Quantum Monte Carlo Methods for Strongly Correlated Electron Systems

Shiwei Zhang

ABSTRACT We review some of the recent development in quantum Monte Carlo (QMC) methods for models of strongly correlated electron systems. QMC is a promising general theoretical tool to study many-body systems, and has been widely applied in areas spanning condensed-matter, high-energy, and nuclear physics. Recent progress has included two new methods, the ground-state and finite-temperature constrained path Monte Carlo methods. These methods significantly improve the capability of numerical approaches to lattice models of correlated electron systems. They allow calculations without any decay of the sign, making possible calculations for large system sizes and low temperatures. The methods are approximate. Benchmark calculations show that accurate results on energy and correlation functions can be obtained. This chapter gives a pedagogical introduction to quantum Monte Carlo, with a focus on the constrained path Monte Carlo methods.

1 Introduction

In order to understand properties of correlated electron systems and their theoretical and technological implications, accurate calculations are necessary to correctly treat microscopic correlations. This crucial need is manifested in the daily demand to solve model systems, which are often intractable analytically, in order to develop and benchmark theory and compare with experiments.

To effectively calculate electron correlation effects is an extremely challenging task. It requires the solution of the Schrödinger equation or evaluation of the density matrix for a many-body system at equilibrium. Explicit numerical approaches are an obvious possibility, which have found wide applications and proved very valuable, e.g., exact diagonalization in the study of lattice models for high-temperature superconductivity and for magnetism, and configuration interaction in quantum chemistry. However, since the dimensionality involved grows exponentially with system size, they all suffer from *exponential* complexity, i.e., exponential scaling of the required computer time with system size or accuracy. As an example, with the most powerful computers we now can treat a lattice of about 50 sites in exact diagonalization of the two-dimensional spin-1/2 Heisenberg model, up from ~ 20 sites twenty years ago. During this period, the peak speed of computers has been doubling roughly every two years.

Monte Carlo methods offer a promising alternative as a numerical approach to study many-body systems. Its required computer time scales algebraically (as opposed to exponentially in exact diagonalization) with system size. Rather than explicitly integrating over phase space, Monte Carlo methods sample it. The central limit theorem dictates that the statistical error in a Monte Carlo calculation decays algebraically with computer time. With the rapid advent of scalable parallel computers, these methods offer the possibility to study large system sizes systematically and extract information about the thermodynamic limit.

For fermion systems, however, Monte Carlo methods suffer from the so-called "sign" problem[1, 2, 3]. In these systems, the Pauli exclusion principle requires that the states be anti-symmetric under interchange of two particles. As a consequence, negative signs appear, which cause cancellations among contributions of the Monte Carlo samples of phase space. In fact, as the temperature is lowered or the system size is increased, such cancellation becomes more and more complete. The net signal thus decays *exponentially*. The algebraic scaling is then lost, and the method breaks down. Clearly the impact of this problem on the study of correlated electron systems is extremely severe.

To date most applications of QMC methods to strongly correlated electron models have either lived with the sign problem or relied on some form of approximation to overcome the exponential scaling. The former has difficulties reaching large system sizes or low temperatures. The latter loses the exactness of QMC and the results are biased by the approximation. Despite these limitations, QMC methods have proved a very useful theoretical tool in the study of strongly correlated electron systems. In many cases they have provided very accurate and reliable numerical results which are sometimes the only ones available for the system in question. Especially recently, progress has been rapid in the development of approximate QMC methods for lattice models, which has led to a growing number of applications. For example, with the constrained path Monte Carlo method (CPMC) discussed below, a system of 220 electrons on a 16×16 lattice in the Hubbard model can be studied with moderate computer time[4]. The dimension of the Hilbert space for this system exceeds 10^{150}.[1]

In this chapter, we review some of the recent progress in the study of models for strongly correlated electron systems with quantum Monte Carlo methods. The chapter is not meant to be a comprehensive review. Our focus is on the ground-state and finite-temperature constrained path Monte Carlo methods[5, 3, 6, 7]. We will, however, include in the Appendix some discussions of variational Monte Carlo and Green's function Monte Carlo (GFMC)[8] methods, which have also seen extensive applications in the

---

[1]The size for a 16×16 system with 109 electrons with up spins and 109 with down spins is about 10^{149}.

study of correlated electron models. We do so to illustrate that, while GFMC is very different from the auxiliary-field-based methods we focus on, the underlying ideas and the nature of the sign problems have much in common. Our goal is to highlight these ideas and show the capabilities that QMC methods in general bring as a theoretical tool, as well as the current algorithmic difficulties.

2   Preliminaries

2.1   Starting point of quantum Monte Carlo (QMC)

Most ground-state QMC methods are based on

$$|\Psi_0\rangle \propto \lim_{\tau \to \infty} e^{-\tau \hat{H}}|\Psi_T\rangle; \tag{1.1}$$

that is, the ground state $|\Psi_0\rangle$ of a many-body Hamiltonian $\hat{H}$ can be projected from any known trial state $|\Psi_T\rangle$ that satisfies $\langle \Psi_T|\Psi_0\rangle \neq 0$. In a numerical method, the limit can be obtained iteratively by

$$|\Psi^{(n+1)}\rangle = e^{-\Delta\tau \hat{H}}|\Psi^{(n)}\rangle, \tag{1.2}$$

where $|\Psi^{(0)}\rangle = |\Psi_T\rangle$. Ground-state expectation $\langle \hat{O}\rangle$ of a physical observable $\hat{O}$ is given by

$$\langle \hat{O}\rangle = \lim_{n \to \infty} \frac{\langle \psi^{(n)}|\hat{O}|\psi^{(n)}\rangle}{\langle \psi^{(n)}|\psi^{(n)}\rangle}. \tag{1.3}$$

For example, the ground-state energy can be obtained by letting $\hat{O} = \hat{H}$. A so-called mixed estimator exists, however, which is exact for the energy (or any other $\hat{O}$ that commutes with $\hat{H}$) and can lead to considerable simplifications in practice:

$$E_0 = \lim_{n \to \infty} \frac{\langle \psi_T|\hat{H}|\psi^{(n)}\rangle}{\langle \psi_T|\psi^{(n)}\rangle}. \tag{1.4}$$

Finite-temperature QMC methods use the density matrix. The expectation value of $\hat{O}$ is:

$$\langle \hat{O}\rangle = \frac{\mathrm{Tr}(\hat{O}e^{-\beta\hat{H}})}{\mathrm{Tr}(e^{-\beta\hat{H}})}, \tag{1.5}$$

where $\beta = 1/kT$ is the inverse temperature. In other words, $\langle \hat{O}\rangle$ is simply a weighted average with respect to the density matrix $e^{-\beta\hat{H}}$. In a numerical method, the partition function in the denominator of Eq. (1.5) is written as

$$Z \equiv \mathrm{Tr}(e^{-\beta\hat{H}}) = \mathrm{Tr}[\underbrace{e^{-\Delta\tau\hat{H}}\cdots e^{-\Delta\tau\hat{H}}e^{-\Delta\tau\hat{H}}}_{L}], \tag{1.6}$$

where $\Delta\tau = \beta/L$ is the "time" step, and L is the number of time slices.

Quantum Monte Carlo methods carry out the iteration in Eq. (1.2) or the trace in Eq. (1.6) — both requiring integration in many-dimensional spaces — by Monte Carlo sampling. The difference in the different classes of methods amounts primarily to the space that is used to represent the wave function or density matrix and to carry out the integration. The ground-state and finite-temperature constrained path Monte Carlo methods work in second quantized representation and in an auxiliary-field space, while Green's function Monte Carlo works in first-quantized representation and in configuration space.

2.2   Basics of Monte Carlo techniques

We list a few key elements from standard Monte Carlo techniques. In addition to serving as a brief reminder to the reader, they will help to introduce several results and notations that will be useful in our discussion of the QMC methods. For extensive discussions of Monte Carlo methods, excellent books exist[9].

Monte Carlo methods are often used to compute many-dimensional integrals of the form:

$$G = \frac{\int_{\Omega_0} f(\vec{x})g(\vec{x})d\vec{x}}{\int_{\Omega_0} f(\vec{x})d\vec{x}}, \tag{1.7}$$

where $\vec{x}$ is a vector in a many-dimensional space and Ω_0 is a domain in this space. We will assume that $f(\vec{x}) \geq 0$ on Ω_0 and that it is normalizable, i.e., the denominator is finite. A familiar example of the integral in Eq. (1.7) comes from classical statistical physics, where $f(\vec{x})$ is the Boltzmann distribution.

To compute G by Monte Carlo, we *sample* $\vec{x}$ from a probability density function (PDF) proportional to $f(\vec{x})$, i.e., the PDF $\bar{f}(\vec{x}) \equiv f(\vec{x})/\int_{\Omega_0} f(\vec{x})d\vec{x}$. This means we generate a sequence $\{\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_i, \cdots\}$ so that the probability that any $\vec{x}_i$ is in the sub-domain $(\vec{x}, \vec{x} + d\vec{x})$ is

$$\text{Prob}\{\vec{x}_i \in (\vec{x}, \vec{x} + d\vec{x})\} = \bar{f}(\vec{x})d\vec{x} \tag{1.8}$$

Below when we refer to sampling a function $f(\vec{x})$, it should be understood as sampling the corresponding PDF $\bar{f}(\vec{x})$. There are different techniques to sample a many-dimensional function $f(\vec{x})$. The most general and perhaps most often used is the Metropolis algorithm, which creates a Markov chain random walk in $\vec{x}$-space whose equilibrium distribution is the desired function.

Given $\mathcal{M}$ independent samples from $f(\vec{x})$, the integral in Eq. (1.7) is estimated by

$$G_\mathcal{M} = \frac{1}{\mathcal{M}} \sum_{i=1}^{\mathcal{M}} g(\vec{x}_i). \tag{1.9}$$

The error in the estimate decays algebraically with the number of samples: $|G_{\mathcal{M}} - G| \propto 1/\sqrt{\mathcal{M}}$.

Using the results above, we can compute

$$G' = \frac{\int_{\Omega_0} f(\vec{x})g(\vec{x})h(\vec{x})d\vec{x}}{\int_{\Omega_0} f(\vec{x})h(\vec{x})d\vec{x}}, \qquad (1.10)$$

if the function $h(\vec{x})$ is such that both the numerator and denominator exist. Formally

$$G'_{\mathcal{M}} = \frac{\sum_{i=1}^{\mathcal{M}} g(\vec{x}_i)h(\vec{x}_i)}{\sum_{i=1}^{\mathcal{M}} h(\vec{x}_i)}, \qquad (1.11)$$

although, as we will see, difficulties arise when $h(\vec{x})$ can change sign and is rapidly oscillating.

Integral equations are another main area of applications of Monte Carlo methods. For example[9], the integral equation

$$\Psi'(\vec{x}) = \int_{\Omega_0} K(\vec{x}, \vec{y}) \; w(\vec{y}) \; \Psi(\vec{y})d\vec{y}, \qquad (1.12)$$

can be viewed in terms of a *random walk* if it has the following properties: $\Psi(\vec{y})$ and $\Psi'(\vec{x})$ can be viewed as PDF's (in the sense of f in Eq. (1.7)), $w(\vec{y}) \geq 0$, and $K(\vec{x}, \vec{y})$ is a PDF for $\vec{x}$ conditional on $\vec{y}$. Then, given an ensemble $\{\vec{y}_i\}$ sampling $\Psi(\vec{y})$, the following two steps will allow us to generate an ensemble that samples $\Psi'(\vec{x})$. First an absorption/branching process is applied to each $\vec{y}_i$ according to $w(\vec{y}_i)$. For example, we can make $\text{int}(w(\vec{y}_i) + \xi)$ copies of $\vec{y}_i$, where ξ is a uniform random number on $(0, 1)$. Second we randomly walk each new $\vec{y}_j$ to an $\vec{x}_j$ by sampling the PDF $K(\vec{x}, \vec{y}_j)$. The resulting $\{\vec{x}_j\}$ are Monte Carlo samples of $\Psi'(\vec{x})$. We emphasize that the purpose of our discussion of random walks here is to illustrate the basic concept, which we will use later. The simple procedure described above is thus not meant as an accurate account of the technical details necessary for an efficient implementation.

2.3   Slater determinant space

In auxiliary-field-based QMC[10, 11, 12, 13] method, the Monte Carlo algorithm works in a space of Slater determinants. The building blocks of Slater determinants are single-particle basis states. The single-particle basis states can be plane waves, or lattice sites in the Hubbard model, or energy eigenstates in a mean-field potential. Often the space of single-particle basis states is truncated. Single-particle wave functions (orbitals) are formed with the basis states. Slater determinants are then built from the single-particle orbitals.

We first define some notations that we will use throughout the discussion of standard auxiliary-field quantum Monte Carlo (AFQMC) and then later the constrained path Monte Carlo (CPMC) methods.

- N: number of single-electron basis states. For example, N can be the number of lattice sites $(L \times L)$ in the two-dimensional Hubbard model.

- $|\chi_i\rangle$: the i^{th} single-particle basis state $(i = 1, 2, \cdots, N)$. For example, $|\chi_{\mathbf{G}}\rangle$ can be a plane wave basis state with $\chi_{\mathbf{G}}(\mathbf{r}) \propto e^{i\mathbf{G}\cdot\mathbf{r}}$, where $\mathbf{r}$ is a real-space co-ordinate.

- $c_i^\dagger$ and c_i: creation and annihilation operators for an electron in $|\chi_i\rangle$. $n_i \equiv c_i^\dagger c_i$ is the corresponding number operator.

- M: number of electrons (if we omit spin index, e.g., if the system is fully polarized). In the more general case, M_σ is the number of electrons with spin σ ($\sigma = \uparrow$ or $\downarrow$). Of course, the choice of N above must ensure that $M_\sigma \leq N$.

- φ_m: single-particle orbital (we include an index m for discussions below to distinguish different single-particle orbitals). A single-particle orbital φ_m, given in terms of the single-particle basis states $\{|\chi_i\rangle\}$ as $\sum_i \varphi_{i,m}|\chi_i\rangle = \sum_i c_i^\dagger \varphi_{i,m}|0\rangle$, can be conveniently expressed as an N-dimensional vector:

$$\begin{pmatrix} \varphi_{1,m} \\ \varphi_{2,m} \\ \vdots \\ \varphi_{N,m} \end{pmatrix}$$

Given M different single-particle orbitals, we form a many-body wave function from their anti-symmetrized product:

$$|\phi\rangle \equiv \hat{\varphi}_1^\dagger \hat{\varphi}_2^\dagger \cdots \hat{\varphi}_M^\dagger |0\rangle \tag{1.13}$$

where the operator

$$\hat{\varphi}_m^\dagger \equiv \sum_i c_i^\dagger \varphi_{i,m} \tag{1.14}$$

creates an electron in the m^{th} single-particle orbital $\{\varphi_{1,m}, \varphi_{2,m}, \cdots, \varphi_{N,m}\}$. The many-body state $|\phi\rangle$ in Eq. (1.13) can be conveniently expressed as an $N \times M$ matrix:

$$\Phi \equiv \begin{pmatrix} \varphi_{1,1} & \varphi_{1,2} & \cdots & \varphi_{1,M} \\ \varphi_{2,1} & \varphi_{2,2} & \cdots & \varphi_{2,M} \\ \vdots & \vdots & & \vdots \\ \varphi_{N,1} & \varphi_{N,2} & \cdots & \varphi_{N,M} \end{pmatrix}$$

Each column of this matrix represents a single-particle orbital that is completely specified by its N-dimensional vector.

If the real-space co-ordinates of the electrons are $R = \{\mathbf{r}_1, \mathbf{r}_2, \cdots, \mathbf{r}_M\}$, the many-body state in Eq. (1.13) gives

$$\langle R|\phi\rangle = \phi(R) = \det \begin{pmatrix} \varphi_1(\mathbf{r}_1) & \varphi_2(\mathbf{r}_1) & \cdots & \varphi_M(\mathbf{r}_1) \\ \varphi_1(\mathbf{r}_2) & \varphi_2(\mathbf{r}_2) & \cdots & \varphi_M(\mathbf{r}_2) \\ \vdots & \vdots & & \vdots \\ \varphi_1(\mathbf{r}_M) & \varphi_2(\mathbf{r}_M) & \cdots & \varphi_M(\mathbf{r}_M) \end{pmatrix},$$

where $\varphi_m(\mathbf{r}) = \sum_i \varphi_{i,m}\chi_i(\mathbf{r})$.

The many-body state $|\phi\rangle$ is known as a Slater determinant. One example of a Slater determinant is the Hartree-Fock (HF) solution $|\phi_{\mathrm{HF}}\rangle = \prod_\sigma |\phi_{\mathrm{HF}}^\sigma\rangle$, where each $|\phi_{\mathrm{HF}}^\sigma\rangle$ is defined by a matrix $\Phi_{\mathrm{HF}}^\sigma$ whose columns are the N_σ lowest HF eigenstates. We can now add the following to our list of notations above:

- $|\phi\rangle$: a many-body wave function which can be written as a Slater determinant.

- Φ: an $N \times M$ matrix which represents the Slater determinant $|\phi\rangle$. Φ_{ij} will denote the matrix element of the matrix Φ in the i^{th} row and j^{th} column. For example, $\Phi_{ij} = \varphi_{i,j}$ above in Φ. Below when a Slater determinant $|\phi\rangle$ is referred to, it will often be helpful to think in terms of the matrix representation Φ operationally.

- $|\Psi\rangle$ (upper case): a many-body wave function which is not necessarily a single Slater determinant, e.g., $|\Psi^{(n)}\rangle$ in Eq. (1.2).

Several properties of the Slater determinant are worth mentioning. For any two real non-orthogonal Slater determinants, $|\phi\rangle$ and $|\phi'\rangle$, it can be shown that their overlap integral is

$$\langle\phi|\phi'\rangle = \det(\Phi^{\mathrm{T}}\Phi'). \tag{1.15}$$

The operation on any Slater determinant by any operator $\hat{B}$ of the form

$$\hat{B} = \exp\left(\sum_{ij} c_i^\dagger U_{ij} c_j\right) \tag{1.16}$$

simply leads to another Slater determinant[14], i.e.,

$$\hat{B}|\phi\rangle = \hat{\phi}_1'^\dagger \hat{\phi}_2'^\dagger \cdots \hat{\phi}_M'^\dagger|0\rangle \equiv |\phi'\rangle \tag{1.17}$$

with $\hat{\phi}_m'^\dagger = \sum_j c_j^\dagger \Phi'_{jm}$ and $\Phi' \equiv e^U\Phi$, where U is a square matrix whose elements are given by U_{ij} and $B \equiv e^U$ is therefore an $N \times N$ square matrix as well. In other words, the operation of $\hat{B}$ on $|\phi\rangle$ simply involves multiplying an $N \times N$ matrix to an $N \times M$ matrix.

The many-body trace of an operator which is an exponential of one-body operators (i.e., of the form $\hat{B}$ in Eq. (1.16)) or a product of exponentials of one-body operators can be conveniently evaluated. The grand-canonical trace, which involves summing over a complete basis for a fixed number of particles M as well as over different numbers of particles (from 0 to N), has a particularly simple form[10, 12]:

$$\text{Tr}(\hat{B}) = \det[I + B], \tag{1.18}$$

where I is the $N \times N$ unit matrix and B once again is the corresponding matrix of the operator $\hat{B}$. The canonical trace, which is for a fixed number of particles M, is given by

$$\text{Tr}_M(\hat{B}) = \frac{1}{M!} \frac{d^M}{d\epsilon^M} \det[I + \epsilon B]\big|_{\epsilon=0}. \tag{1.19}$$

Note that this is simply the sum of all rank-M diagonal minors of the matrix B. As we would expect, the sum over all possible values of M of Eq. (1.19), with the appropriate factor for the chemical potential, recovers Eq. (1.18).

For an operator $\hat{O}$, we can define its expectation with respect to a pair of Slater determinants

$$\overline{\langle \hat{O} \rangle} \equiv \frac{\langle \phi | \hat{O} | \phi' \rangle}{\langle \phi | \phi' \rangle} \tag{1.20}$$

or with respect to a propagator B under the finite-temperature grand-canonical formalism of Eq. (1.18)

$$\overline{\langle \hat{O} \rangle} \equiv \frac{\text{Tr}(\hat{O}\hat{B})}{\text{Tr}(\hat{B})}. \tag{1.21}$$

The "bar" distinguishes these from the true interacting many-body expectations in Eq.'s (1.3) and (1.5). The latter are of course what we wish to compute with QMC.

The simplest example of Eq.'s (1.20) and (1.21) is the single-particle Green's function $G_{ij} \equiv \overline{\langle c_i c_j^\dagger \rangle}$. In the ground-state formalism,

$$G_{ij} \equiv \frac{\langle \phi | c_i c_j^\dagger | \phi' \rangle}{\langle \phi | \phi' \rangle} = \delta_{ij} - [\Phi'(\Phi^{\text{T}}\Phi')^{-1}\Phi^{\text{T}}]_{ij}. \tag{1.22}$$

In the finite-temperature grand-canonical formalism

$$G_{ij} \equiv \frac{\text{Tr}(c_i c_j^\dagger \hat{B})}{\text{Tr}(\hat{B})} = (I + B)_{ij}^{-1}. \tag{1.23}$$

Given the Green's function G, the general expectation defined in Eq.'s (1.20) and (1.21) can be computed for most operators of interest. This is an

important property that will be used in ground-state and finite-temperature QMC calculations. For example, we can calculate the expectation of a general two-body operator, $\hat{O} = \sum_{ijkl} O_{ijkl} c_i^\dagger c_j^\dagger c_k c_l$, under definitions (1.20) and (1.21):

$$\overline{\langle \hat{O} \rangle} = \sum_{ijkl} O_{ijkl}(G'_{jk}G'_{il} - G'_{ik}G'_{jl}), \tag{1.24}$$

where the matrix G' is defined as $G' \equiv I - G$.

2.4 Hubbard-Stratonovich transformation

In order to carry out Eq.'s (1.2) and (1.6) in the Slater-determinant space we have introduced above, we write the many-body propagator $e^{-\Delta\tau\hat{H}}$ in single-particle form. For simplicity we will treat the system as spin-polarized and suppress the spin index in most of the discussions below. It is, however, straightforward to generalize the discussions to include spins. Assuming that the many-body Hamiltonian involves two-body interactions only, we can write it as

$$\hat{H} = \hat{K} + \hat{V} = \sum_{i,j} K_{ij}(c_i^\dagger c_j + c_j^\dagger c_i) + \sum_{ijkl} V_{ijkl}\, c_i^\dagger c_j^\dagger c_k c_l. \tag{1.25}$$

For example, $\hat{K}$ and $\hat{V}$ can be the kinetic and potential energy operators, respectively. With a small $\Delta\tau > 0$, the Trotter approximation can be used:

$$e^{-\Delta\tau\hat{H}} \approx e^{-\Delta\tau\hat{K}} e^{-\Delta\tau\hat{V}}, \tag{1.26}$$

which introduces a Trotter error. For actual implementation of the algorithms we discuss here, higher order Trotter break-ups are often used. The Trotter error can be further reduced with an extrapolation procedure after separate calculations have been done with different values of $\Delta\tau$.

The $\hat{K}$ part of the propagator in (1.26) is the exponential of a one-body operator. The $\hat{V}$ part is not. It is, however, possible to rewrite $e^{-\Delta\tau\hat{V}}$ in this desired form through a so-called Hubbard-Stratonovich transformation[15]. As we will show below, $\hat{V}$ can be written as a sum of terms of the form $\lambda\hat{v}^2/2$, where λ is a constant and $\hat{v}$ is a one-body operator similar to $\hat{K}$. The Hubbard-Stratonovich transformation then allows us to write

$$e^{-\Delta\tau/2\,\lambda\,\hat{v}^2} = \int_{-\infty}^{\infty} dx \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} e^{x\sqrt{-\Delta\tau\lambda}\,\hat{v}}, \tag{1.27}$$

where x is an auxiliary-field variable. The constant in the exponent on the right-hand side can be real or imaginary dependent on the sign of λ. The key is that the quadratic form (in $\hat{v}$) on the left is replaced by a linear one on the right.

We now show one way to write $\hat{V}$ as a sum of $\lambda \hat{v}^2/2$. With the most general form of $\hat{V}$ in Eq. (1.25) we can cast V_{ijkl} in the form of a Hermitian matrix by introducing two indices $\alpha = (i, l)$ and $\beta = (k, j)$ and letting $\mathcal{V}_{\alpha\beta} = \mathcal{V}_{(i,l),(k,j)} = V_{ijkl}$. The Hermitian matrix $\mathcal{V}$ can then be diagonalized and written as $\mathcal{V} = R\Lambda R^\dagger$, where R is a matrix whose columns are the eigenvectors of $\mathcal{V}$ and Λ is a diagonal matrix containing the corresponding eigenvalues λ_α. That is

$$\mathcal{V}_{\alpha\beta} = \sum_\gamma R_{\alpha\gamma} \lambda_\gamma R^\star_{\beta\gamma}. \tag{1.28}$$

The two-body operator $\hat{V}$ can therefore be written as

$$
\begin{aligned}
\hat{V} &= \sum_{ijkl} V_{ijkl} \, c^\dagger_i c_l c^\dagger_j c_k - \sum_{ijkl} V_{ijkl} \, c^\dagger_i c_k \, \delta_{jl} \\
&= \sum_\gamma \lambda_\gamma (\sum_{il} R_{(i,l)\gamma} c^\dagger_i c_l)(\sum_{jk} R^\star_{(k,j)\gamma} c^\dagger_j c_k) - \sum_{ik}(\sum_j V_{ijkj}) c^\dagger_i c_k.
\end{aligned}
$$

Noting that $\hat{V}$ is Hermitian, we can put the above in a more symmetric form

$$\hat{V} = \frac{1}{2} \sum_\gamma \lambda_\gamma \{\hat{\rho}_\gamma, \hat{\rho}^\dagger_\gamma\} + \hat{\rho}_0, \tag{1.29}$$

where the one-body operators are defined as $\hat{\rho}_\gamma \equiv \sum_{il} R_{(i,l)\gamma} c^\dagger_i c_l$ and $\hat{\rho}_0 \equiv -\sum_{ik} [\sum_j (V_{ijkj} + V_{jijk})/2] c^\dagger_i c_k$. Since

$$\{\hat{\rho}_\gamma, \hat{\rho}^\dagger_\gamma\} = \frac{1}{2}[(\hat{\rho}_\gamma + \hat{\rho}^\dagger_\gamma)^2 - (\hat{\rho}_\gamma - \hat{\rho}^\dagger_\gamma)^2], \tag{1.30}$$

we have succeeded in writing $\hat{V}$ in the desired form.

The way to decompose $\hat{V}$ above leads to approximately $2N^2$ auxiliary fields. Often the interaction simplifies $\hat{V}$ and the number of auxiliary fields can be much reduced. In fact, certain type of interactions have particularly simple forms of Hubbard-Stratanovic transformations. For example, for the repulsive on-site interaction $U n_{i\uparrow} n_{i\downarrow}$ ($\uparrow$ and $\downarrow$ denote electron spin) in the Hubbard model, an *exact*, discrete Hubbard-Stratonovich transformation[16] exists:

$$e^{-\Delta\tau U n_{i\uparrow} n_{i\downarrow}} = e^{-\Delta\tau U(n_{i\uparrow} + n_{i\downarrow})/2} \sum_{x_i=\pm 1} \frac{1}{2} e^{\gamma x_i (n_{i\uparrow} - n_{i\downarrow})}, \tag{1.31}$$

where the constant γ is determined by $\cosh(\gamma) = \exp(\Delta\tau U/2)$. Similarly, for an attractive interaction $V n_{i\uparrow} n_{i\downarrow}$ with $V < 0$:

$$e^{-\Delta\tau V n_{i\uparrow} n_{i\downarrow}} = e^{-\Delta\tau V(n_{i\uparrow} + n_{i\downarrow} - 1)/2} \sum_{x_i=\pm 1} \frac{1}{2} e^{\gamma x_i (n_{i\uparrow} + n_{i\downarrow} - 1)}, \tag{1.32}$$

where $\cosh(\gamma) = \exp(\Delta\tau|V|/2)$

If we denote the collection of auxiliary fields by $\vec{x}$ and combine one-body terms from $\hat{K}$ and $\hat{V}$, we obtain the following compact representation of the outcome of the HS transformation:

$$e^{-\Delta\tau\hat{H}} = \int d\vec{x}\, p(\vec{x})\hat{B}(\vec{x}), \tag{1.33}$$

where $p(\vec{x})$ is a probability density function (e.g., a multi-dimensional Gaussian). The propagator $\hat{B}(\vec{x})$ in Eq. (1.33) has a *special form*, namely, it is a product of operators of the type in Eq. (1.16), with U_{ij} depending on the auxiliary field. The matrix representation of $\hat{B}(\vec{x})$ will be denoted by $B(\vec{x})$.

In essence, the HS transformation replaces the two-body interaction by one-body interactions with a set of random external auxiliary fields. In other words, it converts an interacting system into many *non-interacting* systems living in fluctuating external auxiliary-fields. The sum over all configurations of auxiliary fields recovers the interaction.

Different forms of the HS transformation exist[17, 18]. It is reasonable to expect that they can affect the performance of the QMC method. Indeed experience shows that they can not only impact the statistical accuracy, but also lead to different quality of approximations under the constrained path methods that we discuss below. Therefore, although we discuss the algorithms under the generic form of Eq. (1.33), we emphasize that it is worthwhile to explore different forms of the HS transformation in an actual implementation.

3   Standard Auxiliary-Field Quantum Monte Carlo

In this section we briefly describe the standard ground-state[11, 13] and finite-temperature[10, 12] auxiliary-field quantum Monte Carlo (AFQMC) methods. We will rely on the machinery established in the previous section. Our goal is to illustrate the essential ideas, in a way which will facilitate our discussion of the sign problem and help introduce the framework for the constrained path Monte Carlo methods. We will not go into details such as how to efficiently sample the auxiliary fields or how to stabilize the matrices. They are described in the literature. In addition, although the approaches are not identical, we will see these issues manifested in the next section and gain sufficient understanding of them when we discuss the corresponding constrained path Monte Carlo methods.

3.1 Ground-state method

Ground-state expectation $\langle\hat{O}\rangle$ can now by computed with (1.2) and (1.33). The denominator is

$$\langle\psi^{(0)}|e^{-n\Delta\tau\hat{H}}\,e^{-n\Delta\tau\hat{H}}|\psi^{(0)}\rangle$$

$$= \int\langle\psi^{(0)}|\Big[\prod_{l=1}^{2n}d\vec{x}^{(l)}p(\vec{x}^{(l)})\hat{B}(\vec{x}^{(l)})\Big]|\psi^{(0)}\rangle$$

$$= \int\Big[\prod_l d\vec{x}^{(l)}p(\vec{x}^{(l)})\Big]\det\Big([\Psi^{(0)}]^{\mathrm{T}}\prod_l B(\vec{x}^{(l)})\Psi^{(0)}\Big). \qquad (1.34)$$

In the standard ground-state AFQMC method[10], a (sufficiently large) value of n is first chosen and fixed throughout the calculation. If we use X to denote the collection of the auxiliary-fields $X = \{\vec{x}^{(1)}, \vec{x}^{(2)}, \ldots, \vec{x}^{(2n)}\}$ and $D(X)$ to represent the integrand in Eq. (1.34), we can write the expectation value of Eq. (1.3) as

$$\langle\hat{O}\rangle = \frac{\int\overline{\langle\hat{O}\rangle}\,D(X)\,dX}{\int D(X)\,dX} = \frac{\int\overline{\langle\hat{O}\rangle}\,|D(X)|\,s(X)\,dX}{\int|D(X)|\,s(X)\,dX}, \qquad (1.35)$$

where

$$s(X) \equiv D(X)/|D(X)| \qquad (1.36)$$

measures the "sign" of $D(X)$. The non-interacting expectation for a given X is that defined in Eq. (1.20):

$$\overline{\langle\hat{O}\rangle} \equiv \frac{\langle\phi_L|\hat{O}|\phi_R\rangle}{\langle\phi_L|\phi_R\rangle} \qquad (1.37)$$

with

$$\begin{aligned}
\langle\phi_L| &= \langle\psi^{(0)}|\hat{B}(\vec{x}^{(2n)})\hat{B}(\vec{x}^{(2n-1)})\cdots\hat{B}(\vec{x}^{(n+1)}) \\
|\phi_R\rangle &= \hat{B}(\vec{x}^{(n)})\hat{B}(\vec{x}^{(n-1)})\cdots\hat{B}(\vec{x}^{(1)})|\psi^{(0)}\rangle,
\end{aligned}$$

which are both Slater determinants.

$D(X)$ as well as $\langle\phi_L|$ and $|\phi_R\rangle$ are completely determined by the path X in auxiliary-field space. The expectation in Eq. (1.35) is therefore in the form of Eq. (1.10), with $f(X) = |D(X)|$ and $g(X) = \overline{\langle\hat{O}\rangle}$. The important point is that, for each X, $|D(X)|$ is a number and $g(X)$ can be evaluated using Eq.'s (1.22) or (1.24). Often the Metropolis Monte Carlo algorithm[9] is used to sample auxiliary-fields X from $|D(X)|$. Any $\langle\hat{O}\rangle$ can then be computed following the procedure described by Eq. (1.9).

The Metropolis Monte Carlo algorithm allows one to, starting from an (arbitrary) initial configuration of X, create a random walk whose equilibrium distribution is $f(X)$. A kernel $K(X', X)$ must be chosen prior to the

calculation which is a probability density function of X' (conditional on X). The only other condition on K is ergodicity, i.e., it allows the random walk to eventually reach any position from any other position in X-space. For example, it can be a uniform distribution in a hyper-cube of linear dimension Δ centered at X

$$K(X', X) = \begin{cases} 1/\Delta^{2nd}, & \text{if } X' \text{ inside hyper-cube} \\ 0, & \text{otherwise} \end{cases}$$

where Δ is a parameter and d is the dimensionality of the auxiliary-field $\vec{x}$.

The random walk resembles that in Eq. (1.12). A move from X to X' is proposed by sampling an X' according to $K(X', X)$. The key of the Metropolis algorithm is an acceptance probability

$$A(X', X) = \min\left[\frac{K(X, X')f(X')}{K(X', X)f(X)}, 1\right] \qquad (1.38)$$

whose role is similar to that of w in Eq. (1.12). The one important difference is that, if the outcome of $A(X', X)$ is to reject X', X is kept and counted as one new Monte Carlo sample — and hence one new step in the random walk, even though the walker did not actually move.

The choice of K, which has a great deal of freedom, controls the efficiency of the random walk, i.e., how quickly it converges, how correlated successive steps of the random walk are, etc. A choice that is perhaps the most common in practice is to "sweep" through the field one component at a time. A change is proposed for only one component of the many-dimensional vector X, from a kernel $K(\vec{x}_i', \vec{x}_i)$ in similar form to K , while the rest of X is kept the same; the acceptance/rejection procedure is then applied. With the newly generated field, a change is proposed on the next $((l+1)$-th) component and the step above is repeated. The "local" nature of this choice of the kernel can often lead to simplifications in the computation of A.

3.2   Finite-temperature method

The standard finite-temperature auxiliary-field QMC method works in the grand-canonical ensemble. This means that the Hamiltonian contains an additional term $-\mu \sum_i n_i$ which involves the chemical potential μ. The term leads to an additional diagonal one-body operator in the exponent in Eq. (1.16). We will assume that the resulting term has been absorbed in $\hat{B}$ in Eq. (1.33).

Substituting Eq. (1.33) into Eq. (1.6) leads to

$$Z = \int \left[\prod_l d\vec{x}_l p(\vec{x}_l)\right] \text{Tr}[\hat{B}(\vec{x}_L) \cdots \hat{B}(\vec{x}_2)\hat{B}(\vec{x}_1)]. \qquad (1.39)$$

If we apply Eq. (1.18) to the above and again use X to denote a complete

path in auxiliary-field space, $X \equiv \{\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_L\}$, we obtain

$$Z = \int \det[I + B(\vec{x}_L) \cdots B(\vec{x}_2) B(\vec{x}_1)] \, p(X) \, dX. \qquad (1.40)$$

Again representing the integrand by $D(X)$, we see that the finite-temperature expectation in Eq. (1.5) reduces to Eq. (1.35), exactly the same expression as in ground-state AFQMC. The non-interacting expectation $\overline{\langle \hat{O} \rangle}$ is now

$$\overline{\langle \hat{O} \rangle} \equiv \frac{\mathrm{Tr}[\hat{O} \, \hat{B}(\vec{x}_L) \cdots \hat{B}(\vec{x}_2) \hat{B}(\vec{x}_1)]}{\mathrm{Tr}[\hat{B}(\vec{x}_L) \cdots \hat{B}(\vec{x}_2) \hat{B}(\vec{x}_1)]}. \qquad (1.41)$$

Note that, as in the ground-state algorithm, this can be computed for each X. Eq. (1.23) remains valid. The matrix B must be replaced by the ordered product of the corresponding matrices, $B(\vec{x}_L) \cdots B(\vec{x}_2) B(\vec{x}_1)$. "Wrapping" this product (without changing the ordering) gives the equal time Green's function at different imaginary times.

Thus we have unified the ground-state and finite-temperature AFQMC methods with Eq. (1.35). Although the actual forms for $D(X)$ and $\overline{\langle \hat{O} \rangle}$ are different, the finite-temperature algorithm is the same as that of the ground state described earlier. Applying the same procedure, we can sample auxiliary-fields X from $|D(X)|$, and then compute the desired many-body expectation values from Eq. (1.35).

4   Constrained Path Monte Carlo Methods — Ground-State and Finite-Temperature

In this section, we review the ground-state and finite-temperature constrained path Monte Carlo (CPMC) methods. These methods are free of any decay of the average sign, while retaining many of the advantages of the standard AFQMC methods. The methods are approximate, relying on what we will generally refer to as the constrained path approximation. Below, other than presenting a way to make a more formal connection between the ground-state and finite-temperature methods, we will largely follow Ref.'s [5, 7] in discussing these methods.

4.1   Why and how does the sign problem occur?

In the standard AFQMC methods, the sign problem occurs because $D(X)$ is not always positive.[2] The Monte Carlo samples of X therefore have to

---

[2] In fact $\hat{B}(\vec{x})$ is complex for a general HS transformation when the interaction is repulsive. $D(X)$ is therefore complex as well. We will limit our discussions here of the sign problem and the constrained path approximation to the case where $D(X)$ is real, although generalizations are possible.
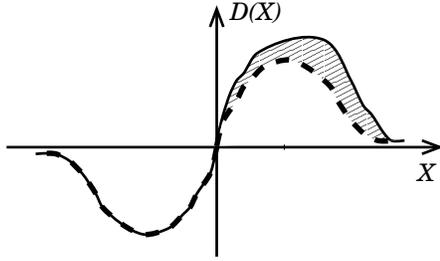
FIGURE 1. Schematic illustration of the sign problem. The X-axis represents an abstraction of the many-dimensional auxiliary-field paths X; each point denotes a collection of X's. The sign problem occurs because the contributing part (shaded area) is exponentially small compared to what is sampled, namely $|D(X)|$.
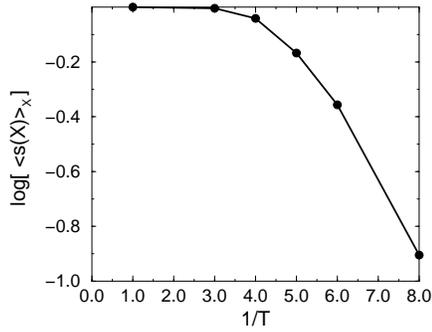

FIGURE 2. Decay of the average sign as a function of inverse temperature. Note the logarithmic scale on the vertical axis. The calculation is for a 4×4 Hubbard model using the finite-temperature formalism. The chemical potential is such that the average electron density is 0.875 (corresponding to 7 ↑ 7 ↓ electrons). The on-site repulsion $U = 4$. (Data courtesy of Richard Scalettar [19].)

be drawn from the probability distribution function defined by $|D(X)|$. As Fig. 1 illustrates, however, $D(X)$ approaches an antisymmetric function exponentially as N or the the length of the path is increased. That is, its average sign, i.e., the denominator in Eq. (1.35), vanishes at large system size or low temperature. Thus the variance of the calculation in Eq. (1.11) grows exponentially, which negates the advantage of the "square-root" behavior of a Monte Carlo calculation (cf. Eq. (1.9)). Fig. (2) shows exponential decay of the average sign $\langle s(X) \rangle$ of the Monte Carlo samples as a function of inverse temperature. This problem was largely uncontrolled in both the ground-state and finite-temperature algorithms, preventing general simulations at low temperatures or large system sizes.

Here we will develop a more detailed picture for the behavior of $D(X)$ and

hence for how the sign problem occurs. Under this picture the constrained path approximation, which will eliminate the decay of the denominator in Eq. (1.35) as a function of n or lattice size, emerges naturally.

For simplicity, we will now use one common symbol L to denote the length of the path X. That is, in the ground-state AFQMC algorithm in the previous section, we let $L = 2n$. $D(X) = D(\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_L)$ will denote the corresponding integrand in either the ground-state or finite-temperature AFQMC algorithm. We recall that the goal is to sample X according to $D(X)$ (although we had to resort to sampling $|D(X)|$ in the previous section).

To gain insight we conduct the following thought experiment[7]. We imagine sampling the complete path X by L successive steps, from $\vec{x}_1$ to $\vec{x}_L$. We consider the contribution in Eq. (1.35) by an individual *partial path* $\{\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_l\}$:

$$\mathcal{P}_l(\{\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_l\}) \equiv \int D(\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_l, \vec{x}_{l+1}, \cdots, \vec{x}_L) \, d\vec{x}_{l+1} \cdots d\vec{x}_L.$$

(1.42)

Note that this is the same as replacing $e^{\Delta\tau\hat{H}}$ by Eq. (1.33) only in the first l time slices. For simplicity we will use $\hat{\mathcal{B}}$ to denote the many-body propagator $e^{\Delta\tau\hat{H}}$ in our discussion below, i.e, $\hat{\mathcal{B}} \equiv e^{\Delta\tau\hat{H}}$. Under the ground-state formalism,

$$\begin{aligned} &\mathcal{P}_l(\{\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_l\}) \\ =\ &\langle\psi^{(0)}|\big[\ \underbrace{\hat{\mathcal{B}}\hat{\mathcal{B}}\cdots\hat{\mathcal{B}}}_{L-l}\ \hat{B}(\vec{x}_l)\cdots\hat{B}(\vec{x}_2)\hat{B}(\vec{x}_1)\big]|\psi^{(0)}\rangle p(\vec{x}_l)\cdots p(\vec{x}_1), \end{aligned}$$ (1.43)

while under the finite-temperature formalism

$$\mathcal{P}_l(\{\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_l\}) = \mathrm{Tr}\big[\underbrace{\hat{\mathcal{B}}\hat{\mathcal{B}}\cdots\hat{\mathcal{B}}}_{L-l}\ \hat{B}(\vec{x}_l)\cdots\hat{B}(\vec{x}_2)\hat{B}(\vec{x}_1)\big]p(\vec{x}_l)\cdots p(\vec{x}_1).$$

(1.44)

We consider the case when $\mathcal{P}_l = 0$. This means that, after the remaining $L-l$ steps are finished, the collective contribution from *all* complete paths that result from this particular partial path will be precisely zero. In other words, all complete paths that have $\{\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_l\}$ as their first l elements collectively make no contribution in the denominator in Eq. (1.35). This is because the path integral over all possible configurations for the rest of the path, $\{\vec{x}_{l+1}, \vec{x}_{l+2}, \cdots, \vec{x}_L\}$, is simply given by Eq. (1.42). In other words, the path integral simply reproduces the $\hat{\mathcal{B}}$'s in Eq. (1.43) or (1.44), leading to $\mathcal{P}_l(\{\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_l\})$ which is zero by assumption.

Thus, in our thought experiment any partial path that reaches the axis in Fig. 3 immediately turns into noise, regardless of what it does at future l's. A complete path which is in contact with the axis at any point belongs to the "antisymmetric" part of $D(X)$ in Fig. 1, whose contributions cancel.
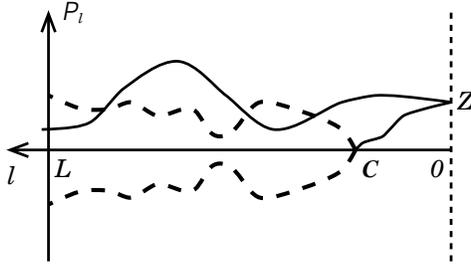
FIGURE 3. Schematic illustration of the boundary condition to control the sign problem in AFQMC. $\mathcal{P}_l$ (Eq. (1.42)) is shown as a function of the length of the partial path, l, for several paths. The remainder of the path integral (of length $L - l$) is assumed to have been evaluated *analytically*. (See Eq. (1.42).) All paths start from the right ($l = 0$) at $\mathcal{P}_0$, which is either $\langle\psi_T|e^{L\Delta\tau\hat{H}}|\psi_T\rangle > 0$ (ground-state) or Z (finite-temperature). When $\mathcal{P}_l$ becomes 0, ensuing paths (dashed lines) collectively contribute zero. Only complete paths with $\mathcal{P}_l > 0$ for all l (solid line) contribute in the denominator in Eq. (1.35).

As L increases, it becomes more and more likely for a path to reach the axis at some point l. Therefore, the "noise" paths become an increasingly larger portion of all paths unless paths are completely prohibited from reaching the axis, such as in the half-filled Hubbard model with repulsive interaction, where particle-hole symmetry makes $D(X)$ positive. A complete path X which crosses the axis at least once, i.e., a noise path, can have a $D(X)$ which is either positive or negative. In the Monte Carlo sampling in AFQMC, we sample X according to $|D(X)|$, which makes no distinction of such paths from the contributing ones that stay completely above the axis. Asymptotically in L, the Monte Carlo samples consist of an equal mixture of positive and negative contributions in $D(X)$. The Monte Carlo signal is lost and the Monte Carlo estimate of Eq. (1.35) becomes increasingly noisy. This is the origin of the sign problem.

4.2   The constrained-path approximation

$\mathcal{P}_0$ is positive, because $\langle\psi_T|e^{L\Delta\tau\hat{H}}|\psi_T\rangle > 0$ (ground-state formalism) and $Z > 0$ (finite-temperature). $\mathcal{P}_l$ changes continuously with l at the limit $\Delta\tau \to 0$. Therefore a complete path contributes if and only if it stays entirely above the axis in Fig. 3. Thus, in our thought experiment, imposition of the boundary condition[5, 7, 20]

$$\mathcal{P}_1(\{\vec{x}_1\}) > 0$$
$$\mathcal{P}_2(\{\vec{x}_1, \vec{x}_2\}) > 0$$
$$\cdots$$
$$\mathcal{P}_L(\{\vec{x}_1, \vec{x}_2, \cdots, \vec{x}_L\}) > 0 \tag{1.45}$$

will ensure all contributing complete paths to be selected while eliminating all noise paths. The axis acts as an infinitely absorbing boundary. A partial path is terminated and discarded as soon as it reaches the boundary. By discarding a path, we eliminate all of its future paths, including the ones that would eventually make positive contributions. The boundary condition makes the distribution of complete paths vanish at the axis, which accomplishes complete cancellation of the negative and the corresponding positive contributions in the antisymmetric part of $D(X)$.

In practice $\hat{\mathcal{B}}$ is of course not known. We replace it by a known trial propagator $\hat{B}_T$. The boundary condition in Eq. (1.45) is then clearly approximate. Results from our calculation will depend on the trial propagator used in the constraint. We expect the quality of the approximation to improve with a better $\hat{B}_T$. This is the basic idea of the constrained path approximation.

For ground-state calculations, let us imagine taking L to infinity. If $\hat{B}_T$ is a mean-field propagator, the wave function

$$\langle \psi_T | = \lim_{L \to \infty} \langle \psi^{(0)} | \hat{B}_T^{L-l} \tag{1.46}$$

can be thought of as the corresponding ground-state solution. The right-hand side of Eq. (1.43) can then be evaluated:

$$\mathcal{P}_l^T \equiv \langle \psi_T | \phi^{(l)} \rangle p(\vec{x}_l) \cdots p(\vec{x}_1), \tag{1.47}$$

where the Slater determinant that results from a particular partial path is defined as

$$|\phi^{(l)}\rangle \equiv \hat{B}(\vec{x}_l) \cdots \hat{B}(\vec{x}_2) \hat{B}(\vec{x}_1) |\psi^{(0)}\rangle. \tag{1.48}$$

The approximate condition in Eq. (1.45) can then be written as

$$\langle \psi_T | \phi^{(l)} \rangle > 0. \tag{1.49}$$

Notice that the constraint now becomes independent of the length of the path l. It is expressed in terms of the overlap of a determinant with a trial wave function $|\psi_T\rangle$. For a choice of $|\psi_T\rangle$ that is a single Slater determinant or a linear combination of Slater determinants, Eq. (1.15) allows straightforward computation of the overlap in Eq. (1.49).

In principle, we could keep L finite and directly implement the constraint in the framework of the standard AFQMC, using the propagator $\hat{B}_T$ instead of $|\psi_T\rangle$. But, as we will see below, the formalism above lends itself more naturally to the random walk realization of the ground-state method.

The constraint is manifested in one other way in the computation of expectations of operators $\hat{O}$ that do not commute with $\hat{H}$. Because we need to insert $\hat{O}$ in the middle in the numerator of Eq. (1.37), we will have to "back-propagate"[5, 3] from the left to obtain a $\langle \phi_L |$. The constraint, however, is implemented in the opposite direction, as we sample the path

from right to left. 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 violate it any even number of times when going from left to right. This sense of direction means that the expectation $\langle \hat{O} \rangle$ we compute is not strictly $\langle \psi_0^c | \hat{O} | \psi_0^c \rangle / \langle \psi_0^c | \psi_0^c \rangle$, with $|\psi_0^c\rangle$ the approximate ground-state wave function under the constraint. This difference is not crucial, however, since it is expected to be of the same order as the error due to the constrained path approximation itself.

For finite-temperature calculations, the right-hand side of Eq. (1.44) can be evaluated if $\hat{B}_T$ is in the form of a single-particle propagator

$$\mathcal{P}_l^T \equiv \det[I + \underbrace{B_T B_T \cdots B_T}_{L-1} B(\vec{x}_1)] p(\vec{x}_l) \cdots p(\vec{x}_1), \qquad (1.50)$$

where, following our convention, B_T is the matrix representation of $\hat{B}_T$. Combining this with Eq. (1.45) leads to the following matrix representation of the approximate boundary condition:

$$\det[I + \underbrace{B_T B_T \cdots B_T}_{L-1} B(\vec{x}_1)] > 0$$

$$\det[I + \underbrace{B_T B_T \cdots B_T}_{L-2} B(\vec{x}_2) B(\vec{x}_1)] > 0$$

$$\cdots$$

$$\det[I + B(\vec{x}_L) \cdots B(\vec{x}_2) B(\vec{x}_1)] > 0 \qquad (1.51)$$

Formally, the constrained path approximation in Eq.'s (1.49) and (1.51) has similarity to the fixed-node[21, 22, 23] (ground-state) or restricted-path[24] (finite-temperature) approximation in configuration space[3] Significant differences exist, however, because of the difference between the real space and a Slater determinant space, which is non-orthogonal.

The goal of the new CPMC methods is to carry out the thought experiment stochastically. We wish to generate Monte Carlo samples of X which *both* satisfy the conditions in (1.49) or (1.51) *and* are distributed according to $D(X)$. The most natural way to accomplish this is perhaps to incorporate the boundary conditions into the standard AFQMC algorithm, for example, as an additional acceptance condition. However, such an approach is likely to be inefficient: The boundary condition is *non-local* and breaks translational invariance in imaginary time, which requires simultaneous updating of the entire path. Without a scheme to propose paths that incorporates information on future contributions, it is difficult to find complete paths which satisfy all the constraints, especially as L increases.

---

[3]See the Appendix for a brief description of the fixed-node approximation for lattice Fermion systems.

We therefore seek to sample paths via a random walk whose time variable corresponds to the imaginary time l. We will see that we must introduce importance sampling schemes to guide the random walks by their projected contribution in $D(X)$. Below we discuss details of such schemes.

4.3 Ground-state constrained path Monte Carlo (CPMC) method

We start from Eq. (1.2). Using (1.33), we write it as

$$|\psi^{(l+1)}\rangle = \int d\vec{x}\, p(\vec{x})\hat{B}(\vec{x})|\psi^{(l)}\rangle. \qquad (1.52)$$

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

$$|\psi^{(l)}\rangle \propto \sum_k w_k^{(l)}|\phi_k^{(l)}\rangle, \qquad (1.53)$$

where k labels the Slater determinants and an overall normalization factor of the wave function has been omitted. A weight factor $w_k^{(l)}$ is introduced for each walker, even though in Eq. (1.52) the kernal p is normalized. This is because single-particle orbitals in a Slater determinants cease to be orthogonal to each other as a result of propagation by $\hat{B}$. When they are re-orthogonalized (see Section 4.5), an overall factor appears, which we will view as the w term in the integral equation Eq. (1.12).

The structure of the random walk now resembles that of Eq. (1.12). For each random walker we sample an auxiliary-field configuration $\vec{x}$ from the probability density function $p(\vec{x})$ and propagate the walker to a new one via $\Phi_k^{(l+1)} = B(\vec{x})\Phi_k^{(l)}$. If necessary, a re-orthogonalization procedure (e.g., modified Gram-Schmidt) is applied to $\Phi_k^{(l)}$ prior to the propagation: $\Phi_k^{(l)} = [\Phi_k^{(l)}]'R$, where R is an $M \times M$ upper-triangular matrix. $[\Phi_k^{(l)}]'$ is then used in the propagation instead; the weight of the new walker is $w_k^{(l+1)} = w_k^{(l)}\det(R)$.

The simple random walk procedure is correct and can be very useful for thinking about many conceptual issues. It is, however, not efficient enough as a practical algorithm in most cases of interest, because the sampling of $\vec{x}$ is completely random with no regard to the potential contribution to $D(X)$. The idea of importance sampling is to iterate a modified equation with a modified wave function, without changing the underlying eigenvalue problem of (1.52). Specifically, for each Slater determinant $|\phi\rangle$, we define an importance function

$$O_T(\phi) \equiv \langle\psi_T|\phi\rangle, \qquad (1.54)$$

which estimates its overlap with the ground-state wave function. We can

then rewrite equation (1.52) as

$$|\tilde{\psi}^{(l+1)}\rangle = \int d\vec{x}_{l+1}\tilde{p}(\vec{x}_{l+1})\hat{B}(\vec{x}_{l+1})|\tilde{\psi}^{(l)}\rangle, \tag{1.55}$$

where the modified "probability density function" is

$$\tilde{p}(\vec{x}_{l+1}) = \frac{O_T(\phi^{(l+1)})}{O_T(\phi^{(l)})}p(\vec{x}_{l+1}). \tag{1.56}$$

It is easy to see that the new kernel $\tilde{p}$ in Eq. (1.56) can be written in terms of the trial partial path contributions in Eq. (1.47):

$$\tilde{p}(\vec{x}_{l+1}) = \frac{\mathcal{P}_{l+1}^T}{\mathcal{P}_l^T}. \tag{1.57}$$

As we will see, this is formally analogous to the importance-sampling kernel in the finite-temperature algorithm we discuss next.

With the new kernel $\tilde{p}$, the probability distribution for $\vec{x}_{l+1}$ vanishes smoothly as $\mathcal{P}_{l+1}^T$ approaches zero, and the constraint is naturally imposed. The auxiliary field $\vec{x}_{l+1}$ is sampled according to our best estimate of the partial path contributions (from the trial propagator/wave function). 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 (1.52) and (1.55) are identical.

To better see the effect of importance sampling, we observe that if $|\psi_T\rangle = |\psi_0\rangle$, the normalization $\int \tilde{p}(\vec{x})d\vec{x}$ becomes a constant. Therefore the weights of walkers remain a constant and the random walk has no fluctuation.

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

$$|\psi^{(l)}\rangle \propto \sum_k w_k^{(l)}|\phi_k^{(l)}\rangle/O_T(\phi_k^{(l)}), \tag{1.58}$$

although in actual measurements it is $|\tilde{\psi}^{(l)}\rangle$ that is needed and division by O_T does not appear. The random walk process is similar to that discussed for Eq. (1.52), 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 $N(\phi_k^{(l)})$ and rewrite the iterative relation as

$$|\phi_k^{(l+1)}\rangle \leftarrow N(\phi_k^{(l)}) \int d\vec{x} \frac{\tilde{p}(\vec{x})}{N(\phi_k^{(l)})}B(\vec{x})|\phi_k^{(l)}\rangle. \tag{1.59}$$

This iteration now forms the basis of the ground-state constrained path Monte Carlo algorithm. For each walker $|\phi_k^{(l)}\rangle$, one step of the random walk consists of:

1. sampling an $\vec{x}$ from the probability density function $\tilde{p}(\vec{x})/N(\phi_k^{(l)})$

2. constructing the corresponding $B(\vec{x})$ and propagating the walker $\Phi_k^{(l)}$ by it to generate a new walker

3. assigning a weight $w_k^{(l+1)} = w_k^{(l)} N(\phi_k^{(l)})$ to the new walker

Note that, in contrast with the primitive algorithm in Eq. (1.52), the weight factor of a walker does not need to be modified here when the re-orthogonalization procedure is applied. This is because the upper-triangular matrix R only contributes to the overlap O_T, which is already represented by the walker weight. After each modified Gram-Schmidt procedure, R can simply be discarded.

4.4 Finite-temperature method

We again construct an algorithm which builds directly into the sampling process both the constraints and some knowledge of the projected contribution. In terms of the trial projected partial contributions $\mathcal{P}_l^T$ defined in Eq. (1.50), the fermion determinant $D(X)$ can be written as

$$D(X) = \frac{\mathcal{P}_L^T}{\mathcal{P}_{L-1}^T} \frac{\mathcal{P}_{L-1}^T}{\mathcal{P}_{L-2}^T} \cdots \frac{\mathcal{P}_2^T}{\mathcal{P}_1^T} \frac{\mathcal{P}_1^T}{\mathcal{P}_0^T} \, \mathcal{P}_0^T. \tag{1.60}$$

We again construct the path X by a random walk of L steps, corresponding to stochastic representations of the L ratios in Eq. (1.60). In Fig. 4 the sampling procedure for one walker is illustrated schematically. At the $(l + 1)$-th step, we sample $\vec{x}_{l+1}$ from the conditional probability density function defined by $\tilde{p}(\vec{x}_{l+1}) = \mathcal{P}_{l+1}^T/\mathcal{P}_l^T$, which is precisely the same as Eq. (1.57) in the ground-state method. Like in the ground-state method or in Green's function Monte Carlo, we simultaneously keep an ensemble of walkers, which undergo branching in their random walks (of L steps).

In the random walk, each walker is a product of B_T's and $B(\vec{x})$'s. They are initialized to $\mathcal{P}_0^T$, with overall weight 1. At the end of the l-th step, each walker has the form

$$\underbrace{B_T B_T \cdots B_T}_{L-l} \, B(\vec{x}_l) \cdots B(\vec{x}_2) B(\vec{x}_1) \tag{1.61}$$

where $\{\vec{x}_l, \cdots, \vec{x}_2, \vec{x}_1\}$ is the partial path already sampled for this walker. Each step for a walker is formally identical to that in the ground-state method:

1. pick an $\vec{x}_{l+1}$ from the probability density function $\tilde{p}(\vec{x}_{l+1})$

2. advance the walker by replacing the B_T next to $B(\vec{x}_l)$ by $B(\vec{x}_{l+1})$
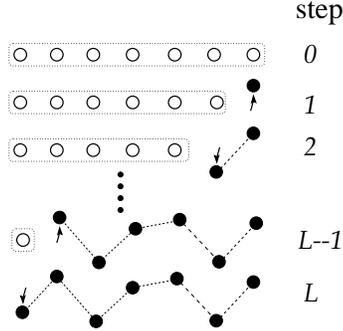
FIGURE 4. Illustration of the sampling procedure in the algorithm. Circles represent auxiliary-fields $\vec{x}_l$. A row shows the field configuration at the corresponding step number shown on the right. Within each row, the imaginary-time index l increases as we move to the left, i.e., the first circle is $\vec{x}_1$ and the last $\vec{x}_L$. Empty circles indicate fields which are not "activated" yet, i.e., $\hat{B}_T$ is still in place of $\hat{B}$. Solid circles indicate fields that have been sampled, with the arrow indicating the one being sampled in the current step.

3. multiply the overall weight of the walker by the normalization factor of $\tilde{p}(\vec{x}_{l+1})$.

The basic idea of the algorithm is, similar to the ground-state method, to select $\vec{x}_l$ according to the best estimate of its projected contribution in Z. Note that the probability distribution for $\vec{x}_l$ again vanishes smoothly as $\mathcal{P}_l^T$ approaches zero. This naturally imposes the boundary condition in Eq. (1.51). The Monte Carlo samples are generated from a probability distribution function of the contributing paths only (solid lines in Fig. 3), which is more efficient than sampling paths from $|D(X)|$ and imposing the boundary condition as an additional acceptance/rejection step.

4.5   Additional technical issues

Implementation of the constraint at finite $\Delta\tau$

In actual simulations $\Delta\tau$ is finite and paths are defined only at a discrete set of imaginary times. The boundary condition on the underlying continuous paths is the same, namely that the probability distribution must vanish at the axis in Fig. 3.

In Fig. 5, we illustrate how the boundary condition is imposed under the discrete representation. The "contact" point with the axis is likely to be between time slices and not well defined, i.e., $\mathcal{P}_l$ may be zero at a non-integer value of l. To the lowest order, we can terminate a partial path when its $\mathcal{P}_l$ first turns negative. That is, we still impose Eq. (1.45) in our thought experiment to generate paths. In Fig. 5, this means terminating the path at $l = n$ (point B) and thereby discarding all its future paths
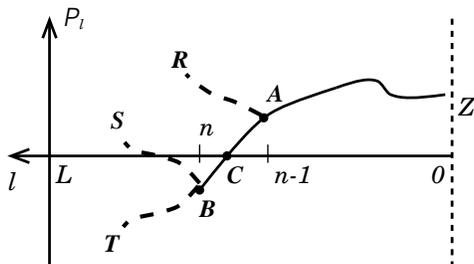
FIGURE 5. Imposition of the boundary condition at finite $\Delta\tau$. Paths are discrete. The point of contact, C (see Fig. 3), must be approximated, either by B (low order algorithm) or by interpolation between B and A (higher order).

(represented by the dashed lines 'BS...' and 'BT...').

We actually use a higher order approximation, by terminating at either $l = n - 1$ or $l = n$, i.e., either point B or point A. The probability for terminating at A is chosen such that it approaches 1 smoothly as $\mathcal{P}_{n-1} \to 0$, for example, $p_\mathrm{A} = 1/[1+\mathcal{P}_{n-1}/|\mathcal{P}_n|]$. If A is chosen, all future paths from A are discarded (represented by 'AR...' and 'AB...'); otherwise we terminate at B as above.

It is important to note that, in both approaches, the finite-$\Delta\tau$ error in imposing the boundary condition vanishes as $\Delta\tau \to 0$. The latter algorithm, which resembles the use of mirror images to construct Green's functions that vanish at a boundary, is a higher order approach, and is sometimes referred as a "mirror correction."

Population control and bias correction

In the CPMC methods walkers carry weights. Often some form of population control is applied. 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], 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].

In the ground-state algorithm, 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.

In finite-temperature calculations, we generally keep all overall rescaling factors throughout the L steps of the random walk. The global product of these rescaling factors are used in the computation of expectation values. Clearly, the overall factor only matters when we wish to combine the results from separate runs (each of L steps). It is advantageous to make the number of walkers as large as possible in finite-temperature calculations so as to reduce the appearance of these overall rescaling factors.

Another approach to eliminate — or at least to reduce — bias is to do several calculations with different (average) population sizes and extrapolate to the infinite population limit. Although this is in general less efficient, it is straightforward to implement.

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 would most likely be more productive to attempt to improve the importance sampling and the algorithm, rather than details of the bias correction scheme.

Stabilizing matrix products

In either the ground-state or finite-temperature method, repeated multiplications of the matrix $B(\vec{x})$ can lead to a numerical instability. Eventually round-off errors will dominate. The product represents an unfaithful propagation of the single-particle propagators along the path. The same instability appears in the standard AFQMC methods. The problem is controlled[12] by a numerical stabilization technique that requires the periodic re-orthonormalization of the product. Our approach in the CPMC method is similar. The differences are a result of the different sampling procedures, namely the random walk formalism here versus Metropolis sampling in the standard methods where the entire path X is stored all at once.

In the ground-state method, we use a modified Gram-Schmidt procedure to stabilize the Slater determinants $|\phi_k^{(l)}\rangle$. The procedure is simpler than

the corresponding ground-state AFQMC. 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 $|\phi\rangle$. As we mentioned at the end of Section 4.3, with importance sampling only Q needs to be kept to represent the walker.

The instability in the ground-state formalism is directly related to the collapse to a bosonic ground state. That is, if we let the propagation continue without any stabilization procedure, all single-particle orbitals each walker would become the same. In Green's function Monte Carlo 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 in configuration space. The procedure to stabilize the Slater determinants is like analytically canceling out the bosonic component. In GFMC it is possible to do such a cancellation stochastically, but the cancellation of walkers of positive and negative signs requires a large density of walkers and does not scale well with system size[26]. With Slater determinants, antisymmetry is automatically imposed in each walker. This would suggest that the sign problem is reduced in this formalism compared to GFMC. One trivial observation that 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.

In the finite-temperature method, the objects of interest in the calculation are of the form in Eq. (1.61). We store the product of B_T's at regular intervals from 1 to L. For each walker, the $B_{(\vec{x})}$ part is stored as a product of three $N \times N$ square matrices: UDV, where U is full, D is diagonal, and V is upper-triangular. As a random walker takes a step, a $B(\vec{x}_l)$ is multiplied to the left of UDV. Periodically, the product UD is rewritten as $U'D'V''$ for each walker. $V''V$ is then combined to obtain a new upper-triangular matrix V'. The process can be carried out at regular imaginary-time intervals as required by the severity of the problem.

5   Illustrative Results

In a variety of benchmark calculations[5, 3, 7, 27], on the Hubbard model in one- and two-dimensions, the constrained path Monte Carlo methods have produced very accurate estimates of the energy as well as expectation values of other observables. Practically, running on current workstations the *systematic* error of CPMC, even for small (6×6) systems, is often smaller than the *statistical* error of the standard AFQMC. The statistical errors in the latter increase rapidly with system size and inverse temperature.

(For example, at 12×12, to achieve the same statistical error as CPMC, ground-state AFQMC would have to run at least four orders of magnitude longer[28, 5]).

The CPMC methods have seen various applications[4, 27, 29] to study the Hubbard model, the periodic Anderson model, and the three-band Hubbard model. It has also been applied[30] to systems including zeolites and organic superconductors, and in nuclear physics. Here, we show a few benchmark results on the two-dimensional Hubbard model. The model provides a good test case, with both its challenging nature and the availability of certain benchmark data. Our goal here is to illustrate the general behavior of the algorithms. Many more results from benchmark studies can be found in the references.

In Table 1.1 we show results from Ref. [5]. It compares ground-state energies from zero-temperature CPMC with available data from other numerical approaches, including stochastic diagonalization (SD) [31], AFQMC, and density-matrix renormalization group (DMRG) [32] methods. The trial Slater determinant $|\psi_T\rangle$ in CPMC is either a non-interacting (free) or an unrestricted Hartree-Fock (uHF) wave function. The SD method uses Monte Carlo methods to construct a basis for approximating the ground state wave function of the system. Since an explicit basis is used, no sign problem occurs; however, an exponential growth in computing time occurs reflecting the increased effort in selecting members of the basis as system size increases. In contrast, the AFQMC method is in principle exact, as we have seen, but suffers from exponential growth in computing time as system size increases because of the sign problem. Finally, the DMRG method is a variational method that has proved very effective for one-dimensional and quasi-one-dimensional systems.

In Table 1.2, we show results for a 4×4, $U = 4$ system. The average density is chosen such that the sign problem is the most severe. (See Fig. 2.) This limits the range of temperatures where accurate calculations can be carried out with the standard AFQMC algorithm. The system hence presents a challenging test case for the CPMC method. At high T, the CPMC algorithm gives results in excellent agreement with AFQMC results[19], which are exact. At low T, it reaches convergence and leads to results in good agreement with those from $T = 0\,\mathrm{K}$ exact diagonalization[34].

6   Summary

In summary, we have reviewed quantum Monte Carlo methods, particularly auxiliary-field-based methods, for strongly correlated electron systems. We have developed a formalism here that unifies the different methods and allows for a systematic understanding of their strengths and weaknesses

TABLE 1.1. Hubbard model ground-state energies from ground-state CPMC simulations compared with available results from other approaches. The first two columns indicate system size and electron filling ($L \times L$ and $M_\uparrow M_\downarrow$, respectively). The interaction strength U is 4. The stochastic diagonalization (SD) results are from [31]; the AFQMC results are from [28]; the density-matrix renormalization group (DMRG) results on two-chains are from [33]. The statistical errors are in the last one or two digits, as indicated.

system		$|\psi_T\rangle$	E_{CPMC}	E_{SD}	E_{AFQMC}
4×4	$5\uparrow\ 5\downarrow$	free	-19.582(5)	-19.58	-19.58(1)
6×6	$13\uparrow\ 13\downarrow$	free	-42.34(2)	-40.77	-42.32(7)
6×6	$14\uparrow\ 14\downarrow$	uHF	-40.17(2)		-40.44(22)
8×8	$25\uparrow\ 25\downarrow$	free	-72.48(2)	-67.00	-72.80(6)
8×8	$27\uparrow\ 27\downarrow$	uHF	-67.46(4)		-67.55(19)
10×10	$41\uparrow\ 41\downarrow$	free	-109.55(3)		-109.7(6)
12×12	$61\uparrow\ 61\downarrow$	free	-153.43(5)		-151.4(1.4)
system		$|\psi_T\rangle$	E_{CPMC}	E_{DMRG}	
2×8	$7\uparrow\ 7\downarrow$	free	-13.067(4)	-13.0664(2)	
2×16	$14\uparrow\ 14\downarrow$	free	-26.87(2)	-26.867(3)	

TABLE 1.2. Comparison of finite-temperature and ground-state CPMC (indicated by FT and GS, respectively) with standard AFQMC and exact diagonalization (ED). $G(\mathbf{l})$ is the average Green's function $\langle c^\dagger_{i+\mathbf{l}\sigma} c_{i\sigma} \rangle$, and $P_d(\mathbf{l})$ the d-wave pairing correlation[35], at separation $\mathbf{l} = (l_x, l_y)$. Numbers in parentheses indicate statistical errors in the last digit.

β		E/N	$G(1,0)$	$G(2,2)$	$P_d(2,1)$
3	CPMC-FT	$-0.9437(8)$	$0.1631(1)$	$-0.0415(1)$	$0.0625(2)$
	AFQMC	$-0.9434(3)$	$0.1631(1)$	$-0.0418(1)$	$0.0630(3)$
6	CPMC-FT	$-0.9648(6)$	$0.1663(3)$	$-0.0470(4)$	$0.077(2)$
	AFQMC	$-0.965(3)$	$0.1662(2)$	$-0.0465(2)$	$0.083(3)$
20	CPMC-FT	$-0.977(2)$	$0.166(1)$	$-0.050(1)$	$0.078(2)$
∞	CPMC-GS	$-0.9831(6)$	$0.167(1)$	$-0.051(1)$	$0.078(2)$
	exact	-0.9838	0.167	-0.051	??

and common features. These methods have been widely applied to study lattice models for electron correlations. In addition, they have seen many applications in high-energy and nuclear physics. The standard auxiliary-field quantum Monte Carlo methods allow essentially exact calculations of ground-state and finite-temperature equilibrium properties of interacting many fermion systems. Their effectiveness, however, is severely limited by the well-known sign problem, which prevents calculations at large system sizes and low temperatures.

The recently developed ground-state and finite-temperature constrained

path Monte Carlo methods allow simulations of fermion systems without any decay of sign. The elimination of the sign decay is possible by studying the behavior of the paths in the general space of Slater determinants. In doing so, we have underlined the common feature of the sign problem under different forms of quantum Monte Carlo methods. The CPMC methods are approximate. Benchmark calculations have shown that accurate results can be obtained with simple choices of $\hat{B}_T$ or $|\psi_T\rangle$. An improved trial propagator or wave function will lead to improved results.

The CPMC methods make possible calculations under the field-theoretical formalism whose required computer time scales algebraically, rather than exponentially, with inverse temperature and system size. With the second-quantized representation, it complements the fixed-node GFMC method and the restricted path-integral Monte Carlo method[24] in real space. The algorithm automatically accounts for particle permutations and allows easy computations of both diagonal and off-diagonal expectations, as well as imaginary-time correlations. While much work is needed to study various forms of $\hat{B}_T$ or $|\psi_T\rangle$, and to understand the subtleties of the methods because of the non-orthogonal and over-complete nature of the Slater determinant space involved, we expect the method and the concept brought forth here to see many applications, and to significantly enhance the applicability of quantum simulations in interacting fermion systems.

Brief review of configuration-space methods

1 Variational Monte Carlo

This is perhaps the simplest way of performing a quantum Monte Carlo calculation for ground-state properties. It is conceptually different from most other QMC approaches. Its theoretical basis is the variational principle. A variational trial wave function $|\Psi_T\rangle$ is first constructed, which has a set of parameters that can be varied. One then optimizes these parameters so as to obtain the lowest possible variational ground-state energy. The only role Monte Carlo plays is in computing the many-dimensional integrals that are necessary for the ground-state energy.

We start with writing the trial state $|\Psi_T\rangle$ in terms of the configurations of the system

$$|\Psi_T\rangle = \sum_R |R\rangle\langle R|\Psi_T\rangle, \qquad (A.1)$$

where the coefficients $\langle R|\Psi_T\rangle = \Psi_T(R)$ depend on the set of variational parameters p_i. Here, $|R\rangle$ is a state in configuration space as defined in Sec. 2.3. For example, for M electrons of the same spin on an $N = L \times L$ two-dimensional square lattice, a configuration could be a state with electrons on lattice sites $[i_1, i_2, \cdots, i_M]$ such that $1 \leq i_k \leq N$ for $1 \leq k \leq M$ and $i_{k1} \neq i_{k2}$ for any pair of particle indices $k1 \neq k2$. The variational ground-state energy in the state $|\Psi_T\rangle$ can be expressed as

$$
\begin{aligned}
E_T &= \frac{\langle\Psi_T|\hat{H}|\Psi_T\rangle}{\langle\Psi_T|\Psi_T\rangle} \\
&= \frac{\sum_R E_L(R)\langle\Psi_T|R\rangle\langle R|\Psi_T\rangle}{\sum_R \langle\Psi_T|R\rangle\langle R|\Psi_T\rangle},
\end{aligned}
\qquad (A.2)
$$

where the local energy in configuration R is defined as

$$E_L(R) = \frac{\langle\Psi_T|\hat{H}|R\rangle}{\langle\Psi_T|R\rangle}. \qquad (A.3)$$

The many-dimensional sum in Eq. (A.2), while in general impossible to evaluate analytically, is in exactly the form of Eq. (1.7) and can therefore be evaluated by Monte Carlo. We can sample the function $f(R) = \langle\Psi_T|R\rangle\langle R|\Psi_T\rangle = |\Psi_T(R)|^2$ by the Metropolis algorithm. Given $\mathcal{M}$ independent Monte Carlo samples, we evaluate the integral for E_T by Eq. (1.9) and obtain a Monte Carlo estimate, E_{VMC}. In addition, we can easily evaluate the statistical error of E_{VMC}: $\delta E = \sigma_E/\sqrt{\mathcal{M}}$. The variance of the

local energy, σ_E, is defined as

$$\sigma_E = \sqrt{\langle E^2 \rangle - \langle E \rangle^2},$$ (A.4)

where $\langle E \rangle = E_{VMC}$ and $\langle E^2 \rangle = \frac{\sum_R E_L^2(R)}{\mathcal{M}}$. It is known[36] that σ_E, which has a lower bound of 0, is a better quantity to optimize than E_{VMC}. We have often used a combination of both quantities as indicators in our search of variational parameters.

We use the Metropolis algorithm to generate a set of $\mathcal{M}$ configurations distributed according to $f(R)$. The general Metropolis algorithm is reviewed briefly in Section 3.1 under the AFQMC algorithm. For a lattice system such as the two-dimensional Hubbard model, working in configuration space allows some further simplification. For example, the kernel K for each electron can simply propose a move uniformly to one of its 4 neighboring sites. The acceptance is then given by

$$A = \min \left[\frac{|\Psi_T(R_{\text{new}})|^2}{|\Psi_T(R_{\text{old}})|^2}, 1 \right],$$ (A.5)

where R_{new} and R_{old} are the new and old configurations, respectively. Clearly, the move would be rejected if the neighboring site is already occupied by an electron of like spins. For most choices of the trial wave function $|\Psi_T\rangle$, which could contain Slater determinants multiplied by two-body (or higher) correlation functions, the evaluation of Eq. (A.5) is simple becuase R_{new} and R_{old} differ only by one electron position. The random walk is repeated sufficient number of steps so that $\mathcal{M}$ independent walkers are obtained.

2 Green's function Monte Carlo (GFMC)

The basic premise of the GFMC method[8] is the same as that of the ground-state CPMC method. That is, it involves random walks based on Eq. (1.2). The random walk formalism utilizes first-quantized representation and is in configuration space.

For a lattice system such as the Hubbard or t-J model, one starts with the operator[37]

$$\hat{F} \equiv C - \hat{H}$$ (A.6)

in place of $e^{-\Delta\tau\hat{H}}$. The operator $\hat{F}$ can be viewed as a first order expansion of the latter. But for discrete systems, projection with $\hat{F}$ leads to the exact ground-state wave function as well, because the system has a spectrum that is bounded. In this case, it is often advantageous to use the operator $\hat{F}$ in the projection, because of the sparseness of its matrix representation. The constant C is chosen to ensure that all diagonal elements are positive and

the spectrum of $\hat{F}$ is positive. Similar to ground-state AFQMC or CPMC, if an initial state $|\Psi^{(0)}\rangle$ has any overlap with the ground state $|\Psi_0\rangle$ of $\hat{H}$, the process

$$|\Psi^{(n)}\rangle = \hat{F}^n |\Psi^{(0)}\rangle \tag{A.7}$$

will lead to $|\Psi_0\rangle$ at large n. The Green's function Monte Carlo method realizes the above process by a Monte Carlo random walk in configuration space.

In order to improve efficiency of the random walk process, one more mathematical manipulation of Eq. (A.7) is necessary. This is done by introducing an operator $\tilde{\hat{F}}$ whose matrix elements are related to those of $\hat{F}$ by a similarity transformation:

$$\tilde{\hat{F}}(R', R) \equiv \langle \Psi_T | R' \rangle \, \langle R' | \hat{F} | R \rangle \, \frac{1}{\langle \Psi_T | R \rangle}. \tag{A.8}$$

The stochastic realization of Eq. (A.7) is actually with $\tilde{\hat{F}}$ instead of $\hat{F}$. While mathematically equivalent, the use of $\tilde{\hat{F}}$ can significantly reduce the fluctuation of the Monte Carlo process if $|\Psi_T\rangle$ is a reasonable approximation of $|\Psi_0\rangle$. As we have seen with the constrained path Monte Carlo methods, this is the idea of importance sampling.

The program is then to start with a set of walkers distributed according to $\langle \Psi_T | R \rangle \langle R | \Psi_T \rangle$. This can be accomplished with a variational Monte Carlo calculation prior to GFMC. The walkers then undergo random walks in R-space. For each walker, denoted by R, a step in the random walk means randomly selecting and moving to a new position R' with probability $\tilde{\hat{F}}(R', R) / \sum_{R'} \tilde{\hat{F}}(R', R)$. Note that the sparsity in the structure of $\hat{H}$ means that there are only $\mathcal{O}(N)$ possible R''s, whose corresponding probabilities can be computed. Because the overall normalization $\sum_{R'} \tilde{\hat{F}}(R', R)$ is not a constant, each walker carries a weight which fluctuates, or a branching scheme is introduced to allow the total number of walkers to fluctuate, or both.

The ground state energy is given exactly by the so-called *mixed estimate* (see Eq. (1.4))

$$
\begin{aligned}
E_0 &= \frac{\langle \Psi_T | \hat{H} | \Psi_0 \rangle}{\langle \Psi_T | \Psi_0 \rangle} \\
&= \frac{\sum_R E_L(R) \langle \Psi_T | R \rangle \langle R | \Psi_0 \rangle}{\sum_R \langle \Psi_T | R \rangle \langle R | \Psi_0 \rangle},
\end{aligned} \tag{A.9}
$$

where $E_L(R)$ is defined in Eq. (A.3). After a sufficient number of steps, the walkers are distributed according to $\langle \Psi_T | R \rangle \langle R | \Psi_0 \rangle$. E_0 can therefore be computed from such walkers as the (weighted) average of $E_L(R)$ with respect to walker positions R. We denote this Monte Carlo estimate of E_0 by E_{GFMC}.

Expectation values of operators other than $\hat{H}$ can also be computed from the mixed estimate. However, if the operator does not commute with $\hat{H}$, this estimate is *not* exact. This is an important distinction which requires careful analysis of the bias in the results due to $|\psi_T\rangle$. There exist ways to improve upon the mixed estimate and to possibly extract exact estimates of expectation values, but we will not discuss them here. The computation of *off-diagonal expectations* in general presents difficulty in GFMC, because it is difficult to sample two groups of random walkers $\{R'\}$ and $\{R\}$ whose overlap is well behaved statistically to allow for computation of matrix elements such as $\langle\Psi_0|R'\rangle\,\langle R'|\hat{O}|R\rangle\,\langle\Psi_0|R\rangle$.

The sign problem occurs in GFMC if the particles are fermions (e.g., the Hubbard model) or any of the matrix elements in Eq. (A.8) are negative (e.g., anti-ferromagnetic coupling of quantum spins). The fixed-node method eliminates the sign decay by imposing the nodal structure of the trial wave function $|\psi_T\rangle$ on the projected wave function. For continuum systems, such as atoms or molecules, the fixed-node approximation is straightforward to implement[21, 22]. One would simply modify the transition probability in Eq. (A.8), setting all negative matrix elements to zero. At the limit of $\Delta\tau \to 0$, this ensures that the random walk process solves the Schrödinger equation under the boundary condition that $\Psi_0^{\text{FN}}(R) = 0$ for all configurations R that satisfy $\Psi_T(R) = 0$, where $|\Psi_0^{\text{FN}}\rangle$ is then the approximate solution for $|\Psi_0\rangle$ under the fixed-node approximation.

For discrete systems, however, the implementation of the fixed-node approximation is more subtle. Unlike in continuum systems, the random walk does *not* become continuous at the limit of $\Delta\tau \to 0$. A method was developed a few years ago[23] which successfully generalized the continuum fixed-node approximation to lattice fermion systems. In addition to setting the transition probability to zero when the matrix element is negative in Eq. (A.8), one also modifies diagonal matrix elements $\tilde{\tilde{F}}(R, R)$ when the corresponding configuration R is in the vicinity of the node defined by $|\Psi_T\rangle$. R is in the vicinity of the trial node if there exists at least one negative $\tilde{\tilde{F}}(R', R)$, i.e., at least one "hop" which would take the walker across the trial node. For each such configuration R, an extra part is added to the diagonal matrix element. The extra part serves as an effective potential. It is defined as $\sum_{R'}^{-} \tilde{\tilde{F}}(R', R)$, where the negative sign means that the sum is only over those configurations R' that lead to negative elements $\tilde{\tilde{F}}(R', R)$. The idea of this effective potential is in fact related to the discussion in the first part of Section 4.5.

3   References

[1] 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).

[2] 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).

[3] Shiwei Zhang, in *Quantum Monte Carlo Methods in Physics and Chemistry*, Ed. M. P. Nightingale and C. J. Umrigar, NATO ASI Series (Kluwer Academic Publishers, 1999).

[4] Shiwei Zhang, J. Carlson, and J. E. Gubernatis, Phys. Rev. Lett. **78**, 4486 (1997).

[5] Shiwei Zhang, J. Carlson, and J. E. Gubernatis, Phys. Rev. Lett. **74**, 3652 (1995); Phys. Rev. B **55**, 7464 (1997).

[6] J. Carlson, J. E. Gubernatis, G. Ortiz, and Shiwei Zhang, Phys. Rev. B, **59**, 12788 (1999).

[7] Shiwei Zhang, Phys. Rev. Lett. **83**, 2777 (1999); *Comp. Phys. Comm.*, **127**, 150 (2000).

[8] 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).

[9] See, e.g., M. H. Kalos and P. A. Whitlock, *Monte Carlo Methods*, Vol I, Wiley, 1986.

[10] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, Phys. Rev. D **24**, 2278 (1981).

[11] G. Sugiyama and S.E. Koonin, Ann. Phys. (NY) **168**, 1 (1986).

[12] 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).

[13] S. Sorella, *et. al.*, Int. Jour. Mod. Phys. B **1** 993 (1988).

[14] See, e.g., D. R. Hamann and S. B. Fahy, Phys. Rev. B **41**, 11352 (1990).

[15] J. Hubbard, Phys. Rev. Lett. **3**, 77 (1959).

[16] J. E. Hirsch, Phys. Rev. B **31**, 4403 (1985).

[17] J. W. Negele and H. Orland, *Quantum many-particle systems* (Adison-Wesley, New York, 1987), chap. 7.

[18] S. Rombouts, K. Heyde, and N. Jachowicz, Phys. Lett. A **242**, 271 (1998).

[19] R. T. Scalettar, private communication.

[20] S. B. Fahy and D. R. Hamann, Phys. Rev. Lett. **65**, 3437 (1990); Phys. Rev. B **43**, 765 (1991).

[21] J. B. Anderson, J. Chem. Phys. **63**, 1499 (1975); **65**, 4122 (1976).

[22] 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).

[23] 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).

[24] 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).

[25] C. J. Umrigar, M. P. Nightingale, and K. J. Runge, J. Chem. Phys. **99**, 2865 (1993); M. P. Nightingale, *Quantum Monte Carlo Methods in Physics and Chemistry*, Ed. M. P. Nightingale and C. J. Umrigar, NATO ASI Series (Kluwer Academic Publishers, 1999).

[26] 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).

[27] Shiwei Zhang and E. C. Allman, in *Computer Simulation Studies in Condensed Matter Physics XII*, Eds. D. P. Landau, S. P. Lewis, and H. B. Schuttler (Springer Verlag, Heidelberg, Berlin, 1999).

[28] N. Furukawa and M. Imada, J. Phys. Soc. Jpn. **61**, 3331 (1992).

[29] See, *e.g.*, M. Guerrero, J.E. Gubernatis, and Shiwei Zhang, Phys. Rev. B **57**, 11980 (1998); J. Bonca and J. E. Gubernatis, Phys. Rev. B **58**, 6992 (1998); M. Guerrero, G. Ortiz, J. E. Gubernatis, Phys. Rev. B **59**, 1706 (1999); Shiwei Zhang and E. C. Allman, in *Computer Simulation Studies in Condensed Matter Physics XII*, Eds. D. P. Landau, S. P. Lewis, and H. B. Schuttler (Springer Verlag, Heidelberg, Berlin, 1999).

[30] Otto F. Sankey, A. A. Demkov, and T. Lenosky, Phys. Rev. B **57**, 15129 (1998); S. Mazumdar, S. Ramasesha, R.T. Clay, D.K. Campbell, Phys. Rev. Lett. **82**, 1522 (1999); K. E. Schmidt and S. Fantoni, Phys. Lett. B **446**, 99 (1999).

[31] H. De Raedt and M. Frick, Phys. Reports, **231**, 109 (1993).

[32] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992); Phys. Rev. B **48**, 10345 (1993).

[33] R. M. Noack, private communication.

[34] 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).

[35] The definition we used for the d-wave pairing correlation is: $P_d(\mathbf{l}) \equiv \langle \Delta_d^\dagger(\mathbf{l})\Delta_d(\mathbf{0})\rangle$, with $\Delta_d(\mathbf{l}) = \sum_\delta f(\delta)\,(c_{\mathbf{l}\uparrow}c_{\mathbf{l}+\delta\,\downarrow} - c_{\mathbf{l}\downarrow}c_{\mathbf{l}+\delta\,\uparrow})$, where δ is $(\pm 1, 0)$ and $(0, \pm 1)$, and $f(\delta)$ is 1 when $\delta = (\pm 1, 0)$, and -1 otherwise.

[36] C. J. Umrigar *et. al.*, *Phys. Rev. Lett.* **60** 1719 (1988).

[37] N. Trivedi and D. M. Ceperley, *Phys. Rev. B* **41**, 4552 (1990)