

Efficient Quantum Monte Carlo Energies for Molecular Dynamics Simulations

Jeffrey C. Grossman

Lawrence Livermore National Laboratory, 7000 East Avenue L-415, Livermore, California 94550, USA

Lubos Mitas

Department of Physics, North Carolina State University, Raleigh, North Carolina 27695-8202, USA

(Received 27 August 2004; published 10 February 2005)

A method is presented to treat electrons within the many-body quantum Monte Carlo (QMC) approach “on-the-fly” throughout a molecular dynamics (MD) simulation. Our approach leverages the large (10–100) ratio of the QMC electron to MD ion motion to couple the stochastic, imaginary-time electronic and real-time ionic trajectories. This continuous evolution of the QMC electrons results in highly accurate total energies for the full dynamical trajectory at a fraction of the cost of conventional, discrete sampling. We show that this can be achieved efficiently for both ground and excited states with only a modest overhead to an *ab initio* MD method. The accuracy of this dynamical QMC approach is demonstrated for a variety of systems, phases, and properties, including optical gaps of hot silicon quantum dots, dissociation energy of a single water molecule, and heat of vaporization of liquid water.

DOI: 10.1103/PhysRevLett.94.056403

PACS numbers: 71.15.Pd, 31.25.-v, 71.15.Nc, 71.15.Mb

Molecular dynamics (MD) combined with density functional theory (DFT) [1] has become a widely accepted and highly effective computational modeling tool to study a broad range of chemical, physical, and biological problems. The DFT approach offers a balance of accuracy and computational efficiency that is well suited for both static and dynamic calculations of numerous properties of systems with hundreds and even thousands of atoms. Yet, in some cases, one requires a more accurate computational approach, such as the quantum Monte Carlo (QMC) approach, that goes beyond the single-particle approximation to solve the Schrödinger equation. Unfortunately, the trade off of speed for accuracy in such higher accuracy methods has prevented their application to molecular dynamics simulations in which a typical trajectory involves thousands or millions of atomic configurations. As a result, *ab initio* molecular dynamics (AIMD) methods have been limited to the accuracy dictated by DFT.

In this Letter, we present a method to treat the electrons within the many-body QMC method “on the fly” throughout the AIMD simulation, providing an improved, significantly more accurate total energy for the full dynamical trajectory. The QMC method relies on a stochastic sampling of the wave function by an ensemble of walkers propagating in electron configuration space. Our approach leverages this unique aspect of the QMC method by coupling the stochastic electronic steps with the AIMD atomic steps so that the sampling ensemble follows the evolving wave function. Since AIMD atomic steps are small enough (10^{-3} to 10^{-4} Å), the QMC samples can be continuously updated and follow the dynamical trajectory with moderate computational overhead (50% to 100%) over AIMD. This approach provides the same energies as conventional, dis-

crete QMC sampling and gives error bars comparable to separate, much longer QMC calculations.

In static (i.e., ions are stationary) diffusion Monte Carlo (DMC) calculations, the Hamiltonian remains fixed while the electronic wave function is sampled stochastically in imaginary time [2]. During a dynamical simulation, *both* the Hamiltonian and the wave function are updated at each MD step. In order to propagate the DMC electrons in concert with the AIMD evolution of ionic positions [6], one needs to modify the DMC algorithm so that the correct wave function and Hamiltonian are being sampled. In particular, when the Hamiltonian changes by changing the atomic positions, the DMC walker population sampling the wave function from the previous dynamical step has to be updated. Let us denote by $f(\mathbf{R}, \tau; \hat{H})$ the probability distribution of DMC walkers at position \mathbf{R} , imaginary-time τ , for the Hamiltonian \hat{H} . Therefore at time $t_{\text{MD}} = 0$ the walkers sample $f(\mathbf{R}, \tau; \hat{H}(0))$, where $\hat{H}(0)$ corresponds to the Hamiltonian at $t_{\text{MD}} = 0$. At the next MD step, $t_{\text{MD}} = 1$, one should be sampling the distribution $f(\mathbf{R}, \tau + n\Delta\tau; \hat{H}(1))$ where $\Delta\tau$ is the DMC step and n is small (ideally, $n = 1$). In order to sample this distribution correctly, one should start the DMC propagation from $f(\mathbf{R}, \tau; \hat{H}(1))$ instead of the original $f(\mathbf{R}, \tau; \hat{H}(0))$. Fortunately, this can be accomplished by the correlated sampling technique [3,7] which has been applied successfully to other problems (e.g., forces, responses). The propagation of walkers is governed by the commonly used DMC approximation for the Green’s function which is a product of Gaussians describing diffusion and drift times a weight with exponentiated local energy. For each step along the MD trajectory, each DMC walker is reweighted by (e.g., at $t_{\text{MD}} = 1$)

$$w(\mathbf{R}(\tau), t_{\text{MD}} = 1) = \frac{G[\mathbf{R}(\tau) \leftarrow \mathbf{R}(\tau - \Delta\tau); \hat{H}(1)]}{G[\mathbf{R}(\tau) \leftarrow \mathbf{R}(\tau - \Delta\tau); \hat{H}(0)]} \cong \frac{\Psi_T^2[\mathbf{R}(\tau); \hat{H}(1)]}{\Psi_T^2[\mathbf{R}(\tau); \hat{H}(0)]} e^{-\Delta\tau(E_{\text{loc}}[\hat{H}(1)] - E_{\text{loc}}[\hat{H}(0)])/2},$$

where E_{loc} is the local energy of the walker evaluated with the given Hamiltonian and corresponding trial function Ψ_T . The ratio of wave functions corrects for the dynamical part of the Green's function (see [7]) while the exponent renormalizes the local energy weight term. For simplicity, the weight is estimated using only a single DMC step propagation history which is straightforward to calculate with little overhead. This approach can be generalized by taking a few more steps from the DMC history.

When adjusting the wave function to the new position of ions, one needs to relax both the weights from local energies and from dynamical parts of the Green's function. In our tests across a range of systems, from a few to 256 electrons, we found that three DMC propagation steps per MD step eliminated any visible bias in the DMC energies, and therefore we used three steps in all applications presented below.

The accuracy of this continuous DMC (CDMC) approach is assessed by comparing the energies obtained with those from separate, discretely sampled MD snapshots. Figure 1 shows this comparison of total energies over a 50 fs time window for SiH_4 , Si_5H_{12} , and $\text{Si}_{14}\text{H}_{20}$ clusters at 1000 K. In each case, the cluster was heated and allowed to equilibrate for several picoseconds before the thermostat was turned off and data recorded. For simplicity, the DMC correlation parameters, initially optimized for each different cluster size, were kept intact during the dynamics since we found only marginal impacts on the DMC energies. As seen in Fig. 1, the agreement between CDMC and discretely sampled DMC energies is excellent for all three cluster sizes. Furthermore, both the correct average energy and a detailed description of the DMC energy path are obtained within the CDMC approach in these cases, all of which were run for 2 ps.




FIG. 1 (color online). DMC total ground state energies vs time for 5 fs time intervals of an AIMD simulation of SiH_4 , Si_5H_{12} , and $\text{Si}_{14}\text{H}_{20}$ at 1000 K. The circles (lines) correspond to discretely sampled (continuous) DMC calculations. The error bars of the discretely sampled data are smaller than the symbol size.

A comparison of efficiencies between the AIMD, CDMC, and discretely sampled DMC methods for these hot silicon clusters illustrates the efficiency of the CDMC approach. As always, the DMC simulation time depends on the desired or acceptable level of statistical noise; timing comparisons are only meaningful once this criterion is determined. For each cluster size, we chose a number of DMC walkers such that the statistical fluctuations in the CDMC energies are one-tenth the size of the variation in total energy as a function of the MD time. This choice allows one to observe the detailed changes in the energy landscape throughout the dynamical simulation and leads to a roughly similar amount of computational time for the CDMC and the AIMD methods. The statistical resolution is then independent of size since both the AIMD and the CDMC methods scale as N^3 with the number of electrons in the system [8]; by fixing the computational time required per MD step for the CDMC and the AIMD methods to be the same for a given cluster size, we are in effect pinning the statistical resolution across different cluster sizes as well. For the AIMD part of the simulations, 12 electronic iterations were found to be sufficient to converge the total energy to within 10^{-6} a.u. at each MD iteration.

These results illustrate that with a modest computational overhead beyond AIMD (i.e., a factor of 2 or less), a very accurate description of the DMC energy landscape is obtained using the CDMC approach. On the other hand, a full DMC calculation performed on individual snapshots from the MD simulation is substantially less efficient. For this comparison, we have carried out a 3000-step DMC calculation every 100 MD steps using the same number of walkers as in the CDMC runs. The resulting error bars are roughly the size of the symbols in Fig. 1. Even with this somewhat coarse sampling of MD snapshots, the discrete sampling DMC approach is an order of magnitude less efficient than the CDMC approach.

Energy differences may also be evaluated on the fly within the CDMC approach. For example, the optical gap can be computed by coupling two DMC electron populations to the AIMD simulation, one for the ground state and one for a given excited state (this doubles the CDMC computational requirements). The excited state population evolves in the same manner as described above, and the ground to excited state total energy difference can be evaluated at each step. Figure 2 shows the DMC singlet optical gap computed within the CDMC approach and the discrete sampling DMC approach for the Si_5H_{12} cluster in a 100 fs time window. For better comparison, we have doubled the number of discretely sampled DMC points. Again, note the excellent agreement between CDMC and the discretely sampled DMC energies throughout.

An improvement in the description of the optical gap during dynamical simulations illustrates an ideal applica-




FIG. 2 (color online). DMC HOMO-LUMO gap (eV, top panel) and difference between DMC and LDA gaps (eV, lower panel) vs time for a 100 fs time interval of an AIMD simulation of Si_5H_{12} at 1000 K. For the DMC gaps, the circles (lines) correspond to discretely sampled (continuous) DMC calculations.

tion of the CDMC approach since DFT gaps are notoriously inaccurate and other methods for improving the gap are too costly to apply during dynamics. The lower panel of Fig. 2 shows the energy difference between the DMC and local density approximation (LDA) highest occupied molecular orbital (HOMO)–lowest unoccupied molecular orbital (LUMO) gaps as a function of MD time. Note that the difference is not constant but rather fluctuates between 0.8 and 1.2 eV over time, indicating that the trends in LDA gaps as a function of strain may be inaccurate. The finite temperature affects both DFT and DMC average gaps which are smaller by ~ 1.0 eV when compared with $T = 0$ calculations.

In electronic structure calculations with ionic motion, one can encounter an occasional change in the occupation and ordering of one-particle orbitals due to temperature fluctuations. This is well known in simulations of metals, at state crossings, and in systems with (near-)degenerate orbitals as in high-symmetry clusters. The orbital swapping can be detected by evaluating overlaps between the relevant subsets of orbitals, e.g., orbitals close to the Fermi energy. Fortunately, the DMC samples enable one to compute these overlaps both rapidly and with sufficient accuracy. Such an event is usually infrequent and when it occurs, as in our calculation of the Si_5H_{12} excited state, we reequilibrate the set of walkers since the change in nodal surfaces alters the spatial distribution of walkers very significantly. Over a 2 ps MD time window, swapping occurred one (zero) times among excited (ground) state orbitals for this case.

In order to test the ability of the CDMC approach to describe a bond breaking process, we compute the O-H bond dissociation energy in a water molecule. This reaction has been studied extensively by numerous theoretical methods [9], and the dissociation energy (125.9 kcal/mol) is well known experimentally. We use constrained AIMD, such that the O-H distance is stretched uniformly by 0.01 \AA

at each step and the molecule is allowed to relax at each given distance constraint.

Figure 3 shows the CDMC energies as a function of the change in O-H bond length from its equilibrium value (ΔR) compared to discretely sampled DMC energies. Note that from the ground state molecule through the transition structure region ($\Delta R \sim 0.5 \text{ \AA}$) to full dissociation ($\Delta R > 1.0 \text{ \AA}$) the CDMC approach accurately reproduces the correct DMC energy pathway. The DMC results give a dissociation energy of 127(2) kcal/mol, in excellent agreement with experiment. The total number of CDMC steps taken to obtain the data shown in Fig. 3 (solid line) is over 30 times less than the number of discrete sampling DMC steps (circles). To describe the bond breaking, we tested both the spin unrestricted single- and restricted two-configuration trial wave functions; the wave function with the double excitation to an antibonding orbital improves the variational energy; however, the DMC energies are within the error bars of CDMC energies in Fig. 3.

In addition to high temperature molecular oscillations and bond dissociation, we apply the CDMC approach to a full liquid simulation of 32-water molecules. The MD trajectory is taken from a simulation at ambient conditions using the TIP4F [10] classical flexible water model, in order to assure that the structure and diffusion are in good agreement with experiment, and to avoid the issues associated with DFT water at 300 K [11]. The system was equilibrated for 100 ps at 300 K using classical MD and subsequent data taken using a time step of 0.25 fs with a weakly coupled Berendsen-type thermostat [12]. Trial functions for CDMC energies were obtained from a DFT [Perdew-Burke-Ernzerhof (PBE) [13]] ground state wave function calculation at the TIP4F geometries. The heat of vaporization of water is computed in our simulation from the liquid binding energy, i.e., the average CDMC energy for the 32-water system minus 32 times the DMC energy of a single water molecule, $U_{\text{liq}} - U_{\text{gas}}$.




FIG. 3 (color online). DMC energies for the dissociation of a single H_2O molecule as a function of the change in O-H distance (ΔR) with respect to the equilibrium value. The circles correspond to discretely sampled DMC energies with error bars approximately of the size of the symbols. The solid line corresponds to CDMC energies.

TABLE I. Total electronic energies per molecule computed within PBE and CDMC simulations averaged over a 0.5 ps MD simulation for 32-water simulation at 300 K (U_{liq} , a.u.), H_2O molecule at 300 K (U_{gas} , a.u.), and corresponding heat of vaporization of water (H_{vap} , kcal/mol).

	U_{liq}	U_{gas}	H_{vap}
PBE	-17.0887(1)	-17.0763(1)	6.2(1)
CDMC	-17.1024(5)	-17.0854(5)	9.1(4)

Table I shows our results for a 0.5 ps CDMC simulation of the 32-water molecule system at 300 K. The electronic energies per molecule are listed with error bars due to thermal fluctuations over the MD time interval for PBE and to both thermal and stochastic fluctuations for a CDMC simulation. The heat of vaporization was computed as the liquid binding energy, with the addition of 1.5 kcal/mol to account for the quantum motion of the hydrogen atoms [14]. The energy of an H_2O molecule at 300 K was computed as the energy of the ground state plus a thermal correction computed from a Gaussian PBE frequency calculation. Despite a rather short simulation time, the CDMC results show a nearly 50% improvement over PBE and provide an estimate very close to the experimental value of 9.92 kcal/mol [15]; to our knowledge the CDMC result represents the best value obtained from an AIMD calculation.

It is not difficult to imagine further development of the kind of coupling introduced in this work. For example, a hybrid MD approach could provide not only improved energies but also an improvement in the trajectories. It is well known that, especially at nonequilibrium geometries such as saddle regions of chemical reactions or bond dissociations, the DFT methods are less reliable. One can therefore envision the following combination of methods: (i) QMC [either variational Monte Carlo or DMC] forces are computed periodically at appropriately frequent checkpoints but not at each MD step, and (ii) whenever significant discrepancies between DFT and QMC forces are observed one switches to the more accurate QMC forces until the discrepancies return below a given threshold. In fact, even a full QMC force or MD with both energies and forces generated within the QMC method may be possible in the near future.

In conclusion, we have proposed and demonstrated an approach for the calculation of accurate energies during molecular dynamics simulations by coupling the stochastic propagation of QMC walkers with the deterministic evolution of the ions. This approach is shown to reproduce the correct DMC energies with a modest overhead to AIMD of a factor of 2 for the systems studied here, which range from room temperature liquid to high temperature clusters and bond dissociation properties. The continuous evolution of DMC energies allows one to take advantage of the MD time step correlation through correlated sampling, and we have shown that it can be applied to classical or quantum molecular dynamics simulations. In the present work, this

approach has been demonstrated for evaluating the DMC energy at atomic positions generated by other methods, as has been the standard technique for most QMC calculations; further possibilities are currently being explored.

We are grateful to G. Galli and E. Schwegler for many stimulating discussions, and to L. Wagner for help with efficient implementation of this method. This work was done under the auspices of the U.S. Department of Energy by the University of California Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48 and further supported by the Grant No. ONR-N00014-01-1-0408 and by NSF Grants No. DMR-0102668 and No. DMR-0121361 with a part of the runs done at the NCSA and PSC facilities.

- [1] R. Car and M. Parrinello, *Phys. Rev. Lett.* **55**, 2471 (1985).
- [2] A trial function is constructed as a product of a single Slater determinant of DFT orbitals times a Jastrow correlation factor [3–5]. The pseudopotential plane-wave orbitals from the AIMD simulation are projected onto a saturated Gaussian basis in order to improve efficiency. A variational Monte Carlo approach is used to optimize the correlated many-body trial function, and a fixed-node diffusion Monte Carlo approach, which removes most of the variational bias, is used for all energies.
- [3] B.L. Hammond, W.A. Lester, Jr., and P.J. Reynolds, *Monte Carlo Methods in Ab Initio Quantum Chemistry* (World Scientific, Singapore, 1994).
- [4] D.M. Ceperley and L. Mitas, *Quantum Monte Carlo Methods in Chemistry, Advances in Chemical Physics* (Wiley, New York, 1996), Vol. XCIII.
- [5] W.M.C. Foulkes, L. Mitas, R.J. Needs, and G. Rajagopal, *Rev. Mod. Phys.* **73**, 33 (2001).
- [6] In our molecular dynamics simulations, the Kohn-Sham energy functional is minimized at each MD step using DFT with the Perdew-Burke-Ernzerhof functional and nonlocal pseudopotentials. We use plane waves with 12 Ryd cutoff, and a time step of 0.075 fs. For molecular calculations, the periodically repeating cubic supercell allows for at least 10 Å between image molecules.
- [7] C. Filippi and C.J. Umrigar, *Phys. Rev. B* **61**, R16291 (2000).
- [8] The DMC method using a localized Gaussian basis scales as N^2 for the computation of total energies. An extra factor of N must be added for the total energies to have the same statistical error bar as a function of size.
- [9] F. Jensen, in *Introduction to Computational Chemistry* (John Wiley and Sons, Chichester, England, 1999), pp. 274–284.
- [10] M.W. Mahoney and W.L. Jorgensen, *J. Chem. Phys.* **115**, 10758 (2001).
- [11] J.C. Grossman *et al.*, *J. Chem. Phys.* **120**, 300 (2004).
- [12] H.J.C. Berendsen *et al.*, *J. Chem. Phys.* **81**, 3684 (1984).
- [13] J.P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **80**, 891 (1998).
- [14] H.A. Stern and B.J. Berne, *J. Chem. Phys.* **115**, 7622 (2001).
- [15] G.S. Kell, *J. Chem. Eng. Data* **20**, 97 (1975).