

Metropolis Methods for Quantum Monte Carlo Simulations

D. M. Ceperley

*NCSA and Dept. of Physics
University of Illinois Urbana-Champaign
1110 W. Green St., Urbana, IL, 61801*

Abstract. Since its first description fifty years ago, the Metropolis Monte Carlo method has been used in a variety of different ways for the simulation of continuum quantum many-body systems. This paper will consider some of the generalizations of the Metropolis algorithm employed in quantum Monte Carlo: Variational Monte Carlo, dynamical methods for projector monte carlo (*i.e.* diffusion Monte Carlo with rejection), multilevel sampling in path integral Monte Carlo, the sampling of permutations, cluster methods for lattice models, the penalty method for coupled electron-ionic systems and the Bayesian analysis of imaginary time correlation functions.

INTRODUCTION

Though the original applications of the Metropolis *et al.* method[1] was to a classical system of hard disks, the algorithm has since been found indispensable for many different applications. In this talk, I will discuss some of these applications of the Metropolis algorithm to quantum many-body problems. This article will be strictly limited to the use of the Metropolis rejection method within quantum Monte Carlo (QMC) and not discuss other aspects of QMC. The richness of the Metropolis algorithm, and the brevity of this article implies that I can only touch briefly on a subset of these developments and must limit myself to a superficial discussion. Others will discuss its use in quantum lattice models, both of condensed matter and in lattice gauge theory, so my focus will be on non-relativistic continuum applications, in particular, on developments requiring generalization of the basic Metropolis algorithm. I will only mention briefly the physics behind these applications, instead referring to review articles.

We define the Metropolis algorithm as follows. Suppose s is a point in a phase space and we wish to sample the distribution function $\pi(s)$. In the simplest algorithm, there is a single transition probability: $T(s \rightarrow s')$. Later we will generalize this to a menu of transition probabilities. One proposes a move with probability $T(s \rightarrow s')$ and then accepts or rejects the move, with an acceptance probability $A(s \rightarrow s')$. Detailed balance and ergodicity are sufficient to ensure that the random walk, after enough iterations, will converge to $\pi(s)$, where by detailed balance we mean that:

$$\pi(s)T(s \rightarrow s')A(s \rightarrow s') = \pi(s')T(s' \rightarrow s)A(s' \rightarrow s). \quad (1)$$

By ergodicity, we mean that there is a non-zero probability of making a move from any state to any other state in a finite number of moves. We refer to this class of algorithms as

Markov Chain Monte Carlo (MCMC) or Metropolis. The generalization of Metropolis to non-uniform transitions was suggested by Hastings[2]. Key defining features of MCMC are the use of detailed balance to drive the distribution to a desired equilibrium state and, in particular, the use of rejections to achieve this.

In the “classic” Metropolis algorithm[1], the state space is the $3N$ vector of coordinates of particles and the distribution to be sampled is the classical Boltzmann distribution, $\pi \propto \exp(-\beta V(s))$. The moves consist of single particle displacements chosen uniformly in a hypercube centered around the current position and the moves are accepted with probability given by $\min[1, \exp(-\beta(V(s') - V(s)))]$. In the hard sphere system, it comes down to determining whether there are overlapping spheres in the proposed move.

How can this simple Metropolis algorithm be generalized? First of all for quantum system not only does the distribution need to be sampled, the many-body wavefunction is unknown. This is done by augmenting the spatial coordinates with an imaginary time, thereby mapping the d dimensional quantum system onto a $d + 1$ dimensional space, which is then sampled with MC. In some situations the space needs further enlarging, for example to add a permutation of particles in taking into account Bose or Fermi statistics. The biggest problem with the Metropolis method is that the approach to equilibrium and the autocorrelation of properties can be very slow. So another important generalization is in finding new, more efficient ways of moving through the sample space. In the case of the penalty method, discussed below, we allow noisy and cheap estimates of the potential energy function. Finally, the Bayesian methods use MCMC to extract the last bit of information from data generated by another MCMC simulation, *i. e.* for data analysis.

There are several themes of this article. First, it is fruitful to generalize MCMC to much more complicated systems. Second, without the ability to generalize, much of quantum many-body physics would still not be accessible to theory. Finally, the algorithmic techniques easily cross disciplinary boundaries, from condensed matter, to high energy, nuclear and chemical physics and more recently into statistics and economics, making a meeting such as this important. The Metropolis technique is indeed perhaps the most powerful computational method (the criterion being that there are no alternative algorithms) and has become the computational common denominator amongst those dealing with many-body systems. Of course, this generality implies that there have been many developments of which I can only touch on a few. Others will be covered elsewhere in this conference.

VARIATIONAL MONTE CARLO

The first major application of Quantum Monte Carlo (QMC) to a many-body system was by McMillan[3] to liquid ^4He . Previous calculations were for only a few particles and are discussed by Kalos elsewhere in this volume. The method McMillan used is now referred to as Variational Monte Carlo and is based on the familiar variational method for solving ground state quantum problems. As often happens, there was a simultaneous calculation by Levesque et al.[4] never published.

Assuming an appropriate trial function $\psi_T(R; a)$ where R are the particle coordinates possibly including spin, and a is a set of parameters, the variational energy,

$$E_V = \frac{\int dR \psi_T^* \hat{H} \psi_T}{\int dR |\psi_T|^2} = \langle \psi_T^{-1} \hat{H} \psi_T \rangle \quad (2)$$

is an upper bound to the ground state energy. In the second equality, $\langle \dots \rangle$ denotes an average over the probability distribution of the trial function $|\psi_T|^2$. It is in generating this distribution, that Metropolis Monte Carlo comes into play. Traditionally, in the variational method the above integrals are done deterministically, severely limiting the form of the trial function. However, using MCMC, any distribution can be sampled including those having explicit particle correlations. In fact, it is a consequence of the large value of the electron-nuclear mass ratio that mean field methods are so pervasive in solid state physics and quantum chemistry. Such methods are much less efficacious for liquid ^4He and other important correlated quantum systems.

McMillan in his pioneering work on the simulation of liquid ^4He used the pair product (Jastrow) trial function:

$$\psi_2(R) = \exp\left[-\sum_{i<j} u(r_{ij}; a)\right]. \quad (3)$$

This form is identical to that of a classical Boltzmann distribution for a system of atoms interacting with a pair potential, such as a rare gas liquid if we make the substitution $u(r) \rightarrow v(r)/k_B T$. His work and many that have followed, have established that this wavefunction provides a physically correct, though not exact, description of the ground state of liquid helium. Though the code that computed the distribution was thereby identical to a classical code, McMillan had to introduce new features in the algorithm in order to calculate properties such as the kinetic energy and momentum distribution and to optimize parameters in the trial function.

To treat fermion systems, such as liquid ^3He , one generalizes the trial function by multiplying by a Slater determinant of one body orbitals, a Slater-Jastrow trial function. There was a twelve year lag in generalizing VMC to fermion systems, primarily because of a psychological barrier. The perception in the community was that fermion MC would be too slow and the algorithm could be non-ergodic if the determinant were to be included in the sampling, so instead, approximate methods were introduced take into account antisymmetry.

Concerning the question of efficiency, it takes order N^3 operations to evaluate a determinant. However, it turns out there is an algorithm, the Sherman-Morrisson formula, discovered when determinants were evaluated by hand, that allows one to change a single column or row in only N^2 operations[5]. Using the update method with Slater determinants, it is possible to do single particle moves within MCMC quite efficiently. In fact, until one reaches several hundred fermions, the time per Monte Carlo step is not much slower than MCMC for the equivalent bosonic system. One implication for the algorithm is that rejections are much “cheaper” than acceptance. As a result efficient VMC fermion calculations have larger step sizes and much smaller average acceptance ratios than usual.

The other psychological barrier concerned the non-classical distribution when a squared determinant is present in the trial function. If the trial function is real, or de-

scribes a closed shell (non-degenerate ground state), the phase space will be divided into regions separated by the nodes of the trial function. It seems possible that the random walk could get trapped in one pocket of phase space, giving a bias to computed properties. However, once a Metropolis random walk with a determinant was attempted, it was observed that the sign of the trial function changed every few steps. The issue of ergodicity was finally settled with the proof of the tiling theorem[6]: for determinants coming from the solution of mean field theory, all the various pockets are equivalent, so that even if the random walk remained in one pocket, the results would be unbiased.

A side effect of attempting to simulate fermion systems was a generalization of MCMC, known as “force-bias” MC. In MCMC there is always a struggle between moves with high acceptance rates that go long distances and moves that are fast¹. Until the update formula allowing much faster single particle moves was employed, N -particle moves were used. However, as N becomes large, one is forced to take increasingly smaller step sizes to get reasonable acceptance rates. Note that in better wavefunctions such as backflow, all rows depend on all of the particle coordinates, so update methods are not useful and non-uniform methods are needed for these more accurate VMC calculations. Kalos suggested[5] that a non-uniform transition probability, $T(s \rightarrow s')$ which locally approximates the equilibrium distribution will have a larger acceptance rate, and hence will allow a larger step size and faster convergence. (Note that this followed the work of Hastings[2].) These “directed” improvements to the classic Metropolis were soon picked up for simulations of classical liquids giving rise to what was called “force-biased” Monte Carlo[7] and “smart” Monte Carlo[8].

In the “smart” Monte Carlo algorithm, a form that will also appear in diffusion MC, one uses an offset Gaussian:

$$T(s \rightarrow s') = C \exp(-(s' - s - D \nabla \ln(\pi))^2 / 4D) \quad (4)$$

where \sqrt{D} is the step size and C a normalization constant. Then the new move is accepted with probability equal to:

$$A(s \rightarrow s') = \min\left[1, \frac{T(s' \rightarrow s)\Pi(s')}{T(s \rightarrow s')\Pi(s)}\right] \quad (5)$$

By making a Taylor expansion about the current positions, one can verify that the acceptance probability deviates from unity only by terms second order in $s' - s$.

Another generalization of the Metropolis algorithm within VMC concerns the properties of a bosonic solid[9]. An accurate trial function for the quantum solid is a symmetrized product of localized functions:

$$\psi_T = \Psi_2(R) \frac{1}{\sqrt{N!}} \sum_P \prod_i \phi(r_i - Z_{P_i}) \quad (6)$$

where P is a permutation and $\phi(r)$ is a localized orbital, *e. g.* $\exp(-cr^2)$, and Z_i is the set of lattice sites. Mathematically the solid wavefunction is a permanent, which is

¹ This illustrates richness of the Metropolis method; one has great freedom in attempting moves. In molecular dynamics, once you specify the Lagrangian you are more or less stuck with the resulting dynamics.

very slow to evaluate explicitly once N gets large with an operation count proportional to $N 2^N$. However, it is easy to add the permutation to the sample space since the trial wavefunction is non-negative for all permutations. To sample the combined space, we need to add a transition move to change the permutation. Pair permutations are sufficient to move through the entire space of $N!$ permutations. This way of generalizing Metropolis had a spin-off into another area of quantum physics: analyzing the debris produced by colliding bosons in particle accelerators where, again, one must symmetrize over assignment of particle labels[10]. The idea of using Metropolis simultaneously in both a continuous and discrete space turned out to be essential in the VMC simulation of nuclei. Lomnitz-Alder et al. [11] sampled the ordering of spin and tensor operators in a VMC calculation of small nuclei by executing a random walk in the operator ordering along with normal moves of the positions of the nucleons.

For inhomogeneous systems, the above form of trial function of a crystal is awkward since it requires prior knowledge of the crystal sites. In the *shadow wavefunction*[12] for each real particle, one adds a complimentary shadow particle, also with a Jastrow correlation and coupled to the real particles. The shadow particle acts like the lattice site, but can move to find its optimal position. This function can be considered a cousin to path integrals, discussed below. The shadow trial function spontaneously orders, so one does not have to specify the lattice beforehand. It gives lower energy at the expense of additional integrals, which don't cost much in VMC anyway. Without the MCMC method, this wave function, which is quite accurate for a variety of problems, would be unlikely to be considered.

One issue that arises for these more complicated trial functions is how to move through the space efficiently. The simplest way to implement the random walk is through a "menu" of moves, where first one takes a particle coordinate move and then a permutational change move (or a shadow move). One generalizes Eq. 1 to enforce detailed balance for each move separately. However, the product of such probabilities, needed to have a Markov process, does not, in general, satisfy detailed balance. But this is not important; it is sufficient that the desired stationary state be an eigenfunction of each of the possible menu items. The freedom to depart from strict detailed balance is invaluable. This was commonly known[2] in the early days though rarely discussed.

BROWNIAN DYNAMICS AND DIFFUSION MONTE CARLO

It is often stated that one does not use MC to study dynamical problems. However, early on, MCMC was used to study kinetic phenomena in the Ising model[13]. Around 1977, we began doing simulations of polymer dynamics[14, 15]. A polymer in a solvent moves by Brownian motion. If hydrodynamic forces are ignored then the master equation for the probability density is the Smoluchowski equation:

$$\frac{df(R,t)}{dt} = D\nabla[\nabla + \beta\nabla V(R)]f(R,t). \quad (7)$$

Here $f(R;t)$ is the distribution function for the system at time t and D is the diffusion constant. Rossky et al. [8] observed that the smart MC algorithm is equivalent to one

step of the Langevin equation assuming that the velocity is randomized at each step. This randomization occurs physically because a large blob of the polymer gets frequently hit by the solvent during each time interval.

A problem we encountered in the simulation of the Brownian dynamics was that it was necessary to take a very small “time step” in order to get the proper equilibrium distribution, because two particles would overlap in a region where the interparticle potential was highly non-linear and the subsequent step they would be thrown into a completely unphysical region. Initially, we would solve this with a kludge; by putting an upper limit on the force. However, after our experience with smart MC with VMC, we decided to enforce detailed balance at each step, by accepting or rejecting à la Metropolis. Detailed balance is a property of the exact solution of Eq. 7, and quite easy to enforce with rejections. This allows better scaling in the number of particles and a convenient way to decide on what time step to take. Of course, the equilibrium distribution will be exactly the Boltzmann distribution for any time step, but to get realistic dynamics, we adjusted the time step to get the average acceptance ratio greater than 90%. One can approximately correct for the effect of rejections by rescaling the time in order to get the exact diffusion constant. This use of rejection in Brownian dynamics was a precursor to hybrid methods latter developed within lattice gauge theory[16]. In hybrid methods, one typically takes multiple dynamical steps before deciding whether to accept the trajectory.

Now returning to quantum mechanics, it is highly desirable to go beyond variational MC, since it is difficult to get more out of the simulation than is put into the trial wavefunction. One needs a more automatic scheme, where the stochastic process generates the distribution. Such an approach, attributed to Fermi and Wigner, who realized that the non-relativistic Schrodinger equation in imaginary time is a random walk, and in the limit of large imaginary time gives the ground state of the quantum system

$$-\frac{d\phi(R;t)}{dt} = [-D\nabla^2 + V(R)]\phi(R;t) \quad (8)$$

where $D = \hbar^2/2m$. The first many-body application of this approach (GFMC) by Kalos *et al.*[17] is not a Metropolis method but a zero time step error method, in which imaginary time is sampled. GFMC is based on the integral equation formulation of the eigenvalue problem. Related methods have been reintroduced in recent years in quantum lattice models under the names of continuous time world-line algorithms[18] as discussed elsewhere in this volume by Troyer.

After working on the dynamics of polymers, the similarity between the GFMC approach and that of Brownian Dynamics was apparent. The connection is made by applying importance sampling to Eq. 8: multiply the equation by a trial function Ψ_T and rewrite in terms of $f = \phi\Psi_T$. The resulting master equation for f is:

$$-\frac{df(R;t)}{dt} = [-D\nabla^2 + 2D\nabla(\nabla \ln |\psi_T|) + \psi_T^{-1}\hat{H}\psi_T]f(r;t). \quad (9)$$

This is the same as Eq. 7 for Brownian dynamics except for the addition of the last term which is a branching process already familiar within GFMC. The process of solving the Schrodinger equation with a drift, diffusion and branching process is known as diffusion

Monte Carlo (DMC). The values of imaginary times are not sampled as in GFMC, so the method does have a time step error; however the concepts of detailed balance and Metropolis rejection are applicable since the exact evolution satisfies:

$$|\Psi_T(R)|^2 f(R \rightarrow R'; t) = |\Psi_T(R')|^2 f(R' \rightarrow R; t). \quad (10)$$

This is the familiar detailed balance equation, however with a subtle difference: $f(R \rightarrow R'; t)$ is not a normalized p.d.f. because of the branching process. Nonetheless, adding a rejection step is both possible and highly desirable and results in faster convergence at large N as compared with GFMC². More importantly it is simpler and easier to integrate with fixed-node method [19] for treating Fermi statistics and is now the almost universal choice for zero temperature quantum problems.

Methods higher order in the time step have been occasionally investigated to simulate the diffusion Monte Carlo equation[20]. However, if they rely on expansions of the trial function or potential, the algorithm will fail badly at non-analytic points, such as when two particles get close together. As a consequence, higher order methods have not been very successful for continuum electronic systems. The approach based on rejections has a different principle than expansion in the timestep, namely detailed balance is put in. The equivalent in deterministic algorithms such as molecular dynamics, is reversibility in time. The crucial question is not the integration order, but how much computer time it takes to get the error to a certain accuracy. By enforcing detailed balance in DMC, one gets the correct result both as the time step goes to zero and as the trial wavefunction gets more accurate.

The DMC method had a spectacular debut[21], in a paper which has the highest citation count of any simulation. The citations were not for introducing the DMC method, but because the energy of a homogenous system of electrons is taken as the reference system in density functional calculations of molecules and solids. The details of the DMC method are given in ref. [22]. The Metropolis algorithm using rejection was crucial in giving accurate results for a variety of densities, phases and particle numbers. A recent review of applications obtained with DMC are described in Foulkes *et al.*[23].

PATH INTEGRAL METHODS

The previous sections described applications at zero temperature. We now consider the finite temperature Path Integral Monte Carlo algorithms. In the same year as the celebrated Metropolis paper, Feynman[24] showed that bosonic systems in equilibrium are mathematically isomorphic to classical “polymer-like” systems. A single classical particle turns into a ring polymer $r(t)$ where t the time index is in the range $0 \leq t \leq \beta = (k_B T)^{-1}$. In discrete-time path integrals, the time index has only discrete values, $t_k = k\tau$ with $1 \leq k \leq M$, τ is the time step, and k the “Trotter” index. The equivalent to the classical potential of the “polymer” is a sum of the “spring” terms (from the quantum

² There is a problem with the combination of branching and rejection: that of persistent configurations. If there is a region of phase space where $qe^{-\tau(E_L - E_T)} > 1$ then the algorithm is unstable.

kinetic energy) and the potential energy and called the action:

$$S(R(T)) = \sum_{i=1}^M \frac{(r_i - r_{i+1})^2}{4D\tau} + \tau V(r_i). \quad (11)$$

An important aspect of equilibrium paths is that they are periodic in imaginary time, which means for distinguishable particles $r(t + \beta) = r(t)$.

The most interesting applications of PIMC involve systems with Bose and Fermi statistics. There one must symmetrize over the particle labels by allowing the paths to close on a permutation of themselves. For Bose systems, all permutations have a positive contribution. In Feynman's theory[24], it was the onset of a macroscopic permutation which was responsible for the phase transition of liquid ^4He at low temperature and the observable properties of Bose condensation and superfluidity. The MCMC simulations were crucial[25] in finally convincing most of the low temperature community that BEC really is responsible for superfluidity and that Feynman had it right after all. Applications to PIMC to helium are discussed at length in ref [26].

There was a twenty-five year lag between the development of Metropolis MC and Feynman path integrals and the large scale computer applications of PIMC. There was considerable small-scale work during this period, much of it unpublished (see Jacucci[27]), but not until the late 1970's did this mature into computational efforts attacking important physical systems, in the lattice gauge community[28], in chemical physics[29] and in lattice models for solid state physics[30, 31, 32]. This 25 year lag was most likely due to a lack of access to sufficiently powerful facilities, combined with use of the basic Metropolis algorithm, which is notoriously inefficient for path integrals.

There were some problems to overcome in using MCMC to simulate bosonic superfluids or indeed any system at a low temperature[33, 34]. The primary problem is similar to that encountered in polymer simulations; namely, that as a polymer gets longer the correlation time of the random walk increases. One can show[26] that the efficiency of any MC with local moves will drop at least as fast as M^{-3} even for free particles. This leads to a bad scaling versus the number of time slices. There are several overlapping solutions to this problem. At the technical level, one wants the best feasible approximation for the action so as to avoid the necessity of many time slices (*i. e.* large M). Though this is an important topic, it is irrelevant for the purposes of this lecture. The other approach is to try to optimize how the state is changed: the transition probability.

Essentially, the problem is how to move a chunk of the path together. In order to get a permutation move accepted of n atoms, in liquid ^4He , it is essential to move a substantial section, typically 8 time slices of the n atoms involved. (Hence one needs to move 24 coordinates at once.) One wants to sample gross features of the change and then sample finer details. The idea is not to waste time computing details until gross features are shown to be reasonable. There were several attempts to make an efficient algorithm including the "staging" method[35] and a method inspired by diffusion Monte Carlo[33]. Here we will sketch the most successful generalization, multilevel sampling[34, 26].

In the multilevel method, a move is partitioned into ℓ levels with an approximate "action" or distribution function for that level $\pi_\ell(s)$ with the requirement that the level action equal the true action at the highest level: $\pi_\ell(s) = \pi(s)$. One samples the trial variables at each level according to some probability distribution, $T(s'_k)$. Then those variables are

tentatively accepted with a generalized Metropolis formula, ensuring detailed balance at each level:

$$A_k = \min\left[1, \frac{T_k(s)\pi_k(s')\pi_{k-1}(s)}{T_k(s')\pi_k(s)\pi_{k-1}(s')}\right]. \quad (12)$$

If a move is rejected at any level, one returns to the lowest level and constructs a completely new move. This algorithm is applied to path integrals by first, sampling the midpoint of the path, and with a certain probability, continue the construction, sampling the midpoints of the midpoints, etc. One gains in efficiency because the most likely rejection will occur at the first level, when only $2^{-\ell}$ of the computational work has been done.

Now consider the problem of how to carry out the bisection; in other words, how to sample the midpoint of a Brownian bridge. Given two fixed end points of the bridge, R_0 at time 0 and R_β at time β , what is the distribution of a point on the bridge R_t at time t with $0 < t < \beta$? For free Brownian motion this was solved by Lévy. For any quantum system, the probability distribution of R_t is:

$$T(R_t) = \frac{\langle R_0 | e^{-t\hat{H}} | R_t \rangle \langle R_t | e^{-(\beta-t)\hat{H}} | R_\beta \rangle}{\langle R_0 | e^{-\beta\hat{H}} | R_\beta \rangle}. \quad (13)$$

In this respect, path integrals are simpler than polymers, since the action only contains terms local in imaginary time. For free particles $T(R_t)$ is a Gaussian with an easily computed mean and variance. For interacting particles, one needs to approximate this by a Gaussian, with a mean and covariance perturbed from the free particle values by the interaction with neighboring particles[26]: an approach similar to smart Monte Carlo.

To start off the multilevel sampling, one first samples the permutation change. In the case of the quantum crystal described in the previous section, we were only interested in small local permutations. However in a superfluid, one is particularly interested in permutation changes which span the entire cell. To get winding number changes, which changes the superfluid density, you need to construct a permutation which can cross the entire system. That is, the cycle length needed to make a change in the winding number is roughly equal to $N^{1/3}$ in 3D. Such a permutation can be found with a random walk in index space[26]. Once the permutation is established, then the actual path is constructed.

The multi-level method has been independently discovered several times. A recent example is the technique of pre-rejection used for classical simulations[36]. Suppose one computes an empirical pair potential first and then a more complicated potential using LDA after the first screening is done. If one has an accurate empirical potential, one can quickly make many large displacement moves, perhaps with a fairly low acceptance probability, and then only on those rare moves that make it through the first level, go to the expense of computing the accurate potential. Multilevel sampling can also be used to improve the efficiency of VMC[37]. There are also similar ideas developed in the polymer world as described in Frenkel's contribution.

For Fermi statistics, one subtracts the sum of the odd permutations from that contribution of even permutations, leading to the infamous fermion sign problem. We are still struggling with this problem today. One approach is to restrict the paths to stay in the positive half of phase space as defined by the fermion density matrix[38]. This would be

a rigorous procedure if we knew how to partition the space, but in practice one needs to make an ansatz for the restriction. A key unsolved problem for these restricted path integrals is that the dynamics appears very slow and non-ergodic at very low temperatures.

A key property for a Bose superfluid is the momentum distribution, or its Fourier transform, the single particle off-diagonal density matrix. This is obtained in path integrals by allowing one path to be open, a linear “polymer.” The two open ends of this polymer can become separated in a superfluid if that atom is a part of a long permutation cycle. Linear polymers have a very efficient way to move through phase space, the reptation motion, (*i.e.* move like a slithering snake), developed for lattice and continuum polymer simulations[39, 40]. The reptation algorithm is obtained by cutting off one end of the polymer and growing that part onto the other end while keeping the body of the snake and the other polymers unchanged. This is a very fast operation both in computer time and in how quickly it refreshes the configuration of the polymer. If one allows the length of the polymer to fluctuate, growth at one end and shrinkage at the other, need not be explicitly coupled. Real world polymers are polydisperse anyway. The reptation algorithm is another example of a fruitful insight in quantum algorithms coming from the polymer world.

One application of the reptation algorithm for quantum simulations is to ground state path integral calculations[41]. In the ground state limit, closing of the paths becomes unimportant, and instead one works with open polymers, closed on the ends with a trial wavefunction. Since they are open, they can move by reptation. This gets around the DMC problem of mixed estimators.

There has been considerable development of MCMC algorithms for quantum lattice models. There are several crucial distinctions between the lattice models and the continuum models, even for bosonic systems. First of all, on the lattice, the action is bounded, leading to other ways of approximating the action. Secondly, there is a finite set of possible local moves, allowing one to use heat bath methods. Suppose, we define the “neighborhood” of a state, as all states that can be reached by a certain class of moves. By heat bath we mean that we directly sample the equilibrium distribution in the neighborhood: $C_s \pi(s)$. Finally, in lattice models the random walks are not continuous trajectories. Some important principles such as fixed-node and winding number estimators were discovered in the continuum because they require continuous trajectories.

PIMC for a lattice model such as the Bose Hubbard model, is known as “world line Monte Carlo.” For the reasons discussed above, it has problems with convergence in the low temperature limit. Progress[42, 43] has made with loop and cluster moves, as described elsewhere in this volume by Troyer. These ideas have given rise to the “meron” methods to solve the sign problem for certain models[44]. In Prokofev’s[45] worm algorithm, one starts with an open polymer and allows the two ends to grow and shrink independently, as we described above with reptation. One gets correct equilibrium statistics (which require closed loops) by taking averages over only those configurations where the head and tail happen to land on the same sites. These new methods have allowed simulations of large lattices and computation of critical properties of quantum phase transitions.

Lattice PIMC for fermion systems is referred to as determinantal MC. There one performs a Stratonovitch-Hubbard transformation of the interaction term, leading to the interaction of a Slater determinant with a random field[31]. Heat bath algorithm and

fermion update formulas are used in the implementation of MCMC in this approach[46]. Aside from the half filled Hubbard model, one has a serious fermion sign problem. Recent progress has been made in developing fixed-node approaches[47, 48] for determinantal Monte Carlo.

METROPOLIS WHEN THE ENERGY FUNCTION IS RANDOM

A significant generalization of the MCMC algorithm is the penalty method. In most of the classical MC applications to date, it is assumed that the energy function is computable in a finite number of operations, and most applications before 1985 used an empirical pair potential, with an occasional more complex functional form. In 1985, Car-Parrinello[49], showed that one can solve the mean field density functional equations at each step in a molecular dynamics simulation. However, to reach the accuracy needed for many practical problems it will be necessary to go far beyond mean field or semi-empirical approaches, greatly increasing computer time.

An approach that we are following[50] is to calculate the electronic Born-Oppenheimer energy at zero temperature using a DMC random walk. The ions are moved at a non-zero temperature with MCMC, possibly using multilevel sampling. However, this means that the energy difference in the Metropolis acceptance formula will not be known precisely but will have a statistical fluctuation. For high accuracy, one will need to reduce these fluctuations to much below $k_B T$ to get reliable results, but how much lower does one need to go? The difference in computer time in going from an error of $k_B T$ to $k_B T/10$ is a factor of 100! Several years ago[51], we raised the question of whether it was possible to take into account these fluctuations in the energy in the Metropolis algorithm. There have been a few suggestions[52, 53] about to handle noisy energy evaluations in the past but without concern about large fluctuations in the energy or of the efficiency of the approach. Suppose that δ is an estimate of ΔE from a known probability distribution, $P(\delta)$. That is $\int d\delta P(\delta)\delta = \Delta E$. Let us require detailed balance on the average:

$$\int d\delta P(\delta) [\pi(s)A(\delta) - \pi(s')A(-\delta)] = 0 \quad (14)$$

where for simplicity, we have assumed a symmetric sampling function, $T(s \rightarrow s') = T(s' \rightarrow s)$. Assuming reasonable conditions on the QMC evaluation of the energy, $P(\delta)$ will approach a normal distribution with a variance σ^2 . Somewhat surprisingly, a nearly optimal solution to the detailed balance equation for a normal distribution has been discovered[51]. To satisfy detailed balance on average one needs to accept a move with probability:

$$A(s \rightarrow s') = \min[1, \exp(-\beta\delta - \beta^2\sigma^2/2)]. \quad (15)$$

One must add a *penalty* equal to $\beta^2\sigma^2/2$ to the energy difference to compensate for the fluctuations. This is much more efficient than simply beating down the error bars. The most efficient simulation is one for which $\beta\sigma > 1$: it is better to take cheap moves, many of which are likely to be rejected, rather than many fewer expensive moves with small values of σ^2 .

Clearly there are other situations aside from quantum simulations where one might wish to evaluate the change in energy only approximately. For example, one might imagine that a classical energy can be split into large short-ranged terms, and slowly varying long-range terms which are slow to evaluate. Those terms can then be sampled. There are recent suggestions on how to use this for protein folding and to deal with non-normally distributed energy fluctuations[54].

BAYESIAN ANALYSIS

In projector and path integral Monte Carlo, the “dynamics” is in imaginary time. By that is meant that we sample matrix elements of $\exp[-\beta\hat{H}]$. An important problem is to extract a maximum amount of information from the imaginary time correlations of these simulations. For example, it is well known that real-time linear response is related by a Laplace transform to the time correlations in PIMC. In this last application of MCMC, we consider how it can be used in the data analysis. This application of MCMC is completely different than the previous ones since the random walk is not in the space of particle coordinates or permutations, but in the space of the real time response.

In fact, we consider a related problem arising from the fermion “sign” problem. The exact transient estimate algorithm[55] allows projection of the ground state for a limited time, but at large time β , the estimates get increasingly noisy. The simple solution is to take the largest time projection as the best estimate of the fermion energy. However, it is clear that there is more statistical information if the earlier values of β are also used[56]. In transient estimate MC, we determine:

$$h(t) = \langle \Psi_T | \exp[-\beta\hat{H}] \Psi_T \rangle \quad (16)$$

as well as the time derivatives of $h(t)$ for a range of times $0 \leq t \leq \beta$. As β gets large, the signal to noise ratio for the energy goes exponentially to zero. But $h(t)$ is related to the spectrum of \hat{H} by:

$$h(t) = \int_{-\infty}^{+\infty} dE c(E) e^{-tE} \quad (17)$$

where the spectral density is

$$c(E) = \sum_i \delta(E - E_i) |\langle \Psi_T | \phi_i \rangle|^2 \quad (18)$$

and (E_i, ϕ_i) are the exact energies and wavefunctions for state i . Some analytic information is known about the spectral density, namely, that it is positive, is identically zero for $E < E_0$ and decays at large E as E^{-k} . Since the “Laplace” transform in Eq.17 is a smoothing operation, the inverse transform needed to find $c(E)$ from $h_{MC}(t)$ is ill-conditioned. Because the evaluation of $h_{MC}(t)$ is stochastic, there is a distribution of $c(E)$ ’s, all of which are consistent with the DMC-determined data. Taking the Bayesian point of view, there is a prior distribution of $Pr_m[c]$ which is conventionally chosen to be an entropic function. Then the posterior distribution of $c(E)$ is given by the Bayesian formula:

$$Pr(c|h_{MC}) \propto Pr_L(h_{MC}|c)Pr_m[c] \quad (19)$$

where $Pr_l(h_{MC}|c)$, the likelihood function, is the probability of obtaining the Monte Carlo data assuming a given spectral density, $c(E)$. By the central limit theorem, the likelihood function is a multivariate Gaussian, whose parameters we can estimate within DMC.

The maximum entropy method, MAXENT[57] consists of finding the most probable value of $c(E)$ and estimating errors by expanding around the maximum. However, MCMC can be put to good use in sampling the distribution of $c(E)$, particularly for cases where the overall distribution $Pr(c|h_{MC})$ is non-gaussian. To do this one represents $c(E)$ on a finite grid, considers moves that change the values of $c(E)$ and accepts or rejects such moves based on how the move changes the value of $Pr(c|h_{MC})$. The structure of this distribution function is different than those that arise in simulations of particle or lattice systems. For example, the moving distance for $c(E)$ should depend on E . Though this analysis takes longer than MAXENT to estimate the spectrum, it typically takes much less time than the original DMC calculation that generated $h_{MC}(t)$. By looking at the output of the MCMC sampling of $c(E)$ one can get a quantitatively precise estimate of which spectral reconstructions are likely. More physical prior functions or analytic insight are easy to put into the distribution function. The density-density response function ($S_k(\omega)$) of liquid ^4He [58] has been calculated using this Metropolis procedure. These statistical estimation methods are increasingly used in computational statistics[59], under the acronym MCMC and are more fully described elsewhere in this volume.

ACKNOWLEDGMENTS

This research was funded by NSF DMR01-04399 and the Dept. of Physics at the University of Illinois. Much of the early work described here was done in collaboration with G. V. Chester and M. H. Kalos. M. H. Kalos is also acknowledged for comments on an early version of this manuscript.

REFERENCES

1. Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A., and Teller, E., *J. Chem. Phys.*, **21**, 1087 (1953).
2. Hastings, W. K., *Biometrika*, **57**, 97–109 (1970).
3. McMillan, W. L., *Phys. Rev. A*, **138**, 442 (1965).
4. Levesque, D., Khiet, T., Schiff, D., and Verlet, L., *Orsay report, unpublished* (1965).
5. Ceperley, D., Chester, G. V., and Kalos, M. H., *Phys. Rev. B*, **16**, 3081 (1977).
6. Ceperley, D. M., *J. Stat. Phys.*, **63**, 1237 (1991).
7. Pangali, C., Rao, M., and Berne, B. J., *Chem. Phys. Lett.*, **55**, 413–17 (1978).
8. Rossky, P., Doll, J. D., and Friedman, H. L., *J. Chem. Phys.*, **69**, 4628–33 (1978).
9. Ceperley, D., Chester, G. V., and Kalos, M. H., *Phys. Rev. B*, **17**, 1070 (1978).
10. Zajc, W. A., *Phys. Rev. D*, **35**, 3396 (1987).
11. Lomnitz-Adler, J., Pandharipande, V. R., and Smith, R. A., *Nucl. Phys.*, pp. 399–411 (1981).
12. Vitiello, S., Runge, K., and Kalos, M. H., *Phys. Rev. Letts.*, **60**, 1970–1972 (1988).
13. Binder, K., and Kalos, M. H., “Monte Carlo Studies of Relaxation Phenomena: Kinetics of Phase Changes and Critical Slowing Down,” in *Monte Carlo methods in Statistical Physics*, edited by

- K. Binder, *Topics in Condensed Matter Physics*, Springer-Verlag, Berlin, 1979.
14. Ceperley, D., Kalos, M. H., and Lebowitz, J. L., *Phys. Rev. Letts.*, **41**, 313 (1978).
 15. Ceperley, D. M., Kalos, M. H., and Lebowitz, J. L., *Macromolecules*, **14**, 1472 (1981).
 16. Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D., *Phys. Letts. B*, **195**, 216–222 (1987).
 17. Kalos, M. H., Levesque, D., and Verlet, L., *Phys. Rev. A*, **9**, 2178–2195 (1974).
 18. Prokof'ev, N. V., Svistunov, B. V., and Tupitsyn, I. S., *JETP Lett.*, **64**, 911–916 (1996).
 19. Anderson, J. B., *J. Chem. Phys.*, **63**, 1499–1503 (1975).
 20. Helfand, E., *Bell Syst. Tech. J.*, **581**, 2289 (1979).
 21. Ceperley, D. M., and Alder, B. J., *Phys. Rev. Lett.*, **45**, 566 (1980).
 22. Reynolds, P. J., Ceperley, D. M., Alder, B. J., and Lester, W. A., *J. Chem. Phys.*, **77**, 5593 (1982).
 23. Foulkes, W. M. C., Mitas, L., Needs, R. J., and Rajagopal, G., *Rev Mod Phys*, **73**, 33–84 (2001).
 24. Feynman, R. P., *Phys. Rev.*, **91**, 1291–1301 (1953).
 25. Griffin, A., *Excitations in a Bose-condensed Liquid*, Cambridge University Press, Cambridge U.K., 1993.
 26. Ceperley, D. M., *Rev. Mod. Phys.*, **67**, 279 (1995).
 27. Jacucci, G., "Path Integral Monte Carlo," in *Monte Carlo methods in Quantum Problems*, edited by M. H. Kalos, NATO ASI Series Vol 125, D. Reidel, Dordrecht, 1984.
 28. Creutz, M., Jacobs, L., and Rebbi, C., *Phys. Rev. Letts.*, **42**, 1390–1393 (1979).
 29. Barker, J. A., *J. Chem. Phys.*, **70**, 2914–18 (1979).
 30. Suzuki, M., Miyashita, S., Kuroda, A., and Kawabata, C., *Phys. Letts.*, **60A**, 478 (1977).
 31. Scalapino, D. J., and Sugar, R. L., *Phys. Rev. Letts.*, **46**, 519–521 (1981).
 32. Raedt, H. D., and Lagendijk, A., *Phys. Rev. Letts.*, **46**, 77–80 (1981).
 33. Ceperley, D. M., and Pollock, E., *Phys. Rev. B*, **39**, 2084 (1989).
 34. Ceperley, D. M., and Pollock, E. L., *Phys. Rev. Lett.*, **56**, 351–354 (1986).
 35. Sprik, M., Klein, M. L., and Chandler, D., *Phys. Rev. B*, **31**, 4234–4244 (1985).
 36. Gelb, L. D., *J. Chem. Phys.*, **118**, 7747–7750 (2003).
 37. Dewing, M., *J. Chem. Phys.*, **113**, 5123–5125 (2000).
 38. 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.
 39. Kron, A., *Polymer Science*, **7**, 1361 (1965).
 40. Webman, I., Lebowitz, J. L., and Kalos, M. H., *J. Physique A*, **41**, 579 (1980).
 41. Baroni, S., and Moroni, S., *Phys. Rev. Lett.*, **82**, 4745–4748 (1999).
 42. Evertz, H. G., Lana, G., and Marcu, M., *Phys. Rev. Letts.*, **70**, 875–879 (1993).
 43. Kawashima, N., and Gubernatis, J. E., *Phys. Rev. E*, **51**, 1547 (1995).
 44. Chandrasekharan, S., and Wiese, U.-J., *Phys. Rev. Lett.*, **83**, 3116 (1999).
 45. N. V. Prokof'ev, B. V. S., and Tupitsyn, I. S., *Phys. Lett. A*, **238**, 253–257 (1998).
 46. Blankenbecler, R., Scalapino, D. J., and Sugar, R. L., *Phys. Rev. D*, **24**, 2278–2286 (1981).
 47. Zhang, S., Carlson, J., and Gubernatis, J. E., *Phys. Rev. B*, **55**, 7464–7477 (1997).
 48. and, S. Z., *Phys. Rev. Letts.*, **90**, 136401 (2003).
 49. Car, R., and Parrinello, M., *Phys. Rev. Lett.*, **55**, 2471–2474 (1985).
 50. Ceperley, D., Dewing, M., and Pierleoni, C., "The Coupled Electronic-Ionic Monte Carlo Simulation Method," in *Bridging Time Scales: Molecular Simulations for the Next Decade*, edited by P. Nielaba, M. Mareschal, and G. Ciccotti, Topics in Condensed Matter Physics, Springer-Verlag, Berlin, 2002.
 51. Ceperley, D. M., and Dewing, M., *J. Chem. Phys.*, **110**, 9812 (1999).
 52. Kennedy, A. D., and Kuti, J., *Phys. Rev. Lett.*, **54**, 2473–76 (1985).
 53. Krajci, M., and Hafner, J., *Phys. Rev. Lett.*, **54**, 5100–03 (1995).
 54. Ball, R. C., Fink, T. M. A., and Bowler, N. E., *preprint server* (2003).
 55. Lee, M., Schmidt, K. E., Kalos, M. H., and Chester, G. V., *Phys. Rev. Lett.*, **46**, 728–731 (1981).
 56. Caffarel, M., and Ceperley, D. M., *J. Chem. Phys.*, **97**, 8415 (1992).
 57. Jarrell, M., and Gubernatis, J. E., *Physics Reports*, **269**, 133–195 (1996).
 58. Boninsegni, M., and Ceperley, D. M., *J. Low Temp. Phys.*, **104**, 339 (1996).
 59. Ming-Hui Chen, Q.-M. S., and Ibrahim, J. G., *Monte Carlo Methods in Bayesian Computation*, Springer-Verlag, 2000.