\{\alpha_n\}, \mathbf{R})]$. The antisymmetric part Ψ_A is usually given by a determinant of a matrix whose elements, A_{ij}, are given by evaluating orbital $\phi_j(\mathbf{r})$ at the position of the i^{th} electron. The orbitals are real or complex functions of a three-dimensional variable, and are expanded in a linear basis. For systems up to about a thousand electrons, the evaluation of the

orbitals incurs the dominant computation cost in a typical QMC simulation. For larger systems, the N^3 scaling of the determinant evaluation becomes dominant.

Electron correlation is included explicitly in the wave function through a parametrized Jastrow correlation factor, $J(\{\alpha_n\}, \mathbf{R})$. The factor includes several terms, the most important of which prevent electrons from coming close too each other. This is critical since the electrostatic repulsion between electrons would make such a configuration energetically unfavorable. In actual calculations these parameters are optimized to obtain the wave function which yields the lowest energy. The form used for the Jastrow factor allows the dominant correlation effects to be captured with relatively few parameters. For example, comparison of wave functions for C_{20} clusters showed that 20 variational parameters in Slater-Jastrow VMC wave function achieved the same accuracy as $\sim 10^6$ variational parameters in a traditional quantum chemistry approach.

2.2. Diffusion Monte Carlo

Once a trial wave function is optimized within VMC, it is further improved through a stochastic projection process known as diffusion Monte Carlo (DMC). It is straightforward to show that for $\tau \to \infty$ the operator $\exp(-\tau\hat{\mathcal{H}})$ projects out the ground state of a given symmetry from any trial function with nonzero overlap. Such a projection operator can be readily implemented in a Monte Carlo approach [11].

A central concept in the DMC algorithm is a single point in the $3N$-dimensional space known as a *walker*. In the algorithm, an ensemble of walkers, termed the *population*, is propagated stochastically from generation to generation. In each propagation step, the walkers are moved through position space. At the end of each step, each walker may reproduce itself, be killed, or remain unchanged, depending on the value of the local energy at that point in space, $E_L(\mathbf{R})$. This *branching* process causes fluctuations in the population size from generation to generation. Averaging the local energy over all walkers and generations yields an energy significantly lower than the VMC energy, but still guaranteed to be an upper bound of the true energy.

Formally, the DMC method is exact provided the *nodes* of the true wave function (i.e. the locus of points for which $\psi(\mathbf{R}) = 0$) are known. In practice, the nodes are assumed to be those of the trial wave function ψ_T. This gives rise to a positive-definite bias known as the *fixed-node error*, which is, in almost all cases, very small. The fixed-node DMC method has proven to be remarkably robust, accurate and effective. For the trial wave functions outlined above, the fixed-node approximation introduces only small errors which allow one to obtain ∼95% of the correlation energy and agreement with experiments within 1-2% for ground state properties such as the lattice constants, bulk modulus, and enthalpies of formation. High accuracies have been found in essentially all systems we have studied: atoms, molecules, clusters and solids. Furthermore, new forms of trial functions and optimization techniques promise to further improve the accuracy.

2.3. The phaseless auxiliary-field QMC method

The recently developed phaseless auxiliary-field (AF) method [12–14] is a complimentary method to DMC. The AF QMC method adds several new dimensions to the QMC methodology. It provides a different route to the sign problem [15, 16] from fixed-node DMC, because its random walks take place in a manifold of Slater determinants [12, 17]. Applications show that this often reduces the reliance of the solution on the quality of the trial wave function. The potential for making QMC even more accurate is of crucial significance to the proposed research, given the central goal of robust algorithms with high predictive power. The AF method also expands the application areas of QMC. It is formulated in a Hilbert space defined by any one-particle basis. In complex systems such as transition metal oxides, the use of an effective basis in which to carry out the random walks could lead to significant gains in efficiency. By truncating

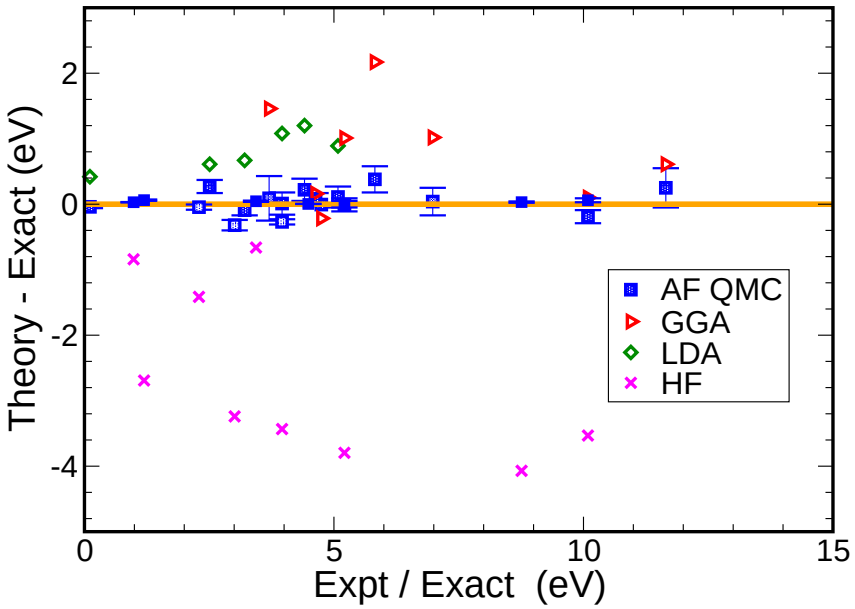

Figure 1. Calculated binding energies of molecules compared with experimental values. The discrepancy between theory and experiment is plotted versus the exact energy. The QMC is fed a trial wave function to start, which is taken directly from DFT [with either local-density approximation (LDA) or generalized-gradient approximation (GGA) functionals] or Hartree-Fock (HF). The corresponding DFT or HF results are also shown. As can be readily observed, the AF QMC results are in excellent agreement with experiment and are significantly more accurate than those from DFT and HF.

the basis (*e.g.*, through a down-folding approach [18]), we can do QMC calculations on *realistic model* Hamiltonians in which higher and irrelevant energy scales can be systematically removed.

To date, the new AF QMC method has been applied to close to 100 systems, including transition metal oxide molecules [19], simple solids [12, 20], post-*d* elements [21], van der Waals systems [22], and in molecules in which bonds are being stretched or broken [23, 24]. In these calculations we have operated largely in an automated mode, inputting only the DFT or Hartree-Fock solutions. The method demonstrated excellent accuracy, consistently able to correct errors in the mean-field trial wave function as summarized in Figure. 1.

3. Breakthrough petascale QMC simulations

We are in an era of exponential growth both in machine performance and in algorithm developments (as far as simulation is concerned). There have been many exciting advances in QMC methods in recent years, including new forms of wave functions (*e.g.* backflow, geminals and pfaffians) [25–28], methods for efficiently optimizing all the parameters in the wave functions [29–33], efficient methods to compute forces and energy differences [34–38], methods to extrapolate finite-size systems to the thermodynamic limit [20, 39–41], (*e.g.* twist average boundary conditions and scaling methods) [42], approaches for coupled electron Monte Carlo and molecular dynamics [3], more accurate methods for treating nonlocal pseudopotentials [43], new techniques to calculate observables [44, 45], algorithms which scale better with the number of electrons (linear or quadratic scaling) [9], and advances in auxiliary-field approaches to do QMC in different manifolds [12–14]. By building on these improvements, we will be able to demonstrate important science accomplishments with petascale facilities.

3.1. Coupled-Electron Monte Carlo (CEIMC) and Molecular Dynamics

During the last ten years we developed a method to perform a simultaneous simulation of electrons and ions, the Coupled Electron-Ion Monte Carlo (CEIMC) method [46], similar in spirit to the Car-Parrinello method [7], but having several advantages over that method. Of most importance is the use of explicit correlation using Reptation Quantum Monte Carlo (similar to DMC), leading to a more accurate Born-Oppenheimer surface, for all configurations of the ions, not just near equilibrium configurations. The CEIMC method can use path integrals for protons, which due to their small mass, are also quantum particles. A remarkable feature of the CEIMC method, is that adding ionic path integrals and twist averaging does not intrinsically slow down the simulation over the use of periodic boundary condition and classical protons. The CEIMC method has been used [3] to study the high-pressure phase diagram of hydrogen with computer time about the same as for corresponding *ab initio* Molecular Dynamics (AIMD) approaches. This opens opportunities for studies of light nuclei quantum effects such as proton tunneling in applications involving hydrogen at finite temperatures.

As an alternative to CEIMC, we have developed, in collaboration with J.C. Grossman from UC Berkeley, a new approach for obtaining QMC energies along the *ab initio* molecular dynamics paths generated in the DFT framework. This enables us to obtain QMC energy for an arbitrary point along the MD ionic trajectories with high efficiency. Our approach leverages the QMC propagation of walkers in the electronic position space by coupling it to the ionic MD steps so that the sampling ensemble follows the evolving wavefunction. By using this coupling and correlation between the QMC sampling points, we can speed up the calculations of energies by a factor 20 to 50 when compared to separate, point-by-point DMC calculations. We have tested this on small hydrogenated silicon clusters at high temperture for a few picoseconds with time step of about 0.5 fs. In addition to the ground state, we also carried out the calculation of excited states along the path, thus gaining insight into the temperature and impact of dynamics on the excitations. Finally, we applied this method to 32 water molecules with periodic boundary conditions, and we were able to calculate the heat of vaporization of water in an excellent agreement with experiment (unlike the DFT methods which were off by about 30%) with statistical error bar of $\sim$1 kcal/mol. This was a significant advance since computing thermodynamic quantities is often orders of magnitude more expensive than static calculations. Currently, we are working on expanding the study to larger systems.

3.2. Transition metal oxides

Transition metal oxides (TMOs) are a class of materials which exhibit a diverse set of unusual behaviors, with applications ranging from high temperature superconductivity to high-density magnetic storage, battery technology, catalysis, and spintronics. The electronic ground state of these systems is determined by several competing energetic contributions, including charge transfer between metallic ions and oxygens, large exchange and correlation energies in the d-electron subshells, crystal field effects, Jahn-Teller and other distortion related effects, etc. Because each of these contributions is of similar size, the application of pressure or chemical doping can drive transitions which change the fundamental properties of the system.

Unfortunately, the same competing energy scales that give rise to rich behaviors in materials have also made them a challenge for theory. Because each contribution to the total energy is small, even relatively small errors introduced by more approximate methods can result in qualitative mispredictions. Very basic characteristics such as equilibrium lattice constants, the energetic ordering of structural phases, electronic gaps, etc., require years of fine-tuning of mean-field theories and/or addition of *ad hoc* parameters. For example, in the LDA+U approach an adjustable parameter, U, is introduced to approximately account for the energy resulting from the double occupancy of a $3d$ orbital [47]. The complexity of many-body effects in these systems calls for new methods and strategies.

The Mott transition in TMOs such as MnO, CoO, and NiO is a prototypical problem. A half century ago, Mott [48] described how electron-electron interactions could induce an insulating state in these materials with a transition to a metallic state under pressure. Evidence of the Mott transition has recently been observed in the pressure-induced metalization of manganese oxide (MnO) [49, 50]. As such, it is an excellent testing case for expanding the applicability of QMC to transition metal solids. At high pressures (about one million atmospheres), the material undergoes a structural phase transition and very close to that also an electronic phase transition to a metal with essentially zero magnetic moment on the Mn ions. The DFT approaches, including extended DFT methods such as LDA+U, have significant difficulties to explain this transition. For example, with proper tuning of the correlation parameter, U, one can reproduce some, but not all, of the experimental observations, indicating that this requires much higher accuracy of the treatment of the electron correlation. Our preliminary QMC calculations provided very encouraging results for the estimation of the band gap and cohesive energy of the MnO solid at the ambient conditions. Therefore we are well poised to address the challenge of describing correctly the electronic structure of two different phases of TMO systems.

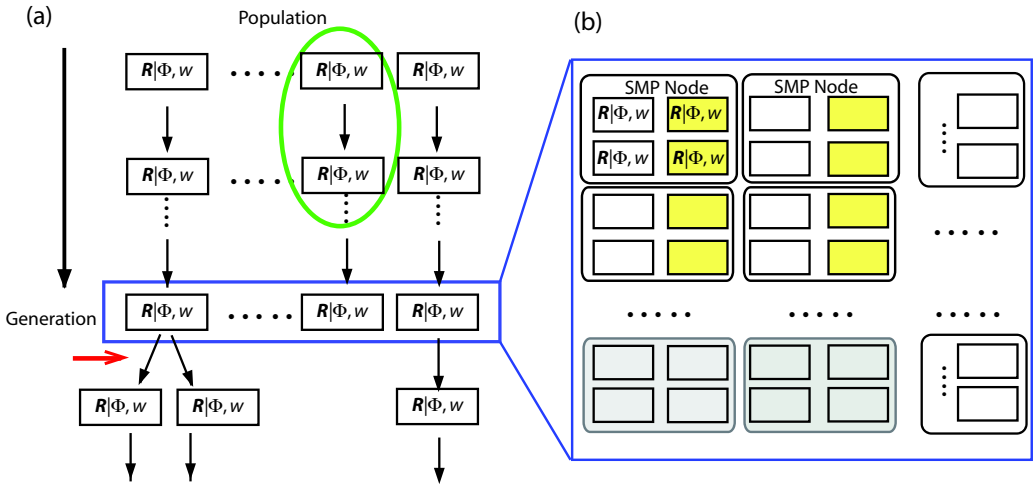

Figure 2. (a) Computational structure of DMC and AF QMC. Each column is a random walk stream. The number of columns gives QMC population size. The red arrow indicates a population control, where weights $\{w\}$ are adjusted and birth/death can occur. Shown in green is one QMC step for one walker. Each walker is a configuration of an N-particle system for DMC (shown as $\mathbf{R}$) or a Slater determinant which looks like a DFT solution for AF QMC (shown as $|\Phi\rangle$). (b) The parallel processor view of a generation after a population control. The population is distributed among SMP nodes to balance computational loads between the population control steps. Each SMP node can utilize multiple threads, e.g., two walkers for each thread on a dual-core node of this example. A group of SMP nodes as indicated by grey boxes can share single-particle orbitals whose storage needs exceed the physical memory of a SMP node.

4. Quantum Monte Carlo: Algorithms, Scalability, and Performance on current HPC systems

The vast majority of the computational time in an accurate QMC simulation is spent in diffusion Monte Carlo (DMC) and phaseless auxiliary-field QMC (AF AMC) methods. A variational optimization of the wave function is required to make the DMC run efficiently, but this step

is almost trivially parallelizable and can be treated as a special case of DMC. For this reason, we will focus our discussion on the algorithm and parallelization of DMC and AF QMC and analyze their performances as implemented in current QMC codes.

Both DMC and AF QMC sample a multi-dimensional space by propagating an ensemble of walkers (population) stochastically from generation to generation. This common algorithm is illustrated in Figure 2(a). The key difference of the two methods lies in what a walker represents. In DMC, a walker is a $3N$-dimensional configuration of an N-particle system, while a walker in AF DMC is a Slater determinant in a Hilbert space. A QMC step is carried out independently over the walkers. This allows QMC methods to be efficiently parallelized over many processors by distributing the walkers to each processor. At the end of a generation, the branching process results in a fluctuating population and the walkers are redistributed so that the next QMC step can be executed with a good load-balance. A typical number of walkers is $10^3 - 10^4$ for each population. A feedback mechanism is used to keep the average number of walkers at the target population size.

Figure 2(b) describes the current parallelization scheme of DMC/AF QMC algorithms for a given generation. Each virtual processor, typically a MPI node on a Symmetric Memory Processor (SMP) node, holds in its memory a small number of walkers and proceeds with a propagation of local walkers. Current QMC codes, which employ simple load balancing schemes and a lock-step propagation of each generation, scale well to a few thousand processors. As long as each processing element, *e.g.*, a SMP node of Fig. 2(b), is responsible for several walkers, transient load imbalances and communication time result in a few percent loss in efficiency. Figure 3 shows a pilot AF QMC calculation on the Cray XT3 at ORNL. The parallel speedup is shown for a model system of a P_2 molecule in a supercell. The calculation has a population size of $\mathcal{M}$ =10,000. We see that a parallel efficiency of over $\simeq$85% is achieved on up to 2000 MPI nodes, when the average number of walkers per node is 5. The deviation from perfect scaling is due mostly to load imbalance at the limit of few walkers per node. A similar parallel efficiency is achieved for DMC calculations of a wide range of systems. While DMC and AF QMC share the common parallelization and load-balance schemes in propagating a population over the generations, the operations for each QMC step are fundamentally different.

The computational effort of DMC is dominated by that of the wave functions. Accurate functional forms are essential to achieve high quality results with small statistical error. The dominant computation for a large system is set by the antisymmetric part Ψ_A of the trial wave function: i) N^3 scaling to update Slater determinants; and ii) $N - N^3$ scaling to fill in the Slater determinants with one-particle orbitals. An example is given in Fig. 4 (a).

An analysis of computational kernels [53] is presented in Fig. 4 (b). Of three alternative

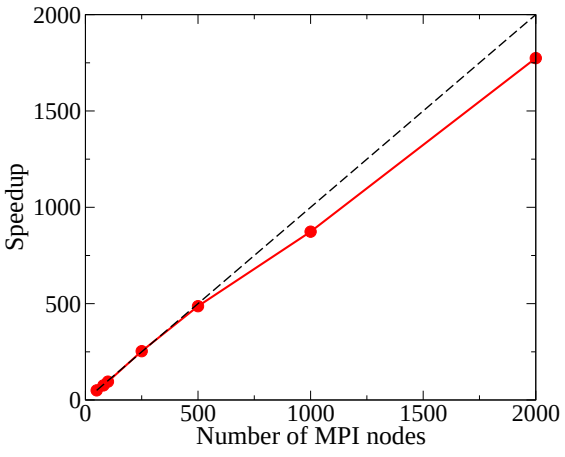

Figure 3. The parallel efficiency of plane-wave AF QMC for a model system, the P_2 molecule in a supercell with a single-particle basis of 9939 plane-waves. The speedup on a Cray XT3 (dual-core AMD Opteron Opteron 2.6 GHz CPUs) is plotted as a function of the number of MPI nodes. "Hard scaling" is shown, with the number of walkers fixed at ∼10,000 (the total fluctuates due to population control). The dashed line represents perfect scaling. Efficiencies of $\simeq$85% are achieved up to 2000 nodes, where there are only ∼5 walkers per node on average.

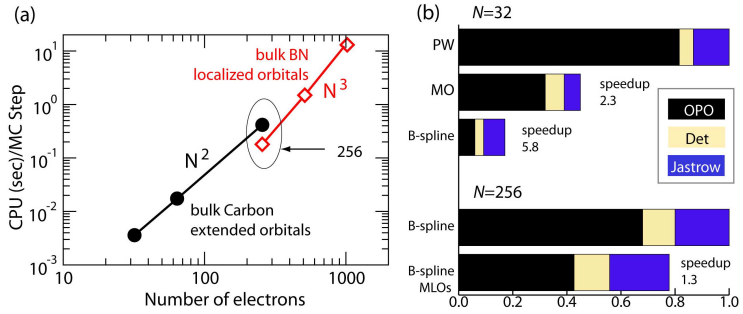

Figure 4. (a) CPU time per Monte Carlo (MC) step for bulk carbon (diamond) and BN (zincblend) [51]. Each MC step involves random moves of N electrons. (b) Breakup of many-body wave function computations for 32- and 256-electron systems using one-particle orbitals in a plane-wave (PW), molecular-orbital (MO) basis and on a real-space grid with B-spline methods [52].

representations of the one-particle orbitals, the use of B-spline interpolation of orbitals on a grid speeds up the computation sixfold over use of a plane-wave (PW) basis set, while the cost with molecular-orbital (MO) basis set falls between B-spline and PW. The gain in computation time of real-space methods become increasingly large as the system size grows. The time to evaluate one-particle orbitals can be further reduced by using maximally localized orbitals (MLOs) [54]. The speedup at $N = 256$ is a direct consequence of using the localized orbitals. For large systems, the performance of determinant updates and Jastrow function evaluations, which typically scale as N^2, becomes as important as one-particle orbital evaluations.

In AF QMC, each walker is a Slater determinant: $|\phi\rangle = |\varphi_1, \varphi_2, \cdots, \varphi_N\rangle$ [55], where N is the number of electrons, and each electron orbital is given by $\varphi_i = \sum_{j=1}^{M} c_{j,i}|\chi_j\rangle$, where $|\chi_j\rangle$ is a one-electron basis state, and M is the size of the basis. The main operations include (i) propagating a walker and (ii) calculating the so-called local energy. A walker propagates by

$$e^{\hat{h}(\sigma)}|\phi\rangle \rightarrow |\phi'\rangle, \tag{5}$$

where $\hat{h}(\sigma)$ is an independent-electron Hamiltonian which depends on the external auxiliary fields, σ. The fields are sampled by the random walk and fluctuate both spatially and *temporally*. The local energy is given by

$$E_L(\phi) = \frac{\langle \Psi_T|\hat{H}|\phi\rangle}{\langle \Psi_T|\phi\rangle}, \tag{6}$$

where $|\Psi_T\rangle$ is a Slater determinant (i.e., of the same form as $|\phi\rangle$).

There are presently two main flavors of the phaseless AF QMC method, from two different choices of the one-electron basis: (i) plane-wave and norm-conserving pseudopotential, and (ii) Gaussian basis sets. In (i), a simulation supercell or twist-average boundary conditions. A basis function is specified by $|\chi_j\rangle \propto \exp(i\mathbf{G}_j \cdot \mathbf{r})$, where $\mathbf{G}_j$ is a reciprocal lattice vector. As in plane-wave DFT calculations, our basis set consists of plane waves with $|\mathbf{G}_j|^2/2 \leq E_{\rm cut}$, where the cutoff energy, $E_{\rm cut}$, controls the size of the basis, $M \propto E_{\rm cut}^{3/2}$. The dimension of the AF σ is $\mathcal{O}(8\,M)$. The computational cost of each MC step (advancing all components of σ) scales roughly as $N^2\,M\ln M$. In (ii), the size of the basis is much smaller, but the cost of each step currently scales as M^3. The calculation of the local energy is similar to the calculation of the variational energy in a Hartree-Fock calculation.

5. Transition from the Terascale to the Petascale QMC.

As demonstrated in Figure 3, current QMC algorithms scale well to a few thousand processors provided that each processing element has a sufficient number of local walkers. They also highlight the limitation of the current QMC implementation for the petascale systems of 10^4-10^5

compute cores. There are several critical impediments in changing from a terascale performance to a petascale one. Consider two types of computations of interest, neither possible on today's computers:

First, consider the computation of a system of ten thousand electrons versus the current limitation of about one thousand electrons. Given a computation time scaling of between the second and third power of the number of electrons, each iteration would take between a hundred and a thousand times as long, assuming minimal change in processor speed. However, the number of needed walkers would not necessarily increase, so the strategy of simply increasing the number of walkers would vastly increase the needed wall clock time to reach convergence. The memory requirements for storing the trial functions would increase in the same way [56].

Second, consider our grand goal, namely the ability to perform classical MD using quantum forces. The utility of these MD calculations will be determined by the amount of wall clock time needed to advance a single time step; *i.e.* to calculate the forces to sufficient accuracy. This will require walkers to advance enough steps so that they become decorrelated with their starting values. Adding additional walkers beyond what is needed would not be effective because this critical number of steps is fixed by the physics of the system, not the number of processors.

Both of these examples argue that we must exploit the shared memory capability of a single node. First, having the processors on a given node share the wave function reduces tje memory requirement by the number of cores per node. Second, by partitioning the wave function calculation of a single walker across the cores on a node, the wall clock time of a single iteration will be reduced by the same amount.

5.1. Multithreading

The architecture of current and proposed petascale systems is based on multicore Symmetric Memory Processor (SMP) nodes. While the memory per core is likely to remain at the current size of few gigabytes, the number of cores per SMP node is expected to grow steadily over the years and hence also the shared memory per SMP node. The current QMC codes can reach a petascale performance without major redesign, simply by utilizing the increased shared memory capacity and improved performance of multi-threaded software. We have implemented the simplest loop-parallelization over the walkers on a node using an MPI/OpenMP hybrid scheme as depicted by the white and yellow boxes in Figure 2(b). The parallel efficiency remains high for QMC as we increase the number of OpenMP threads while fixing the average number of walkers on a SMP node. In fact, DMC scales nearly perfectly with respect to the number of threads: the additional parallelism from multithreading allows more walkers per node and improves the overall parallel efficiency and load balance among SMP nodes within the hierarchical partitioning. Evaluations of different terms of a trial wave function, single particle orbitals and dense/sparse updates of determinants can be readily parallelized by OpenMP threads, although care must be taken to schedule multiple tasks with varying computation loads to maximize the parallelism while minimizing data dependency and overhead of thread management.

5.2. Multi-level algorithmic partitioning

In many applications, components of the existing QMC algorithms can be cast into averages over a set of variables or parameters, providing opportunities for efficient subdivision. Such possibilities enable grouping of processors and assigning each group an instance of the variable or parameter subset. For example, in calculations of systems with periodic boundary conditions, especially for metals, carrying out calculations with different twist angles (*i.e.*, Bloch **k**-vectors) [42] reduces the statistical errors and the finite system-size errors. A group of processors operate on a single twist-angle but averages are taken over all the groups. In the CEIMC method which uses QMC to compute electronic energy differences and Path Integral Monte Carlo to perform simulation of the protonic (or light ionic) degrees of freedom, every protonic time slice

can be assigned to a group of processors, and the overall average of the energy difference used to decide on movement of the entire path, thus allowing much lower temperatures than one can do on current platforms.

6. Challenges for QMC at the petascale and beyond

6.1. Load balancing in DMC and AF QMC

Current QMC codes, which employ simple load balancing schemes and a lock-step propagation of each generation, scale well to a few thousand processors. As long as each processing element is responsible for several walkers, transient load imbalances and communication time, result in a few percent loss in efficiency. A finite number of generations is required for the population to reach equilibrium, prior to which no useful data may be collected. Since the total effort required to reach equilibrium scales with the population size, there is an advantage to keeping the population as near to the minimum required population as possible. Thus, it would be highly desirable to have an algorithm which would allow the simulation to approach the *dilute limit* of one to two walkers per node without sacrificing efficiency due to load imbalance.

Development of just such a method is underway. It aims to achieve near perfect efficiency despite load imbalance by relaxing the requirement that all walkers be at the same generation number at all time. This removes the need for barrier synchronizations and blocking communication and allows the communication required for walker redistribution to be overlaid with computation. A simple, but effective, implementation of such a scheme could be done with a manager/worker model. At set wall-clock intervals, each node would send a packed message containing the number of walkers on each node and the physical properties of those walkers which are the output of the simulation. They would then continue propagating their local population of walkers. Meanwhile, the manager would decide on an optimal walker redistribution to be effected by pointwise transfers of walkers between pairs of nodes. It would send messages to each node that needs to participate in communications. Workers would occasionally check for such posted messages, and comply by sending or receiving walkers as appropriate. If the scale requires, this scheme can be efficiently generalized to a two-level hierarchy.

To allow scaling to simulations containing $\mathcal{O}(10^4)$ electrons, the work required to propagate a single walker may be distributed with threads across cores sharing a memory address space. Present indications suggest that a typical petascale machine will have on the order 32 cores in an SMP node. If care is taken to avoid threading at too fine a scale, it should be possible to achieve efficient distribution of single-walker propagation over this number of cores.

6.2. Fault tolerance and reproducibility

A significant concern for the petascale drive in general is the issue of fault tolerance. The sheer number of components involved in a single calculation on the order of 10^5 to 10^6 processing cores may reduce the mean time between failures (MTBF) to below the average run time for a job. To avoid wasting large amounts of CPU time, it will thus be necessary to have codes which are robust to failure of individual processors. The manager/worker architecture introduced for efficient load balancing can easily be designed to accommodate the failure of individual nodes. In the simplest implementation, a node which has failed to communicate for some time beyond the designated wall-clock interval will be invalidated and its walkers removed from the pool. Since this removal is triggered by a random event (i.e. hardware failure), it should not introduce bias into the walker population. Furthermore, the proposed load balancing scheme should allow the remainder of the nodes to proceed without loss of efficiency.

6.3. Very large systems and orbital storage

The memory usage of quantum Monte Carlo for complex materials is dominated by the storage of single-particle orbitals. In the simplest implementation, this usage scales quadratically with

system size. Just as in density functional calculations, this scaling stems from the delocalized character of the orbitals and can be reduced to linear scaling for large systems using localized orbitals [9, 54]. Furthermore, we have developed several new methods to reduce the memory required for orbital storage through the use of nonuniform meshes and by exploiting symmetry. We anticipate that the memory per SMP node on upcoming petascale machines will be sufficient for a wide range of problems, since the complete *read-only* orbital data can be shared by multiple threads on a SMP node.

For a sufficiently large, disordered system, the memory required for orbital storage may exceed what is available on a single SMP node. In this case, it should be possible to distribute the orbitals among groups of tightly coupled nodes. Within each group of a few nodes, electron positions would be broadcast in an all-to-all fashion. Each node would then evaluate the orbitals stored in its local, shared memory at the required coordinates. The nodes would then send the data back to the necessary nodes in an operation akin to a small, distributed matrix transpose. This communication may be overlaid with the computation required for other parts of wave function evaluation, and should thus incur very little performance penalty.

Using an efficient B-spline basis, complex orbital evaluations can be performed at a rate of approximately 4 million per second per core on current processors, thus generating data at a rate of 64 MB/s per core. For this scheme to work well, the point-to-point bandwidth should exceed 64 N_{core} MB/s. This requirement is significantly exceeded on most present terascale and sub-petascale machines, and will continue to be if the present ratio of interconnect bandwidth to FLOPS is maintained in future machines.

In the limit of a large system, AF QMC will parallelize the calculation by splitting each walker, a Slater determinant, across multiple processors. If we have $\mathcal{M}$ walkers, the calculation will look like $\mathcal{M}$ loosely coupled DFT calculations. Each walker, *i.e.*, each "DFT stream", will reside on n processors. Typically $\mathcal{M}$ is larger than, say, 300, to ensure that population control bias is small. This means that $n < 350$ for the expected petascale machine with more than 100,000 processors. Current implementation of DFT is well into this regime. For example, a plane-wave DFT calculation on an 807-atom FePt nanoparticle has been performed on the Cray XT3 by team member Kent and co-workers, achieving a parallel efficiency of about 75% on 512 processors, and scaling to at least 2048 processors. This scaling is already sufficient for the proposed AF QMC petascale calculations. Because population controls occur infrequently in the QMC (every $\mathcal{O}(10)$ steps), the grain size will be exceedingly large relative to load-balancing cost. In other words, the QMC efficiency should closely follow that of the DFT with respect to the number of processors, and be very weakly dependent on the total number of processors. In developing the plane-wave AF QMC kernel, we will need to address several additional technical issues, such as the calculation of the local energy $E_L(\Phi)$ and the force bias [12, 57], but the road-map is clear.

6.4. Pseudorandom number generation

At the petascale, Monte Carlo simulations will require one to two orders of magnitude more pseudorandom numbers than have ever been used to date. Since the statistical quality of the random streams has never been tested at this scale, it is critical to verify the integrity of this fundamental input. Work is currently underway to test the Scalable Parallel Random Number Generator (SPRNG) suite of generators, utilizing Jaguar, the Cray XT4 at ORNL with over 30,000 cores.

We tested the Multiplicative Lagged Fibonacci Generator in SPRNG. This generator passed all the statistical tests in the SPRNG test suite, both tests of sequential random number generators and tests of parallel random number generators. Around 28,000 random number streams were used in these tests, with a total of around 10^{13} random numbers per test. We also tested this generator with a 2D Ising model using the Metropolis algorithm. In this test,

256 random number streams are used on a 16×16 lattice. Such computation is repeated with 512 different sets of streams. We verified the accuracy of each set of streams, and then verified the cumulative results of the entire sets of streams. We believe that this is the largest test of random number generators, both in terms of the number of streams used and the total number of random numbers consumed. These tests give us confidence that even at the petascale, our simulations will be free of bias from correlated random numbers.

7. Summary

We have outlined some of the most relevant ideas which will enable to fulfill the potential of petaflop machines for solving breakthrough many-body quantum problems. We plan to test these devlopments on some of the current problems which are at the forefront of attention of computational practitioners in this field. The role of flagship supercomputers in science is that of *time machines*: whatever the computational scientists do on supercomputers today will be available to scientists at large in 10-15 years. The goal of our scientific thrust is to prepare and develop the underlying methodologies, algorithms and implementations for the benefit of broad scientific communities in physics, chemistry, materials science and other fields.

Acknowledgments

The QMC Endstation project is supported by the U.S. Department of Energy (DOE) under Contract No. DOE-DE-FG05-08OR23336. This research used resources of the National Center for Computational Sciences and the Center for Nanophase Materials Sciences, which are sponsored by the respective facilities divisions of the offices of Advanced Scientific Computing Research and Basic Energy Sciences of DOE under Contract No. DE-AC05-00OR22725.

References

[1] D M Ceperley and B J Alder 1980 *Phys. Rev. Lett.* **45** 566
[2] D M Ceperley and B J Alder 1987 *Phys. Rev.* B **36** 2092–2106
[3] Pierleoni C, Ceperley D M and Holzmann M 2004 *Phys. Rev. Lett.* **93** 146402
[4] L Mitas, J C Grossman, I Stich, J Tobik 2000 *Phys. Rev. Lett.* **84** 1479
[5] Maezono R, Ma A, Towler M and Needs R 2007 *Phys. Rev. Lett.* **98** 025701
[6] N Metropolis, A W Rosenbluth, M N Rosenbluth, A H Teller and E Teller 1953 *J. Chem. Phys.* **21** 1087–1092
[7] Car R and Parrinello M 1985 *Phys. Rev. Lett.* **55** 2471–2474
[8] Kalos M H, Levesque D and Verlet L 1974 *Phys. Rev.* A **9** 2178–2195
[9] A J Williamson, R Q Hood and J C Grossman 2001 *Phys. Rev. Lett.* **87** 246406
[10] J C Grossman and L Mitas 2005 *Phys. Rev. Lett.* **94** 056403
[11] W M C Foulkes, L Mitas, R J Needs and G Rajagopal 2001 *Rev. Mod. Phys.* **73** 33
[12] Zhang S and Krakauer H 2003 *Phys. Rev. Lett.* **90** 136401
[13] Al-Saidi W A, Zhang S and Krakauer H 2006 *J. Chem. Phys.* **124** 224101
[14] Suewattana M, Purwanto W, Zhang S, Krakauer H and Walter E J 2007 *Phys. Rev.* B **75** 245123 (pages 12) URL `http://link.aps.org/abstract/PRB/v75/e245123`
[15] D M Ceperley and B J Alder, J Chem Phys **81**, 5833 (1984); Shiwei Zhang and M H Kalos, Phys Rev Lett **67**, 3074 (1991); J. B. Anderson in *Quantum Monte Carlo: Atoms, Molecules, Clusters, Liquids and Solids*, Reviews in Computational Chemistry, Vol. 13, ed. by Kenny B. Lipkowitz and Donald B. Boyd (1999)
[16] Zhang S 1999 *Quantum Monte Carlo Methods in Physics and Chemistry* ed Nightingale M P and Umrigar C J (Kluwer Academic Publishers) cond-mat/9909090
[17] Zhang S, Carlson J and Gubernatis J E 1997 *Phys. Rev.* B **55** 7464
[18] Andersen O K and Saha-Dasgupta T 2000 *Phys. Rev.* B **62** R16219
[19] Al-Saidi W A, Krakauer H and Zhang S 2006 *Phys. Rev. B* **73** 075103
[20] Kwee H, Zhang S and Krakauer H 2008 *Phys. Rev. Lett.* In press.
[21] Al-Saidi W A, Krakauer H and Zhang S 2006 *J. Chem. Phys.* **in press** physics/0608298
[22] Al-Saidi W A, Krakauer H and Zhang S 2007 *The Journal of Chemical Physics* **126** 194105 (pages 8) URL `http://link.aip.org/link/?JCP/126/194105/1`
[23] Al-Saidi W A, Zhang S and Krakauer H 2007 *The Journal of Chemical Physics* **127** 144101 (pages 8) URL `http://link.aip.org/link/?JCP/127/144101/1`

[24] Purwanto W, Al-Saidi W A, Krakauer H and Zhang S 2008 *The Journal of Chemical Physics* **128** 114309 (pages 7) URL `http://link.aip.org/link/?JCP/128/114309/1`
[25] P Lopez Rios, A Ma, N D Drummond, M D Towler and R J Needs 2006 *Phys. Rev.* E **74** 066701
[26] Casula M and Sorella S 2003 *The Journal of Chemical Physics* **119** 6500–6511 URL `http://link.aip.org/link/?JCP/119/6500/1`
[27] Casula M, Attaccalite C and Sorella S 2004 *J. Chem. Phys.* **121** 7110
[28] Bajdich M, Mitas L, Drobny G, Wagner L and Schmidt K E 2006 *Phys. Rev. Lett.* **96** 130201
[29] Umrigar C J and Filippi C 2005 *Phys. Rev. Lett.* **94** 150201
[30] Umrigar C J, Toulouse J, Filippi C, Sorella S and Hennig R G 2007 *Phys. Rev. Lett.* **98** 110201
[31] Toulouse J and Umrigar C J 2007 *J. Chem. Phys.* **126** 084102
[32] Sorella S, Casula M and Rocca D 2007 *J. Chem. Phys.* **127** 014105
[33] Toulouse J and Umrigar C J 2008 *J. Chem. Phys.* **128** 174101
[34] Filippi C and Umrigar C J 2000 *Phys. Rev.* B **61** R16291
[35] Filippi C and Umrigar C J 2002 *Recent Advances in Quantum Monte Carlo Methods, Vol. III* (World Publishing: ed. by W. A. Lester and S. M. Rothstein)
[36] R Assaraf and M Caffarel 2003 *J. Chem. Phys.* **119** 10536
[37] Mosé Casalegno M M and Rappe A M 2003 *J. Chem. Phys.* **118** 7193
[38] S Chiesa, D M Ceperley, and S Zhang 2005 *Phys. Rev. Lett.* **94** 036404
[39] L M Fraser, W M C Foulkes, G Rajagopal, R J Needs, S D Kenny and A J Williamson 1996 *Phys. Rev.* B **53** 1814
[40] Williamson A J, Rajagopal G, Needs R J, Fraser L M, Foulkes W M C, Wang Y and Chou M Y 1997 *Phys. Rev.* B **55** 4851
[41] Chiesa S, Ceperley D M, Martin R M and Holzmann M 2006 *Phys. Rev. Lett.* **97** 076404
[42] C Lin F H Z and Ceperley D M 2001 *Phys. Rev.* E **64** 016702
[43] Casula M 2006 *Phys. Rev.* B **74** 161102
[44] Assaraf R and Caffarel M 1999 *Phys. Rev. Lett.* **83** 4682
[45] Toulouse J, Assaraf R and Umrigar C J 2007 *J. Chem. Phys.* **126** 244112
[46] Ceperley D M and Dewing M 1999 *J. Chem. Phys.* **110** 9812
[47] Kasinathan D, Kunes J, Koepernik K, Diaconu C V, Martin R L, Prodan I D, Scuseria G E, Spaldin N, Petit L, Schulthess T C and Pickett W E 2006 *Phys. Rev.* B **74** 195110 URL `http://link.aps.org/abstract/PRB/v74/e195110`
[48] Mott N F 1949 *Proceedings of the Physical Society. Section A* **62** 416
[49] Patterson J R, Aracne C M, Jackson D D, Malba V, Weir S T, Baker P A and Vohra Y K 2004 *Phys. Rev.* B **69** 220101 (pages 4)
[50] Yoo C S, Maddox B, Klepeis J H P, Iota V, Evans W, McMahan A, Hu M Y, Chow P, Somayazulu M, Hausermann D, Scalettar R T and Pickett W E 2005 *Phys. Rev. Lett.* **94** 115502 (pages 4)
[51] Dell PowerEdge 1955 Blade Server Cluster, Intel64 (Clovertown, 2.33GHz, quad core) and InfiniBand (IB) interface.
[52] Performance analysis by PerfSuite on 1.6 GHz Itanium 2 (6.4 GFLOPS peak).
[53] Kufrin R perfSuite, http://perfsuite.ncsa.uiuc.edu/
[54] Reboredo F A and Williamson A J 2005 *Phys. Rev.* B **71** 121105(R)
[55] Shiwei Zhang 2003 *Theoretical Methods for Strongly Correlated Electron Systems* (Springer-Verlag) chap 2, pp 39–74
[56] For the one-particle orbitals in a plane-wave and molecular-orbital basis, the memory usage scales N^2 for the system size N. The use of localized orbitals can lead to a linear scaling
[57] Purwanto W and Zhang S 2004 *Phys. Rev.* E **70** 056702