CONSTRAINED PATH MONTE CARLO FOR FERMIONS

SHIWEI ZHANG
Department of Physics and Department of Applied Science
College of William & Mary, Williamsburg, VA 23187, USA

1. Introduction

In these lectures we describe the constrained path Monte Carlo (CPMC) method for quantum many-fermion systems. We will focus on the ground-state CPMC algorithm[1] and its applications[2], but will also briefly describe work in progress on a finite-temperature extension[3].

In previous lectures we have seen both the auxiliary-field quantum Monte Carlo (AFQMC)[4, 5] and Green's function Monte Carlo (GFMC) [6] methods. We will see here how the ground-state fermion CPMC algorithm combines the concepts of Hubbard-Stratonovich transformation and Slater determinants with branching random walks and importance sampling. We will demonstrate the origin of the fermion sign problem [7, 8] in this context and then discuss an approximate solution, the constrained path approximation. CPMC is free of any sign decay and its computing time scales algebraically with system size. The Hubbard model will be used as an example to illustrate the basic idea and actual implementation of the algorithm, although the formalism of the algorithm is more general.

Having introduced the Hubbard model and discussed the ground-state CPMC method, we will then describe in some detail an application of CPMC. The Hubbard model has been the subject of intense theoretical effort for the past decade, to study its ground-state properties and understand whether it contains the relevant electron correlations to describe high-Tc superconductivity[9, 10]. CPMC enabled, for the first time, ground-state calculations on large system sizes. On the other hand, the task of characterizing the electron pairing correlations is an extremely difficult one and requires great accuracy. This application illustrates both the promise and the challenge facing our current fermion methods.

In the last part we will briefly discuss some preliminary results on a finite-temperature ($T > 0\,\mathrm{K}$) CPMC method. The goal is to have an algorithm which enables calculations at finite temperatures free of any sign decay, while maintaining much of the grand-canonical formalism of the standard Blankenbecler, Scalapino, and Sugar (BSS) algorithm[4].

Why do we need another approximate fermion method? After all, we have fixed-node[11, 12] Green's function Monte Carlo (GFMC) (including diffusion Monte Carlo (DMC)[11, 13]) and projector quantum Monte Carlo[5, 14], for ground states, and path integral Monte Carlo[15, 16] and auxiliary-field quantum Monte Carlo[4] for finite-temperature properties. The short answer is that there are *many* applications that demand it. In fact the ground-state CPMC algorithm has already made possible various calculations and is seeing an increasing number of applications in different areas. We will also give a long answer to this question by making a rough table of the existing algorithms, as shown on the next page. We *loosely* categorize these fermion algorithms into **continuum** and **lattice** methods. The table is a synopsis to capture some basic ideas, which will provide a very brief review of some of the materials we have learned, and also help set the reference frame for the algorithms we will discuss. (Readers unfamiliar with AFQMC may wish to first read Section 2.1.)

While cross-fertilization certainly has occurred, the division line between the two classes of methods remains quite visible. Each has distinct strengths, which have to a large extent maintained their separation in application areas. Although each has seen a great deal of success, limitations do exist in both, as we see from the table. In this summer institute, we have seen considerably more "mixing" of the two, and we can hope that this will stimulate even more such efforts.

2. Ground-state ($T = 0\,\mathrm{K}$) CPMC algorithm

The ground-state constrained path Monte Carlo (CPMC) algorithm filled in the void (indicated by the first set of question marks in Table 1.) of a ground-state auxiliary-field algorithm without exponential scaling. CPMC is free of any "sign decay", which is the signature of the fermion sign problem. The required computation time therefore scales algebraically with system size. It is approximate, with results dependent on the trial wave function that is used as constraint. Similar to the fixed-node approximation [11, 12] in configuration space, the process of CPMC solves the Schrodinger equation under some boundary condition defined by the trial wave function. But, differently from fixed-node, this boundary condition is not limited to real space, and can be a more "global" one, as we will discuss in detail below. CPMC also retains (at least to a large degree) an important advantage of AFQMC compared to continuum methods, namely the ability to compute easily many expectation values, including off-diagonal ones such

TABLE 1. **Synopsis of standard fermion quantum Monte Carlo (QMC) algorithms.**

	Continuum	**Lattice**
Applications	electronic structure, quantum chemistry, ^{3}He, few-body nuclear physics	correlated electron models, nuclear shell model, quantum field theory
Basis	configuration space: $R = \{\vec{r}_1, \vec{r}_2, \ldots, \vec{r}_N\}$, where $\vec{r}_i$ is position of the i-th of N fermions	Slater determinant space: $\lvert\Phi\rangle = \{\phi_1, \phi_2, \ldots, \phi_N\}$, where each ϕ_i is a *complete* single-particle orbital.
	GROUND-STATE:	
Algorithm	Green's Function MC (GFMC)[6, 13]	Projector QMC (PQMC)[5,14]
Description $P \equiv e^{-\Delta\tau H}$ $(\Delta\tau > 0)$ $\lvert\Psi^{(0)}\rangle$: known initial w.f.	$\langle\Psi^{(0)}\lvert PP \ldots P \ldots \to$ ground state. 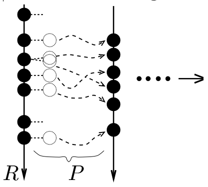 Wave function amplitude is given by distribution of a collection of R's. Each application of P is realized by random walk for each R: branching + diffusion.	Compute $\dfrac{\langle\Psi^{(0)}\lvert P \ldots O \ldots P\,\lvert\Psi^{(0)}\rangle}{\langle\Psi^{(0)}\lvert P \ldots P\,\lvert\Psi^{(0)}\rangle}$ 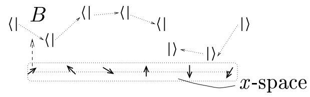 P can be written as $\int dx\, p(x) B(x)$. $B(x)$ is *1-electron* operator and depends on auxiliary-field x. MC samples x ("spins") to evaluate integrals. $\langle\,\rvert$ and $\lvert\,\rangle$ denote Slater determinants. $\langle\,\rvert B \to \langle\,\rvert$; $\langle\,\rvert\cdot\lvert\,\rangle \to$ number.
Sign problem	Caused by R's that represent + and − amplitudes of the wave function collapsing to the *same* asymptotic distribution — a symmetric one (bosonic).	Caused by cancellations: numbers of paths that lead to − and + contributions become increasingly *equal* as path grows longer (lower temperature).
	fixed-node approximation[11,12] available to achieve algebraic scaling	**???**
Observables	*Hard* to compute (R's are orthogonal).	*Easy*
	FINITE-TEMPERATURE:	
Algorithm	Path-Integral MC (PIMC)[15]	(Grand-Canonical) QMC[4]
Description	Feynman path integrals allow mapping to a classical problem of a ring-polymer in R-space, whose length is $\propto 1/T$. MC is used to sample the "classical" partition function (to compute observables).	Similar to above. However, $\lvert\Psi^{(0)}\rangle$ is implicit: N is *not* fixed; trace over all possible values can be evaluated *analytically* which yields expression involving only x. Integral over all x is again done by MC. (Length of path $\propto 1/T$.)
Sign problem	*restricted path-integral appr.* [16]	**???**

as electron pairing correlations. This is an important point since in many applications, particularly those in the study of strongly correlated models, it is crucial to have information on properties beyond the total energy.

The ground-state CPMC algorithm has two main components: The *first* is to formulate the projection of the ground state as open-ended random walks with importance sampling, as in GFMC, although the random walks take place in a space of Slater determinants. The *second* is to constrain the paths of the random walks so that any Slater determinant generated maintains a positive overlap with a known trial wave function $|\psi_T\rangle$. Below we describe them in sections 2.2 and 2.3 respectively. The first is an alternative way to do projector QMC which in many cases can be more efficient than the latter, because of the use of importance sampling. Like projector QMC, it is exact, but suffers from the sign problem. The second component is the constrained path approximation, which eliminates the sign decay, but introduces a systematic error in the algorithm. These two components are independent of each other, and can be used separately. Their combination is what we are calling the ground-state CPMC algorithm.

In describing the CPMC algorithm, we will use the Hubbard model and often use a trivial system in our illustration. We do so to make the description more pedagogical. However, it is important to think about how the method generalizes. This part of the notes should be used to complement Ref [1], in which more formal and rigorous discussions can be found. In addition, details of the algorithm that we will not be able to fully cover here can be found in those papers.

2.1. BACKGROUND

Here we introduce some background materials and illustrate them with a simple example. Some of these have been discussed by A. Muramatsu in his lecture[17].

2.1.1. *The Hubbard Hamiltonian*

The one-band Hubbard model is a simple paradigm of a system of interacting electrons. Its Hamiltonian is given by

$$H = K + V = -t \sum_{\langle ij\rangle\sigma} (c_{i\sigma}^\dagger c_{j\sigma} + c_{j\sigma}^\dagger c_{i\sigma}) + U \sum_i n_{i\uparrow} n_{i\downarrow}, \qquad (1)$$

where t is the hopping matrix element, and $c_{i\sigma}^\dagger$ and $c_{i\sigma}$ are electron creation and destruction operators, respectively, of spin σ on site i. The notation $\langle\ \rangle$ indicates near-neighbors. The on-site Coulomb repulsion is $U > 0$, and $n_{i\sigma} = c_{i\sigma}^\dagger c_{i\sigma}$ is the electron number operator. We will denote the number of lattice sites by N, and the linear dimension by L. The numbers of electrons with spin $\sigma = \uparrow, \downarrow$, which we denote by $N_\uparrow$ and $N_\downarrow$, will be fixed in each calculation. That is, we will work in the canonical ensemble. We will set $t = 1$ and impose periodic boundary conditions.

2.1.2. *The Hubbard-Stratonovich transformation*
We now introduce Hirsch's discrete Hubbard-Stratonovich (HS) transformation[18]. For any positive $\Delta\tau$, which we will choose to be small, the following identity holds:

$$e^{-\Delta\tau U n_{i\uparrow} n_{i\downarrow}} = e^{-\Delta\tau U (n_{i\uparrow}+n_{i\downarrow})/2} \sum_{x_i=\pm 1} p(x_i) e^{\gamma x_i (n_{i\uparrow}-n_{i\downarrow})}, \tag{2}$$

where γ is determined by $\cosh(\gamma) = \exp(\Delta\tau U/2)$. For reasons that will become clear in the next section, we have inserted a function $p(x_i)$ in place of the constant factor 1/2 and will interpret $p(x_i)$ as a (discrete) probability density function of x_i, with $x_i = \pm 1$.

The essence of the HS transformation is the conversion of an interacting system into many *non-interacting* systems living in fluctuating external auxiliary-fields, and the summation over all such auxiliary-field configurations recovers the many-body interactions. More specifically, the goal is to write the many-body operator $e^{-\Delta\tau H}$ in terms of one-electron operators. In Equation (2), the exponent on the left, which comes from the interaction term V, is quadratic in n, indicating the interaction of *two* electrons. The exponents on the right, on the other hand, are linear in n, indicating two non-interacting electrons in a (common) external field characterized by x_i.

The HS transformation above is one special case for the Hubbard interaction. Since our purpose is to illustrate the CPMC algorithm with an example, we will not dwell on more general forms of HS transformations[19] here, except to state that they exist for various other forms of interactions. In general, after an HS transformation is applied, the propagator $e^{-\Delta\tau H}$ can be written as

$$e^{-\Delta\tau H} = \int d\vec{x}\, p(\vec{x}) B(\vec{x}), \tag{3}$$

where $p(\vec{x})$ is a probability density function and $B(\vec{x})$ is a one-electron operator whose matrix elements are functions of the many-dimensional auxiliary-fields $\vec{x}$.

2.1.3. *A specific system to illustrate things*
We now look in some details into the ground-state of a simple 2×2 system of the Hubbard model to help explain the "language" we will use when describing CPMC later. We will consider $N_\uparrow = 2$ and $N_\downarrow = 1$. We label the four sites 1 thru 4 such that site 1 is near-neighbors with sites 2 and 3.

First let us examine the trivial case of free electrons, i.e., we set $U = 0$. We can write down the *1-electron* Hamiltonian matrix, which is of dimension 4×4:

$$H = \begin{pmatrix} 0 & -1 & -1 & 0 \\ -1 & 0 & 0 & -1 \\ -1 & 0 & 0 & -1 \\ 0 & -1 & -1 & 0 \end{pmatrix}.$$

The eigenstates of H can be obtained by direct diagonalization. With these eigenstates, we immediately obtain the ground-state wave function $|\psi_0\rangle$ of the 3-electron system from the Pauli exclusion principle:

$$|\psi_0\rangle = \begin{pmatrix} 0.5 & -0.58 \\ 0.5 & 0.40 \\ 0.5 & 0.40 \\ 0.5 & -0.58 \end{pmatrix} \otimes \begin{pmatrix} 0.5 \\ 0.5 \\ 0.5 \\ 0.5 \end{pmatrix},$$

where the first matrix contains two single-particle orbitals (two columns) for the two $\uparrow$ electrons and the second matrix contains one single-electron orbital for the one $\downarrow$ electron. Each single-electron orbital is an eigenvector of H. Note the second and third lowest single-electron states are degenerate, causing the 3-electron system to be open shell and $|\psi_0\rangle$ to be degenerate. We have simply picked one particular linear combination for the second $\uparrow$-electron in $|\psi_0\rangle$ above.

An object of the form of $|\psi_0\rangle$ is of course nothing more than a Slater determinant. For example, the amplitude of the configuration $|R\rangle = |\downarrow 0 \uparrow\uparrow\rangle$, i.e., two $\uparrow$ electrons on sites 3 and 4 and the one $\downarrow$ electron on site 1, is given by

$$\langle R|\psi_0\rangle = \det \begin{pmatrix} 0.5 & 0.40 \\ 0.5 & -0.58 \end{pmatrix} \cdot \det \begin{pmatrix} 0.5 \end{pmatrix}.$$

That is, more formally,

$$|R\rangle = \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

and $\langle R|\psi_0\rangle = \det(R^\dagger \cdot \Psi_0)$, where R and Ψ_0 denote the matrices corresponding to $|R\rangle$ and $|\psi_0\rangle$, respectively. In general, the overlap of two Slater determinants

$$\langle \phi|\phi'\rangle = \det(\Phi^\dagger \cdot \Phi') \tag{4}$$

is a number.

For this non-interacting system, an alternative (albeit indirect and indeed circular) way of obtaining $|\psi_0\rangle$ is by the power method. From the eigenvalues and eigenvectors of K, we can easily construct the matrix for $e^{-\Delta\tau H}$, which has the same structure as H above (i.e., a 4×4 matrix). Denote this matrix by B_K. With an arbitrarily chosen initial Slater determinant $|\psi^{(0)}\rangle$ (with non-zero overlap with $|\psi_0\rangle$), we can then repeatedly apply $e^{-\Delta\tau H}$, which means multiplying both the $4\times 2\uparrow$ matrix and the 4×1 $\downarrow$ matrix by B_K. As indicated in Table 1, this leads to $|\psi_0\rangle$ asymptotically, i.e.,

$$|\psi^{(n+1)}\rangle \propto e^{-\Delta\tau H}|\psi^{(n)}\rangle \tag{5}$$

gives $|\psi_0\rangle$ as $n\to\infty$.

Now suppose we turn on the interaction U. The first approach of directly diagonalizing H is the method of exact diagonalization, which scales exponentially. The power method of Eq. (5), on the other hand, can still apply if we can write $e^{-\Delta\tau H}$ in some one-electron form. The HS transformation does just that. Assuming $\Delta\tau$ is small and applying the Trotter break-up, we have

$$\begin{aligned} e^{-\Delta\tau H} = \sum_{\vec{x}} p(\vec{x}) \quad & \begin{pmatrix} e^{\gamma x_1} & 0 & 0 & 0 \\ 0 & e^{\gamma x_2} & 0 & 0 \\ 0 & 0 & e^{\gamma x_3} & 0 \\ 0 & 0 & 0 & e^{\gamma x_4} \end{pmatrix} \cdot B_K \\ \otimes & \begin{pmatrix} e^{-\gamma x_1} & 0 & 0 & 0 \\ 0 & e^{-\gamma x_2} & 0 & 0 \\ 0 & 0 & e^{-\gamma x_3} & 0 \\ 0 & 0 & 0 & e^{-\gamma x_4} \end{pmatrix} \cdot B_K, \end{aligned}$$

where $\vec{x} = \{x_1, x_2, x_3, x_4\}$. This is just Eq. (3). Note that $B(\vec{x})$ has an $\uparrow$ and a $\downarrow$ component, each of which is a 4×4 matrix. Applying each $B(\vec{x})$ to a Slater determinant means precisely the same as in the non-interacting case (with $B_K \otimes B_K$). In other words, $B(\vec{x})$ operating on any Slater determinant $|\phi\rangle$ simply involves matrix multiplications for the $\uparrow$ and $\downarrow$ components separately, leading to another Slater determinant $|\phi'\rangle$:

$$|\phi'\rangle = B(\vec{x})|\phi\rangle. \tag{6}$$

We now have all the basics to understand the projector QMC algorithm described in Table 1, and to move on to discuss CPMC.

2.2. RANDOM WALK IN SLATER DETERMINANT SPACE

The first component of the CPMC algorithm is the reformulation of the projection process as branching, open-ended random walks (a la GFMC) in Slater determinant space.

We start from the projection process in Eq. (5). Using (3), we write it

$$|\psi^{(n+1)}\rangle = \int d\vec{x}\, p(\vec{x}) B(\vec{x}) |\psi^{(n)}\rangle. \tag{7}$$

In the Monte Carlo realization of this iteration, we represent the wave function at each stage by a finite ensemble of Slater determinants, i.e.,

$$|\psi^{(n)}\rangle \propto \sum_k |\phi_k^{(n)}\rangle, \tag{8}$$

where k labels the Slater determinants and an overall normalization factor of the wave function has been omitted. Analogous to GFMC, the Slater determinants will be referred to as *random walkers*. We note that, contrary to the walkers in GFMC, the walkers here are non-orthogonal. Further, their characterization is not as simple as that of a point, because various combinations of single-particle orbitals can lead to the same Slater determinant.

The iteration in (7) is achieved stochastically by Monte Carlo (MC) sampling of $\vec{x}$:

$$|\phi_k^{(n+1)}\rangle \leftarrow \int d\vec{x}\, p(\vec{x}) B(\vec{x}) |\phi_k^{(n)}\rangle; \tag{9}$$

that is, for each random walker we choose an auxiliary-field configuration $\vec{x}$ from the probability density function $p(\vec{x})$ and propagate the walker to a new one via $|\phi_k^{(n+1)}\rangle = B(\vec{x})|\phi_k^{(n)}\rangle$. We repeat this procedure for *all* walkers in the population. These operations accomplish one step of the random walk. The new population represents $|\psi^{(n+1)}\rangle$ in the sense of (8), i.e., $|\psi^{(n+1)}\rangle \propto \sum_k |\phi_k^{(n+1)}\rangle$. These steps are iterated indefinitely. After an equilibration phase, all walkers thereon are MC samples of the ground-state wave function $|\psi_0\rangle$ and ground-state properties can be computed.

The algorithm we have described so far, while correct, is not efficient, because the sampling of $\vec{x}$ is completely random with no regard to the potential contribution to the ground-state wave function. In order to improve the efficiency of (7) and make it a practical algorithm, an importance sampling[6] scheme is required. To further motivate the need for importance sampling and develop some insights on how to proceed, we consider the so-called mixed estimate of the ground-state energy. As in GFMC, this is given by $E_0 \equiv \langle\psi_T|H|\psi_0\rangle/\langle\psi_T|\psi_0\rangle$. Estimating E_0 requires estimating the denominator by $\sum_\phi \langle\psi_T|\phi\rangle$, in which $|\phi\rangle$ denotes random walkers after equilibration. Hence, the term $\langle\psi_T|\phi\rangle$ plays the role of the Boltzmann distribution in statistical physics, while the denominator resembles the partition function. Our current way of sampling corresponds to sampling completely randomly configurations from the ensemble, which would clearly lead to large statistical fluctuations in the evaluation of the partition function. To improve the situation, it is natural to draw the analogy with the standard practice of Monte Carlo in statistical physics and try to sample $|\phi\rangle$ according to $\langle\psi_T|\phi\rangle$. This way, each walker would ideally contribute equally to the estimate, and E_0 would be given by the average of the quantity $\langle\psi_T|H|\phi\rangle/\langle\psi_T|\phi\rangle$ with respect to the sampled walkers. For a good $|\psi_T\rangle$, fluctuations in this quantity are small and we expect the statistical error to be much reduced.

With importance sampling, we iterate a modified equation with a modified wave function, without changing the underlying eigenvalue problem of (7). Specifically, for each Slater determinant $|\phi\rangle$, we define an importance function

$$O_T(\phi) \equiv \langle\psi_T|\phi\rangle, \tag{10}$$

which estimates its overlap with the ground-state wave function. We can then rewrite equation (7) as

$$|\tilde{\psi}^{(n+1)}\rangle = \int d\vec{x}\tilde{p}(\vec{x})B(\vec{x})|\tilde{\psi}^{(n)}\rangle, \tag{11}$$

where the modified "probability density function" is

$$\tilde{p}(\vec{x}) = \frac{O_T(\phi^{(n+1)})}{O_T(\phi^{(n)})}p(\vec{x}). \tag{12}$$

As expected, $\tilde{p}(\vec{x})$ is a function of both the current and future positions in Slater-determinant space. Further, it modifies $p(\vec{x})$ such that the probability is increased when $\vec{x}$ leads to a determinant with larger overlap and is decreased otherwise. It is trivially verified that equations (7) and (11) are identical.

In the random walk, the ensemble of walkers $\{\,|\phi_k^{(n)}\rangle\,\}$ now represents the modified wave function: $|\tilde{\psi}^{(n)}\rangle \propto \sum_k |\phi_k^{(n)}\rangle$. The true wave function is then given formally by

$$|\psi^{(n)}\rangle \propto \sum_k |\phi_k^{(n)}\rangle/O_T(\phi_k^{(n)}), \tag{13}$$

although in actual measurements it is $|\tilde{\psi}^{(n)}\rangle$ that is needed and division by O_T does not appear. The iterative relation for each walker is again given by (9), but with $p(\vec{x})$ replaced by $\tilde{p}(\vec{x})$. The latter is in general not a normalized probability density function, and we denote the normalization constant for walker k by $N(\phi_k^{(n)})$ and rewrite (9) as

$$|\phi_k^{(n+1)}\rangle \leftarrow N(\phi_k^{(n)}) \int d\vec{x}\, \frac{\tilde{p}(\vec{x})}{N(\phi_k^{(n)})}B(\vec{x})|\phi_k^{(n)}\rangle. \tag{14}$$

This iteration now forms the basis of the CPMC algorithm. As in the DMC method, it is convenient to associate a weight $w_k^{(n)}$ with each walker, which can be initialized to unity. With the introduction of the weight, $|\tilde{\psi}^{(n)}\rangle \propto \sum_k w_k^{(n)}|\phi_k^{(n)}\rangle$. For each walker $|\phi_k^{(n)}\rangle$, one step of the algorithm is then as follows:

1. sample a $\vec{x}$ from the probability density function $\tilde{p}(\vec{x})/N(\phi_k^{(n)})$,
2. propagate the walker by $B(\vec{x})$ to generate a new walker,
3. compute a weight $w_k^{(n+1)} = w_k^{(n)}N(\phi_k^{(n)})$ for the new walker.

To better see the effect of importance sampling, we observe that if $|\psi_T\rangle = |\psi_0\rangle$, the normalization $\int \tilde{p}(\vec{x})d\vec{x}$ is constant. Therefore the weights of walkers remain a constant and the random walk has no fluctuation. Furthermore, we refer again to the estimator for E_0. With importance sampling, the denominator becomes the sum of weights w, while the numerator is $\sum_\phi \langle\psi_T|H|\phi\rangle w_\phi/\langle\psi_T|\phi\rangle$, where again $|\phi\rangle$ denotes walkers after equilibration. As $|\psi_T\rangle$ approaches $|\psi_0\rangle$, all walkers contribute with equal weights to the estimator and the variance approaches zero.

We stress again that the importance sampling transformation does not alter the underlying iterative equation in (7). The results of the calculations are not affected; different choices of the importance function only affects the efficiency of the algorithm and thus the statistical error.

2.3. THE SIGN PROBLEM AND THE CONSTRAINED PATH APPROXIMATION

The second component of the CPMC algorithm is the constrained path approximation to eliminate the sign decay.

As a side-product, the reformulation of the projection process described in 2.2. provides a clear picture of the origin of the sign problem. Below we first elaborate on this picture, from which the constrained path approximation emerges naturally.

The sign problem occurs because of the fundamental symmetry between the fermion ground-state $|\psi_0\rangle$ and its negative $-|\psi_0\rangle$[20, 21]. For any ensemble of Slater determinants $\{|\phi\rangle\}$ which gives a Monte Carlo representation of the ground-state wave function, this symmetry implies that there exists another ensemble $\{-|\phi\rangle\}$ which is also a correct representation. In other words, the Slater determinant space can be divided into two degenerate halves (+ and $-$) whose bounding surface $\mathcal{N}$ is defined by $\langle\psi_0|\phi\rangle = 0$. This surface is in general *unknown*.

In some special cases, such as the particle-hole symmetric, half-filled one-band Hubbard model, symmetry prohibits any crossing of $\mathcal{N}$ in the random walk. The calculation is then free of the sign problem[18]. In more general cases, however, walkers do cross $\mathcal{N}$ in their propagation by $e^{-\Delta\tau H}$. The sign problem then invariably occurs.

In Fig. 1, we illustrate the space of Slater determinants by a one-dimensional (horizontal) line. The node $\mathcal{N}$ is the (red) dot in the middle. Imaginary time (or n) is in the vertical direction and increases as the arrow suggests. That is, as the walker moves in the horizontal line, we stretch out continuously "snapshots" of its position along the vertical direction. Now we follow an initial Slater determinant. With no loss of generality, we assume it has a positive overlap with $|\psi_0\rangle$. At time $n = 0$ it is indicated by the (green) dot on the right. As the random walk evolves, the walker can reach the node, which is the (red) vertical line. At the instant it lands on $\mathcal{N}$, the walker will make no further contribution to the representation of the ground state, since

$$\langle\psi_0|\phi\rangle = 0 \;\Rightarrow\; \langle\psi_0|e^{-\tau H}|\phi\rangle = 0 \text{ for any } \tau. \tag{15}$$

Paths that result from such a walker have equal probability of being in either half of the Slater determinant space. A few of these possible paths are shown by dashed lines. Computed analytically, they would cancel and not make any contribution in the ground-state wave function, as indicated by their symmetric placement with respect to the node line. But since the random walk has no knowledge of $\mathcal{N}$, these paths continue to be sampled (randomly) in the random walk and become Monte Carlo noise. Only paths that are completely confined to the right-hand side, as shown by the solid (green) line, will lead to contributions to the ground state, but the relative number of such confined paths decreases exponentially with n. Asymptotically in n, the Monte Carlo representation of the ground-state wave function consists of an *equal* mixture of the + and − walkers, regardless of where the random walks originated. The Monte Carlo signal is therefore lost. The decay of the signal-to-noise ratio, i.e. the decay of the average sign of $\langle\psi_T|\phi\rangle$, occurs at an exponential rate with imaginary time.

To eliminate the decay of the signal-to-noise ratio, we impose the constrained-path approximation. Fahy and Hamann first used[20] such a constraint in the framework of the standard AFQMC method. However, the non-local nature of such a constraint proved difficult to implement efficiently in the "path-integral-like" projector QMC scheme. Here, with the random walk formalism, the constraint only needs to be imposed one time-step at a time and is extremely simple to implement. It requires that each random walker at each step have a positive overlap with the trial wave function $|\psi_T\rangle$:

$$\langle\psi_T|\phi_k^{(n)}\rangle > 0. \tag{16}$$

This yields an approximate solution to the ground-state wave function, $|\psi_0^c\rangle = \sum_\phi |\phi\rangle$, in which all Slater determinants $|\phi\rangle$ satisfy (16).

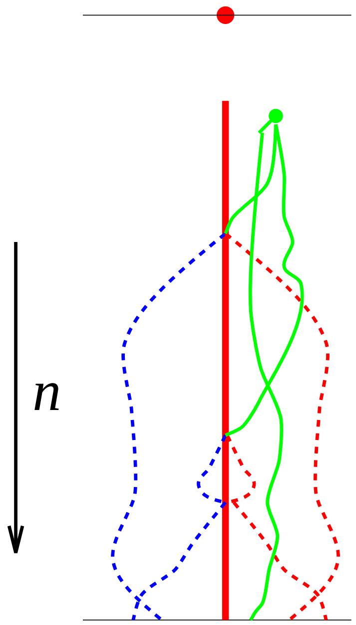

Figure 1. Schematic illustration of the sign problem. The top line represents Slater determinant space; the dot represents the "node" $\mathcal{N}$, where a determinant is orthogonal to the ground state $|\psi_0\rangle$. As the projection continues (increasing n), Slater determinants undergo random walks, tracing out "paths" as shown. When a walker reaches $\mathcal{N}$, its future paths will *collectively* cancel in their contribution to $|\psi_0\rangle$, indicated by the symmetric distribution of dashed paths about the nodal line. The Monte Carlo sampling, with no knowledge of this cancellation, continues to sample such paths randomly. The relative number of paths with real contributions (solid paths) decreases *exponentially* as n increases.

From (15), it follows that the constrained path approximation becomes *exact* for an exact trial wave function $|\psi_T\rangle = |\psi_0\rangle$. The overall normalization of walkers remains a constant on the average, since the loss of walkers at $\mathcal{N}$ is compensated by the branching of walkers elsewhere; that is, the eigenvalue problem with $\mathcal{N}$ as boundary has a stable solution.

To implement the constrained-path approximation in the random walk, we redefine the importance function by (10):

$$O_T(\phi) \equiv \max\{\langle\psi_T|\phi\rangle, 0\}. \tag{17}$$

This prevents walkers from crossing the trial nodal surface $\mathcal{N}$ and entering the "$-$" half-space as defined by $|\psi_T\rangle$. In the limit $\Delta\tau \to 0$, (17) ensures that the walker distribution vanishes smoothly at $\mathcal{N}$ and the constrained-path approximation is properly imposed. With a finite $\Delta\tau$, however, $\tilde{p}(\vec{x})$ has a discontinuity at $\mathcal{N}$ and the distribution does not vanish. We have found this effect to be very small for reasonably small imaginary-time steps $\Delta\tau$. Nonetheless, we correct for it by modifying $\tilde{p}(\vec{x})$ near $\mathcal{N}$ so that it approaches zero smoothly at $\mathcal{N}$.

2.4. ADDITIONAL ALGORITHMIC ISSUES

2.4.1. *Computing expectation values*

Given the ground-state wave function $|\psi_0^c\rangle$ from CPMC, where the superscript "c' indicates that the wave function is obtained under the constrained-path approximation, we can compute the ground-state energy with the mixed estimate[6]:

$$E_{\rm CPMC}^{\rm mixed} = \frac{\langle\psi_T|H|\psi_0^c\rangle}{\langle\psi_T|\psi_0^c\rangle}. \tag{18}$$

This can be implemented in a way similar to GFMC. However, contrary to fixed-node in GFMC, the ground-state energy computed from the mixed estimate is *not* an upper bound[22]. (The argument given in Ref. [1] on variational properties is incorrect.) The origin for the disappearance of the upper bound property in the mixed-estimate is the non-orthogonal nature of the Slater determinant space. In contrast with configuration space in fixed-node GFMC, where points on the node have no overlap with either $|\psi_T\rangle$ or the fixed-node solution, a point in the Slater determinant space which has no overlap with $|\psi_T\rangle$ is not necessarily orthogonal with $|\psi_0^c\rangle$. Eq. (18) is *not* equivalent to the estimator $\langle\psi_0^c|H|\psi_0^c\rangle/\langle\psi_0^c|\psi_0^c\rangle$, which is variational.

In the Hubbard model, the effect of this appears to be small[22]. In Table I, we show several cases we could find where the computed ground-state energies by CPMC fell below the exact results. The system is 4×4, with $5\uparrow$ and $5\downarrow$ electrons, which corresponds to a special case of a closed-shell system in the non-interacting limit. The free-electron wave function was used as the trial $|\psi_T\rangle$. As we see, the amount by which the mixed estimates fall below the exact results is no more than 0.1% across a wide range of interactions. For $U=8$, an "almost" free-electron wave function was used as $|\psi_T\rangle$ in Ref. [1], namely an unrestricted Hartree-Fock wave function generated with a small U. That $|\psi_T\rangle$ led to a ground-state energy *above* the true value[1].

Several possibilities exist to correct for the non-variational nature of the mixed estimate, at least in principle, and make it an upper bound[22]. For example, we could compute $\langle\psi_0^c|H|\psi_0^c\rangle/\langle\psi_0^c|\psi_0^c\rangle$ directly with a method similar to forward walking[23, 24] in GFMC. That is, we can create a separate walk for the left-hand wave function $\langle\psi_0^c|$, which propagates from $\langle\psi_T|$ with the constraint properly implemented, and then compute the overlap with populations in the regular (right-hand) walk. It would clearly be inefficient

TABLE 2. Examples showing the lack of upper bound property in the *mixed-estimate* of CPMC. The system is 4×4, with 5 $\uparrow$ and 5 $\downarrow$ electrons. Calculations were done for three values of the interaction strength U, 4, 8, and 20. The CPMC results are obtained from the mixed estimate; statistical errors are in the last digit and are indicated in parentheses. These are compared with results from exact diagonalization.

U	4	8	20
CPMC	-19.582(5)	-17.519(2)	-15.460(3)
exact	-19.580	-17.510	-15.452

to *randomly* sample the left-hand walks, whose overlap with the regular walks would fluctuate greatly. We have devised an "importance sampling" scheme that both imposes the constraint with $|\psi_T\rangle$ and attempts to incorporate the necessary correlations between the left-hand random walker $\langle\phi'|$ and its matching right-hand walker. We take an importance function $O_T^L(\phi') = |\langle\phi'|\phi_k^{(n)}\rangle| + \alpha\langle\phi'|\psi_T\rangle$, where α is a parameter and $|\phi_k^{(n)}\rangle$ is the walker with which $\langle\phi'|$ will be matched in the regular population representing the right-hand wave function. While this offered a significant improvement over naive sampling, we found it still prone to large fluctuations. With the variational correction term to the mixed estimate being very small, the statistical accuracy obtained by this scheme did not appear adequate to be of practical use for large systems. We note that, although this way of forward walking is somewhat similar to the back-propagation technique we discuss below, they are not equivalent. The latter tends to be much more accurate statistically. However, as we see below, it does not maintain the correct sense of direction in imposing the constraint and does not yield strictly $\langle\psi_0^c|$.

The back-propagation (BP) technique enables the computation of the expectation value of an operator $\mathcal{O}$ that does not commute with H. As we have mentioned, the essence of BP also comes from the forward walking technique in the GFMC method:

$$\langle\mathcal{O}\rangle_{\rm BP} = \lim_{\tau\to\infty} \frac{\langle\psi_T \exp(-\tau H_c)|\mathcal{O}|\psi_0^c\rangle}{\langle\psi_T \exp(-\tau H_c)|\psi_0^c\rangle}. \tag{19}$$

A subtle distinction, however, exists between back-propagation and forward walking. In back-propagation, $\langle\psi_T \exp(-\tau H_c)| = \langle\psi_T|\exp(-\tau H_c)$ is restricted to "constrained" paths, i.e., those paths that do not violate the constraint in the *original forward direction* $\exp(-\Delta\tau H_c)|\psi_0^c\rangle$. In fixed-node GFMC a path in configuration space has no sense of direction with respect to the nodal surface. In the CPMC method, however, there is a sense of direction: a set of determinants along the path of a random walk which does not violate the constraint at any step when going from right to left may indeed violate it any even number of times when going from left to right. Because of this sense of direction, expression (19) may not yield the true expectation value $\langle\psi_0^c|\mathcal{O}|\psi_0^c\rangle/\langle\psi_0^c|\psi_0^c\rangle$. However, since $|\psi_0^c\rangle$ is itself approximate, this issue is not crucial. $\langle\mathcal{O}\rangle_{\rm BP}$ does have the correct limiting behavior, remaining exact for an exact trial wave function.

It is relatively simple to implement the back-propagation scheme on top of a regular CPMC calculation. The key is that once a path has been sampled, the single-particle propagator can be easily constructed. Propagating a Slater determinant with it, either forward or backward, simply means matrix multiplications. We choose an iteration n and store the entire population $\{\,|\phi_k^{(n)}\rangle\,\}$. As the random walk proceeds, we keep track of the following two items for each new walker: (1) the sampled auxiliary-field variables that led to the new walker from its parent walker and (2) an integer that labels the parent. After an additional m iterations, we carry out the back-propagation: For each walker l in the $(n+m)^{\rm th}$ (current) population, we initiate a determinant $\langle\psi_T|$ and act on it with the corresponding propagators, but taken in reverse order. The m successive propagators are constructed from the stored items, with $\exp(-\Delta\tau K/2)$ inserted where necessary. The resulting determinants $\langle\bar{\phi}_l^{(m)}|$ are combined with its parent from iteration n to compute $\langle\mathcal{O}\rangle_{\rm BP}$, in a way similar to the mixed estimator (18). The weights are given correctly by $w_l^{(n+m)}$, due to importance sampling in the regular walk. Starting from another iteration n', this process can be repeated and the results accumulated.

We note again that, because the random walkers in CPMC are full Slater determinants, the algorithm lends itself well to the calculation of expectation values. The non-orthogonal nature of the determinant space ensures that overlaps between two walkers can be computed. This is to be contrasted with "continuum" methods in Table 1, where off-diagonal expectations are in general difficult to compute, because it is difficult to sample two group of random walkers whose overlap is well behaved statistically.

2.4.2. *population control and bias correction*

In CPMC walkers carry weights or, equivalently, they branch. The procedure we use to control the population is similar to that used in many GFMC calculations. First, a branching (or birth/death) scheme is applied, in which walkers with large weights are replicated and ones with small weights are eliminated by probabilities defined by the weights. There exist various ways to do this[25, 24, 26], with the guideline being that the process should not change the distribution statistically. In general, how this step is done only affects the efficiency of the algorithm, but does not introduce any bias.

Branching allows the total number of walkers to fluctuate and possibly become too large or too small. Thus as a second step, the population size is adjusted, if necessary, by rescaling the weights with an overall factor. Re-adjusting the population size, i.e., changing the *overall* normalization of the population, does introduce a bias[25, 24]. We correct for this bias by carrying the m most recent overall rescaling factors and including them in the estimators when computing expectation values. In the calculation we keep a stack which stores the m latest factors, $f^{(j)}$ $(j = 1, m)$. Suppose that at the current step the total number of walkers exceeds the pre-set upper bound. We modify the weight of each walker by a constant factor $f < 1$ which reduces the population size to near the expected number. We then replace the oldest $f^{(j)}$ in the stack by f. Whenever we compute expectation values, we multiply the weight of each walker by $1/\prod_j f^{(j)}$. In our calculations on the Hubbard model, m is typically between 5 and 10. As we include more such factors, i.e., increasing m, the bias is reduced, but the statistical error increases. On the other hand, as we reduce m, the statistical error becomes smaller, but the bias increases. The choice of m is thus a compromise between these two. Another common approach to eliminate or reduce bias, which is in general less efficient, is to do several calculations with different (average) population sizes and extrapolate to the infinite population limit.

Clearly, the schemes we have described to correct for bias have some arbitrariness and further improvements are possible. But this should not be a major factor in the calculation. In cases where the bias is *substantially* larger than can be handled by correction schemes in this spirit, it is likely more productive to attempt to improve the importance sampling and the algorithm, rather than details of the bias correction scheme.

2.4.3. *re-orthogonalization a Slater determinant*

Repeated multiplications of $B(\vec{x})$ to a Slater determinant, either in the regular walk or in BP, lead to a numerical instability, such that round-off errors dominate and $|\phi_k^{(n)}\rangle$ represents an unfaithful propagation of $|\phi_k^{(0)}\rangle$. This instability is well-known in the AFQMC method and is controlled[27, 28] by a numerical stabilization technique that requires the periodic re-orthonormalization of the single particle orbitals in $|\phi_k^{(n)}\rangle$.

We use the *modified* Gram-Schmidt procedure[28, 27] to stabilize the Slater determinants. The procedure is further simplified because of the random walk formalism in CPMC. For each walker $|\phi\rangle$, we factor the matrix Φ as $\Phi = QR$, where Q is a matrix whose columns are a set of orthonormal vectors representing the re-orthogonalized single-particle orbitals. R is a triangular matrix. Note that Q contains all the information about the walker $|\phi\rangle$. R, on the other hand, only contributes to the overlap of $|\phi\rangle$. Since the overlap is already represented by the walker weight, R does not need to be carried along and can be discarded.

It is interesting to note that this instability originates from the tendency for each walker (Slater determinant) to collapse to a bosonic ground state. In GFMC this actually happens and is the first cause for the sign problem. Here the instability can be eliminated by the re-orthonormalization procedure because a walker is explicitly a Slater determinant, not a single point. The procedure to stabilize the Slater determinants is like analytically canceling walkers in GFMC[21], which requires a large density of walkers and does not scale well with system size. CPMC, like AFQMC, automatically imposes the antisymmetry. This would suggest that the sign problem is reduced in this formalism compared to GFMC. One trivial observation which is consistent with this is the case of a non-interacting system, where CPMC does not suffer from the sign problem and has zero-variance, while GFMC does have a sign problem and still requires the fixed-node approximation.

2.5. ILLUSTRATIVE BENCHMARK RESULTS

Here we show a few benchmark results to illustrate characteristics of the CPMC algorithm. These results are for the two-dimensional Hubbard model. Ref. [1] contains more results and a more extensive account of the benchmark calculations that have been performed.

In Figure 2 we show the computed E_0/N at $U = 4$ for various lattice sizes and electron filling $\langle n \rangle = (N_\uparrow + N_\downarrow)/N$. Also shown is available AFQMC data [29]. We see that the CPMC results are in good agreement with the AFQMC data. For example, in the case of 8×8 lattice with $25 \uparrow$ and $25 \downarrow$ electrons, the CPMC result lies approximately $0.4\% \pm 0.2\%$ above the AFQMC result. The behaviors of the statistical errors in these two algorithms are of course completely different. In AFQMC the sign problem causes the variance to increase *exponentially* with N and the projection time $n\Delta\tau$, while the CPMC method is free of any sign decay and exhibits power law scaling with N. For large systems (10×10 and beyond), the CPMC error bars are 30-50 times smaller than those of AFQMC, as indicated for the 12×12 system.

A more stringent test of the algorithm is the calculation of correlation functions. In Table 3, we show computed expectation values for the system of 4×4 with $7 \uparrow$ $7 \downarrow$ electrons, which is an open-shell case with a severe sign problem. The one-body density matrix is the expectation value of the Green's function element: $\rho(\mathbf{l}) = \langle c_0^\dagger c_\mathbf{l} \rangle$, where the vector $\mathbf{l} = (l_x, l_y)$ denotes

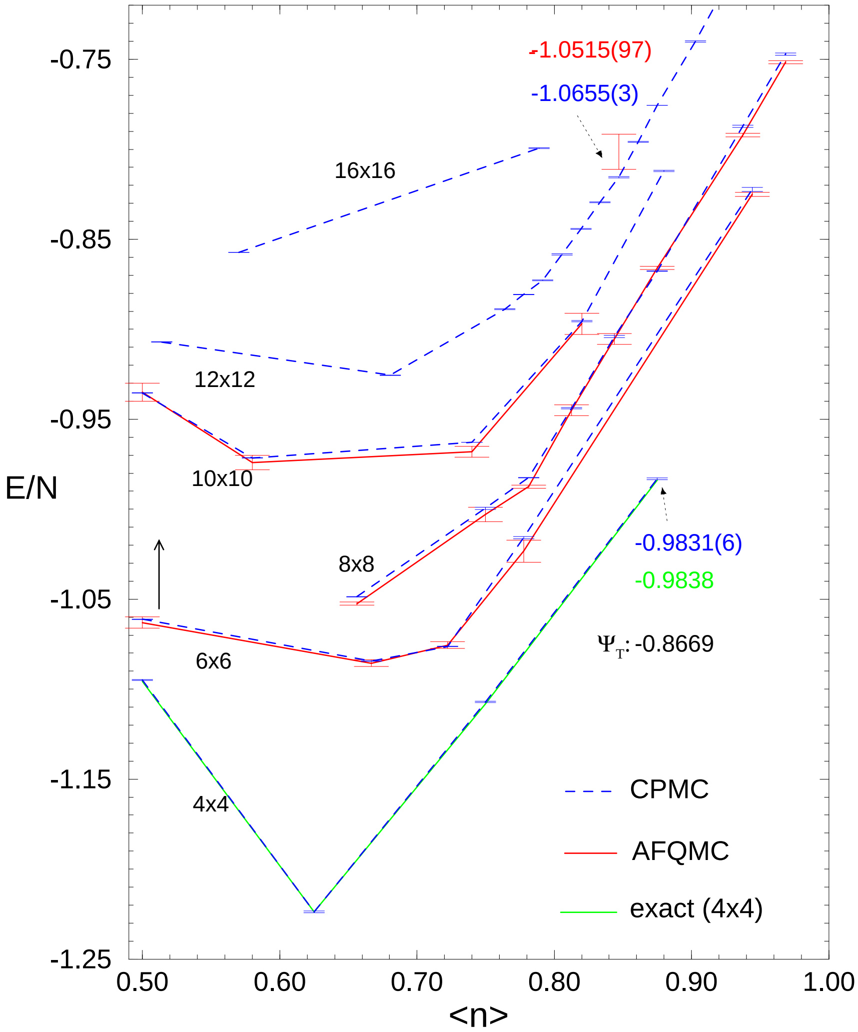

Figure 2. Computed energies per site vs. electron fillings from CPMC, compared with available AFQMC and exact diagonalization data. The lines are to aid the eye. Curves for $L > 4$ are shifted up. Half-filling is $\langle n \rangle = 1$. For 4×4, the solid curve is from exact diagonalization[30]. At $\langle n \rangle = 0.875$ ($7 \uparrow \ 7 \downarrow$), the CPMC (with error bar in the last digit) and exact numbers are shown, together with the variational energy from the trial wave function $|\psi_T\rangle$ in CPMC. Numbers are also shown for the AFQMC (lone point) and CPMC values for 12×12 with $61 \uparrow \ 61 \downarrow$. Note that the CPMC error bar is smaller by a factor of 30 in this case.

the position of a site (vs. i, which we have used so far and which is the site label). The spin density structure factor is

$$S(k_x, k_y) = S(\mathbf{k}) = 1/N \sum_{\mathbf{l}} \exp(i\mathbf{k} \cdot \mathbf{l}) \langle \mathbf{s_0 s_l} \rangle, \tag{20}$$

where $\mathbf{s_l} = n_{\mathbf{l}\uparrow} - n_{\mathbf{l}\downarrow}$ is the spin at site $\mathbf{l}$. We show results from CPMC simulations with two different trial wave functions. Both are unrestricted Hartree-Fock wave functions, but $|\psi_{T1}\rangle$ was obtained with a U of 0.1, while $|\psi_{T2}\rangle$ with $U = 4$. The calculation with $|\psi_{T1}\rangle$ has much less fluctuation, even though $|\psi_{T2}\rangle$ has a lower variational energy[1]. We see that the two trial wave functions yield very different variational estimates, but their CPMC results are consistent and in good agreement with exact results. For example, in the free-electron-like $|\psi_{T1}\rangle$, the momentum distribution is a step function, and the $\mathbf{k} = (1, 0)$ state is completely occupied ($n_k = 1$), but even with this trial wave function the CPMC method still gives the correct occupation of 0.92(1). The spin structure factor $S(\pi, \pi)$ is the only case where CPMC does not produce the exact result. However, even with the free-electron-like $|\psi_{T1}\rangle$, which severely underestimates the anti-ferromagnetic correlation, CPMC correctly recovers the physics and predicts the presence of a strong peak of $S(\mathbf{k})$ at (π, π).

The fact that CPMC shows consistency with different choices of $|\psi_T\rangle$ is reassuring. Such consistency checks are crucial in real applications where benchmark data is scarce. The fact that CPMC appears to be able to accurately calculate correlation functions in these systems is significant. Correlation functions are often much more difficult to compute than the energy. It is here that the difference in scaling behaviors between CPMC and AFQMC becomes much more magnified. The AFQMC algorithm tends to break down for system sizes much smaller than indicated in Fig. 2. The fixed-node GFMC also encounters difficulties in computing certain correlation functions because forward walking does not yield proper overlaps, as we discussed earlier.

In Table 4, we show a comparison of the computed ground-state energies from CPMC and from lattice fixed-node GFMC[31]. We see that, in addition to the ability to compute correlation functions conveniently, CPMC does better even for the energy. (The particular system is an easy case for CPMC and the results may not be typical of the Hubbard model.) At small U, CPMC is expected to perform better because the linear combination of mean-field states provides an effective description of the system. On the other hand, at very large U, we expect that, at least in the current way of Hubbard-Stratonovich transformation, the statistical error in CPMC will

[1]This trend appears to be rather general: free-electron-like wave functions tend to be better importance functions than unrestricted Hartree-Fock wave functions. This likely reflects to what extent translational invariance is maintained.

TABLE 3. Computed expectation values and correlation functions from CPMC for a 4×4 lattice with $7 \uparrow$ $7 \downarrow$ electrons and $U = 4$, compared with exact results. Results are shown for two different trial wave functions $|\psi_{T1}\rangle$ and $|\psi_{T2}\rangle$. Both are unrestricted Hartree-Fock wave functions, but $|\psi_{T1}\rangle$ was obtained with a U of 0.1, while $|\psi_{T2}\rangle$ with $U = 4$. Exact diagonalization results are from Ref. [30]; numbers in parentheses indicate either the range of values due to the ground-state degeneracy or uncertainties in extracting numbers from a graph. Statistical errors on the last digit of the CPMC results are in parentheses. E_k is the kinetic energy, $\rho(l_x, l_y)$ the one-body density matrix, S the spin density structure factor, and n_k the momentum distribution.

		E_k	$\rho(2,2)$	$S(\pi,\pi)$	$n_k(\pi/2,0)$
variational	$\|\psi_{T1}\rangle$	-24.0	1.654	0.625	1.0
	$\|\psi_{T2}\rangle$	-21.88	-0.0602	4.39	0.941
CPMC	$\|\psi_{T1}\rangle$	-21.44(2)	-0.051(1)	2.90(1)	0.92(1)
	$\|\psi_{T2}\rangle$	-21.39(8)	-0.049(1)	2.92(2)	0.92(1)
exact		-21.39(1)	-0.051	2.16(2)	0.93(1)

be larger than GFMC. More complete studies are necessary to characterize where the "cross-over" occurs.

TABLE 4. Comparison of CPMC with lattice fixed-node (FN) GFMC for 4×4, with $5 \uparrow$ and $5 \downarrow$ electrons. Shown are computed ground-state energies per site. The CPMC result was obtained with a free-electron $|\psi_T\rangle$, as was the first number for fixed-node GFMC. The FN GFMC result "Gutzwiller" was obtained with a better $|\psi_T\rangle$ which included a Gutzwiller factor. Statistical errors are in the last digit and are indicated in parentheses. Fixed-node results are from Ref. [32].

FN GFMC	FN GFMC (Gutzwiller)	CPMC	exact
-1.2186(4)	-1.2201(4)	-1.2239(3)	-1.2238

3. Application: Is the Hubbard model the right one for high-T_c?

Here we describe an application of the CPMC algorithm to study the ground state of the one-band Hubbard Hamiltonian of Eq. (1). In particular, we will focus on the electron pairing correlation in the $d_{x^2-y^2}$-wave channel. Most of the data in this section is from Ref. [2].

The Hubbard model is a simple microscopic description of a many-electron system which includes both band-structure and electron interactions. Since the discovery of high-T_c superconductivity, it has been widely believed that this model contains essential features relevant to the properties of the CuO_2 planes in the cuprate superconductors. Viewed as the basis for further progress on understanding the mechanism for high-T_c, the model has been the subject of intense theoretical activity[9]. Various calculations predict[10] that the doped model exhibits an attractive interaction between pairs; the extended s- and $d_{x^2-y^2}$- symmetries of this attraction are consistent with the likely symmetries of the experimentally measured superconducting gap [10]. Yet unobserved, however, is convincing evidence that the attractive interaction leads to a ground state with off-diagonal long-range order[9, 10]. In other words, we do not have the complete answer to whether the ground state of this model has a condensate of electron pairs and, if yes, what their nature is.

Numerical approaches have been a promising tool for answering some of the questions on the Hubbard model. The AFQMC method in particular has been applied extensively and has seen a great deal of success[33, 9]. However, due to the sign problem, it has typically been limited to relatively small system sizes, high temperatures, and selected electron fillings. CPMC largely eliminates these difficulties, although we must bear in mind the approximate nature of the constrained path approximation.

The pairing correlation function we computed is defined as

$$D_\alpha(\mathbf{l}) = \langle \Delta_\alpha^\dagger(\mathbf{l})\Delta_\alpha(\mathbf{0})\rangle, \tag{21}$$

where α indicates the nature of pairing. The pair-field operator at site $\mathbf{l}$ is $\Delta_\alpha(\mathbf{l}) = \sum_\delta f_\alpha(\delta)[c_{\mathbf{l}\uparrow}c_{\mathbf{l}+\delta\downarrow} - c_{\mathbf{l}\downarrow}c_{\mathbf{l}+\delta\uparrow}]$, where δ is $(\pm 1, 0)$ and $(0, \pm 1)$. For the extended s-wave ($s^\star$-), $f_{s^\star}(\delta) = 1$. For d-wave, $f_d(\delta)$ is 1 when $\delta = (\pm 1, 0)$ and -1 otherwise.

Because the algorithm has an uncontrolled approximation and because the electron pairing correlations are extremely difficult quantities to characterize, we have carried out benchmark calculations[2, 1] and many self-consistency checks. Due to the limited availability of data for benchmark, the latter becomes even more important. In Fig. 3, we show another such example, for a fairly difficult case of 8×8 with $27\uparrow$ $27\downarrow$ electrons.

Fig. 4 shows the $d_{x^2-y^2}$-wave pairing correlations for 12×12 and 16×16 lattices at $U = 2$ and 4. The electron filling is 0.85, which corresponds to closed-shell cases, with $N_\uparrow = N_\downarrow = 61$ for 12×12 and $N_\uparrow = N_\downarrow = 109$ for 16×16. The free-electron wave function was used as $|\Psi_T\rangle$ in CPMC. Also shown in each graph is the $D_d(\mathbf{l})$ for the non-interacting system. We see that the computed $d_{x^2-y^2}$-wave pairing correlations at both values of U are bounded from above by the $U = 0$ result. Further, consistent with this

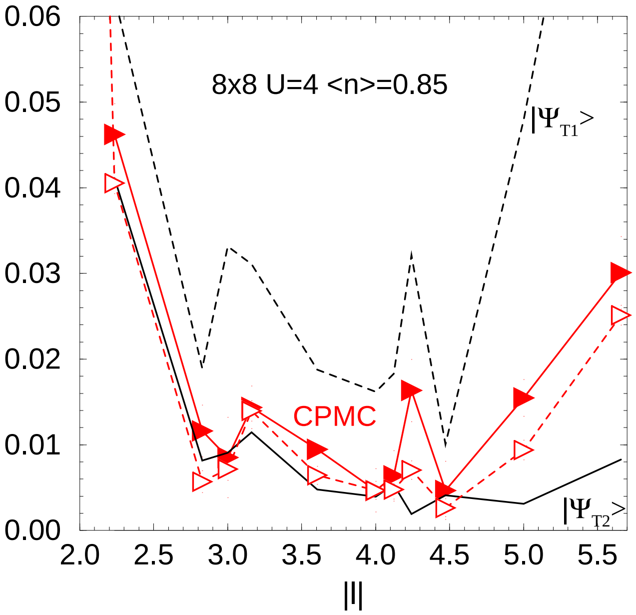

Figure 3. Sensitivity of CPMC results to the choice of trial wave function $|\Psi_T\rangle$. Shown is the near-neighbor d-wave pairing correlation as a function of pair separation $|\mathbf{l}|$. The short distance portion is omitted. This system corresponds to a difficult open-shell case. The two trial wave functions are both unrestricted Hartree-Fock wave functions. The first, $|\psi_{T1}\rangle$, was obtained with a very small interaction strength U, resembling a free-electron wave function, while $|\psi_{T2}\rangle$ was generated with the true U of 4. The CPMC pairng correlations, with $|\psi_{T1}\rangle$ and $|\psi_{T2}\rangle$ as trial wave functions respectively, are shown (curves with symbols (red)), with the dashed one from $|\psi_{T1}\rangle$. The mean-field results that $|\psi_{T1}\rangle$ and $|\psi_{T2}\rangle$ yield are also shown (no symbols). We see that the two sets of CPMC results, although not in exact agreement, over-correct the mean-field results and show consistency. The statistical error bars on the CPMC results are indicated by small dots, the maximum being several times the symbol size.

observation, we see that at large distances $D_d(\mathbf{l})$ decreases as U increases from 2 to 4. In the 16×16 case, the long-range correlations vanish as the interaction strength is increased to 4. The statistics is enough to discern in this case the irregular oscillations of $D_d(\mathbf{l})$ around zero.

Calculation was also done for the 12×12 system at $U = 8$ with the same (free-electron) trial wave function. The error bars become substantially larger. $D_d(\mathbf{l})$ oscillates around zero and shows no indication of enhancement at large distances[2].

We also note that the behavior of $D_d(\mathbf{l})$ in Fig. 4 at short distances is just the opposite of that at large distances, i.e., $D_d(\mathbf{l})$ increases as U is increased. Due to the large discrepancy in magnitude between short and long distances, the integrated value of $D_d(\mathbf{l})$ would therefore not give the correct trend even at system sizes as large as these.

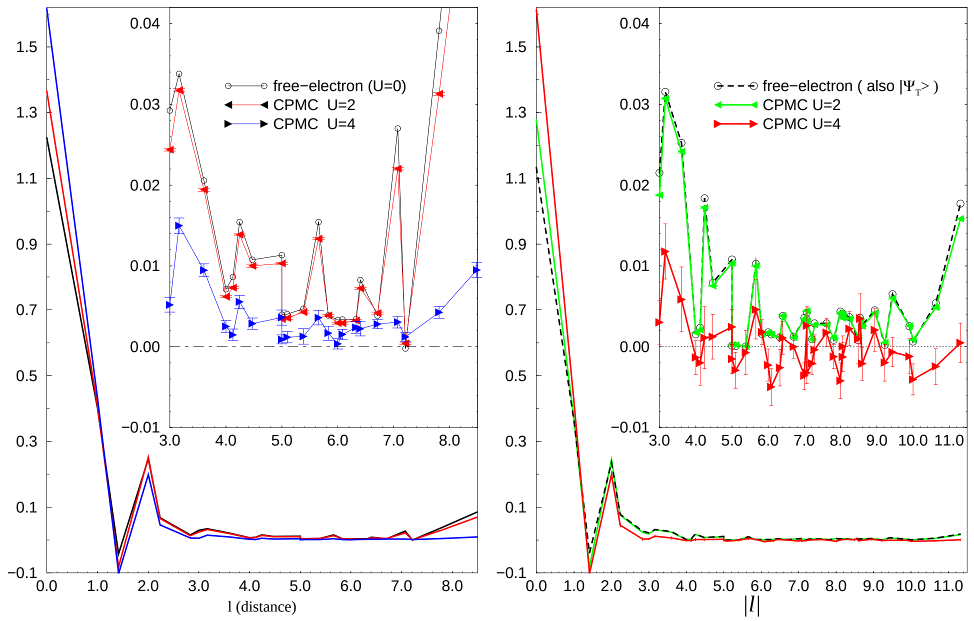

Figure 4. $d_{x^2-y^2}$-wave pairing correlation $D_d(\mathbf{l})$ vs. distance $l = |\mathbf{l}|$. LEFT: $D_d(\mathbf{l})$ for a 0.85-filled 12×12 lattice at interaction $U = 2$ and 4. The inset shows the the portion with $l \geq 3$. Shown together is the ground-state pairing correlation for the non-interacting system, whose wave function is also the trial wave function used in the CPMC calculations for both values of U. At large distances (inset) $D_d(\mathbf{l})$ systematically decays as U is increased. Also note that at the irrelevant (short) distances the behavior is the opposite. RIGHT: Same plot, for a 16×16 system. Same trend as that of 12×12 continues, but now the long-range correlation vanishes at $U = 4$.

In summary, the near-neighbor $d_{x^2-y^2}$-wave electron pairing correlation seems to disappear as the interaction strength increases and no indication of enhancement is found over the non-interacting system. Indeed, the behavior of the pairing correlation appears similar across a fairly broad range of electron fillings, including the half-filled case, which exhibits pairing correlations of similar magnitude[2].

We emphasize again that, due to the approximate nature of the CPMC method, the possibility can not be completely ruled out that the results were biased such that the pairing correlation was suppressed. However, the many benchmark calculations, and particularly the variety of self-consistency checks, indicate strongly that the results are accurate and robust.

4. Finite-temperature CPMC algorithm

Extension of the CPMC algorithm to finite temperatures will allow us to remove the second set of question marks in Table 1.

While the standard method of Blankenbecler, Scalapino, and Sugar (BSS)[4, 27] has been widely applied, the exponential scaling has hindered its ability to effectively study true phase transitions. The difficulty with developing a finite-temperature counterpart of CPMC lies in the implicit nature of the path-integral picture in the BSS formalism. As discussed in Table 1, paths in the grand-canonical formulation, contrary to ground-state projector QMC, do not originate or end at a single explicit point in Slater determinant space. Instead they involve many points, due to the trace over N (the number of particles, not the number of sites). Indeed the points do not even have the same "dimension"[4]. A naive application of the current way of imposing the constrained path approximation, which would require separate constraining wave functions and conditions for paths with different end points, would therefore not be practical.

We have developed a mechanism to incorporate the constraint into the analytic evaluation of the trace[3], which we are currently testing. In Table 5, we show preliminary results on total energies per site, for a simple case of the 4×4 system. The chemical potential is adjusted so that the average filling fraction corresponds to $5\uparrow$ and $5\downarrow$ electrons. As high temperatures, the new algorithm reproduces the results from AFQMC, while at a low temperature of $T = 0.0625$ ($\beta = 16$), it is in good agreement with the ground-state result from exact diagonalization, as well as that from the $T = 0\,\mathrm{K}$ CPMC algorithm (cf Table 1).

TABLE 5. Preliminary data from the new finite-temperature algorithm. The computed total energies per site, as a function of temperature, are compared with results from AFQMC[34] or exact diagonalization. The system is 4×4 at $U = 4$. The chemical potential is adjusted so that the average filling fraction corresponds to $5\uparrow$ and $5\downarrow$ electrons. Statistical errors are in the last digit and are indicated in parentheses. The * indicates the exact *ground-state* energy from exact diagonalization.

T	1	0.25	0.0625
CPMC	-0.8459(2)	-1.1947(6)	-1.2237(6)
AFQMC	-0.8457(1)	-1.1947(2)	-1.2238*

Acknowledgements

This work was supported by the US National Science Foundation under grant DMR-9734041 and by an award from Research Corporation.

References

1. Shiwei Zhang, J. Carlson, and J. E. Gubernatis, Phys. Rev. Lett. **74**, 3652 (1995); Phys. Rev. B **55**, 7464 (1997).
2. Shiwei Zhang, J. Carlson, and J. E. Gubernatis, Phys. Rev. Lett. **78**, 4486 (1997).
3. Shiwei Zhang, to be published.
4. R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, Phys. Rev. D **24**, 2278 (1981).
5. G. Sugiyama and S.E. Koonin, Ann. Phys. (NY) **168**, 1 (1986).
6. M. H. Kalos, D. Levesque, and L. Verlet, Phys. Rev. **A9**, 2178 (1974); D. M. Ceperley and M. H. Kalos, in *Monte Carlo Methods in Statistical Physics*, edited by K. Binder (Springer-Verlag, Heidelberg, 1979).
7. E. Y. Loh Jr., J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, Phys. Rev. B **41**, 9301 (1990).
8. K. E. Schmidt and M. H. Kalos, in *Applications of the Monte Carlo Method in Statistical Physics*, edited by K. Binder (Springer Verlag, Heidelberg, 1984).
9. See, e.g., E. Dagotto, Rev. Mod. Phys. **66**, 763 (1994) and references therein.
10. See, for example, D. J. Scalapino, Phys. Reports **250**, 329 (1995), and references therein.
11. J. B. Anderson, J. Chem. Phys. **63**, 1499 (1975); **65**, 4122 (1976).
12. J. W. Moskowitz, K. E. Schmidt, M. A. Lee, and M. H. Kalos, J. Chem. Phys. **77**, 349 (1982); P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester, J. Chem. Phys. **77**, 5593 (1982).
13. For a review of the DMC method, see B. L. Hammond, W. A. Lester, Jr., and P. J. Reynolds, *Monte Carlo methods in ab initio quantum chemistry*, World Scientific, Singapore, 1994).
14. S. Sorella, *et. al.*, Int. Jour. Mod. Phys. B **1** 993 (1988).
15. D. M. Ceperley, Rev. Mod. Phys. **67**, 279 (1995).
16. D. M. Ceperley, Phys. Rev. Lett. **69**, 331 (1992); in *Simulation in Condensed Matter Physics and Chemistry*, Ed. by K. Binder and G. Ciccotti (1996).
17. A. Muramatsu, this volume.
18. J. E. Hirsch, Phys. Rev. B **31**, 4403 (1985).
19. J. Hubbard, Phys. Rev. Lett. **3**, 77 (1959).
20. S. B. Fahy and D. R. Hamann, Phys. Rev. Lett. **65**, 3437 (1990); Phys. Rev. B **43**, 765 (1991).
21. Shiwei Zhang and M. H. Kalos, Phys. Rev. Lett., **67**, 3074 (1991); J. B. Anderson, C. A. Traynor, and B. M. Boghosian, J. Chem. Phys. **95**, 7418 (1991).
22. J. Carlson, J. E. Gubernatis, G. Ortiz, and Shiwei Zhang, in preparation.
23. M. H. Kalos, J. Comp. Phys. **2**, 257 (1967).
24. M. P. Nightingale, this volume.
25. C. J. Umrigar, M. P. Nightingale, and K. J. Runge, J. Chem. Phys. **99**, 2865 (1993).
26. J. H. Hetherington, Phys. Rev. A **30**, 2713 (1984).
27. S. R. White, D. J. Scalapino, R. L. Sugar, E. Y. Loh Jr., J. E. Gubernatis, R. T. Scalettar,, Phys. Rev. B **40**, 506 (1989).
28. See E. Y. Loh Jr. and J. E. Gubernatis, in *Electrons Phase Transitions*, edited by W. Hanke and Y. V. Kopaev (Elsevier, New York, 1990), and references therein.
29. N. Furukawa and M. Imada, J. Phys. Soc. Jpn. **61**, 3331 (1992).
30. A. Parola, S. Sorella, S. Baroni, R. Car, M. Parrinello, and E. Tosatti, Physica C **162-164**, 771 (1989); A. Parola, S. Sorella, M. Parrinello, and E. Tosatti, Phys. Rev. B **43**, 6190 (1991); G. Fano, F. Ortolani, and A. Parola, Phys. Rev. B **42**,

6877 (1990).
31. D. F. B. ten Haaf, H. J. M. van Bemmel, J. M. J. van Leeuwen, W. van Saarloos, and D. M. Ceperley, Phys. Rev. B **51**, 13039 (1995); H. J. M. van Bemmel, D. F. B. ten Haaf, W. van Saarloos, J. M. J. van Leeuwen, and G. An, Phys. Rev. Lett., **72**, 2442 (1994).
32. D. F. B. ten Haaf and J. M. J. van Leeuwen, unpublished.
33. D. J. Scalapino, in *High Temperature Superconducting Proceedings*, ed. by K. S. Bedell, D. Coffey, D. E. Meltzer, D. Pines, and J. R. Schrieffer (Addison Wesley, 1990), p314.
34. R. T. Scalettar, private communication.