

Quantum Monte Carlo Methods in the Study of Nanostructures

J. Shumway* and D. M. Ceperley†

*Dept. of Physics and Astronomy, Arizona State University, Tempe AZ, 85287

†Dept. of Physics, University of Illinois, 1110 W. Green St., Urbana IL 61801

Abstract. We describe quantum Monte Carlo methods for simulating quantum systems. These techniques are particularly well-suited for nanomaterials where quantum effects often fall outside conventional approaches. We review recent applications to quantum dots, buckyballs, and superfluid helium nanoclusters. (For Handbook of Computational Nanotechnology Encyclopedia of Nanoscience and Nanotechnology)

CONTENTS

1	Introduction	2
2	Common models for nanostructures	3
2.1	The <i>ab initio</i> Hamiltonian	4
2.2	Effective mass or quasiparticle models	5
2.3	Interatomic potentials	5
2.4	Lattice models and tight-binding Hamiltonians	6
3	Zero temperature methods	6
3.1	Metropolis Monte Carlo	7
3.2	Variational Monte Carlo	8
3.2.1	Slater-Jastrow trial functions	9
3.2.2	Improved determinants	10
3.2.3	Order (N) methods for wavefunction evaluation	10
3.3	Sampling with drifted Gaussians	11
3.4	Diffusion Monte Carlo	12
3.4.1	Estimators in diffusion Monte Carlo	14
3.5	The fermion sign problem	15
3.5.1	The fixed node approximation	15
3.5.2	Release node methods	16
3.6	Excited states	16
3.7	Pseudopotentials	18
3.8	Position-dependent mass and pseudo-Hamiltonians	18
3.9	Other Hamiltonians	19
3.10	Fixed phase and generalized boundary conditions	19
3.11	Zero temperature examples and applications	20

3.11.1	Fullerenes	20
3.11.2	Colloidal quantum dots	20
3.11.3	Effective mass models of quantum structures	21
4	Path integral methods	23
4.1	Basic theory	23
4.2	The action	25
4.3	Sampling the paths	30
4.4	Bosons, superfluidity and Bose condensation	34
4.5	Nanodroplets	39
4.6	Excitation spectra using MaxEnt	40
4.7	Fermion PIMC: Fixed-node and fixed-phase methods	41
4.8	Ground state path integrals	44
4.9	Path integral simulations of quantum dots	45
5	Conclusions and outlook	46

1. INTRODUCTION

Monte Carlo methods are computational approaches which, one way or another, involve averaging over stochastic or random processes. Computer simulations are the only known general approaches for exactly treating high-dimensional systems; that is, if the number of degrees of freedom is greater than about ten. In classical statistical mechanics one uses a type of Monte Carlo, namely random walks, to determine thermodynamic properties of many-body systems, such as the phase diagram or the energy as a function of temperature. The correspondence principle says that quantum mechanics reduces to classical mechanics in the limits of low density and high temperature, so it is natural to expect that general methods for calculating properties of quantum many-body systems must also involve random walks. Quantum simulations are more complicated: in classical mechanics one only has to do integrals over the Boltzmann distribution, while in quantum mechanics one must also determine the quantum density matrix or, at low temperature, the wave function. Exact methods which can treat many-particle systems are only known in a few cases. However, even with their imperfections, the Quantum Monte Carlo (QMC) methods are the most powerful computational techniques available to treat quantum systems.

The first quantum Monte Carlo methods were by Donsker and Kacs in the early 1950's [1]. In the 1960's Kalos developed methods for bosonic systems, leading to the solution of the problem of the hard sphere system in 1972 [2, 3]. Anderson pioneered the use of methods for small molecules in 1975 [4, 5]. Ceperley and others developed methods for fermion systems, in particular the fixed-node importance sampled diffusion Monte Carlo method [6]. Its application to the electron gas was important in improving density functionals [7]. The realization of Feynman's theory of superfluid systems was developed by Ceperley and Pollock in the 1980's [8, 9, 10]. In the 1990's methods for more realistic systems including pseudopotentials were developed by Fahy [11, 12], Mitas [13, 14], and others [15, 16]. In this period, generalization of fixed-node methods

for finite temperature [17] and for complex trial functions have been developed [18, 19]. There has also been a great deal of progress in regards to lattice calculations not described in this article.

The abilities of QMC make it a particularly useful technique for studying nanomaterials. A lot of experimental research is aimed at realizing novel quantum states in nanostructures. Mean field theories often fail to describe these quantum states, and QMC is a useful alternative. Many of the Hamiltonians used to model nanostructures cannot be treated with conventional quantum methods, such as density functional theory and the use of variation methods with specialized basis sets. In these cases, QMC is often more flexible and reliable than those approaches. Thermal effects can be important in nanomaterials, and can be naturally included in QMC simulations. For example, magnetic transitions at finite temperature are often studied with Monte Carlo Ising simulations, so one would expect similar properties of quantum nanomagnets will require QMC simulations. The nature of simulations also embraces the complexity present in many nanosystem, in contrast to more reductionist analytic techniques.

Quantum Monte Carlo often scales linearly or as a low-order polynomial with system size, usually no worse than order(N^3). This is in contrast to exact diagonalization and other many body techniques that are limited to just a few particles because of poor size scaling. With the exponential growth of computer power described by Moore's law, QMC is able to directly tackle nanoscale structures containing hundreds or thousands of particles. Experimental techniques now allow manipulation of nanostructures with tens to thousands of particles; this progress is meeting the exponentially increasing capabilities of QMC simulations. Recently, the proliferation of low-cost computer clusters, built from commodity personal computer parts, has given a further edge to Monte Carlo simulations since QMC calculations are naturally parallel. One simply uses different random number streams on different processors and averages across the independent simulations. Parallelism also allows direct averaging over model variables, such as alloy fluctuations or generalized boundary conditions, by averaging over an ensemble of simulations, each run on its own processor.

This article is organized as follows. QMC algorithms are applicable to a wide variety of models, which we summarize in Section 2. In Section 3 we describe ground state QMC techniques. These are the best developed QMC algorithms, and have been very successful at calculating properties quantum dots, molecules, and nanocrystals. In Section 4 we describe the finite temperature techniques. In these path-integral-based algorithms, the correspondence with classical simulations is most clear. We review applications of path integrals to nanoscale helium droplets and quantum dots. We conclude in Section 5 with a discussion of current trends and important challenges for QMC.

2. COMMON MODELS FOR NANOSTRUCTURES

There are many types of nanostructures, and we illustrate a few commonly studied systems in Fig. 1. While there is much variety, all nanostructures contain many atoms arranged on scales between a nanometer and a micron. A variety of models are commonly used to describe nanostructures.

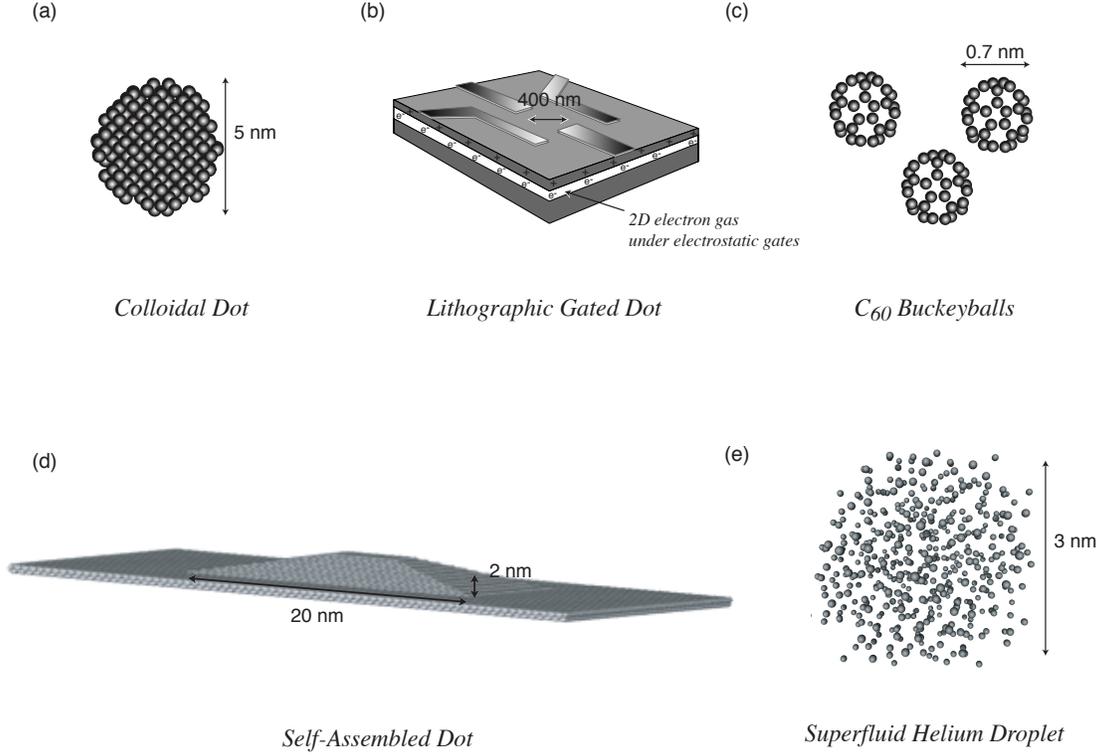


FIGURE 1. Illustration of the nanostructures considered in this review.

2.1. The *ab initio* Hamiltonian

The non-relativistic many-body Hamiltonian,

$$H = \sum_{i=1}^{N_e} \frac{p_i^2}{2m_e} + \sum_{i=1}^{N_{\text{ion}}} \frac{P_i^2}{2M_i} + \sum_{\substack{i<j \\ \text{electrons}}} \frac{e^2}{r_{ij}} + \sum_{\substack{i<j \\ \text{ions}}} \frac{Z_i Z_j e^2}{R_{ij}} - \sum_{i=1}^{N_e} \sum_{j=1}^{N_{\text{ion}}} \frac{Z_j e^2}{|r_i - R_j|}, \quad (1)$$

is the starting point for most *ab initio* work on atoms, molecules, and solids. In this article R_i refers to the coordinate of the i^{th} ion and r_i to the i^{th} electron. Although Eq. 1 is general, determination of exact numerical solutions quickly becomes intractable due to the many interacting degrees of freedom and the required fermion antisymmetry of the wavefunction. One often resorts to single particle approximate theories, such as Hartree-Fock or density functional theory. The QMC techniques described in this article allow numerical many-particle quantum simulations and are directly applicable to the *ab initio* Hamiltonian.

A common approximation to the *ab initio* Hamiltonian is the use of pseudopotentials to replace core electrons with a (non-local) ionic pseudopotential,

$$H = \sum_{i=1}^{N_e} \frac{p_i^2}{2m_e} + \sum_{i=1}^{N_{\text{ion}}} \frac{P_i^2}{2M_i} + \sum_{\substack{i<j \\ \text{electrons}}} \frac{e^2}{r_{ij}} + \sum_{\substack{i<j \\ \text{ions}}} \frac{Z_i Z_j e^2}{R_{ij}} + \sum_{i=1}^{N_e} \sum_{j=1}^{N_{\text{ion}}} \sum_{\ell=0}^{\infty} \hat{V}_{\ell}^{(j)}(|r_i - R_j|). \quad (2)$$

Here $\hat{V}_{\ell}^{(j)}$ is a non-local operator that acts only on the subspace with angular momentum ℓ about atom j .

2.2. Effective mass or quasiparticle models

Ab initio techniques are currently limited to a system of roughly one thousand electrons, though increases due to improved methods and hardware can be anticipated. Typical nanostructures are much larger than this, *e. g.* a typical self-assembled quantum dot contains about 10,000 atoms. To model these structures, it may be necessary to formulate the electronic system in terms of quasiparticle excitations. In semiconductors, these quasiparticles are conduction electrons and valence holes. In the simplest model the conduction-band electrons are assumed to be a single, parabolic band,

$$H = \sum_{i=1}^N \frac{p^2}{2m^*} + \sum_{i=1}^N V_e(\mathbf{r}_i) + \sum_{i<j}^N \frac{e^2}{\epsilon r}, \quad (3)$$

where $V_e(\mathbf{r})$ is an effective confining potential arising from strain, conduction-band edges of heterostructures, and electric potentials, m^* is the electron effective mass, and ϵ is the dielectric constant.

The simplest and most-studied model for a quantum dot is a two-dimensional model with parabolic confinement and a magnetic field, $\mathbf{B} = \nabla \times \mathbf{A}$,

$$H = \sum_{i=1}^N \frac{1}{2m^*} \left(\mathbf{p}_i - \frac{|e|\hbar}{c} \mathbf{A}_i \right)^2 + \sum_{i=1}^N \frac{1}{2} m^* \omega^2 r_i^2 + \sum_{i<j}^N \frac{e^2}{\epsilon r_{ij}} + \sum_{i=1}^N \frac{e}{c} \mathbf{B}_i \cdot \boldsymbol{\sigma}_i. \quad (4)$$

The sum over N particles now refers to the N conduction band electron quasiparticles, which are far fewer than the N_e electrons in the *ab initio* or pseudopotential Hamiltonians, Eq. (1) and Eq. (2). The single-particle states in the absence of a magnetic field are the solutions to the two-dimensional harmonic oscillator, with energies $E_{n_x n_y} = \hbar\omega(n_x + n_y + 1)$. With a magnetic field, the single-particle problem is still analytically solvable and gives the well-known Fock-Darwin states [20]. The interacting system is commonly studied with QMC many-body simulations.

2.3. Interatomic potentials

To study spatial arrangement of the ions in a structure, a common starting point is an empirical interatomic potential that replaces the full Born-Oppenheimer potential. The

simplest approximation, appropriate to rare gas atoms is a pair potential:

$$H = \sum_{i=1}^{N_{\text{atoms}}} \frac{P_i^2}{2M_i} + \sum_{\substack{i < j \\ \text{atoms}}} V_{ij}(R_{ij}). \quad (5)$$

Quantum Monte Carlo techniques can simulate the quantum behavior of these atoms. For example, the effect of zero point motion of light atoms in nanostructures can be studied. Also, Bose condensates, such as superfluid He⁴ nanodrops, can be numerically simulated with QMC [21] using an empirical interatomic potential.

2.4. Lattice models and tight-binding Hamiltonians

Lattice models, such as the Hubbard model, have been extensively studied with quantum Monte Carlo. However, these models have seen little application to nanostructures. Very realistic models of single-electron properties in semiconductors have been constructed using tight-binding Hamiltonians [22, 23, 24, 25]. While QMC can be applied to tight-binding models, the number of atoms and electrons in typical semiconductor nanostructures makes QMC too costly. Instead, many-body effects in tight-binding models have been included with Slater-determinant, or configuration interaction (CI), expansions [26]. The recent work of Zhang and Krakauer [27] using lattice QMC in a space of Slater determinants may be the best approach for future QMC applications to tight-binding Hamiltonians.

Another area where lattice models may become relevant is the study of coupled nanostructures. While single-particle physics may suffice to describe individual quantum dots, arrays of coupled dots may form highly correlated structures, such as a Hubbard lattice. This has already been seen for crystalline C₆₀, where the lattice of buckeyballs is described well by a three-band Hubbard model. Effective electronic coupling parameters have been extracted from atomistic calculations for coupled C₆₀ molecules [28, 29] and for coupled self-assembled dots [30]. These parameters are input for lattice Hamiltonians that can be studied with lattice QMC [31].

3. ZERO TEMPERATURE METHODS

Although absolute zero is unattainable in the real world, an essential starting place is the quantum ground state. One of the reasons for this is that for many physical systems, room temperature (300 K \approx 25 meV) is low compared to the Fermi temperature of a typical metal (5×10^4 K \approx 5 eV) or the binding energy of a hydrogen atom (1.5×10^5 K \approx 13.6 eV). The zero point motion of quantum mechanics makes even the ground state of a quantum system quite interesting.

The ground state of an N -particle system may be described by a wavefunction in $3N$ dimensions. (For the most part we neglect the spin degrees of freedom for simplicity of presentation.) Traditional numerical methods for solving partial differential equations, such as basis set expansions or the use of a grid, are unable to handle the high dimension-

ality of the wavefunction for even a few particles. Monte Carlo techniques, on the other hand, scale very well to high dimensional systems.

In this section, several ground state Monte Carlo algorithms are described. The first of these, variational Monte Carlo (VMC), is perhaps easiest to describe because of its direct analogy with classical thermodynamic simulations. As its name implies, it is a variational method and is simply a technique for numerically evaluating the energy and other properties of rather realistic trial wavefunctions. Another method, diffusion Monte Carlo (DMC), is able to project properties of the true ground state using a stochastic process. This section concludes with a discussion of difficulties posed by fermionic systems and excited states, and presents several common approaches to these problems. For a review of ground state quantum Monte Carlo to molecules and solids, see Ref. [32].

3.1. Metropolis Monte Carlo

Before discussing quantum Monte Carlo methods, we briefly review the Metropolis (or Markov chain) Monte Carlo method as applied to classical thermodynamic simulations. Thermodynamic quantities may be written as averages over the Boltzmann distribution,

$$\langle \mathcal{O} \rangle = \frac{\int dS \pi(S) \mathcal{O}(S)}{\int dS \pi(S)}, \quad (6)$$

where $\pi(S)$ is the weight or probability of a configuration S in the ensemble and $\mathcal{O}(S)$ is the estimator for property \mathcal{O} . The configuration S may refer to the positions of the electrons, or of the ions, or of electron spin, or in the case of path integrals the particle permutations; the formalism is identical in all of these cases.

Metropolis *et al.* [33] proposed the following Markov process for evaluating the ratio of integrals in Eq. (6):

1. Start with an initial (arbitrary) configuration S_0 .
2. To advance the configuration from S_i to S_{i+1} :
 - (a) Choose a trial configuration S'_i from a transition probability $T(S_i \rightarrow S'_i)$.
 - (b) Evaluate the acceptance ratio, \mathcal{A} ,

$$\mathcal{A}(S_i \rightarrow S'_i) = \min \left\{ 1, \frac{T(S'_i \rightarrow S_i) \pi(S'_i)}{T(S_i \rightarrow S'_i) \pi(S_i)} \right\} \quad (7)$$

- (c) Compare $\mathcal{A}(S_i \rightarrow S'_i)$ to a random variate $\zeta \in [0, 1)$. If $\zeta < \mathcal{A}$, then $S_{i+1} = S'_i$, otherwise $S_{i+1} = S_i$.

The Markov process S_0, S_1, S_2, \dots generated by the Metropolis algorithm eventually samples the distribution $\pi(S)$. The quantity \mathcal{O} is estimated by averaging over the random walk:

$$\langle \mathcal{O} \rangle \approx \frac{1}{n} \sum_{i=i_{\text{first}}}^{i_{\text{last}}} \mathcal{O}(S_i), \quad (8)$$

VARIATIONAL MONTE CARLO CODE

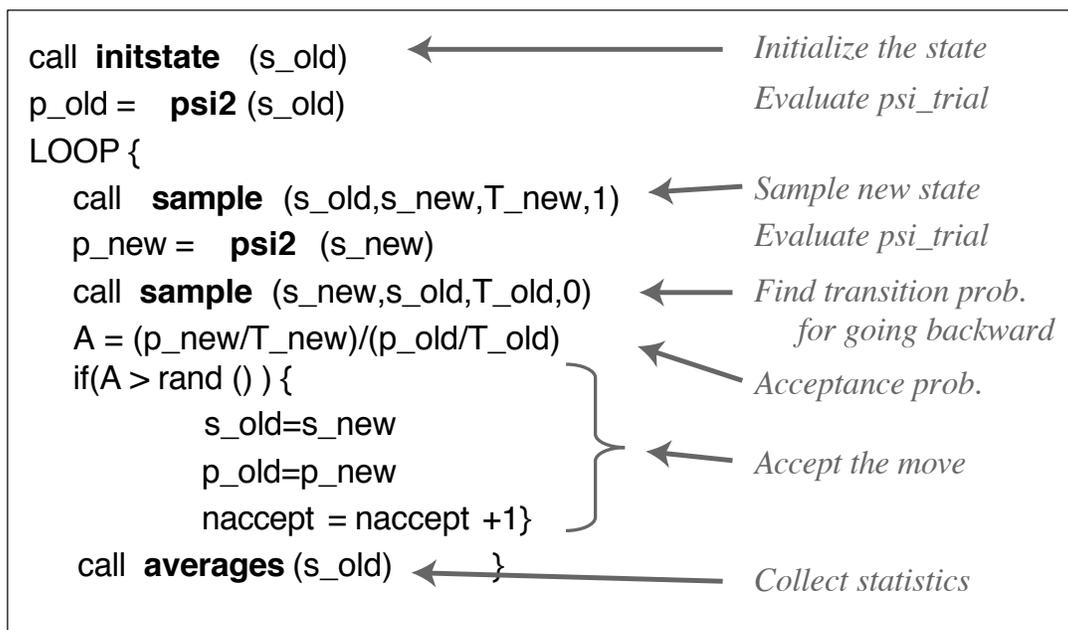


FIGURE 2. Outline of a VMC code. The calculated average and variance of the energy can be used to optimize a parameterized trial function.

where $n = i_{\text{last}} - i_{\text{first}} + 1$ is the number of steps in the average, i_{first} is determined by looking at the time series \mathcal{O} . Discussion of how to perform the statistical analysis of the correlated data from a Markov chain can be found in Allen and Tildesley [34].

3.2. Variational Monte Carlo

The essence of variational Monte Carlo is the use of Metropolis Monte Carlo methods, commonly used in classical thermodynamics, to evaluate expectation values for a trial wavefunction ψ_T by sampling the density $\pi(R) = |\psi_T(R)|^2$. We show the outline of a simple VMC code in Fig. 2. As in the case of analytic variation methods, the trial wavefunction is parameterized, and the parameters are optimized to minimize the energy. (One can also minimize the variance of the energy estimator [35].) The advantage of using Monte Carlo is that much more realistic wavefunctions may be tried when Monte Carlo integration is used. Better wavefunctions lead to stricter upper bounds on the energy and better estimates of other expectation values.

To calculate expectations of quantum mechanical operators $\mathcal{O}(R)$ diagonal in the position basis R we can use the expression in Eq. (6). With an assumed many-body trial wavefunction $\Psi_T(R)$, we take $\pi(R) = |\psi(R)|^2$ and $\mathcal{O}(R) = \psi_T(R)^{-1} \mathcal{O} \psi(R)$. For

example, the energy expectation value for the trial wavefunction is,

$$\langle E \rangle = \frac{\int dR \psi_T^*(R) H \psi_T(R)}{\int dR \psi_T^*(R) \psi_T(R)} = \frac{\int dR |\psi_T(R)|^2 E_L(R)}{\int dR |\psi_T(R)|^2} \geq E_0 \quad (9)$$

where $E_L(R) = \psi_T(R)^{-1} H \psi_T(R)$ is the “local energy” of the trial wavefunction and E_0 is the unknown exact energy. Note that only if $\psi_T = \Phi_0$, then $E_L = E_0$ everywhere, and the variance of the integrand goes to zero. This is the “zero variance” property of the energy estimator.

3.2.1. Slater-Jastrow trial functions

The most common trial wavefunction for fermionic systems is a Slater-Jastrow pair-product wavefunction,

$$\psi_T(R) = \det |\phi_i(r_j)| \exp \left[- \sum_{i < j} u(r_{ij}) \right]. \quad (10)$$

Here the functions ϕ_i are single particle orbitals, and the “Jastrow” factors $u(r_{ij})$ add correlations to the (exchange only) Slater determinant. Both the orbitals ϕ and the functions u can be parameterized. One can determine exact conditions at small distances [36] and at large distances [37]. A particular advantage for charged systems is the accuracy of analytic Jastrow factors [37]. In some cases a single determinant is insufficient and the wavefunction must be written as a sum of determinants. An alternative is to modify the orbitals to include many-body effects.

There are efficient procedures [38] for evaluating the local energy of a trial wavefunction. The local energy is a sum of the potential energy and the kinetic energies of each particle:

$$E_L(R) = V(R) - \sum_{i=1}^N e^u \nabla_i \lambda_i \nabla_i e^{-u}, \quad (11)$$

where ∇_i is the gradient operator for the coordinates of particle i , λ_i is $\hbar^2/2m_i$, the inverse mass of particle i and $u(R) = -\log \psi_T(R)$ is negative of the log of trial wavefunction. This expression may be written in terms of the derivatives of u .

$$E_L(R) = V(R) + \sum_{i=1}^N [\nabla_i \lambda_i \nabla_i u - (\nabla_i u) \lambda_i (\nabla_i u)]. \quad (12)$$

In the effective mass formulation, it is straightforward for λ to describe a position-dependent effective mass tensor.

To evaluate derivatives of Slater determinants, one uses the identity,

$$\mathcal{D}_i \det M = \det M (M^T)^{-1} \mathcal{D}_i M, \quad (13)$$

where M is a matrix and \mathcal{D}_i is any derivative operator acting on particle i , such as $\mathcal{D}_i = \nabla_i$ or $\mathcal{D}_i = \nabla_i^2$. Thus, to evaluate the local energy of a Slater determinant it is sufficient to evaluate the gradient and Laplacian of the single particle orbitals and the inverse of the Slater matrix.

To speed evaluation of Slater determinants, update formulas can be used as follows. When a particle j is moved, the new determinant value is quickly found from a vector multiply,

$$\det M' = \det M \sum_{i=1}^N M_{ji}^{-1} \phi_i(r'_j). \quad (14)$$

If the move of particle j is accepted, the Slater matrix inverse must be updated,

$$(M'_{ki})^{-1} = \begin{cases} \left(\frac{\det M}{\det M'} \right) M_{ki}; & k = j, \\ M_{ki}^{-1} - (M'_{ji})^{-1} \sum_{m=1}^N M_{km} \phi_m(r'_j); & k \neq j. \end{cases} \quad (15)$$

3.2.2. Improved determinants

As discussed below, the location of the nodes of the trial function are the only uncontrolled approximation in ground state quantum Monte Carlo, so improving the nodal surfaces is crucial to achieving higher precision. One successful approach to going beyond a Slater determinant of one body orbitals is an approach based on analyzing the effect of electron-electron interaction as a perturbation to the mean field solution. The Slater Jastrow trial function contains first order terms in a perturbation analysis. The second order terms contain backflow terms as well as three body (bosonic) effects. It is the former which puts correlation in the nodal surfaces. The orbitals are perturbed as

$$\phi_k = \phi_k(\mathbf{r} + \nabla Y(R)) \quad (16)$$

where $Y(R)$ is the backflow potential, usually taken as a sum of renormalized pair interactions [39].

3.2.3. Order (N) methods for wavefunction evaluation

The computer time needed to repeatedly evaluate Slater-Jastrow wavefunctions is the bottleneck that currently limits system size. In most applications, the Slater inverse matrix elements are updated incrementally, so that order(N^3) operations are required to move all N particles. This scaling results because the Slater determinant is dense, and each orbital must be evaluated for each electron position. For insulating systems, it is well-known that the orbitals may be transformed into localized (or Wannier) functions, without affecting the Slater determinant. With these localized functions, the Slater determinant becomes sparse, and each electron overlaps only a small number of orbitals, independent of system size. This “order(N)” technique has recently been implemented

by Williamson, Hood, and Grossman [40]. They have used this for benchmark calculations of bandgaps of silicon quantum dots [41, 42]. The Jastrow functions also involve pairwise sums, which scale as order(N^2). In a further approximation, the Jastrow functions are taken to be short range, so the pairwise sums can be kept to be order(N) as well.

3.3. Sampling with drifted Gaussians

As a way of introducing the Diffusion Monte Carlo algorithm, we discuss an efficient method for moving the particle in Metropolis Monte Carlo and thereby evaluating the expectation values in Eq. (6). Consider a particle diffusing in a potential $U(R)$. Thermodynamics tells us the time-averaged distribution of the particle will be proportional to $\exp[-\beta U(R)]$. We take βU to be a fictitious quantity, which will be chosen to sample the desired probability $\pi(R) = |\psi_T(R)|^2$. The “potential” U must then be determined by our wavefunction, so $\beta U(R) = -2 \log |\psi_T(R)|$.

Now consider the dynamics. A Brownian particle obeys a stochastic equation of motion, which in the overdamped case (inertial mass $\rightarrow 0$) takes the form,

$$dR = (BF)dt + dw, \quad (17)$$

where $F = -\nabla U$ is the force on the particle, B is the mobility, and dw is a random fluctuation with mean 0 and variance λdt . For classical systems, the fluctuation-dissipation theorem tells us that the temperature, (T) the mobility, and the amplitude of the fluctuations are related by $\lambda = Bk_B T$. With this relation, we find that the equation of motion is

$$dR = 2\lambda dt \nabla \log \psi_T + dw. \quad (18)$$

Note that λ sets relation between the fluctuations and the potential — larger λ leads to more fluctuations. In the next section λ will be identified with the inverse mass $\hbar^2/2m$ of particles in the diffusion Monte Carlo algorithm. Since the mobility B can be a symmetric tensor, λ can also be a tensor, an identification that is necessary for application of diffusion Monte Carlo to particles with anisotropic mass tensors.

In VMC calculations, the “time step” $\tau = dt$ is not really taken to be very small. Rather, Eq. (18) is used as a motivation for a transition probability $T(R \rightarrow R')$ of the form of a drifting Gaussian with variance $\lambda \tau$ and drift $2\lambda \tau \nabla \log \psi$, and used with the usual Metropolis algorithm. The rejection step in Metropolis ensures that the proper distribution is sampled for any value of τ . The rejection step allows any transition probability T , so the drifting Gaussian may be further modified (provided the reverse probability remains finite). For example, the maximum “drift” of the Gaussian is often limited to stabilize the algorithm in case $\nabla \log \psi_T$ can become arbitrarily large. For small τ the rejections become very rare and Langevin dynamics are recovered.

3.4. Diffusion Monte Carlo

Although VMC is useful for providing bounds to the energy and investigating forms for trial wavefunctions, it is desirable to have an automatic procedure for getting the true ground state energy. Diffusion Monte Carlo (DMC) is such a procedure. In DMC, the exact ground state $\Phi_0(R)$ is projected using the imaginary time propagator, $e^{-\beta H}$, where H is the many-body Hamiltonian and β is an interval in imaginary time. This property can be understood by expanding an initial trial function ψ_T in terms of the (unknown) exact eigenstates Φ_i and eigenvalues E_i of H . The propagated state is then,

$$e^{-\beta H} \psi_T = \sum_{i=0} e^{-\beta E_i} \Phi_i \langle \Phi_i | \psi_T \rangle. \quad (19)$$

For large β , the $i = 0$ term dominates and we have:

$$\lim_{\beta \rightarrow \infty} e^{-\beta H} \psi_T = \Phi_0 e^{-\beta E_0} \langle \Phi_0 | \psi_T \rangle. \quad (20)$$

For a long-but-finite projection time the second leading term has a relative coefficient $e^{-\beta(E_1 - E_0)}$. Therefore, the systematic error for a finite projection time is exponentially decaying with a decay rate of $(E_1 - E_0)$, with a leading factor that is proportional to the relative contribution of the first excited state in the trial wavefunction.

To sample the ground state wavefunction, factor the projection operator into many small steps,

$$\Phi_0(R_M) \propto \left[\prod_{i=1}^M e^{-\tau H} \right] \psi_T. \quad (21)$$

For small τ the kinetic and potential contributions of the Hamiltonian factor (this breakup is known as Trotter's theorem and is exact as $\tau \rightarrow 0$ [43]), and we have

$$\Phi_0 \propto \lim_{M \rightarrow \infty} \left[\prod_{i=1}^M \frac{1}{(4\pi\lambda\tau)^{Nd/2}} \int dR_{i-1} e^{-\frac{(R_i - R_{i-1})^2}{4\lambda\tau}} e^{-\tau V(R_{i-1})} \right] \psi_T(R_0). \quad (22)$$

If we sample the propagator at each step, this becomes a Markov process that samples the true ground state Φ_0 , after for a sufficiently long projection time β .

This method is known as diffusion Monte Carlo (DMC) because of its similarity to classical diffusion described in the last section. We show an outline of a simple DMC code in Fig. 3. There are two problems with Eq. (22): (i) we would rather sample $\Phi_0^* \Phi_0$, and (ii) the exponential factor $e^{-\tau V}$ is very unstable for many physical problems. One solution these problems is to work with path integrals, as discussed in the next section, Sec. 4. (Most closely related to this discussion is the variational path integral approach of Sec. 4.8, in which a trial density matrix $\rho_T = \psi_T \psi_T^*$ is projected to sample $\Phi_0^* \Phi_0$ [21]).

Importance sampling [7, 44] is needed to make the algorithm efficient. With importance sampling, the propagator $\langle R | \exp(-\beta H) | R' \rangle$ gets multiplied by the trial function:

$$P_{\psi_T}(R \rightarrow R') = \langle R | \exp(-\beta H) | R' \rangle \Psi_T(R)^{-1} \Psi_T(R') \quad (23)$$

DIFFUSION MONTE CARLO CODE

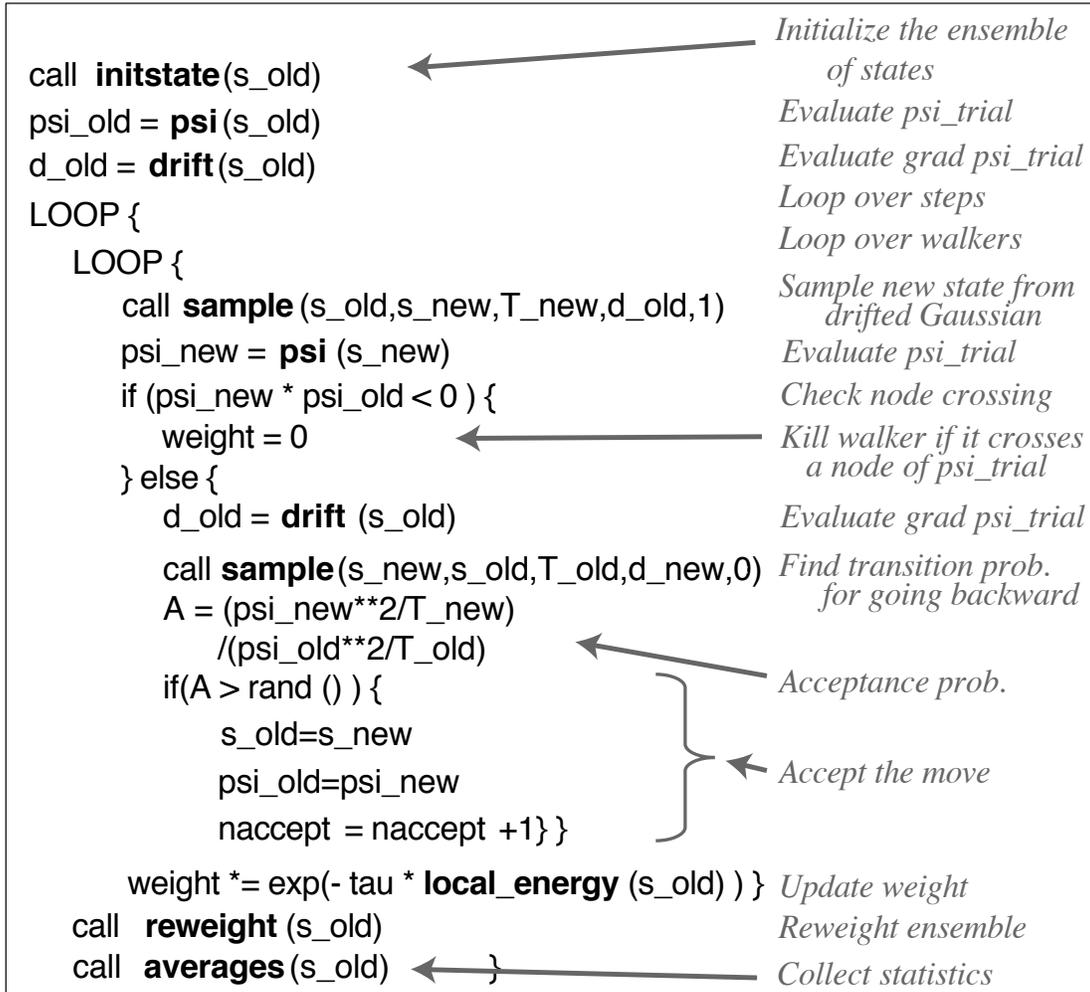


FIGURE 3. Outline of a fixed-node DMC code. The algorithm works with an ensemble of walkers, each with its own state and weight.

In the small time step limit, this propagator can be shown to be given by:

$$P_{\psi_T}(R \rightarrow R') = (4\pi\lambda\tau)^{-\frac{Nd}{2}} \exp \left[-\frac{(R' - R - \lambda\tau\nabla \ln |\Psi_T|^2)^2}{4\lambda\tau} - \tau E_L(R) \right]. \quad (24)$$

Remarkably, this is simply the drifted-Gaussian VMC propagator, Eq. (18), with an additional weight factor proportional to the local energy. Note that the Hamiltonian only enters through E_L in the weight factor; if you turn off the weight you actually sample the many-body Hamiltonian that has ψ_T for its ground state, recovering VMC. In practice, a constant term, the trial energy ϵ_0 , is subtracted from the local energy in the weight term. If $\epsilon_0 = E_0$, the true ground state energy, then the weights will be close to one.

The weight may be handled in different ways. One method, pure diffusion Monte Carlo, uses a large number of configurations, each carrying a cumulative weight. However, this procedure is inefficient, since eventually one configuration accumulates all of the weight. It is more efficient to have branching and death processes for the walkers to keep the weights close to one. At each step, the walker continues as N_{copies} new walkers, where

$$N_{\text{copies}} = \text{floor}[e^{-\tau(E_L - \epsilon_0)} + \zeta], \quad (25)$$

and ζ is a random variate, $0 \leq \zeta < 1$. If $N_{\text{copies}} = 0$ then the walker dies. It is necessary to provide some weak feedback to ϵ_0 in order to stabilize the population. The feedback introduces a bias, so it should be kept small and a large population $N_{\text{walkers}} \gg 1$ of walkers should be maintained. Detailed discussion of the algorithms with a view to reducing the time step error is given in Refs. [6, 45]

An alternative to branching is the use of transient estimates where a cumulative weight is carried along for each walker, but only for a finite projection time. Implementation of this method can be done with a cyclic buffer so that the old contributions to the weight can be discarded as the walker evolves. This method is more stable, however the projection time must be long enough to project out any excited states. Because of the weights and the added overhead of carrying along the weights, the efficiency is lower than with branching methods.

3.4.1. Estimators in diffusion Monte Carlo

Diffusion Monte Carlo has a serious complication in computing most physical observables. Since the quantity being sampled is $\psi_T^* \Phi_0$, averages (called mixed estimators) are of the form

$$\langle \mathcal{O} \rangle_{\text{DMC}} = \frac{\int dR \psi_T^*(R) \mathcal{O} \Phi_0(R)}{\int dR \psi_T^*(R) \Phi_0(R)}. \quad (26)$$

Physical observables are ground state expectations (*i. e.* $|\Phi_0|^2$), rather than the unphysical mixed estimators.

Operators \mathcal{O} that commute with the Hamiltonian, are unbiased. In particular, DMC has no trouble calculating total energies. Rewriting Φ_0 in terms of the projected trial state $\lim_{\beta \rightarrow \infty} e^{-\beta H} \psi_T$,

$$\begin{aligned} \langle \mathcal{O} \rangle_{\text{DMC}} &= \lim_{\beta \rightarrow \infty} \int dR \psi_T^*(R) \mathcal{O} e^{-\beta H} \psi_T(R) \\ &= \lim_{\beta \rightarrow \infty} \int dR \psi_T^*(R) e^{-\beta H/2} \mathcal{O} e^{-\beta H/2} \psi_T(R) \\ &= \int dR \Phi_0^*(R) \mathcal{O} \Phi_0(R), \end{aligned} \quad (27)$$

where commutation of \mathcal{O} and H was necessary to move the operator \mathcal{O} between the two halves of the projector.

Operators such as correlation functions, structure factors, the probability density, and even components of the total energy do not commute. To correct the mixed estimator, we assume that the trial wavefunction is close to the true ground state and extrapolate [36] the ground state estimate using the VMC estimate,

$$\langle \mathcal{O} \rangle = 2\langle \mathcal{O} \rangle_{\text{DMC}} - \langle \mathcal{O} \rangle_{\text{VMC}}. \quad (28)$$

This extrapolation is, of course, an uncontrolled approximation, and should be used with caution.

A better way to deal with these operators in DMC is to use forward walking algorithms [46, 47]. One calculates projected expectation values

$$\langle \mathcal{O} \rangle_{t_{\text{proj}}} = \lim_{\beta \rightarrow \infty} \langle \Psi_T | e^{-H t_{\text{proj}}} \mathcal{O} e^{-\beta H} | \Psi_T \rangle, \quad (29)$$

where t_{proj} is the projection time. Large values of t_{proj} give estimates $\langle \mathcal{O} \rangle_{t_{\text{proj}}}$ that approach the true ground state value, at the cost of an increasing variance. To implement this technique, store the estimator values for each walker in a cyclic buffer. As the random walk proceeds, the properties are copied along with the walkers. The measurements stored on the walkers automatically “age” through DMC weighting and branching processes. It is useful to record estimates for many different projection times during the simulation to determine whether the projection has converged before the variance becomes too large.

3.5. The fermion sign problem

Fermion statistics introduce a serious problem in the diffusion Monte Carlo algorithm. The diffusion propagator must be antisymmetrized introducing negative signs, which cannot be interpreted as classical probabilities. To calculate properties of general quantum states, existing algorithms are either approximate or scale exponentially with the number of particles. For example, Variational Monte Carlo has no fermion sign problem since $|\Psi_T|^2$ is always positive, however the results are biased by the assumed trial function. We first discuss an approximate procedure, then an exact procedure for handling fermion antisymmetry.

3.5.1. The fixed node approximation

It is always possible to take the ground state wavefunction to be real since if a complex wavefunction is a solution so are its real and imaginary parts. Its *nodes* are the locations where the trial function vanishes in the $3N$ dimensional coordinate space. In the fixed node (FN) approximation [4, 5], one requires that the projected wavefunction is restricted to have the same nodes as a real trial function Ψ_N . Effectively, an infinite potential barrier is placed on the nodes of the trial function in the Nd dimensional configuration space, and with this barrier the Bosonic DMC algorithm is applied. The fixed-node constraint ensures that the sampled probability density remains non-negative.

The fixed-node energy is the lowest energy consistent with the assumed nodal surfaces. As such, it must lie between the variational energy and the exact energy. The exact nodes give the true ground state energy, and all other choices give higher energies. Typically 70% to 90% of the error in the variational energy is recovered with the fixed-node approximation [48]. This is accomplished with no additional work in finding better trial functions, simply implementing the branching procedure on top of a VMC algorithm.

In calculations on strongly confined quantum dots, the nodal surfaces are often very well approximated by the trial wavefunction, and FN DMC can obtain very accurate ground state energies.

3.5.2. Release node methods

For systems with smaller number of particles, or when the need for high precision merits longer computer times, there is an automatic way to improve the FN results, known as released node (RN) Monte Carlo [49]. The idea is to let the fermion sign back into the problem, but in a controlled way. Released node is very similar to transient estimate DMC. Both methods are exact for fermions, although released node is tailored to take advantage of the FN DMC algorithm.

When a walker crosses a node, it is allowed to continue (with the proper change of sign) for a certain predefined lifetime, T_R , after which time it dies. If T_R is zero, this is simply the FN DMC algorithm. The estimated energy converges monotonically and exponentially fast to the exact energy. It is desirable to project as long as possible, subject to the exponential growth of the statistical noise as T_R increases.

In the implementation of RN DMC, one samples many different choices of T_R with the same simulation run. Each walker gets a timer t that is initialized to $t = 0$. The first time a walker crosses a node, its timer is started. To sample an estimator with projection time T_R , only those walkers with timers $t < T_R$ are included in the average. Whenever a walkers timer exceeds the largest projection time being examined the walker is killed.

The expectation values may be plotted versus projection time to check for convergence. In this way, it is possible to get exact ground state energy estimates from long simulations of finite fermion systems. Of course, the fermion sign problem is still present, so extension to larger systems or longer projection times T_R eventually becomes exponentially more difficult.

3.6. Excited states

A QMC method to calculate excited state energies is presented in Refs. [50, 51]. A summary follows. Suppose we choose a set of m trial wavefunctions, f_α , $\alpha = 1, \dots, m$ and ask for the lowest energy in the space spanned by this set. The result is the general-

ized eigenvalue problem [50]

$$\sum_{\beta=1}^m [H_{\alpha\beta} - \Lambda_k N_{\alpha\beta}] d_{k\beta} = 0, \quad (30)$$

where $\phi_k = \sum_{\beta} d_{k\beta} f_{\beta}$ is the k^{th} eigenvector with eigenvalue Λ_k and the matrices N and H are the overlap and Hamiltonian matrices in our trial basis:

$$N_{\alpha\beta} = \int dR_1 dR_2 f_{\alpha}(R_2) f_{\beta}(R_1), \quad (31)$$

$$H_{\alpha\beta} = \int dR_1 dR_2 H f_{\alpha}(R_2) f_{\beta}(R_1). \quad (32)$$

McDonald's theorem ensures that $\Lambda_k \geq E_k$ for each k . The matrix elements can be computed by a Metropolis random walk, sampling some trial function $|\psi|^2$. A good choice for the sampling function is $|\psi|^2 = \sum_{k=1}^m |\phi_k|^2$, where ϕ_k is an approximate eigenfunction. With a single random walk, all the needed matrix elements are computed at once.

This procedure can be improved using projection on the basis with the operator e^{-tH} as was done in diffusion Monte Carlo. In the limit of large t the eigenvalues Λ_k approach the low energy eigenvalues E_k within the Hilbert space spanned by the trial functions $\{f_{\alpha}\}$. To sample the matrix elements of N and H , a random walk is performed using the guiding function ψ , as described in Section 3.3. At each step in the walk, the coordinates R_{i+1} of the particles are updated using the guiding function,

$$R^{(i+1)} = R^{(i)} + \tau \lambda \nabla \log |\psi(R^{(i)})|^2 + (2\tau\lambda)^{\frac{1}{2}} \chi_i, \quad (33)$$

where χ_i is a normally distributed random with zero mean and unit variance and τ is the time step. The matrix element of e^{-tH} may be estimated by integrating the local energy of the guiding function, $E_{L\psi}(R)$, along the block,

$$W_{n,n+k} = \exp \left\{ -\tau \sum_{j=n}^{n+k-1} \frac{1}{2} [E_{L\psi}(R_j) + E_{L\psi}(R_{j+1})] \right\}, \quad (34)$$

where the local energy of the guiding function $E_{L\psi}$ is given by

$$E_{L\psi} = \psi^{-1}(R) H \psi(R). \quad (35)$$

Thus, the weight $W_{n,n+k}$ is the weight of the walker at step $n+k$, given by the product of the weights the previous k steps.

The matrices N and H may be estimated by

$$n_{\alpha\beta}(k\tau) = \frac{1}{P} \sum_{i=1}^P F_{\alpha}(R_i) W_{i,i+k} F_{\beta}(R_{i+k}) \quad (36)$$

and

$$h_{\alpha\beta}(k\tau) = \frac{1}{P} \sum_{i=1}^P F_{\alpha}(R_i) W_{i,i+k} F_{\beta}(R_{i+k}) E_{L\psi}(R_{i+k}), \quad (37)$$

where $F_\alpha = f_\alpha/\psi$ and $E_{L\beta} = f_\beta^{-1}(R)Hf_\beta(R)$ is the local energy of the trial wavefunction.

The excited state technique has been used to calculate scattering lengths for exciton-exciton scattering in semiconductors [52], discussed in Sec. 3.11.3. The technique may be also useful for calculating tunneling rates or lifetimes in nanostructures.

3.7. Pseudopotentials

As described in section 2.1, *ab initio* calculations often use pseudopotentials to remove core electrons. For local potentials, the diffusion Monte Carlo technique is straightforward to apply. Non-local pseudopotentials, found in the pseudopotential Hamiltonian of Eq. (2), lead to a non-local contributions to the propagator, $\exp[-\tau\hat{V}_\ell^{(j)}(|r_i - R_j|)]$. This non-local propagation could be implemented sampling the angular distribution, so that electron walks discontinuously hop around the core while acquiring a phase and a weight. An exact evaluation of the non-local pseudopotential propagator would lead to a sign problem that seriously reduces the efficiency. To avoid this sign problem, pseudopotentials are commonly handled with the frozen core approximation [53, 54, 32]. The pseudopotential requires extra work to evaluate angular integrals of the non-local pseudopotentials. This approach is successful for systems containing atoms beyond the first row of the periodic table [32].

3.8. Position-dependent mass and pseudo-Hamiltonians

In some models, the Hamiltonian contains a position-dependent mass. For example, effective mass models of semiconductor heterostructures have a mass tensor $M(\mathbf{r})$ that depends on the local material and strain. The generalization of the kinetic energy must be Hermitian in order to satisfy the continuity equation, thus the tensor mass function $\lambda(\mathbf{r}) = \hbar^2 M^{-1}(\mathbf{r})/2$ gets sandwiched between the gradient operators,

$$H_{\text{KE}}\psi(\mathbf{r}) = -\nabla \cdot \lambda(\mathbf{r}) \cdot \nabla \psi(\mathbf{r}). \quad (38)$$

Such a Hamiltonian term also occurs in the pseudo-Hamiltonian technique for removing core electrons [55, 56]. The propagator for Eq. (38) involves derivatives of the mass tensor,

$$\mathcal{P}(r' \rightarrow r) = \frac{1}{\sqrt{(4\pi\tau)^3 \det|\lambda|}} \exp \left[-\frac{(\Delta\mathbf{r}) \cdot \lambda^{-1} \cdot (\Delta\mathbf{r})}{4\tau} - \tau E_L \right], \quad (39)$$

where

$$\Delta\mathbf{r} = \mathbf{r} - \mathbf{r}' - \lambda \tau \nabla |\ln \Psi_T|^2 - \tau \nabla \lambda. \quad (40)$$

3.9. Other Hamiltonians

In effective-mass semiconductor physics, one must consider coupling of multiple bands. For example, the Luttinger Hamiltonian [57],

$$H_L = \frac{\hbar^2}{2m} \left[\left(\gamma_1 + \frac{5}{2} \gamma_2 \right) \nabla^2 - 2\gamma_2 (\nabla \cdot \mathbf{J})^2 + 2(\gamma_3 - \gamma_2) (\nabla_x^2 J_x^2 + \nabla_y^2 J_y^2 + \nabla_z^2 J_z^2) \right], \quad (41)$$

describes valence band mixing of light and heavy holes in cubic semiconductors. Here γ_1 , γ_2 , and γ_3 are the Kohn-Luttinger parameters and \mathbf{J} is an operator that acts of the four valence bands like the angular momentum operator acts on an atomic $J = 3/2$ state. These multiple band Hamiltonians are similar to model Hamiltonians used in nuclear QMC calculations [58, 59]. To our knowledge, multiband QMC has not been attempted for semiconductor nanostructures.

3.10. Fixed phase and generalized boundary conditions

There are many physical situations which lead to complex valued ground state wavefunctions. For example, one may also wish to study excited states with definite linear or angular momenta. For electrons in a magnetic field with vector potential \mathbf{A} , the Hamiltonian is

$$H = \sum_i \frac{1}{2m} \left(\frac{\hbar}{i} \nabla_i - \frac{e}{c} \mathbf{A}(\mathbf{r}_i) \right)^2 + V(R). \quad (42)$$

leading to a complex-valued Hamiltonian. While the equations in previous sections may be generalized to complex wavefunctions, the propagators pick up phases that make Monte Carlo sampling inefficient. Ortiz *et al.* [18] introduced a fixed-phase approximation that allows projection to the lowest energy state for these cases.

Starting from a complex trial function $\Psi_T(R)$, we vary the modulus to minimize the total energy, while keeping the phase fixed. Writing the wavefunction as $\Psi_F(R) = F(R)\Psi_T(R)$, where $F(R)$ is a real positive function, its energy is

$$E[F] = \int \Psi_T^*(R) F^*(R) H F(R) \Psi_T(R) dR. \quad (43)$$

By requiring the functional derivative of $E[F]$ with respect to $F(R)$ to be zero, subject to the normalization constraint we find a Schrödinger-like equation for $|\Psi_T|^2 F$,

$$\left[- \sum_i \frac{\hbar^2}{2m_i} \nabla_i^2 + \sum_i \frac{\hbar^2}{2m_i} \nabla_i \cdot \left(\frac{\nabla_i |\Psi_T|^2}{|\Psi_T|^2} \right) + \Re E_L(R) \right] |\Psi_T|^2 F = \varepsilon |\Psi_T|^2 F, \quad (44)$$

where $E_L(R) = \Psi_T^{-1}(R) H \Psi_T(R)$ is the local energy of the trial wavefunction. For the Hamiltonian, Eq. (42), the local energy is

$$E_L = \left[\sum_i - \frac{\hbar^2}{2m} \left(\frac{\nabla_i^2 \Psi_T}{\Psi_T} \right) + \frac{ie\hbar \mathbf{A}(\mathbf{r}_i)}{mc} \cdot \left(\frac{\nabla_i \Psi_T}{\Psi_T} \right) + V(R) - \frac{e^2 A^2(\mathbf{r}_i)}{2mc^2} \right]. \quad (45)$$

Thus the fixed-phase approximation may be implemented by simply working with complex trial wavefunctions and using the real value of the drift and local energy in the DMC algorithm.

3.11. Zero temperature examples and applications

Applications of zero temperature quantum Monte Carlo algorithms to nanostructures have fallen into two categories: *ab initio* pseudopotential calculations on fullerenes and colloidal dots and effective mass calculations on semiconductor quantum dots.

3.11.1. Fullerenes

Spherical carbon molecules, or fullerenes, have been an active field of research since the discovery of C_{60} in 1985. Quantum Monte Carlo calculations have been aimed at determining the stability and structure of these carbon cages. All these studies have used the DMC algorithm with fixed node approximation and frozen-core-approximated non-local pseudopotentials. Studies by Grossman and Mitas established that the computing relative energies of C_{20} isomers requires the high accuracy of QMC [60]. Surprisingly, very different energetic orderings of ring, bowl, and cage structures are predicted by Hartree-Fock, local density approximation, generalized gradient approximation, and QMC calculations. The prediction by QMC that the closed C_{20} cage is energetically unfavorable was independently confirmed by later quantum chemistry calculations [61].

More recent calculations by Kent, Needs, and Rajagopal have studied carbon clusters near the the crossover to fullerene stability [62]. Torelli and Mitas have looked at the Peierls instability in carbon rings [63]. Using QMC methods adapted to larger systems, Williamson, Hood, and Grossman [40] used DMC to calculate binding energies of a series of fullerene cages: C_{20} , C_{36} , C_{60} , C_{80} , and C_{180} . These calculations illustrate that future DMC studies will be able tackle these larger and more common fullerenes as well as carbon nanotubes.

3.11.2. Colloidal quantum dots

Hydrogen passivated silicon nanoclusters have become a common system for the theoretical study of band gaps. Starting from the silane molecule, SiH_4 , the optical gap of these clusters have been computed with many methods, including: density functional theory, time-dependent density functional theory, *GW* calculations, empirical pseudopotential methods, and fixed-node quantum Monte Carlo.

Grossman and Mitas studied very small ($n \leq 20$) unpassivated clusters [64]. These calculations focused on the role of correlation in the total energies of different isomers. Later calculations by Mitas *et al* combined LDA, DMC, and *ab initio* molecular dynamics to study structure and stability of Si_{20} and Si_{25} clusters [65].

Calculations on hydrogen passivated clusters have taken advantage of linear-scaling algorithms to study clusters with over 200 Si atoms [40]. Williamson *et al.* [41] used DMC to calculate optical gaps of silicon quantum dots with diameters up to 1.5 nm. Since DMC is a ground state technique, they had to take advantage of the nodal constraint to study the first excited state. First they calculated the ground state energy, using atomic positions and nodes from a density functional calculation. To get the energy of the first excited state, they formed excited state nodes by replacing the highest occupied molecular orbital (HOMO) with the lowest unoccupied molecular orbital (LUMO) in the Slater determinant. They then approximated the optical gap as the difference of this DMC-calculated first excited state energy and the DMC ground state energy and compared with time-dependent density functional theory, tight binding, various density functionals, and the *GW* approximation. The QMC optical gaps are significantly higher than those of most other methods, and the good agreement between *GW* and QMC at small sizes suggests QMC gives accurate predications for optical gaps in these systems.

Other publications have combined QMC, density functional theory, and *ab initio* molecular dynamics to study the effects of surface chemistry [66], synthesis conditions [42], and surface reconstructions [67] on the stability and optical properties of silicon nanocrystals.

3.11.3. *Effective mass models of quantum structures*

The effective mass models lend themselves to efficient QMC simulations. There have been a number of applications to several different nanostructures:

Spherical quantum dots. The first QMC calculations in quantum dots are for spherical models. Although many experimental dots have a spherical shape, the presence of multiple valence or conduction bands often limits single-band effective mass models to descriptions of correlation trends, rather than detailed quantitative predictions [68]. Austin determined the binding energy of excitons in a spherical dot as a function of dot radius [69] with DMC. For a GaAs dot embedded in AlGaAs, it was shown that the exciton binding energy is maximum for an intermediate radius of about 30 Å. Pollock *et al.* used a path integral technique to calculate energies for excitons and biexcitons as a function of radius [70, 71], using an infinite confining potential. These calculations also included temperature dependence, as discussed in Sec. 4.9.

Shumway *et al.* performed a systematic study of the role of correlation in finite spherical dots with QMC, and checked a commonly-used truncation of configuration interaction expansions [68]. A finite confining potential was used, with parameters taken to represent typical colloidal and embedded dots. These calculations showed a peak in the exciton binding energy at an intermediate dot radius and a much less pronounced peak in the biexciton binding energy. The relative contributions of perturbation theory, self-consistency, and correlation were quantified, showing that biexciton binding arises entirely from correlation. The comparison to configuration interaction calculations revealed that the usual practice of not including continuum states in the wave expansion misses about half the biexciton binding energy. However, the ability of QMC to capture

all the correlation must be weighed against detailed spectra that result from configuration interaction calculations. Other calculations on spherical models include DMC calculation of exciton and biexciton energies by Tsuchiya [72], and stochastic variational calculations by Varga *et al.* [73].

Self-assembled and gated quantum dots. The experimental interest in self-assembled and gated quantum dots led to much theoretical work on two-dimensional parabolic dots [20]. These models are often applied to electrons, since holes are more complicated to treat because of several near-degenerate valence bands. Strain lifts this degeneracy, so two-component (electron and hole) systems could be studied, but most QMC calculations do not consider this. There have been several DMC calculations on interacting electrons in parabolic dots [74, 75, 76, 77], ranging from two electrons up to as many as twenty. Most of these calculations focus exclusively on total energies and comparison to Hartree-Fock and LSDA. Shumway *et al.* [78] used DMC to calculate total energies for electron addition to a realistic single-band model of a strained InAs/GaAs pyramidal dot, and found LSDA captured the essential correlation energy of electrons in the strongly confined regime. The general trend is that while correlation is important to multiexciton energies in all size dots, charged dots containing only electrons are highly correlated only for very low densities.

Bolton [74] used a fixed-phase approximation [18, 74] to study two, three, and four harmonically confined electrons in a magnetic field. Special trial wavefunctions were used for these low numbers of electrons, and the resulting fixed-phase compared well with exact diagonalization. A general trial wavefunction was proposed that would work for larger numbers of electrons, and the resulting energies agreed reasonably well for the special cases $N \leq 4$. This work preceded the released-phase algorithm of Jones, Ortiz, and Ceperley [19], so released phase was not attempted on this system.

Quantum rings. For certain growth conditions, the centers of large self-assembled dots may be depleted, leading to ring-shaped InAs/GaAs nanostructures. There has been interest in the electronic structure of these rings, as they can trap magnetic flux in their cores, and also resemble one-dimensional wires with periodic boundary conditions. These ring structures have been studied with QMC [79, 80].

Related calculations: exciton-exciton scattering. Excitons are bound states of an electron and a hole, and obey Bose statistics. A low density exciton gas is an experimentally realizable dilute Bose gas. The scattering length a_s of a dilute Bose gas determines its properties, but is difficult to calculate. Shumway and Ceperley [52] performed an essentially exact QMC treatment of exciton-exciton scattering. By fixing nodes in the extended scattering states, the energies of scattering states were calculated from excited state DMC [50, 52]. The dependence of the scattering state energies on node positions determines the phase shifts and scattering lengths. A similar QMC technique has been used much earlier by Carlson and Pandharipande to calculate nuclear cross sections [81]. The exciton calculations found scattering lengths for different spin orientations of the excitons and a significant triplet-triplet to singlet-singlet scattering process [52]. At some mass ratios m_h/m_e , the scattering lengths diverge in conjunction with the appearance of biexciton vibrational states, an effect not found by earlier perturbative treatments.

While these scattering calculations were performed for bulk excitons and did not include nanostructures, the technique should be adaptable to scattering in quantum wells and wires or scattering from bound electrons in quantum dots.

4. PATH INTEGRAL METHODS

Path integral method allow for simulations at non-zero temperatures, without the necessity for considering individual excited states and without construction of trial wave functions. Because of this they are particularly suited for systems undergoing phase transitions. They are based on the same projection operator $\exp(-\beta H)$ as is diffusion Monte Carlo, but the initial and final conditions of the state are treated differently. For more details in this section see the review article by Ceperley [21].

4.1. Basic theory

Static and dynamical properties of a quantum system in thermal equilibrium are obtainable from the thermal density matrix. Again, let us denote the exact eigenvalues and eigenfunctions of a Hamiltonian \mathcal{H} by ϕ_i and E_i . In thermal equilibrium, the probability of a given state i being occupied is $e^{-E_i/k_B T}$, with T the temperature. The equilibrium value of an operator O is

$$\langle O \rangle = Z^{-1} \sum_i \langle \phi_i | O | \phi_i \rangle e^{-\beta E_i}, \quad (46)$$

where the partition function is

$$Z = \sum_i e^{-\beta E_i}, \quad (47)$$

and $\beta = 1/k_B T$. The operator, $e^{-\beta \mathcal{H}}$ is the density matrix. Although the above traces can be carried out in any complete basis, in continuum PIMC one almost always works in a *position basis* because then all of the elements of the density matrix are non-negative and can be interpreted as a probability. The position-space density matrix is

$$\rho(R, R'; \beta) = \langle R | e^{-\beta \mathcal{H}} | R' \rangle = \sum_i \phi_i^*(R) \phi_i(R') e^{-\beta E_i}, \quad (48)$$

where: $R = \{\mathbf{r}_1, \dots, \mathbf{r}_N\}$ and \mathbf{r}_i is the position of the i^{th} particle.

The path-integral formula for the many-body density matrix is arrived at by using the product property of Green's functions M times, giving an expression for the density matrix at a temperature T , in terms of density matrices at a temperature MT .

$$e^{-\beta \mathcal{H}} = \left(e^{-\tau \mathcal{H}} \right)^M, \quad (49)$$

where the *time step* is $\tau = \beta/M$. Written in the position representation,

$$\rho(R_0, R_M; \beta) = \int \dots \int dR_1 dR_2 \dots dR_{M-1} \rho(R_0, R_1; \tau) \rho(R_1, R_2; \tau) \dots \rho(R_{M-1}, R_M; \tau). \quad (50)$$

For MT large enough, we can write down a sufficiently accurate approximation to the density matrix giving an explicit form for the low-temperature density matrix. Suppose the Hamiltonian is split into the kinetic and potential energy, $\mathcal{H} = \mathcal{T} + \mathcal{V}$. The *primitive* approximation to the density matrix neglects their commutator:

$$e^{-\tau(\mathcal{T}+\mathcal{V})} \approx e^{-\tau\mathcal{T}} e^{-\tau\mathcal{V}}. \quad (51)$$

One might worry that this will lead to an error in the limit as $M \rightarrow \infty$ with small errors building up to a finite error. According to the Trotter formula [43] given earlier, one does not have to worry:

$$e^{-\beta(\mathcal{T}+\mathcal{V})} = \lim_{M \rightarrow \infty} \left[e^{-\tau\mathcal{T}} e^{-\tau\mathcal{V}} \right]^M. \quad (52)$$

The Trotter formula holds if the three operators \mathcal{T} , \mathcal{V} and $\mathcal{T} + \mathcal{V}$ are self-adjoint and make sense separately, for example, if their spectrum is bounded below [82]. The kinetic matrix can be evaluated using the eigenfunction expansion of \mathcal{T} ,

$$\langle R_0 | e^{-\tau\mathcal{T}} | R_1 \rangle = \sum_{\mathbf{n}} L^{-3N} e^{-\tau\lambda K_n^2 - iK_n(R_0 - R_1)} \quad (53)$$

$$= (4\pi\lambda\tau)^{-3N/2} \exp\left[-\frac{(R_0 - R_1)^2}{4\lambda\tau}\right]. \quad (54)$$

Equation (54) is obtained by approximating the sum by an integral, appropriate only if the thermal wavelength of one step is much less than the size of the box.

Using Eqs. (50), (52), and (54), we arrive at the discrete path-integral expression for the density matrix in the primitive approximation:

$$\rho(R_0, R_M; \beta) = \int dR_1 \dots dR_{M-1} (4\pi\lambda\tau)^{-3NM/2} \exp\left(-\sum_{m=1}^M \left[\frac{(R_{m-1} - R_m)^2}{4\lambda\tau} + \tau V(R_m)\right]\right). \quad (55)$$

This expression relates the quantum density matrix at any temperature to integrals over the path $R_1 \dots R_{M-1}$ of something that is like a classical Maxwell-Boltzmann distribution function. This is the famous mapping from a quantum system to a classical system. The Feynman-Kacs formula is obtained by taking the limit $M \rightarrow \infty$, making a continuous path.

A single R_k is referred to as the k^{th} *time slice* or a bead. The *path* is the sequence of points $\{R_0, R_1, \dots, R_{M-1}, R_M\}$. The *time* associated with the point R_k is defined as $t_k = k\tau$. A *link* m is a pair of time slices (R_{m-1}, R_m) separated by time τ . The *action* of a link is defined as *minus* the logarithm of the density matrix:

$$S^m \equiv S(R_{m-1}, R_m; \tau) \equiv -\ln[\rho(R_{m-1}, R_m; \tau)]. \quad (56)$$

We can interpret the path-integral expression, Eq. (55), as a classical configuration integral; the action is analogous to a classical potential-energy function divided by $k_B T$. In the classical analog, the kinetic link action corresponds to a spring potential connecting beads representing the same atom in successive time slices. The classical system is a chain of beads connected with springs. We call such a chain a *polymer*. In fact, the bead-spring model of real-life polymers is a useful picture. The potential action represents interactions between beads of different atoms, keeping the polymers out of each other's way (for a repulsive potential). The potential is represented by an interpolymeric, potential which is peculiar from the classical point of view in that it interacts only at the same "time" and only between beads on different chains.

Thermodynamical properties, or static properties diagonal in configuration space, are determined by the trace of the density matrix, *i.e.*, the integral of Eq. (55) over R_0 with $R_0 = R_M$. The formula for diagonal elements of the density matrix then involves a path that returns to its starting place after M steps: a *ring polymer*. Because the partition function of the quantum system is equal to the partition function of the classical system, there is an exact, systematic procedure for understanding many properties of quantum systems purely in terms of classical statistical mechanics. Anything about a quantum system that can be written as matrix elements of the density matrix, has a classical statistical-mechanical analog.

Shown in Fig. 4 are pictures of two-dimensional path integrals. What is usually done is to plot the *trace* of the paths in the x-y plane, projecting out the "time" coordinate. Shown are six distinguishable helium atoms in a periodic square at 2 K with 80 time slices. The filled circles are markers for the beginning of imaginary time, *i.e.*, R_1 . There is nothing special about that value of imaginary time, but it is useful to place a mark at a common bead for each polymer, since the potential acts only at equal "times." Path integrals are like directed polymers: the paths have a direction.

Shown in Fig. 5 are the same paths projected onto the (x,t) and (y,t) planes; the vertical axis is imaginary time. This is the *world-line* perspective. Both real polymers and these paths are "fractals."

One should always keep in mind that all these pictures show are points sampled from a product of thermal density matrices. They have only an indirect relationship to real-time dynamics or paths. They tell us mathematically precise information, but it must be correctly interpreted. Correlations along the paths are the Laplace transform of real-time linear-response functions. Even though the imaginary-time dynamics is not directly real-time dynamics, it is very important. For example, it is impossible to tell from a single time slice whether or not a liquid is a superfluid, but one can recognize a superfluid by examining the connection of the paths in imaginary time.

4.2. The action

It is clearly desirable to make a good but cheap approximation to the exact link action. The better we can make the individual link action, the fewer the number of time slices and the shorter the "polymer." The sampling becomes much easier as the paths have fewer links and the estimation of various quantities such as the kinetic energy have

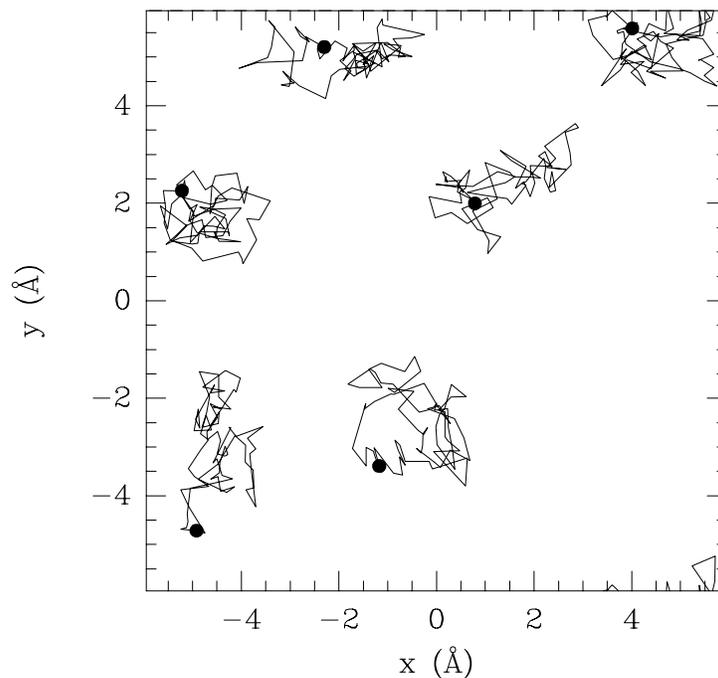


FIGURE 4. The trace of the paths of six helium atoms at a temperature of 2 K with 80 steps on the paths. The filled circles are markers for the (arbitrary) beginning of the path. Paths that exit on one side of the square reenter on the other side. Successive beads are connected with straight lines.

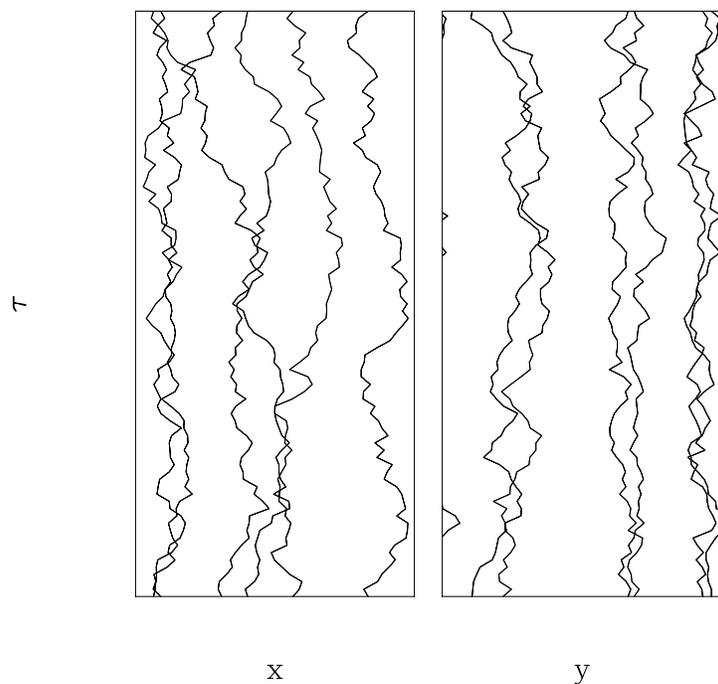


FIGURE 5. A world-line diagram of the paths shown in Fig 4. The vertical axis is imaginary time.

smaller statistical fluctuations.

The Feynman-Kac (FK) formula is a guide to better actions. Using it, we can write the exact action as

$$e^{-U(R_0, R_F; \tau)} = \left\langle \exp \left[- \int_0^\tau V(R(t)) dt \right] \right\rangle_{RW}, \quad (57)$$

where $\langle \dots \rangle_{RW}$ means to average over all Gaussian random walks (bridges) from R_0 to R_F in a “time” τ . One can sample Brownian paths and to evaluate the exponential in the FK formula. One gets into trouble for many-body systems or at low temperatures, since the exponential can fluctuate wildly. It is possible to use direct sampling to estimate the exact action in special situations. The results can be used either to suggest new forms or to adjust parameters in an existing form.

The task of finding a good action is different from that of finding a good integrator for an ordinary differential equation because of the fractal nature of the paths. Since paths do not have continuous derivatives, predictor-corrector or leapfrog methods are not as useful. On the other hand, statistical methods for improving the convergence work very well. Several possible strategies for choosing the high-temperature density matrix will now be examined. All of them converge to the correct answer at sufficiently high temperature, but their rates of convergence depend on the potential energy.

An easily computed property of the exact action is its behavior as two particles approach each other. It can be shown that the divergent part of the action should approach a two-particle form. For a Coulomb interaction, this condition leads to a cusp condition on the action at $r = 0$:

$$\lim_{r_{ij} \rightarrow 0} \frac{dU(R, R'; \tau)}{dr_{ij}} = - \frac{e_i e_j}{(d-1)(\lambda_i + \lambda_j)}. \quad (58)$$

Here e_i is the charge on particle i and $\lambda_i = \hbar^2/2m_i$. The path averaging in the FK formula smooths the potential so that the action is finite at the origin instead of having a r^{-1} singularity. For a Lennard-Jones, r^{-12} potential, the action at small r must diverge as r^{-5} . One can also derive exact properties of the action at large distances by considering the action as a function of the fourier transform of the density and then going to the long-wavelength limit.

Semiclassical methods for improving the action rely on the fact that, at very high temperatures, the major contribution to the FK path-integral comes from paths neighboring a single “classical” path. For small imaginary-times the optimal path is a straight line connecting R_0 and R_F , and the action is the integral of the potential energy along this straight-line path,

$$U_{SC}(R_0, R_F; \tau) = \tau \int_0^1 ds V(R_0 + (R_F - R_0)s). \quad (59)$$

In another approach, we assume that the potential is quadratic in the region near the 2 endpoints. The exact density matrix for a quadratic potential is a Gaussian [83]. The

result to order λ is

$$U_H(R_0, R_F; \tau) = \tau V(R^*) + \frac{\tau^2 \lambda}{6} \nabla^2 V(R^*) \quad (60)$$

$$- \frac{\tau}{12} (R_F - R_0) \nabla \nabla V(R^*) (R_F - R_0) - \frac{\tau^3 \lambda}{12} [\nabla V(R^*)]^2.$$

This expansion is equivalent to the Wigner-Kirkwood or \hbar expansion of the action if we make the choice of $R^* = R_0$. See, for example, the derivation for the diagonal element in Landau and Lifshitz [84]. Other choices are possible.

Although one is picking up higher-order contributions with the harmonic expansion, it does not uniformly converge for a hard potential. At large r , where the potential is small, the expansion is adequate, but at small r , where quantum effects are very important, all terms in the expansion are large. Suppose the potential goes as r^{-12} at small r . Then the second and third terms will diverge as r^{-14} while the last term will diverge as r^{-26} at small r . In fact, quantum diffraction causes the exact action to diverge only as r^{-5} . For the coulomb potential, the cusp condition is not arrived at with this expansion.

One can try to choose the harmonic reference system better. One strategy is to choose it to minimize the free energy of the trial action; see, for example Ref. [85]. One can also evaluate the exact harmonic action instead of keeping only terms of order λ . These improvements could work in a system that was nearly harmonic, for example, a nearly classical solid or a molecule with small zero-point vibrations. However, the improvements will not help much in regions where the potential is highly anharmonic.

If the potential energy is sufficiently smooth, one can take the averaging process of the FK formula up into the exponent. The cumulant expansion [86] of the exponential of a random variate,

$$\langle e^x \rangle = \exp[\langle x \rangle + \frac{1}{2}(\langle x^2 \rangle - \langle x \rangle^2) + \dots]. \quad (61)$$

Here x is some random variate, the brackets represent an average, and higher-order cumulants have been neglected. If x is normally distributed, as it would if it were the sum of many uncorrelated terms, the higher-order cumulants are zero. Now let $x = \int_0^\tau dt V(R(t))$, the total potential of a path, and average over all free-particle paths from R_0 to R_F . The first-order cumulant action is,

$$U_C(R_0, R_F; \tau) = \int_0^\tau dt \langle V(R(t)) \rangle_\mu, \quad (62)$$

where we average over the total density of points, $\mu(R)$, the *cumulant density*, sampled by the random walk in going from R_0 to R_F .

A better approach for a hard-sphere-like system is to determine the exact action for two atoms and then to use that to construct a many-body action [87].

$$V(R) = \sum_{i < j} v(r_i - r_j). \quad (63)$$

Now apply the Feynman-Kacs formula, Eq. (57). What enters is the integral of the potential energy along a path. Let x_{ij} be the exponentiated integral of the pair energy

along a random walk,

$$x_{ij} = \exp \left[- \int_0^\tau dt v(\mathbf{r}_{ij}(t)) \right]. \quad (64)$$

Then x_{ij} is a random variable drawn from some distribution function that depends on the end points (R_0, R_F). If the variables x_{ij} are uncorrelated with each other, we can interchange the product and averaging operation,

$$e^{-U} = \left\langle \prod_{i<j} x_{ij} \right\rangle \approx \prod_{i<j} \langle x_{ij} \rangle. \quad (65)$$

The average is the interacting part of the *exact* action for a *pair* of atoms. The *pair-product* action is

$$U_2(R, R'; \tau) = \sum_{i<j} u_2(\mathbf{r}_{ij}, \mathbf{r}'_{ij}; \tau), \quad (66)$$

where $u_2(\mathbf{r}_{ij}, \mathbf{r}'_{ij}; \tau)$ is the exact action for a pair of atoms.

This approximation has several advantages: the pair product will be correct to lowest order in a density expansion of the action, since it is only when we have three atoms in close proximity that we make a substantial error. The pair density approximation is perfectly well defined for all potentials, either Coulombic or hard-sphere-like, in contrast to the cumulant action, which exists only for integrable potentials. At low temperatures the pair action approaches the pair-product or Jastrow ground-state wave function discussed earlier.

The *end-point* approximation to the pair action,

$$u_{2E}(\mathbf{r}_{ij}, \mathbf{r}'_{ij}; \tau) = \frac{1}{2} [u_2(\mathbf{r}_{ij}, \mathbf{r}_{ij}; \tau) + u_2(\mathbf{r}'_{ij}, \mathbf{r}'_{ij}; \tau)], \quad (67)$$

is exact on the diagonal ($\mathbf{r} = \mathbf{r}'$) and is symmetrical in R and R' . This approximation is very convenient computationally, since the primitive action simply gets replaced by an effective potential and the action depends only on the radial variables r_{ij} . Hence once $u_2(r; \tau)$ has been computed and saved, its use takes the same amount of computer time as the primitive action (assuming potentials are calculated with look-up tables). However, most of the accuracy of the pair action is lost by the end-point approximation.

The exact pair action can be calculated efficiently by the matrix-squaring method [88]. First, the pair density matrix is factorized into a center-of-mass term that is free-particle like and a term that is a function of the relative coordinates. One then expands the relative pair density matrix in partial-waves:

$$\rho(\mathbf{r}, \mathbf{r}'; \tau) = \begin{cases} \frac{1}{2\pi\sqrt{rr'}} \sum_{l=-\infty}^{\infty} \rho_l(r, r'; \tau) e^{il\theta}, & 2\text{D} \\ \frac{1}{4\pi rr'} \sum_{l=0}^{\infty} (2l+1) \rho_l(r, r'; \tau) P_l(\cos \theta), & 3\text{D} \end{cases} \quad (68)$$

where θ is the angle between \mathbf{r} and \mathbf{r}' . Each partial-wave satisfies the 1D recursion formula,

$$\rho_l(r, r'; \tau) = \int_0^\infty dr'' \rho_l(r, r''; \tau/2) \rho_l(r'', r'; \tau/2). \quad (69)$$

If we square the density matrix k times, it will result in a lowering of the temperature by a factor of 2^k . Additional details are given in Ref. [21].

4.3. Sampling the paths

Here we consider how to sample the integrand. The total configuration space to be integrated and summed over is made of elements: $s = [R_1, \dots, R_M]$ where $R_k = \{\mathbf{r}_{1k} \dots \mathbf{r}_{Nk}\}$ are the path variables. We need to sample the distribution

$$\pi(s) = \frac{\exp[-\sum_{k=1}^M S^k]}{Z}, \quad (70)$$

where S^k is the action of the k^{th} link. The partition function Z normalizes the function π in this space. This distribution is different from that of a simple liquid because the points on the path are linked together by the kinetic springs, which can cause the convergence of simple simulation techniques to become slow.

One can sample the distribution using a molecular-dynamics (MD) method [89]. There are two major concerns with MD methods: ergodicity and efficiency. It is easy to see that a free-particle path-integral system will never come into equilibrium because a collection of uncoupled harmonic oscillators will never exchange energy with each other. Even with an interparticle interaction, if the time step τ is small enough, ergodicity is still a worry. However, by using a Nose thermostat, or by resampling the velocities (hybrid Monte Carlo) each dynamical step, one is guaranteed to get convergence to the correct distribution. Tuckerman *et al.* [90] have introduced methods for speeding convergence by separating the slow and fast dynamical scales. They have shown in some cases that path-integral molecular dynamics (PIMD) can be almost as efficient as PIMC.

These results are interesting because there are situations in which dynamical techniques are advantageous. In Monte Carlo methods, an elementary move is typically of one or several particles at several time steps, while in molecular dynamics, one moves all particles at all time steps simultaneously. In some cases, this can be computationally much more efficient. For example, one might want to do a PIMD calculation of a system where the Born-Oppenheimer electronic energy is calculated with a dynamical method such as the Car-Parrinello [91] local-density-functional algorithm. The electronic calculation must be done for all the ions simultaneously so one may as well move all the path variables as well. Another example of a case in which dynamical methods may prove useful is provided by fermion path-integral MC.

However, the chief difficulty with dynamical methods is that it is not possible for the permutation to change continuously, since it is a discrete variable. Dynamical methods by themselves cannot treat problems in which quantum statistics are important.

Now, turning to MC methods, consider the problem of how best to sample a single point R at time τ . This point is connected to two fixed end points, R_1 and R_2 , with

imaginary-time coordinates, 0 and 2τ , respectively. In the simplest choice for the transition probability, a single particle at a single time slice is displaced uniformly inside a cube of side Δ , adjusted to achieve 50% acceptance; Δ must be on the order of, or smaller than, the thermal de Broglie wavelength for a slice, $\Delta \approx \sqrt{\lambda \tau}$.

The heat-bath transition rule will have the smallest correlation time among all transition rules that sample a distribution for R equal to

$$T^*(R) \propto \rho(R_1, R) \rho(R, R_2). \quad (71)$$

To approximate T^* , we drop factors independent of R and factor out the free-particle action. A repulsive potential will cut holes in the free-particle Gaussian distribution where a non-moving atom is present. To incorporate these, one can sample from a correlated Gaussian:

$$T_S(R) = \sqrt{(2\pi)^3 |\mathbf{A}|} e^{-(R-\bar{R}) \cdot (2\mathbf{A})^{-1} \cdot (R-\bar{R})}, \quad (72)$$

where the 3×3 positive-definite covariance matrix \mathbf{A} and the mean position vector \bar{R} are free parameters of the sampling. Let us define the sampling potential \tilde{u} by

$$e^{-\tilde{u}(\mathbf{r}')} = \int d\mathbf{r} \exp \left[-\frac{(\mathbf{r} - \mathbf{r}')^2}{2\sigma} - u(\mathbf{r}, \mathbf{r}; \tau) \right]. \quad (73)$$

Correlated sampling proceeds as follows. First the midpoint $R_m = (R_1 + R_2)/2$ for the move is determined. Then \bar{R} and \mathbf{A} are determined using the following equations with $\tilde{U}(R) = \sum_{i < j} \tilde{u}(r_{ij})$,

$$\bar{R} = \frac{\int dR R T^*(R)}{W} = R_m - \sigma \frac{\partial \tilde{U}(R_m)}{\partial R_m}, \quad (74)$$

$$\mathbf{A} = \frac{\int dR (R - \bar{R}) (R - \bar{R}) T^*(R)}{W} = \sigma \left(\mathbf{I} - \sigma \frac{\partial^2 \tilde{U}(R_m)}{\partial R_m \partial R_m} \right), \quad (75)$$

where \mathbf{I} is the unit tensor. To sample the multivariate normal distribution, one Cholesky-factorizes the covariance matrix as $A = SS^T$, where S is an upper triangular matrix. Then if χ is a vector of Gaussian random numbers with zero mean and unit variance, $S\chi + R_m$ has the desired mean and variance. The effect of atomic interactions on the free sampling is to push the mean position of an atom away from its free-particle mean, if another particle is nearby. Similarly, the free-particle variance is changed by interactions with neighboring particles. In directions where the curvature of the potential is positive, the cage of surrounding atoms results in a narrower Gaussians being sampled.

No matter how well single-bead sampling has been optimized, as the value of τ decreases, the random walk will diffuse through configuration space more and more slowly since the computer time needed to change the overall shape of a single path scales as M^3 . This scaling law, in conjunction with the use of the primitive action, necessitates very large values of M , has caused many path-integral studies to be prohibitively expensive. To achieve faster convergence, one must go beyond moves of a single time slice.

The simplest multiple-slice move is a *displacement* move, in which the entire chain is translated by an amount δ . Moves will be rejected if the chain ends up overlapping with

another chain. The displacement will not change the internal shape of the path; there has to be a different kind of move to do that. Since the move will take $\mathcal{O}(M)$ operations, a displacement should be attempted much less frequently than other kinds of quicker moves.

To generalize the displacement move to the internal degrees of freedom of the paths one can use the normal modes of the kinetic action. These are obtained by a discrete Fourier transform along the “time direction” [92]. We define the normal-mode coordinate by:

$$Q_k = \sum_{l=1}^M R_l e^{2\pi i k l / M}. \quad (76)$$

The total kinetic action is decoupled in normal modes,

$$K = \frac{1}{4\lambda\tau} \sum_{i=1}^M (R_i - R_{i-1})^2 = \frac{1}{\lambda\beta} \sum_k \sin^2(\pi k / M) |Q_k|^2. \quad (77)$$

Each of the $3NM$ normal-mode variables Q_k is independent of the others and has a Gaussian distribution. There are two quite different ways of using normal modes. First, in *normal-mode sampling* one uses this form of the kinetic action to construct a transition move [93]. One samples one or more Q_k from some transition probability, such as a Gaussian distribution with squared width, $\lambda\beta/[2\sin^2(\pi k/m)]$. Then the new path coordinates are determined by the inverse Fourier transform and the move is accepted or rejected based on the change in action and the ratios of transition probabilities. When a potential is present, only the large k modes can be sampled directly from the free-particle Gaussian, since they cause a small movement of the path. In contrast, the low k modes are moved only a small amount, say $|Q'_k - Q_k| < \gamma_k$, with γ_k adjusted to get 50% acceptances. The center-of-mass mode ($k = 0$) is just the displacement move described above.

The second approach is to work directly with the normal-mode variables by rewriting the path integrals as integrals over Q_k instead of R_k . This is the method of *Fourier path integrals* [94]. In using the Feynman-Kacs formula, rather than discretizing the random walk in M time steps, one instead discretizes in M normal modes. The action is now $\int_0^\beta dt V(R(t))$, where

$$R(t) = R_0 + \sum_{k=1}^M Q_k e^{-2\pi i t k / \beta}. \quad (78)$$

One has to worry about approximating the effect of the neglected modes, $k > M$. Their effect can be accounted for by using the cumulant action (the partial averaging technique). See [95] for a systematic comparisons of efficiency of the Fourier-based methods versus the discrete-time path integrals.

Multilevel Monte Carlo (MMC) is another general sampling method [8] which can efficiently make multi-slice, many-particle moves. Suppose the full configuration s is partitioned at the beginning of a Monte Carlo step into $l + 1$ levels $s = (s_0, s_1, \dots, s_l)$, where the coordinates s_0 are unchanged by the move, s_1 are sampled in the first level, s_2 are sampled in the second level, etc. The primed coordinates (s'_1, \dots, s'_l) are the new trial positions in the sense of a Metropolis rejection method. Suppose one has an

PATH INTEGRAL MONTE CARLO CODE

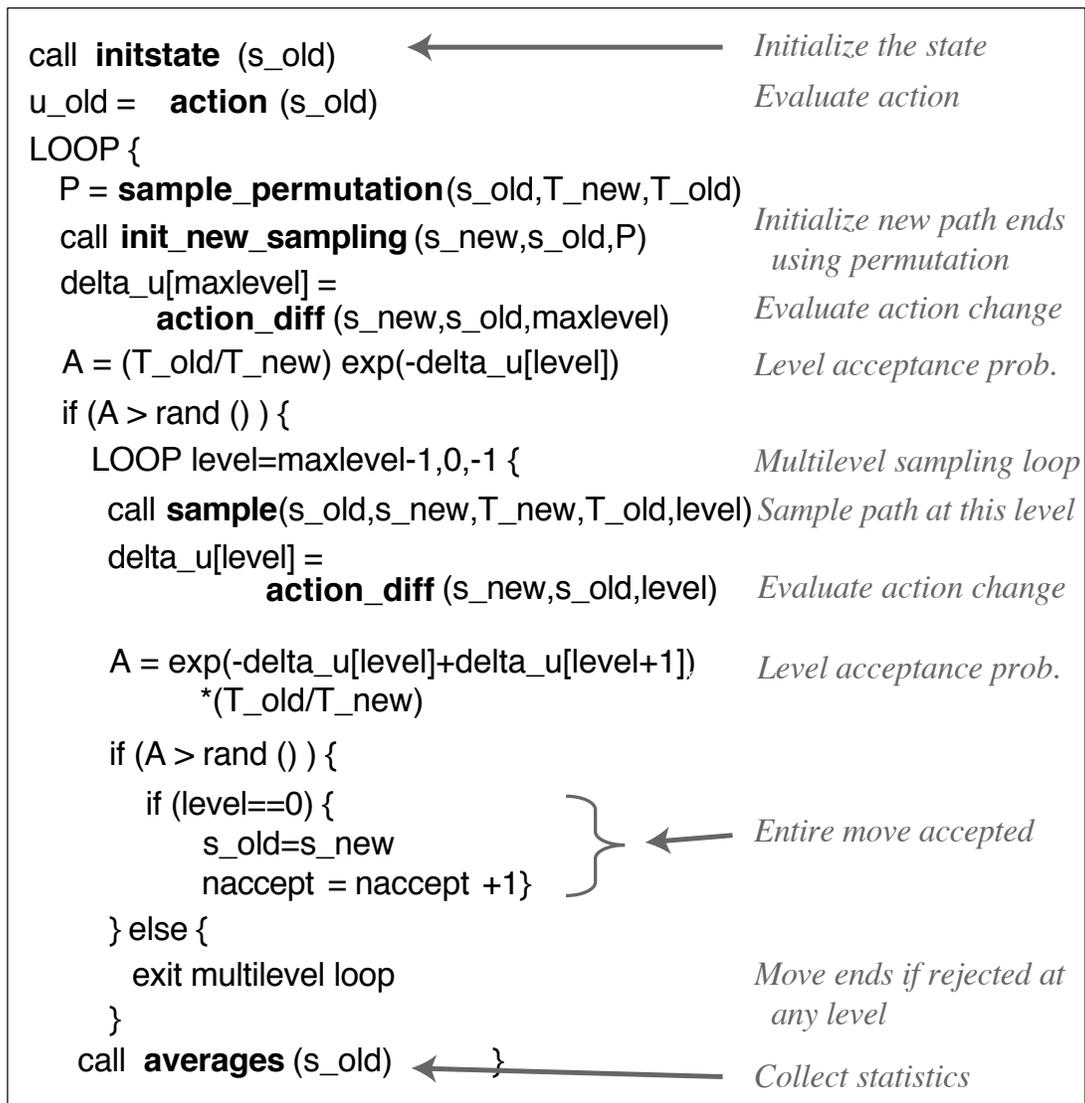


FIGURE 6. Outline of a PIMC code with multilevel sampling. In multilevel Monte Carlo, there are multiple accept/reject steps for each step in the Markov chain. The multilevel algorithm allows large sections of the path to be sampled with a Lévy construction of a Brownian bridge.

approximation to the action as a function of variables in that level and in previous levels. It is used to decide whether the sampling of the path should continue beyond the current level. We require that the action at the last level be exact,

$$\pi_l(s_0, s_1, \dots, s_l) = \pi(s). \quad (79)$$

Now, we choose a sampling rule for s_k contingent on the levels already sampled. The sampling then proceeds past level k with probability

$$A_k(s') = \min \left[1, \frac{T_k(s_k)\pi_k(s')\pi_{k-1}(s)}{T_k(s'_k)\pi_k(s)\pi_{k-1}(s')} \right]. \quad (80)$$

MMC gains in efficiency because the coarsest degrees of freedom are sampled and accepted or rejected before the finer movements are even constructed.

In the Lévy construction of a Brownian bridge, one begins with two fixed end points (say R_0 at “time” 0 and R_β at “time” β) and samples a bisecting point at time $\beta/2$, exactly as was described with free-particle sampling. The bisecting point R of the interval $(\mathcal{R}_1, \mathcal{R}_2)$ is

$$R = \frac{R_1 + R_2}{2} + \eta, \quad (81)$$

where η is a normally distributed random vector with mean zero and covariance, $\sqrt{\lambda|t_2 - t_1|/2}$. Having sampled $R = R_{\beta/2}$, one now bisects the two new intervals $(0, \beta/2)$ and $(\beta/2, \beta)$, generating points $R_{\beta/4}$ and $R_{3\beta/4}$ with the same algorithm. One continues recursively, doubling the number of sampled points at each level, stopping only when the “time” difference of the intervals is τ .

This method is used with the multilevel method to accomplish early rejection. Suppose a single-particle or many-particle path consisting of $m = 2^l - 1$ time slices is “clipped out” where l is the *level*. The fixed end points are R_i and R_{i+m} . The new points to be sampled will have the coordinates: $R_{i+1}, \dots, R_{i+m-1}$. The coordinates are partitioned into levels as in the Lévy construction. First the midpoint is sampled. Then the same algorithm is used to find the midpoints of the two remaining intervals, etc. One needs to approximate an action at the k^{th} level, but the detailed form of the level action does not matter very much since one is looking for a quick and dirty way of deciding whether to continue the bisection procedure or to reject and start over. For example, in helium simulations, any function that detects overlapping atoms will accomplish this.

4.4. Bosons, superfluidity and Bose condensation

The path integral discussion up to this point have been appropriate to distinguishable (Boltzmann) particle statistics. Quantum statistics was not taken into account. For Bose systems only totally symmetric eigenfunctions contribute to the density matrix; those such that $\phi_i(PR) = \phi_i(R)$ where P is a permutation of particle labels, *i.e.*, $PR = (\mathbf{r}_{P_1}, \mathbf{r}_{P_2} \dots \mathbf{r}_{P_N})$. By symmetrizing the density matrix with respect to particle labels, we obtain the bosonic density matrix

$$\rho_B(R_0, R_1; \beta) = \frac{1}{N!} \sum_P \rho(R_0, PR_1; \beta), \quad (82)$$

where ρ is the boltzmann density matrix. A straightforward evaluation of the sum is out of the question once N gets large, since there will be $N!$ terms. Each term in the sum is positive, so we can *sample* the permutations in the sum. A bosonic simulation consists

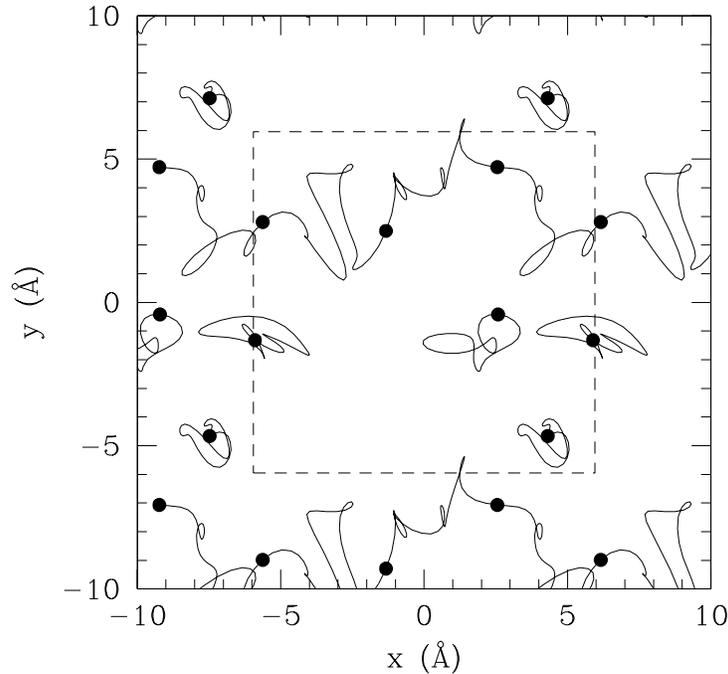


FIGURE 7. The extended trace of six ^4He atoms at a temperature of 0.75 K and with 53 time slices. The dashed square represents the periodic boundary conditions. Three of the atoms are involved in an exchange which winds around the boundary in the x direction.

of a random walk through the path space *and* the permutation space. For Fermions special techniques are employed to deal with the sign as we discuss below.

The partition function for a Bose system has the form

$$Z_B = \frac{1}{N!} \sum_P \int dR_0 \dots dR_{M-1} \exp \left(- \sum_{m=1}^M S^m \right), \quad (83)$$

with the path closure: $PR_m = R_0$. Paths are allowed to close on any permutation of their starting positions. In the classical isomorphic system, ring polymers can “cross-link.” Fig. 7 shows the trace perspective (the projection onto the xy plane) of the same paths. Three of the particles are involved in a cyclic exchange, which wraps around the periodic boundary conditions. This path wrapping around the boundaries is called a *winding* path and is a direct manifestation of superfluidity. In Eq. (82), it appeared that R_0 was singled out to receive the permutation. By change of variables one can make the permutation occur at any time slice. All time slices are still equivalent, and they can all be used to calculate averages. Whenever one constructs an estimator for an observable, one should consider whether Bose symmetry has been taken into account properly.

Now consider the sampling of permutations. The simplest type of move would be to insert a pair permutation of two atoms without moving the coordinates of the path. This type of move can fail to sample the permutation space for a system with repulsive interactions because the trial move is improbable. Instead, one can apply the multilevel Metropolis methods to the joint sampling of permutations and path moves. A set of time

slices are first selected for the path move. A local permutation move consists of applying a cyclic exchange of n atoms to an existing path. Since permutation space is discrete, we can use the heat-bath method to sample a permutational change from the distribution $T^*(P) \propto \rho(R_i, PR_{i+m})$ where P ranges over all cyclic permutations involving n atoms. If we make the end-point approximation for the density matrix, terms involving the interaction will drop out since they are symmetric under particle interchange

$$T^*(P) = \frac{1}{C_I} \exp \left[- \sum_{j=1}^n (\mathbf{r}_{j,i} - \mathbf{r}_{Pj,i+m})^2 / 4n\lambda^* \tau \right], \quad (84)$$

where C_I is a normalization factor. The cyclic permutation can either be explicitly sampled from a precomputed table or it can be implicitly sampled with a walk through particle labels.

According to Feynman's theory [96], the superfluid transition is represented in the classical system by the formation of macroscopic polymers, *i.e.* those stretching across an entire system and involving on the order of N atoms. Now we explain the properties directly related to superfluidity; the specific heat, the momentum distribution, and the superfluid density.

It was the shape of the specific heat curve of liquid ^4He and its singularity at 2.17 K that gave rise to the name of the lambda transition. Feynman explained how macroscopic bosonic exchange gives rise to this peak. The specific heat is proportional to the mean-squared fluctuation of the exchange distance. At the critical point, the specific heat diverges because there are both long and short exchanges present.

London, in 1938, supposed the superfluid transition in liquid ^4He to be the analog of the transition that occurs in an ideal Bose gas, where below the transition, a finite fraction of particles occupy the zero-momentum state. Penrose and Onsager [97] defined *Bose condensation* in an interacting system as the macroscopic occupation of a single-particle state, namely the state of zero momentum. The condensate fraction has a simple meaning in terms of path-integrals. The probability density of observing a single atom with momentum k is the Fourier transform of the single body density matrix,

$$n_{\mathbf{k}} = \frac{1}{\Omega(2\pi)^d} \int d\mathbf{r}_1 d\mathbf{r}'_1 e^{-i\mathbf{k}(\mathbf{r}_1 - \mathbf{r}'_1)} n(\mathbf{r}_1, \mathbf{r}'_1), \quad (85)$$

where $n(\mathbf{r}_1, \mathbf{r}'_1)$ is defined as:

$$n(\mathbf{r}_1, \mathbf{r}'_1) = \frac{\Omega}{Z} \int d\mathbf{r}_2 \dots d\mathbf{r}_N \rho(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N, \mathbf{r}'_1, \mathbf{r}_2, \dots, \mathbf{r}_N; \beta). \quad (86)$$

The paths that we have been discussing up to this point, each ending at the start of another particle's path, cannot be used to calculate the momentum distribution. To get the momentum distribution, one removes the restriction on one of the atoms that it returns to its starting position. At high temperature, where there are no particle exchanges, $n(\mathbf{r}, \mathbf{r}')$ will be almost free-particle like with a Gaussian end-to-end distribution: $n(\mathbf{r}, \mathbf{r}') \propto \exp[-(\mathbf{r} - \mathbf{r}')^2 / (4\lambda\beta)]$. Taking the Fourier transform, we end up with the Maxwellian momentum distribution with a width $k_B T$.

If a macroscopic exchange is present, as is the case in the superfluid state, the two ends can become separated by much more than a thermal wavelength if they are attached to a macroscopic exchange. For a 3D bulk liquid the single particle density matrix in the superfluid state goes to a constant at large $|\mathbf{r} - \mathbf{r}'|$. The momentum distribution, which is the Fourier transform of the single particle density matrix, will then have a delta function at the origin. The condensate fraction is the probability of finding an atom with precisely zero momentum,

$$\tilde{n}_0 = \frac{(2\pi)^3}{\Omega} n_0 = \frac{1}{\Omega^2} \int d\mathbf{r} d\mathbf{r}' n(\mathbf{r}, \mathbf{r}') = \lim_{r \rightarrow \infty} n(r). \quad (87)$$

To find the condensate fraction of an inhomogenous system, one must diagonalize the single particle density matrix, thereby determining the natural orbitals and their occupation numbers.

Superfluidity is experimentally characterized by the response of a system to movements of its boundaries. The rotating bucket experiment was first discussed by Landau in 1941, who predicted that superfluid helium would show an abnormal relation between the energy it takes to spin a bucket and its moment of inertia. A normal fluid in equilibrium will rotate rigidly with the walls. The work done is $E = \frac{1}{2} I \omega^2$, where I is the momentum of inertia and ω is the angular rotation rate. On the other hand, a superfluid will stay at rest if the walls rotate slowly, so that a smaller energy is needed to spin up the container. The liquid that stays at rest is the superfluid. Experiments by Andronikashvili in 1946 confirmed this prediction.

The effective moment of inertia is defined as

$$I = \left. \frac{dF}{d\omega^2} \right|_{\omega=0} = \left. \frac{d\langle \mathcal{L}_z \rangle}{d\omega} \right|_{\omega=0}, \quad (88)$$

where F is the free energy, and \mathcal{L}_z is the total angular momentum operator in the \hat{z} direction. The ratio of the effective moment of inertia to the classical one I_c is defined as the normal density; what is missing is the superfluid density:

$$\frac{\rho_n}{\rho} = 1 - \frac{\rho_s}{\rho} = \frac{I}{I_c}. \quad (89)$$

Thus the superfluid density is the linear response to an imposed rotation, just as the electrical conductivity is the response to an imposed voltage. This can be calculated with path integrals as:

$$\frac{\rho_s}{\rho} = \frac{2m \langle \mathcal{A}_z^2 \rangle}{\beta \lambda I_c}, \quad (90)$$

where the projected area is

$$\vec{\mathcal{A}} = \frac{1}{2} \sum_{i,j} \mathbf{r}_{i,j} \times \mathbf{r}_{i,j+1}. \quad (91)$$

Note that the area of a path is a vector. For rotations about the \hat{z} axis we need only the \hat{z} component of the area. The superfluid density is proportional to the mean-squared area of paths sampled for a container at rest.

Although superfluid response (like conductivity) is not itself a local property, it is possible [98] to calculate a local contribution to the total response and thereby quantify places that contribute to the superflow response. To define a local superfluid density we write:

$$\rho_s(\mathbf{r}) = \frac{2mN}{\beta\lambda I_c} \left\langle \int d\mathbf{r}' \mathcal{A}(\mathbf{r}) \mathcal{A}(\mathbf{r}') \right\rangle = \frac{2mN}{\beta\lambda I_c} \langle \mathcal{A}(\mathbf{r}) \cdot \mathcal{A} \rangle, \quad (92)$$

where $\mathcal{A}(\mathbf{r})$ is defined as

$$\vec{\mathcal{A}}(\mathbf{r}) \equiv \frac{1}{2} \sum_{i,k} (\mathbf{r}_{i,k} \times \mathbf{r}_{i,k+1}) \delta(\mathbf{r} - \mathbf{r}_{i,k}). \quad (93)$$

Since $\vec{\mathcal{A}}(\mathbf{r})$ integrates to the total area, the local superfluid response exactly integrates to the total response. Two types of contributions to the local superfluid density can be distinguished based on the connectivity of the instantaneous paths: contributions of particles on the same chain, which are positive, and contributions of particles on different exchange cycles. By reversing the order of one exchange cycle, the contribution from different cycles will change sign; if the cycles are spatially separated, the magnitude of the contribution will be unaffected, so that their contribution will be much smaller.

A uniform magnetic field acting on charged particles is equivalent to a rotation, and one can use these superfluid formulas for charged particles in magnetic fields. The analogous phenomenon to superfluidity is the Meisner effect of superconductors. The Hamiltonian in a magnetic field is

$$H = \frac{1}{2m} \left[p - \frac{e\mathbf{A}}{c} \right]^2 + V(R). \quad (94)$$

In a constant magnetic field in the Coulomb gauge, the lowest-order change in the Hamiltonian due to the magnetic field is proportional to the angular momentum operator, $\frac{e}{mc} \mathbf{p} \cdot \mathbf{A} = \frac{eB}{2mc} \mathbf{L}_z$, hence the previous discussion applies if we replace the rotation rate with $\frac{eB}{2mc}$. Then the zero-field diamagnetic susceptibility is

$$\chi = -\frac{\langle \mathcal{A}^2 \rangle e^2}{\beta \hbar^2 c^2}. \quad (95)$$

For details see Ref. [99].

Periodic boundary conditions are more convenient for simulations, since no surfaces appear and there is no curvature in making a loop around the boundaries. For the superfluid density we get

$$\frac{\rho_s}{\rho} = \frac{\langle W^2 \rangle}{2\lambda \beta N}, \quad (96)$$

where the winding number is defined by

$$\mathbf{W} = \sum_{i=1}^N \int_0^\beta dt \left[\frac{d\mathbf{r}_i(t)}{dt} \right]. \quad (97)$$

The winding number is a topological invariant of a given path. Paths with a nonzero winding are the signal for superfluidity. This justifies the claim made earlier, that the identification of a Bose superfluid requires the full imaginary-time paths. Static correlation functions are not enough; one needs to know how the paths are connected up.

Macroscopic exchange is necessary to have both superfluidity and momentum condensation. However, neither property is simply proportional to the number of macroscopic exchanges. In 3D systems they go together; in 2D there is no condensation, but the system is still superfluid.

4.5. Nanodroplets

Molecules confined in ^4He nanodroplets have been shown to exhibit excitation spectra with clearly resolved rotational fine structure consistent with that of a free rotor, though with an increased moment of inertia. Toennies *et al.* [100, 101] have shown that only a few layers of ^4He coating the molecule are required to decouple the impurity from the droplet and achieve free rotation. Nanodroplets provide a very interesting interface between the chemistry of solvation and nanoscale superfluidity. Importantly for quantum simulation, they provide both an excellent test bed for methods, as well as many important applications of the methods, because essentially exact calculations can be done of the equilibrium properties. There is no sign problem since Bose statistics are applicable. Finite size effects, far from being a nuisance, are part of the interest, and temperature effects are relevant. The theoretical uncertainties have to do with whether the droplets are in equilibrium and what is the precise interaction between the helium atoms and the impurity.

While one can do simulations of helium droplets with either ground state methods, such as DMC [102] or reptation Monte Carlo [103], here we will discuss the PIMC approach, as it gives the cleanest physical interpretation. Other theoretical approaches are discussed in the review volume [104]. Simulations have been performed with pure ^4He droplets [105], predating the experiment and showing that the droplets of 64 atoms are essentially superfluid, based on the area estimator Eq. 90. It was also found the droplets of H_2 droplets are superfluid [106], as long as the droplets are sufficiently small. Larger clusters freeze into localized crystals, thus preventing bosonic exchange. Such behavior has recently been found [107] for compound systems made from an OCS molecule, a few hydrogen molecules, and surrounded by a droplet of ^4He . PIMC has been used [108] to understand the details of the molecular exchange responsible for the observations.

PIMC calculations of ^4He droplets containing an impurity are relatively straightforward and allow unambiguous calculation of the helium density and the fraction of it that is superfluid. Because helium is strongly attracted to most impurities, a marked layered structure is formed. The first 2 layers are relatively isolated from the rest, those atoms exchange infrequently; the superfluid response in the first layer is reduced due to the anisotropy of the helium density in that region. Only a small fraction of the particles in the first layer are localized (not permuting) at low temperatures. Below 1 K, the majority of the atoms in the droplet are part of exchange cycles with atoms in both the first

layer and the rest of the droplet. Thus, while the effect of confinement on the thermal excitations resembles that of a two-dimensional superfluid, this system is fundamentally different from an isolated helium film. In Ref. [98] it is shown that the superfluid response in the first layer is also significantly reduced by thermal excitations at the physical temperatures of the drops, 0.40 K.

4.6. Excitation spectra using MaxEnt

Most experimental probes of nanostructures are dynamical, for example, a measurement of the electro-magnetic response. Here we discuss how to use the imaginary time correlation functions from PIMC (or DMC) to derive dynamical information. The dynamic structure factor $S_k(\omega)$ is related to the electromagnetic response and is a crucial correlation function to calculate or observe, since through it one can see the excitations. The dynamical structure factor is defined as

$$S_k(\omega) = \frac{1}{2\pi N} \int_{-\infty}^{\infty} dt e^{i\omega t} \langle \rho_k(t) \rho_{-k}(0) \rangle \quad (98)$$

$$= \frac{1}{NZ} \sum_{mn} \delta(\omega - E_m + E_n) e^{-\beta E_n} |\langle m | \rho_k | n \rangle|^2. \quad (99)$$

Here $\rho_k = \sum_i e^{ikr_i}$ is the Fourier transform of the density. The sum in the second equation is over all pairs of exact states $|n\rangle$ with energy E_n .

The imaginary-time density-density response function, defined as

$$F_k(t) = \frac{1}{NZ} \text{Tr} \left\{ \rho_{-k} e^{-t\mathcal{H}} \rho_k e^{-(\beta-t)\mathcal{H}} \right\}, \quad (100)$$

is straightforward to calculate with path integrals or with the ground-state methods discussed below. However, there is little in the way of structure in $F_k(t)$; to the eye it is a featureless exponential decay up to $\beta/2$, and then rises again because of the periodicity in imaginary time; it is related to $S_k(\omega)$ by a Laplace transform,

$$F_k(t) = \int_{-\infty}^{\infty} d\omega e^{-t\omega} S_k(\omega). \quad (101)$$

Mathematically, $F_k(t)$ and $S_k(\omega)$ are equivalent, since an ‘‘inverse Laplace transform’’ (or equivalently analytic continuation from imaginary to real time) is well defined. The presence of statistical noise, however, rules out a direct inversion, since very small features in $F_k(t)$ come from large features in $S_k(\omega)$. The noise destroys the information needed to do the inversion. The numerical inversion of a Laplace transform is a classic ill-conditioned problem.

Silver, Sivia, and Gubernatis [109] have shown that Bayesian, or maximum-entropy, techniques aid in this inversion by making the problem better conditioned. In addition to using the Monte Carlo data for the inverse Laplace transform, one can use *prior* information about $S_k(\omega)$. In some models, the maximum entropy method has given very

satisfactory inversions. Jarrell and Gubernatis [110] have reviewed the fundamentals and practical details of this approach.

Bayes' theorem can be written in the following form:

$$P[S(\omega)|F(t)] \propto P_L[F(t)|S(\omega)]P_p[S(\omega)]. \quad (102)$$

The probability of $S(\omega)$, given our PIMC data and the theoretical input [the *posterior probability*], equals the probability of the PIMC data given $S(\omega)$ [the *likelihood*] times the prior knowledge of $S(\omega)$ [the *prior function*].

The central-limit theorem guarantees that the noise in $F(t)$ from a well-converged PIMC is normally distributed

$$P_L[F(t)|S(\omega)] = \exp \left[-\frac{1}{2} \sum_{t,t'} \delta F(t) \sigma(t,t')^{-1} \delta F(t') \right], \quad (103)$$

where $\delta F(t) = F(t) - \langle F(t) \rangle$.

One knows a few things about any physical $S(\omega)$: it is everywhere nonnegative, it satisfies detailed balance and various sum rules, and it has certain asymptotic behaviors at large and small ω . We need to construct a prior function so that any $S(\omega)$ not satisfying these conditions will have a zero prior function. For image reconstruction and many other applications the entropic prior has worked very well,

$$P_p[S(\omega)] \propto \exp \left[\alpha \sum_{\omega} S(\omega) \log(S(\omega)/m(\omega)) \right], \quad (104)$$

where α is an adjustable parameter and $m(\omega)$ is the *default model*. However, there is no proof that the entropic prior is correct for the "distribution" that describes $S(\omega)$ and it violates the principle that quantum energy levels repel each other.

Having chosen the likelihood and prior function, one can take either of two approaches to using Bayes' theorem. In the *maximum-entropy* approach one finds the $S(\omega)$ that is most likely, the one which maximizes $P[S(\omega)|F(t)]$. This is a good procedure when the probability distribution is narrow. Finding the maximum is very fast. A second, more rigorous approach, *average entropy*, is to sample $S(\omega)$ with a probability equal to the posterior probability. This can easily be done with Metropolis Monte Carlo. This approach is slower (but not nearly as slow as generating the original data) but does not rely on any assumptions about how narrow or skewed the probability function is. Some features of $S(\omega)$ may be tightly constrained by the PIMC data, while other features are not constrained. To compute errors, one simply looks at the fluctuations of $S(\omega)$ coming from the Markov chain that samples $S(\omega)$. In general, maximum entropy gives spectra that are too broad [111]; lines that are sharp become broad features.

4.7. Fermion PIMC: Fixed-node and fixed-phase methods

We now discuss the generalization of the path integral methods to fermion systems. In the *direct fermion method* one sums over permutations just as for bosonic systems.

Odd permutations then contribute with a negative weight. The direct method has a major problem because of the cancellation of positive and negative permutations; the efficiency is:

$$\xi = \left[\frac{Z_F}{Z_B} \right]^2 = \exp[-2\beta(F_F - F_B)] \quad (105)$$

where Z_F and Z_B refer to partition function and F_F and F_B to the total free energies for Fermi and Bose statistics respectively. These are proportional to the number of particles.

The restricted path identity asserts that the nodes of the exact density matrix determine the rule by which one can take only paths with the same sign. Suppose ρ_F is the density matrix corresponding to some set of quantum numbers which is obtained by using the projection operator \hat{a} on the distinguishable particle density matrix. Then,

$$\rho_F(R_\beta, R_*; \beta) = \int dR_0 \rho_F(R_0, R_*; 0) \oint_{R_0 \rightarrow R_\beta \in \Upsilon(R_*)} dR_t e^{-S[R_t]}, \quad (106)$$

where the subscript means that we restrict the path integration to paths starting at R_0 , ending at R_β and are node-avoiding (those for which $\rho_F(R_t, R_*; t) \neq 0$ for all $0 < t \leq \beta$.) The weight of the walk is $\rho_F(R_0, R_*; 0)$. It is clear that the contribution of all the paths for a single element of the density matrix will be of the same sign; positive if $\rho_F(R_0, R_*; 0) > 0$, negative otherwise. In particular, on the diagonal all contributions must be positive.

Underlying the identity is the assumption that the random walk is a continuous trajectory so we can say definitively that if sign of the density matrix changed, it had to have crossed the node at some point. This presents a technical problem for discrete time paths since we must decide whether the path crosses and crosses back in between the sampled times. Lattice models or non-local Hamiltonians do not have continuous trajectories so this method is not as straightforward for those models.

The ‘‘bosonic’’ path integral formulation can be applied to fermion path integrals; all that is required is to take a subset of the bosonic paths. We shall see that it is important to allow non-trivial, even permutations since they are directly related to Fermi liquid behavior.

The problem with using Eq. (106) is that the unknown density matrix appears both on the left-hand side and on the right-hand side (in the criterion of node-avoiding.) To apply the formula directly, we would have to self-consistently determine the density matrix, a nearly impossible task in high dimensions. In practice, what we need to do is to make an *ansatz* for the fermion density matrix. The trial density matrix is used to define node-avoiding. Using trial nodes we generate a better approximation to the density matrix. Hence, $\rho(R', R; \beta)$ is a solution to the Bloch equation inside the trial nodal cells, and it obeys the correct initial conditions. It is not an exact solution to the Bloch equation (unless the nodes of ρ_T are correct) because it can have gradient discontinuities at the trial nodal surfaces.

Consider the case where one has spin $\frac{1}{2}$ fermions. Assume that there is no magnetic field so that the Hamiltonian is independent of spin. Then we can quantize the spin axis in the \hat{z} direction. Suppose we want to calculate the partition function Z_m in an ensemble

where S_z is fixed to be m . It is not hard to see that this equals:

$$Z_m = \frac{1}{N!} \text{Tr}_\sigma \left\{ \sum_P (-1)^P \delta^N(\sigma, P\sigma) \delta(\sum_i \sigma_i, m) \int dR \rho_D(R, PR) \right\}. \quad (107)$$

Now relabel the coordinate indices in the integration so that particles 1 through $N_\uparrow = N/2 + m$ have up spin and particles $N_\uparrow + 1$ through N have down spin. All arrangements of the spin variable in the trace contribute the same amount to the partition function and we can just calculate one term and multiply by the number of ways of assigning spins: $N!/(N_\uparrow!N_\downarrow!)$. The following operator antisymmetrizes over the up spins and the down spins individually.

$$\mathcal{A} = \frac{1}{N_\uparrow!N_\downarrow!} \sum_{P_\uparrow, P_\downarrow} (-1)^{P_\uparrow + P_\downarrow} P_\uparrow \otimes P_\downarrow, \quad (108)$$

where P_\uparrow (P_\downarrow) operates on the up (down) spin coordinates only. With this one can calculate spin-independent correlation functions in the fixed S_z ensemble.

If we want to calculate the total partition function, $Z = \sum_m Z_m$, the magnetization must be sampled. To average over all spins one could occasionally attempt a spin flip from one fixed m ensemble to $m \pm 1$ and accept or reject that flip, depending on the path is now legal. We can recover the fixed total spin ensemble by using the relationship:

$$Z_S = (2S + 1)(Z_m - Z_{m+1}). \quad (109)$$

In an actual calculation, one does not make a geometric interpretation of the nodes. Instead one computes a *trial density matrix*, typically a determinant, and uses that to decide whether a given path is to be allowed. Paths for which the trial density matrix are negative or too close to zero are rejected. Let $v_e(\mathbf{r})$ be a single-particle external potential. The distinguishable particle density matrix is then a product of solutions of the single-particle Bloch equation:

$$-\frac{d}{dt} g(\mathbf{r}, t) = [-\lambda \nabla^2 + v_e(\mathbf{r})] g(\mathbf{r}, t), \quad (110)$$

with the boundary condition,

$$g(\mathbf{r}, \mathbf{r}_*; 0) = \delta(\mathbf{r} - \mathbf{r}_*). \quad (111)$$

Then using the antisymmetric projection operator and the definition of a determinant, we find for the spinless case:

$$\rho_F(R, R_*; t) = \frac{1}{N!} \det[g(\mathbf{r}_i, \mathbf{r}_{j,*}; t)]. \quad (112)$$

In the case where the external potential is zero (or a constant), the single-particle density matrix is a Gaussian.

$$g(\mathbf{r}, \mathbf{r}_*; t) \propto \exp \left[-\frac{(\mathbf{r} - \mathbf{r}_*)^2}{4\lambda t} \right]. \quad (113)$$

Miltzer and Pollock have generalized this to the Hartree-Fock case [112].

To calculate the momentum distribution, the fermion sign comes back in, as it must. Consider the ideal fermi-gas momentum distribution (for spinless fermions).

$$n_{\mathbf{k}} = \begin{cases} 1/(2\pi^3\rho) & \text{for } k < k_F \\ 0 & \text{for } k > k_F. \end{cases} \quad (114)$$

Then

$$n(\mathbf{r}) = \frac{3}{(k_F r)^3} [\sin(k_F r) - k_F r \cos(k_F r)], \quad (115)$$

where the Fermi wavevector for spinless fermion is related to the density by $k_F = (6\pi^2\rho)^{1/3}$. Since $n(r)$ is often negative, even with restricted paths we must have negative weights entering. The momentum distribution has a discontinuity at the Fermi wavevector \mathbf{k}_F . As a consequence the single particle density matrix must decay at large distance as r^{-2} . We can get such long-range behavior only if there are macroscopic exchanges. Hence, the existence of any kind of non-analytic behavior (we mean a discontinuity in $n_{\mathbf{k}}$ or in any of its derivatives) implies that the restricted paths have important macroscopic permutation cycles. Recent calculations with this procedure have been reported by Miltzer *et al.* [113]. See the review article [114] for more details concerning restricted path integrals.

4.8. Ground state path integrals

The ground state path integral method is a way of using the path integral method to perform the averages appearing in diffusion Monte Carlo for the ground state. There are two big advantages this approach gives for ground state properties: properties can be computed without having to rely on the forward-walking or extrapolated procedure and an explicit probability density for the path is obtained, making it easier to get such quantities as forces or energy differences. For a description of reptation and path integral methods see Refs. [21, 115, 116].

A trial function $\Psi_0(R)$ is projected as $\Psi_{\beta/2}(R) = e^{-\frac{\beta}{2}H}\Psi_0$ and will converge to the exact ground state. Let

$$Z_{\beta} = \langle \Psi_0 \exp(-\beta H) \Psi_0 \rangle. \quad (116)$$

The variational energy is:

$$E(\beta) = -\frac{1}{Z} \frac{dZ_{\beta}}{d\beta} = \langle E_L(0) \rangle. \quad (117)$$

This derivative is implemented as the average local energy at either end of the path where the local energy is defined as $E_L(R) = \Re(\Psi_0^{-1}H\Psi_0(R))$. As a function of β , $E(\beta)$ is an upper bound that converges to the exact energy at large β and is equivalent to the transient estimate method. The β derivative of the energy is strictly negative and is the variance of the projected trial function.

To calculate the density matrix, we divide imaginary time into P time slices and make an approximation for $\exp(-\tau H)$ where the timestep is $\tau = \beta/P$. Then the probability distribution function for a path is given by:

$$\Pi(s) = \exp \left[U(R_1) + U(R_p) + \sum_{i=1}^{p-1} S(R_{i+1}, R_i) \right]. \quad (118)$$

Here $U = \Re(\ln(\Psi_0))$ and S is the link action. One can use the type of action used in PIMC: namely a pair product action closed at the end with the wavefunction. The disadvantage is in the treatment of the fermion terms, done with an approximate image potential. The trial function does not appear in the action. With a good trial function the local energy is nearly flat and breakup errors will come from just the commutator of the Laplacian with drift. Using the expression for the importance-sampled Green's function from Ref. [7], we find another expression for the link action:

$$\langle R | e^{-\tau H} | R' \rangle = c \psi(R) / \psi(R') e^{-\tau E_L(R) - (R' - R - 2\lambda \tau F(R))^2 / (4\lambda \tau)}. \quad (119)$$

Here $F(R) = \Re(\nabla U(R))$. This expression should be symmetrized with respect to R and R' .

For efficiency it is necessary to prevent changes in sign of the trial function. This can be done within the fixed-node (real) or fixed phase (complex) approximation [18] as discussed earlier within DMC. We modify the Hamiltonian by adding a term $\lambda \nabla (\Im(\log(\Psi_0(R))))^2$ to the potential. Instead of the exact energy we get an upper bound, the best upper bound consistent with the assumed phase.

There are several ways to move the ground state path integral. In the reptation procedure, one of the two ends is sampled at random to be the growth end. A move consists of sampling a new point near the growth end from a Gaussian distribution, exactly as is done with DMC. The move is accepted or rejected based on the usual Metropolis formula based on the reversed move. The old tail position is then discarded if the move is accepted. Another type of move is an internal sampling move. An internal slice is picked, and the midpoint between the adjacent slices is found and a new slice is sampled from a Gaussian centered on that midpoint. The path moves less because it is constrained in two directions but the move may be useful in getting rid of persistent configurations in the middle. The sampling methods discussed earlier could be applied here to improve the convergence.

One can also take a large step made of from several reptation moves (say m) in one direction, a “block move”. A better approach is the bounce algorithm: we reverse the direction on rejection. The “bounce” algorithm has several advantages over the “block move” algorithm; the average number of steps in one direction is set automatically by the algorithm, averages are taken every step and bad moves are rejected early.

4.9. Path integral simulations of quantum dots

There have been several path integral simulations of quantum dots. The earliest of these was a study of excitons and biexcitons by Pollock and Koch [71, 70]. They used

a simple effective mass model of a spherical dot with an infinite confining potential, with parameters chosen to resemble CdS. With the thermal path integral, they calculated the exciton and biexciton energies and charge distributions for temperatures from 6 K to 200 K. They found an interesting change in radial charge distribution as a function of temperature. At low temperatures, coulomb attraction caused the heavier hole to sit near the center of the dot, somewhat more localized than the lighter electron. At high temperatures the coulomb interactions are secondary, and the heavier hole was able to more uniformly fill the dot, especially near the surface, due to its smaller thermal wavelength. They also found increasing exciton and biexciton binding energy as the dot size decreased due to their hard confining potential. This is in contrast to Refs. [69, 68], where finite confining potentials lead to an optimal radius that maximizes exciton and biexciton binding. It should be noted that Pollock and Koch's thermal calculations do not include thermal spin excitations; the results are for a fixed $S_z = 0$ ensemble, which gives less weight to spin polarized states than a true thermal ensemble.

Another type of path integral simulation was a low temperature (ground state) calculation of recombination rates in dots [117]. The radiative recombination rates of interacting electron-hole pairs in a quantum dot are strongly affected by quantum correlations among electrons and holes in the dot. Recent measurements of biexciton recombination rates in self-assembled dots lie in the theoretically predicted range of one to two times the rate of single exciton recombination, depending on the degree of correlation. Theoretical approaches using effective mass wave functions have been unable to accurately determine biexciton recombination rates for a realistic, material specific model. They have developed a Feynman path-integral formulation of the problem that allows the direct evaluation of the recombination rate, including thermal and many-body effects. Using real-space Monte Carlo integration, they evaluate the path-integral expressions for realistic three-dimensional models of InGaAs/GaAs, CdSe/ZnSe, and InP/InGaP dots, including anisotropic effective masses.

5. CONCLUSIONS AND OUTLOOK

Quantum Monte Carlo methods are among the most general algorithms for quantum many-body systems. They have come into the forefront of nanostructure simulations with the availability of large computational resources. Future trends are for more accurate calculations, larger number of particles, simulations of spin and magnetic effects, and the development of methods for more complicated nanostructures with realistic effects such as variable masses, dielectric functions, band effects, and disorder. A number of problems remain for the future, in fact one can assume that the methods are only in their infancy. Research on the fermion sign problem is progressing but unresolved in practice. Work on the general problem of quantum dynamics is at an even earlier phase. Only a handful of properties have been computed within QMC routinely. Fast accurate methods for computing a wide spectrum of properties is needed. Within classical simulation and for lattice models, substantial progress has been made in using various ensembles such as grand canonical, constant pressure and so forth and algorithms that improve the convergence of the simulation. Because of the mapping from

quantum mechanics into classical mechanics, many of these methods are generalizable to quantum systems. Finally, QMC can deal with quite general systems, use sophisticated mathematical methods (trial functions, sampling schemes estimators) and take full use of distributed and parallel processing, but only if reliable software is available to the research community.

Acknowledgments. The authors wish to acknowledge support from DARPA/ONR N00014-01-1-1062, NSF DMR-0239819, and NSF DMR-0104399 during the preparation of this review.

REFERENCES

1. Donsker, M. D., and Kac, M., *J. Res. Natl. Bur. Std.*, **44**, 551 (1950).
2. Kalos, M. H., *Phys. Rev.*, **128**, 1791–1795 (1962).
3. Kalos, M. H., Levesque, D., and Verlet, L., *Phys. Rev. A*, **9**, 2178–2195 (1974).
4. Anderson, J. B., *J. Chem. Phys.*, **63**, 1499–1503 (1975).
5. Anderson, J. B., *J. Chem. Phys.*, **65**, 4122–4127 (1976).
6. Reynolds, P. J., Ceperley, D. M., Alder, B. J., and Lester, W. A., *J. Chem. Phys.*, **77**, 5593 (1982).
7. Ceperley, D. M., and Alder, B. J., *Phys. Rev. Lett.*, **45**, 566 (1980).
8. Ceperley, D. M., and Pollock, E. L., *Phys. Rev. Lett.*, **56**, 351–354 (1986).
9. Ceperley, D. M., and Pollock, E., *Phys. Rev. B*, **39**, 2084 (1989).
10. Ceperley, D. M., *Rev. Mod. Phys.*, **67**, 279–355 (1995).
11. Fahy, S., Wang, X. W., and Louie, S. G., *Phys. Rev. Lett.*, **61**, 1631–1634 (1988).
12. Fahy, S., Wang, X. W., and Louie, S. G., *Phys. Rev. Lett.*, **65**, 1478–1481 (1990).
13. Mitas, L., Shirley, E. L., and Ceperley, D. M., *J. Chem. Phys.*, **95**, 3467 (1991).
14. Martin, R. M., Kwon, Y., Li, X. P., Mitas, L., Natoli, V., Shirley, E. L., and Ceperley, D. M., “Quantum Monte Carlo calculations on real materials.,” in *Interatomic potential and structural stability: Proceedings of the Taniguchi Symposium*, edited by K. Terakura and H. Akai, Springer-Verlag, Berlin, 1993, p. 191.
15. Greeff, C. W., Lester, W. A., and Hammond, E. L., *J. Chem. Phys.*, **104**, 1973–1978 (1996).
16. Hood, R., Chou, M.-Y., Williamson, A., Rajagopal, G., Needs, R., and Foulkes, W., *Phys. Rev. Lett.*, **78**, 3350–3353 (1997).
17. Ceperley, D. M., *Phys. Rev. Lett.*, **69**, 331 (1992).
18. Ortiz, G., Ceperley, D. M., and Martin, R. M., *Phys. Rev. Lett.*, **71**, 2777–2780 (1993).
19. Jones, M. D., Ortiz, G., and Ceperley, D. M., *Phys. Rev. E*, **55**, 6202–6210 (1997).
20. Jacak, L., Wójs, A., and Harylack, P., *Quantum Dots*, Springer-Verlag, Berlin, 1998.
21. Ceperley, D. M., *Rev. Mod. Phys.*, **67**, 279 (1995).
22. Allan, G., Niquet, Y. M., and Delerue, C., *Appl. Phys. Lett.*, **77**, 639–641 (2000).
23. Niquet, Y. M., Delerue, C., Lannoo, M., and Allan, G., *Phys. Rev. B*, **64**, 113305 (2001).
24. Lee, S., Jönsson, L., Wilkins, J. W., Bryant, G. W., and Klimeck, G., *Phys. Rev. B*, **63**, 235307 (2002).
25. Lee, S., Kim, J., Jönsson, L., Wilkins, J. W., Bryant, G. W., and Klimeck, G., *Phys. Rev. B*, **66**, 235307 (2002).
26. Xie, R.-H., Bryant, G. W., Lee, S., and Jaskoólski, W., *Phys. Rev. B*, **65**, 235306 (2002).
27. Zhang, S., and Krakauer, H., *Phys. Rev. Lett.*, **90**, 136401 (2003).
28. Gunnarsson, O., Satpathy, S., Jepsen, O., and Andersen, O. K., *Phys. Rev. Lett.*, **67**, 3002–3005 (1991).
29. Satpathy, S., Antropov, V. P., Jepsen, O. K. A. O., Gunnarsson, O., and Liechtenstein, A. I., *Phys. Rev. B*, **46**, 1773–1793 (1992).
30. Bester, G., Shumway, J., and Zunger, A., *Theory of excitonic spectra and entanglement engineering in dot molecules* (2004), submitted.
31. Gunnarsson, O., Koch, E., and Martin, R. M., *Phys. Rev. B*, **54**, R11026–11029 (1996).
32. Foulkes, W. M. C., Mitas, L., Needs, R. J., and Rajagopal, G., *Rev. Mod. Phys.*, **73**, 33–83 (2001).

33. Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. M., and Teller, E., *J. Chem. Phys.*, **21**, 1087–1092 (1953).
34. Allen, M. P., and Tildesley, D. J., *Computer Simulation of Liquids*, Clarendon, Oxford, 1987.
35. C. Umrigar, K. W., and Wilkins, J. W., *Phys. Rev. Letts.*, **60**, 1719 (1990).
36. Ceperley, D. M., and Kalos, M. H., “Quantum many-body problems,” in *Monte Carlo methods in statistical physics*, edited by K. Binder, Topics in Condensed Matter Physics, Springer-Verlag, Berlin, 1986.
37. Ceperley, D., *Phys. Rev. B*, **18**, 3126 (1978).
38. Ceperley, D., Chester, G. V., and Kalos, M. H., *Phys. Rev. B*, **16**, 3081 (1977).
39. Holzmann, M., Ceperley, D. M., Pierleoni, C., and Esler, K., *Phys. Rev. E*, **68**, 046707:1–15 (2003).
40. Williamson, A. J., Hood, R. Q., and Grossman, J. C., *Phys. Rev. Lett.*, **87**, 246406 (2001).
41. Williamson, A. J., Grossman, J. C., Hood, R. Q., Puzder, A., and Galli, G., *Phys. Rev. Lett.*, **89**, 196803 (2002).
42. Draeger, E. W., Grossman, J. C., Williamson, A. J., and Galli, G., *Phys. Rev. Lett.*, **90**, 167402 (2003).
43. Trotter, H. F., *Proc. Am. Math. Soc.*, **10**, 545 (1959).
44. Kalos, M. H., Levesque, D., and Verlet, L., *Phys. Rev. A*, **9**, 2178–2195 (1974).
45. C. J. Umrigar, M. P. N., and Runge, K. J., *J. Chem. Phys.*, **99**, 2865–2890 (1993).
46. Reynolds, P. J., Hammond, B. L., Hammond, B. L., and Lester, Jr., W. A., *J. Stat. Phys.*, **43**, 1017–1026 (1986).
47. Hammond, B. L., Lester, Jr., W. A., and Reynolds, P. J., *Monte Carlo Methods in Ab Initio Quantum Chemistry*, World Scientific, 1994.
48. Kwon, Y., Ceperley, D. M., and Martin, R. M., *Phys. Rev. B*, **48**, 12037 (1993).
49. Ceperley, D. M., and Alder, B. J., *J. Chem. Phys.*, **81**, 5833 (1984).
50. Ceperley, D. M., and Bernu, B., *J. Chem. Phys.*, **89**, 6316 (1988).
51. Ceperley, B. B. D. M., and Lester, W. A., *J. Chem. Phys.*, **93**, 552 (1990).
52. Shumway, J., and Ceperley, D. M., *Phys. Rev. E* (2000).
53. Hurley, M. M., and Christiansen, P. A., *J. Chem. Phys.*, **86**, 1069–1070 (1987).
54. Hammond, B. L., Reynolds, P. J., and Lester, W. A., Jr., *J. Chem. Phys.*, **87**, 1130–1136 (1987).
55. Bachelet, G. B., Ceperley, D. M., and Chiocchetti, M., *Phys. Rev. Lett.*, **62**, 2088–2091 (1989).
56. Li, X. P., Ceperley, D. M., and Martin, R. M., *Phys. Rev. B Rapid Commun.*, **44**, 10929–10932 (1991).
57. Luttinger, J. M., *Phys. Rev.*, **102**, 1030–1041 (1956).
58. Schmidta, K. E., and Fantoni, S., *Phys. Lett. B*, **446**, 99–103 (1999).
59. Sarsa, A., Fantoni, S., Schmidt, K. E., and Pederiva, F., *Phys. Rev. C*, **68**, 024308 (2003).
60. Grossman, J. C., Mitas, L., and Raghavachari, K., *Phys. Rev. Lett.*, **75**, 3870–3873 (1995).
61. Murphy, R. B., and Friesner, R. A., *Chem. Phys. Lett.*, **288**, 403 (1998).
62. Kent, P. R. C., Towler, M. D., Needs, R. J., and Rajagopal, G., *Phys. Rev. B*, **62**, 15394–15397 (2000).
63. Torelli, T., and Mitas, L., *Phys. Rev. Lett.*, **85**, 1702–1705 (2000).
64. Grossman, J. C., and Mitás, L., *Phys. Rev. Lett.*, **74**, 1323–1323 (1995).
65. Mitas, L., Grossman, J. C., Stich, I., and Tobik, J., *Phys. Rev. Lett.*, **84**, 1479–1482 (2000).
66. Puzder, A., Williamson, A. J., Grossman, J. C., and Galli, G., *Phys. Rev. Lett.*, **88**, 097401 (2002).
67. Puzder, A., Williamson, A. J., Reboredo, F. A., and Galli, G., *Phys. Rev. Lett.*, **91**, 157405 (2003).
68. Shumway, J., Franceschetti, A., and Zunger, A., *Phys. Rev. B*, **63**, 155316 (2001).
69. Austin, E. J., *Semicond. Sci. Technol.*, **3**, 960–962 (1988).
70. Hu, Y. Z., Koch, S. W., Lindberg, M., Peyghambarian, N., Pollock, E. L., and Abraham, F. F., *Phys. Rev. Lett.*, **64**, 1805–1807 (1990).
71. Pollock, E. L., and Koch, S. W., *J. Chem. Phys.*, **94**, 6776–6781 (1991).
72. Tsuchiya, T., *Physica E*, **7**, 470–474 (2000).
73. Varga, K., Navratil, P., Usukura, J., and Suzuki, Y., *Phys. Rev. B*, **63**, 205308 (2001).
74. Bolton, F., *Phys. Rev. B*, **54**, 4780–4793 (1996).
75. Lee, E., Puzder, A., Chou, M. Y., Uzer, T., and Farrelly, D., *Phys. Rev. B*, **57**, 12281–12284 (1998).
76. Luczak, F., Brosens, F., Devreese, J. T., and Lemmens, L. F., Many body diffusion and interacting electrons in harmonic confinement (2000), available online, cont-mat/0002343.
77. Harting, J., Mülken, O., and Borrmann, P., *Phys. Rev. B*, **62**, 10207–10211 (2000).

78. Shumway, J., Fonseca, L. R. C., Leburton, J. P., Martin, R. M., and Ceperley, D. M., *Physica E*, **8**, 260–268 (2000).
79. Borrmann, P., and Harting, J., *Phys. Rev. Lett.*, **86**, 3120–3123 (2001).
80. Pederiva, F., Emperador, A., and Lipparini, E., *Phys. Rev. B*, **66**, 165314 (2002).
81. Carlson, J., and Pandharipande, V. R., *Nuc. Phys. A*, **424**, 47–59 (1984).
82. Simon, B., Academic, 1979.
83. Feynman, R. P., *Statistical Mechanics*, Addison-Wesley, Reading, MA, 1972.
84. Landau, L., and Lifshitz, E. M., *Quantum Mechanics (Non-relativistic Theory)*, Butterworth Heinemann, London, 1977, third edn.
85. Feynman, R. P., and Kleinert, H., *Phys. Rev. B*, **34**, 5080 (1986).
86. van Kampen, N. G., North Holland, 1981.
87. Barker, J. A., *J. Chem. Phys.*, **70**, 2914 (1979).
88. Storer, R. G., *J. Math. Phys.*, **9**, 964 (1968).
89. Parrinello, M., and Rahman, A., *J. Chem. Phys.*, **80**, 860 (1984).
90. Tuckerman, M. E., Berne, B. J., Martyna, G. J., and Klein, M. L., *J. Chem. Phys.*, **99**, 2796 (1993).
91. Car, R., and Parrinello, M., *Phys. Rev. Lett.*, **55**, 2471–2474 (1985).
92. Feynman, R. P., and Hibbs, A. R., *Quantum Mechanics and Path Integrals*, McGraw-Hill, New York, 1965.
93. Takahashi, M., and Imada, M., *J. Phys. Soc. Jpn.*, **53**, 963 (1984).
94. Coalsen, R. D., Freeman, D. L., and Doll, J. D., *J. Chem. Phys.*, **85**, 4567 (1985).
95. Chakravarty, C., Gordillo, M. C., and Ceperley, D. M., *J. Chem. Phys.*, **109**, 2123 (1998).
96. Feynman, R. P., *Phys. Rev.*, **91**, 1291–1301 (1953).
97. Penrose, O., and Onsager, L., *Phys. Rev.*, **104**, 576 (1956).
98. Draeger, E. W., and Ceperley, D. M., *Phys. Rev. Lett.*, **90**, 065301 (2003).
99. Pollock, E. L., and Runge, K. J., *J. Chem. Phys.*, **96**, 674 (1992).
100. Toennies, J. P., and Vilesov, A. F., *Annu. Rev. Phys. Chem.*, **49**, 1 (1998).
101. Grebenev, S., Toennies, J. P., and Vilesov, A. F., *Science*, **279**, 2083 (1998).
102. Viel, A., Patel, M. V., Niyaz, P., and Whaley, K. B., *Comp. Phys. Comm.*, **145**, 24 (2002).
103. Moroni, S., Sarsa, A., Fantoni, S., Schmidt, K. E., and Baroni, S., *Phys. Rev. Lett.*, **90**, 143401 (2003).
104. Northby, J. A., *J. Chem. Phys.*, **115**, 10065–10565 (2001).
105. Sindzingre, P., Klein, M. L., and Ceperley, D. M., *Phys. Rev. Lett.*, **63**, 1601–1604 (1989).
106. Sindzingre, P., Ceperley, D. M., and Klein, M. L., *Phys. Rev. Lett.*, **67**, 1871–1874 (1991).
107. Grebenev, S., Sartakov, B., Toennies, J. P., and Vilesov, A. F., *Science*, **289**, 1532–1535 (2000).
108. Kwon, Y., and Whaley, K. B., *J. Chem. Phys.*, **119**, 1986 (2003).
109. R. N. Silver, D. D. S., and Gubernatis, J. E., *Phys. Rev. B*, **41**, 2380 (1990).
110. Jarrell, M., and Gubernatis, J. E., *Physics Reports*, **269**, 134–195 (1996).
111. Boninsegni, M., and Ceperley, D. M., *J. Low Temp. Phys.*, **104**, 339 (1996).
112. Militzer, B., and Pollock, E. L., *Phys. Rev. E*, **61**, 3470 (2000).
113. Militzer, B., Pollock, E. L., and Ceperley, D. M., *Phys. Rev. B* (2003), submitted.
114. Ceperley, D., “Path Integral Monte Carlo Methods for Fermions,” in *Monte Carlo and Molecular Dynamics of Condensed Matter Systems*, edited by K. Binder and G. Ciccotti, Editrice Compositori, Bologna, 1996.
115. Baroni, S., and Moroni, S., *Phys. Rev. Lett.*, **82**, 4745–4748 (1999).
116. Sarsa, A., Schmidt, K. E., and Magro, W. R., *J. Chem. Phys.*, **113**, 1366–1371 (2000).
117. Wimmer, M., Shumway, J., and Nair, S. V. (2004), submitted.