

Algorithm for branching and population control in correlated sampling

Siyuan Chen ^{1,*}, Yiqi Yang ^{1,‡}, Miguel Morales,² and Shiwei Zhang ²

¹Department of Physics, College of William & Mary, Williamsburg, Virginia 23187, USA

²Center for Computational Quantum Physics, Flatiron Institute, 162 5th Avenue, New York, New York, USA



(Received 4 August 2023; accepted 9 November 2023; published 22 November 2023)

Correlated sampling has wide-ranging applications in Monte Carlo calculations. When branching random walks are involved, as commonly found in many algorithms in quantum physics and electronic structure, population control is typically not applied with correlated sampling due to technical difficulties. This hinders the stability and efficiency of correlated sampling. In this paper, we study schemes for allowing birth/death in correlated sampling and propose an algorithm for population control. The algorithm can be realized in several variants depending on the application. One variant is a static method that creates a reference run and allows other correlated calculations to be added *a posteriori*. Another optimizes the population control for a set of correlated, concurrent runs dynamically. These approaches are tested in different applications in quantum systems, including both the Hubbard model and electronic structure calculations in real materials.

DOI: [10.1103/PhysRevResearch.5.043169](https://doi.org/10.1103/PhysRevResearch.5.043169)

I. INTRODUCTION

Monte Carlo (MC) methods [1] are widely used in engineering [2], physics [3–7], reliability theory [8], queuing theory [9], finance [10], etc. In many applications, MC methods are used to solve the underlying differential or integral equations of the processes or systems in a stochastic way. This is often achieved by random walks that are typically constructed by Markov chains, where the transition probability is designed to reflect the characteristic of the system.

In a broad class of algorithms, a random walker gains a weight that fluctuates as the random walk proceeds, leading to an exponentially increasing variance [11]. To solve this problem, the random walk is designed to be carried out by multiple walkers simultaneously with branching [1,11–13]. Branching random walks duplicate the walkers with large weights and eliminate walkers with small weights under certain probabilities. In addition, to avoid the unbounded fluctuation of the total population, population control is generally applied [11,13,14]. During the random walk, samples are collected, and the weighted average of all the samples gives a MC representation of the target quantity. Branching random walk algorithms are widely applied in the study of ground-state [4,15,16] (and even finite-temperature [17]) properties of interacting quantum many-body systems.

In quantum physics or chemistry, the difference between two closely related systems is often of great interest and

importance. Examples include binding energies and gaps, redox potential, and reactions. In addition, other observables, such as forces in molecules and solids [18,19] and order parameters in lattice models [20], can be obtained by computing energy derivatives based on finite difference. Quantities computed from MC methods naturally contain statistical errors. When taking their differences, the statistical errors are propagated into the target quantity. If the differences are small, the signal-to-noise ratio would be low.

Correlated sampling [1,21] is generally applied in the MC calculation of the energy differences and gradients to reduce the statistical noise. A considerable amount of work has been done to develop and apply correlated sampling to various forms of quantum MC methods, including variational MC (VMC) calculation of potential energy curves [22] and particle-hole excitations [23]; VMC optimization of trial wave functions [24,25]; Green's function MC (GFMC) and diffusion MC (DMC) calculation of potential energy curves and bond lengths [26], H_3 cations [27], the dipole moment of LiH [28], and forces as well as polarizabilities [29]; and phaseless auxiliary field quantum MC (Ph-AFQMC) calculation of bond dissociation energies, ionization potentials, and electron affinities [30]. Recent advances with automatic differentiation (AD) offers the promise for computing gradients more accurately and conveniently without correlated sampling (see, e.g., Refs. [31,32]). However, there are important situations which do not lend themselves readily to AD, for example, in the computation of ionization potentials or gaps or excitations.

Combining branching random walks and correlated sampling is not straightforward. The difficulty originates from branching decorrelating the random walks in different systems. Even when the systems are very close, branching can eventually cause the random walks to deviate and the correlations between them to then quickly deteriorate. To prolong correlation, typically correlated sampling calculations forgo branching and population control [30], instead simply keeping

*These two authors contributed equally to this work.

†schen24@wm.edu

‡yyang25@wm.edu

the weights of walkers. This exacerbates the asymptotic instability of the calculation. To our knowledge, population control algorithms have not been discussed in correlated sampling calculations, nor has there been a systematic study of how it can affect the efficiency, accuracy, and stability of the calculation.

In this paper, we propose an algorithm for branching and population control in correlated sampling, which ensures stability and thus prolongs the correlation time. Our correlated sampling schemes with branching and population control are then tested in ground-state AFQMC calculations, in both the Hubbard model and a periodic solid system. The statistical fluctuation is found to be significantly reduced compared with that of the runs without correlated sampling, and the correlations between the runs are sustained for significantly longer compared with correlated sampling without population control. Although our population control algorithms are implemented and studied within AFQMC, they apply to any correlated sampling calculations that involve branching random walks, as further discussed below.

The rest of the paper is structured as follows. In Sec. II, we provide a brief overview of correlated sampling as well as branching and population control. This is followed in Sec. III by a description of our population control algorithms, together with discussions on the metrics to quantify the performance of correlated sampling. The test results in the Hubbard model and real materials are shown in Sec. IV. Then we conclude in Sec. V.

II. MOTIVATION

The expectation value of a quantity $A(x)$ measured from a probability distribution of x is expressed by an integral over the probability distribution function (PDF) $p(x)$:

$$\langle A \rangle = \frac{\int p(x)A(x)dx}{\int p(x)dx}, \quad (1)$$

where $p(x)$ can be, for example, the partition function \mathcal{Z} in statistical physics or the probability density $|\Psi|^2$ for a quantum state. The variance of A describes the spread of A for different x values:

$$\text{var}(A) = \frac{\int p(x)[A(x) - \langle A \rangle]^2 dx}{\int p(x)dx}. \quad (2)$$

If the integral in Eq. (1) is in high-dimensional space, one typically uses the MC method to evaluate it.

MC samples form a stochastic representation of the PDF $p(x) \rightarrow \sum_i w_i \delta(x - x_i)$. The MC estimation of quantity A is then given by

$$A_{\text{MC}} = \frac{\sum_{n=1}^N w_n A(x_n)}{\sum_{i=1}^N w_n}, \quad (3)$$

which approaches the expectation value of A as the number of samples N increases: $|\langle A \rangle - A_{\text{MC}}| \propto 1/\sqrt{N}$ [3,33] according to the law of large numbers.

Sometimes we are interested in the difference between two quantities A and B . When computing the difference $A - B$ via the MC method, its variance $\text{var}(A - B)$ can be expressed as follows [1,34]:

$$\text{var}(A - B) = \text{var}(A) + \text{var}(B) - 2\text{Cov}(A, B). \quad (4)$$

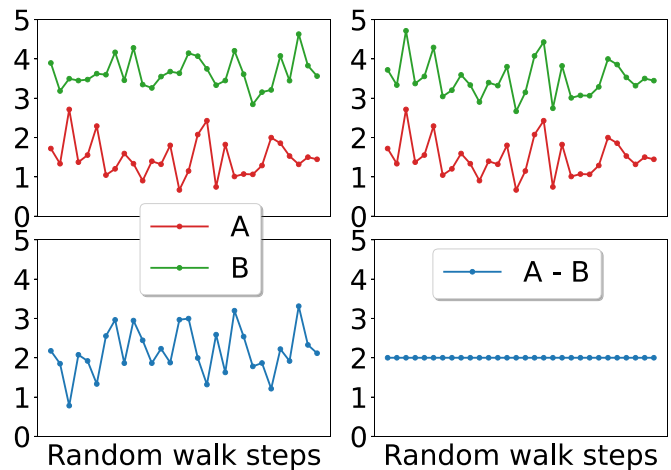


FIG. 1. Schematic illustration of the effect of correlated sampling in computing the differences. The left panels show uncorrelated calculations, while the right panels show correlated sampling. Horizontal axes give Monte Carlo time, i.e., index of the branching random walk steps. The top panels give snapshot measurements in two separate calculations for A and B , while the bottom panels show the computed difference.

Here,

$$\text{Cov}(A, B) = \langle AB \rangle - \langle A \rangle \langle B \rangle \quad (5)$$

denotes the covariance of the two quantities. If the two quantities are independently sampled—in other words, their covariance is zero—the variance of the difference would be the sum of the variance of A and B , as illustrated in the left column of Fig. 1. However, if one can correlate the sampling of the two quantities, for example, by using the same random number stream, to maximize the covariance, the statistical error of their difference will be minimized. The ideal case occurs when correlation is perfect, $2\text{Cov}(A, B) = \text{var}(A) + \text{var}(B)$, and $A - B$ is a constant so that the difference is obtained with no statistical fluctuation, as shown in the right panel of Fig. 1. This motivates the idea of correlated sampling.

In random walks that involve branching and population control, for instance, the AFQMC method [16], the implementation of correlated sampling is not straightforward. This is due to an inherent conflict between correlated sampling and branching random walks, as illustrated in Fig. 2. Panel (b) shows the energy fluctuations for three correlated runs, without population control. Early on in the calculations (a small number of random walk steps), the correlation is strong, and the runs lead to much reduced fluctuations in the estimated differences. This mode has been very useful in quantum chemistry calculations [30,35]. As time goes on, however, the weights will fluctuate a lot more in the absence of branching and population control, causing larger statistical noise in each measured result from Eq. (3), hence also larger noise in the differences. Doing individual and separate branching and population control in each run will break down the correlation rapidly, as shown in panel (c). Since the systems being correlated are not identical, their random walks branch with different rates. As soon as population control leads to different branching decisions (e.g., extra walker or elimination of a

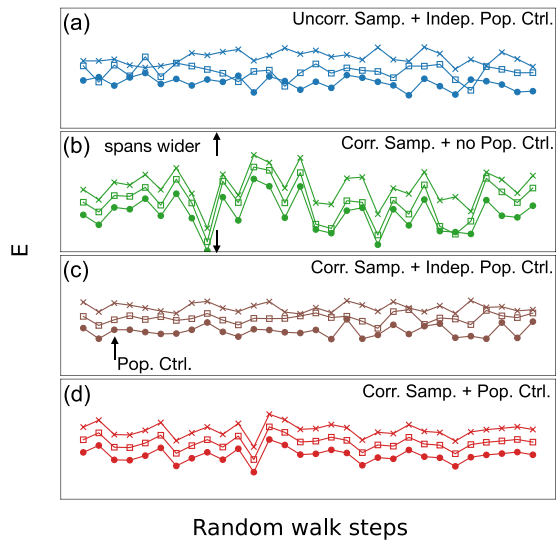


FIG. 2. Schematic illustration of the conflict between correlated sampling and branching and population control. (a) Uncorrelated calculations of three closely related systems incur independent noise and hence large fluctuations in estimating the differences. (b) In correlated sampling without population control, runs remain correlated as random walk proceeds, but large fluctuations in weights make the statistical error grow rapidly. (c) In correlated sampling with independent population control, separate branching can cause the runs to decorrelate quickly. (d) Algorithm for branching and population control in correlated sampling aims to maximize efficiency and prolong correlation.

walker in one run), a correlation between the different runs will be lost. The idea of this paper is to incorporate branching and population control in a correlated manner between the

different runs [panel (d)], which allows the calculations to stay well correlated for a much longer time.

III. BRANCHING AND POPULATION CONTROL IN CORRELATED SAMPLING

In this section, we introduce our algorithm for branching and population control in correlated sampling. The algorithm is described in a flowchart in Fig. 3. We will assume that M systems (runs) are correlated, each with a population size N . The weights of run m are given by $\{w_{n,m}\}$ with walker index $n = 1, \dots, N$. Walkers from different runs with index n are correlated. As the walkers are propagated and branched, this set of correlated walkers must follow the same branching path to delay decorrelation.

To implement this, a reference weight is derived from each set of correlated walkers through some user-defined function:

$$\tilde{w}_n \equiv f(w_{n,1}, w_{n,2}, \dots, w_{n,M}). \tag{6}$$

A universal choice of f is taken and applied to all n , but as we further discuss below, the form of the function is flexible. For instance, f could be the max function to take the maximum weight in each correlated walker set, or the average function, or it could take the weight of a specific system (fixed m , but see below), etc. The reference weights $\{\tilde{w}_n\}$ are then fed into a normal branching and population control routine, i.e., one that is conventionally used in the branching random walk without correlated sampling. This outputs a set of branching and population control decisions on the reference weights, in the form of a new set of weights $\{\tilde{w}'_n\}$. We then update the actual walker weights in each run by keeping constant the ratio

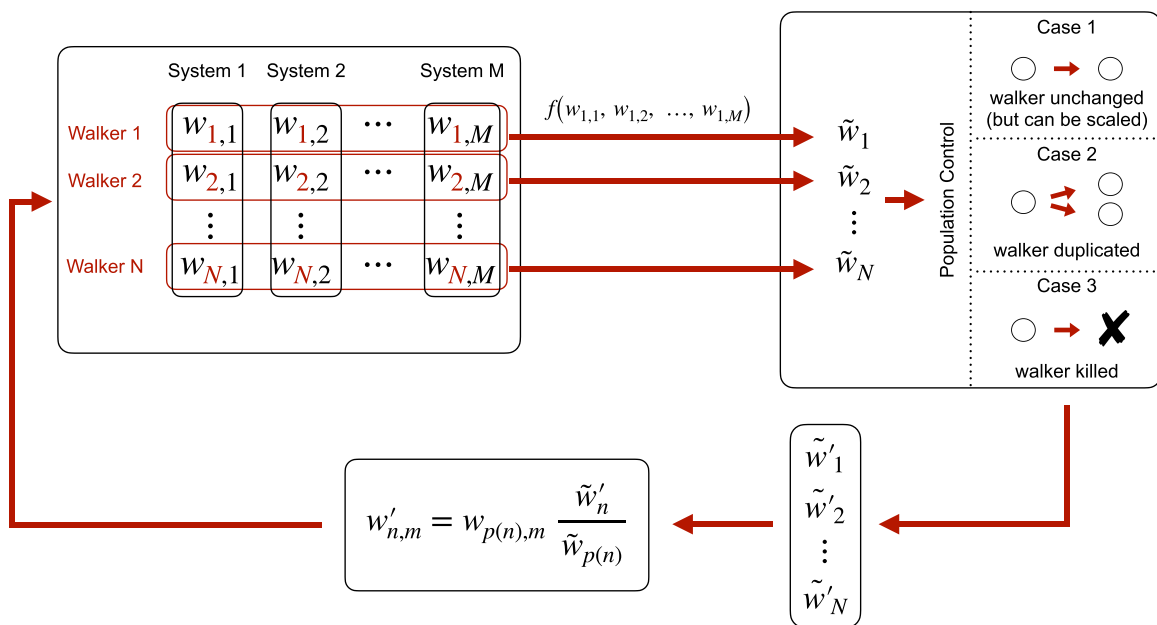


FIG. 3. Flowchart of the branching and population control algorithm in correlated sampling. Walker weights prior (subsequent) to population control are denoted by $w_{p(n),m}$ ($w'_{n,m}$), where $p(n)$ is the index of the parent walker of the n th walker from the population control decision, with M correlated systems/runs and N walkers for each system. The reference weights $\{\tilde{w}_{p(n)}\}$ are obtained via f and fed into population control, which gives output decisions represented by a new set of reference weights $\{\tilde{w}'_n\}$. The individual weights are then updated accordingly.

TABLE I. Characteristics of the three implementations of the branching and population control algorithm in correlated sampling.

	Add runs	Synchronous run	No. reference runs
Dynamic	N	Y	>1
Static	Y	N	1
Semidynamic	Y	Y	>1

between the new and old weights:

$$w'_{n,m} = w_{p(n),m} \frac{\tilde{w}'_n}{\tilde{w}_{p(n)}}, \quad (7)$$

where $p(n)$ is the index of the parent walker for n from the population control decision.

Various realizations exist for the branching and population control algorithm, in part depending on the choice of f , which can lead to different implementations that emphasize different aspects of the computation or programming and can have relative advantages and disadvantages:

(1) Carry out all the correlated runs concurrently, where the reference walker weights are derived from each correlated walker set on the fly. In this way, the branching and population control are optimized dynamically, and we refer to this as the *dynamic* realization. All the choices for f mentioned above are possible within this scheme.

(2) Conduct a single MC run as the reference run and record the reference weights, which store the branching and population control decisions made therein. Other correlated runs (child runs) are then added freely afterward, following the reference weight and branching decisions in the reference run. We refer to this as the *static* realization. In this scheme, the function f must be based on a single run index m which is the reference run.

(3) Combining (1) and (2), one can record the branching and population control decisions from a dynamic run, then add subsequent runs where the decisions are applied. We refer to this as a *semidynamic* realization.

We summarize the characteristics of the three methods in Table I.

In the static realization in (2), the function f has limited choices since we only have knowledge of the walker weights in the reference system. We also need to consider that walkers with zero (or very small) weights need to be carried through and not population controlled because other unknown runs might have nonzero (or much larger) weights. Our choice for the reference weight for the static realization of population control is

$$\tilde{w}_n = \begin{cases} w_{n,1}, & w_{n,1} > w_{\text{th}} \\ w_{\text{safe}}, & \text{otherwise,} \end{cases} \quad (8)$$

where $m = 1$ labels the reference run, w_{th} is a small positive number, and w_{safe} is a positive number that ensures a walker with $w_{n,1} = w_{\text{th}}$ or smaller will yield exactly one copy in the population control. Alternatively, for general population control algorithms, one can simply use $\tilde{w}_n = w_{n,1}$ and manually prevent branching in all runs for this walker when $w_{n,1} \leq w_{\text{th}}$.

The goal of a correlated branching and population control algorithm is to prolong the time during which the systems remain well correlated. Even with branching and population control, the correlation will degrade eventually for each set of correlated walkers and for the entire systems. Small differences in walker weights in a set of correlated walkers accumulate over time to become large differences, and the probability of a rare event (e.g., one walker in a correlated run is killed while others retain a significant weight) will also increase.

We monitor the level of correlation with a number of metrics. For example, we can take the standard deviation of the walker weight for each correlated walker group and normalize it with respect to its mean as a measure of the relative weight fluctuation. The average across the entire population:

$$Q = \sum_{n=1}^N \sqrt{\frac{M \sum_m w_{n,m}^2}{(\sum_m w_{n,m})^2} - 1} / N, \quad (9)$$

then gives a simple global indicator. A Q value larger than a few, for instance, would indicate a significant dispersion of the weight among the correlated runs. This way we could examine the quality of the correlation on the fly (as a function of random walk steps) and stop the calculation when it reaches a certain level. The walker weights are the only quantities needed in this process.

To maximize utilization of the well-correlated period, we minimize the equilibration time by starting correlated sampling with an equilibrated walker set from one of the correlated systems. This is called the *preliminary equilibration scheme* and was proposed by Shee *et al.* [30] previously. To improve statistics and reduce error bars, it is typically more efficient to repeat the entire procedure multiple times with different random seeds. If the correlation between the correlated runs remains significant (e.g., as indicated by a small Q metric value) beyond the autocorrelation time of the measurements in the individual runs, then we can make multiple measurement blocks to collect more data.

In the dynamic realization, instead of the preliminary equilibration scheme, we reset the concurrent runs periodically to repeat the sampling, achieving an effective restart at every reset which accomplishes the same goal. Note that the sampling needs to be equilibrated again (re-equilibration) after each reset. Specifically, a reset period is chosen before the calculation; once the propagation imaginary time reaches this specified reset period, we reset the run by copying the walkers of one of the correlated systems (we call it the reference system, and it can be chosen arbitrarily) to other systems. In Fig. 4, we show a schematic of the reset method for dynamic branching and population control in correlated sampling, using Ph-AFQMC as an example. Note that, in both schemes, a reweighting is needed, as the underlying importance functions are typically different between the two systems. (In Ph-AFQMC, this is reflected by the different trial wave functions $|\Psi_{T1}\rangle$ and $|\Psi_{T2}\rangle$ for Systems 1 and 2, respectively.) When the population of System 2 is replaced by the surrogate population from System 1, the probability density is modified, hence the jumps in the computed expectations in the




FIG. 4. Schematic illustration of the reset method used in dynamic correlated sampling. The computed energies are shown for two systems as a function of random walk time steps (projection time τ in Ph-AFQMC) starting from equilibration. System 1 is chosen (arbitrarily) as the reference. At reset, the population of System 2 is replaced by that of System 1, which causes the jumps in the computed expectation value in System 2. The jump magnitude reaches a constant when both systems have equilibrated. Extra steps are needed for re-equilibrating System 2 after every reset.

illustration. The dynamic reset method provides a natural way to continue the correlated sampling runs to improve statistics, more in the spirit of regular branching random walks.

IV. RESULTS

In this section, we show test results in two very different classes of problems. We use the AFQMC method [15,16] applied to the ground state of the Hubbard model and to solid silicon. They provide a spectrum of tests in real-world problems which are sufficiently challenging and of strong current research interest in quantum physics. Both static and dynamic realizations of the population control algorithm will be tested. The different types of quantum MC algorithms (a lattice problem with local interactions in the former and a continuum problem with long-range interactions in the latter) allow us to do so under different, general conditions. Rather than obtaining new physical results, our focus here is on testing, illustrating, and studying the population control algorithm and correlated sampling. The physical quantities we will compute are simple and well understood, so that we can benchmark the results straightforwardly. To cleanly separate out the effects of correlated sampling and population control, we use the same type of trial wave functions throughout the set of computations

The implementation details are very different for the two classes of problems we consider here as well as the behaviors of the algorithms. We refer readers to the literature for further discussions. For understanding the tests below, it suffices to recall that the AFQMC method [3,34,36] projects out the ground state of a Hamiltonian \hat{H} by applying $e^{-\tau\hat{H}}$, to yield

$$|\psi(\tau)\rangle \propto \sum_n w_n^{(\tau)} |\phi_n^{(\tau)}\rangle / \langle \psi_T | \phi_n^{(\tau)} \rangle, \quad (10)$$

where τ denotes projection time and $|\psi(\tau)\rangle \xrightarrow{\tau \rightarrow \infty} |\psi_G\rangle$ the desired ground state. The projection is realized by a branching random walk whose time step is represented by τ , in which n is the index of the random walker, and w_n is the weight. The random walker $|\phi_n\rangle$ lives in a space of Slater determinants determined by the details of the problem, and $\langle \psi_T | \phi_n \rangle$ is an importance function defined by a known trial wave function $|\psi_T\rangle$. The random walkers are propagated in the manifold of Slater determinants via a set of random auxiliary fields, \mathbf{x} : $B(\mathbf{x})|\phi^{(\tau)}\rangle \rightarrow |\phi^{(\tau')}\rangle$ (with $\tau' = \tau + \Delta\tau$ being the next step in the random walk). In this framework, it is convenient and straightforward to think of correlated sampling as correlating the multidimensional auxiliary fields \mathbf{x} . (See Ref. [15] for a way to cast DMC [4] in this framework.) We will correlate multiple runs for different but related Hamiltonians \hat{H}_m , with different predefined $|\psi_{T,m}\rangle$, producing $|\psi_m(\tau)\rangle$. It is important to note that, in the context of the population control algorithm, the weights $w_{n,m}^{(\tau)}$ are the only objects we will need to deal with.

A. Hubbard model

We study the ground state of the Hubbard model on a square lattice [37], one of the fundamental models in quantum many-body physics:

$$\hat{H} = -t \sum_{\langle ij \rangle \sigma} \hat{c}_{i\sigma}^\dagger \hat{c}_{j\sigma} + U \sum_i \hat{n}_{i\uparrow} \hat{n}_{i\downarrow}, \quad (11)$$

where t is the hopping parameter, U the on-site interaction, $\hat{c}_{i\sigma}^\dagger$ ($\hat{c}_{j\sigma}$) creates (annihilates) an electron with spin σ at site i (j), and $\langle \dots \rangle$ refers to the nearest-neighbor hopping. We will focus on computing the double occupancy

$$D \equiv \langle \psi_G | \sum_i \hat{n}_{i\uparrow} \hat{n}_{i\downarrow} | \psi_G \rangle, \quad (12)$$

which is directly proportional to the interaction energy and which provides an important measure to the nature of electron correlation. (Note that the physical double occupancy on each lattice site is given by D divided by the number of lattice sites.) Using the Hellman-Feynman theorem:

$$D = \langle \psi_G | \frac{d\hat{H}}{dU} | \psi_G \rangle = \frac{dE_G}{dU}, \quad (13)$$

where the derivative on the right-hand side can be evaluated by finite difference:

$$\left. \frac{dE_G}{dU} \right|_U \doteq \frac{E_G(U + \delta U) - E_G(U - \delta U)}{2\delta U}. \quad (14)$$

Thus, the double occupancy D for \hat{H} can be computed by the ground-state energy difference between two Hamiltonians \hat{H}_1 and \hat{H}_2 , with U in Eq. (11) replaced by $U_1 = U + \delta U$ and $U_2 = U - \delta U$, respectively.

We test the dynamic realization for branching and population control in correlated sampling to compute the double occupancy D . We first study a small system of lattice size 4×4 , $U/t = 4.0$, at half-filling with periodic boundary conditions along both x and y directions. The trial wave functions were generated from the generalized Hartree-Fock method with effective on-site interaction U_{eff} [38,39]. The benchmark

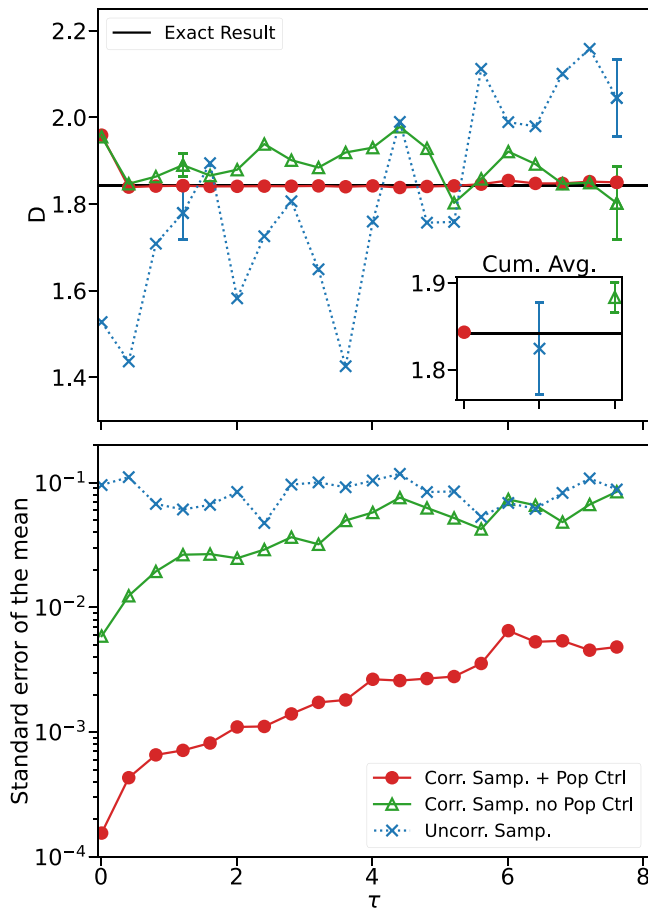


FIG. 5. Double occupancy D as a function of projection time obtained from uncorrelated, correlated sampling with no population control and correlated sampling with population control calculations. These error bars are estimated from 40 resets (or 40 repeated runs for the uncorrelated calculations). In each reset or each run, a (re-)equilibration time of $\tau = 1.2$ is discarded and not shown in the plot. The random walk step size is $\tau = 0.01$. In the upper panel, the statistical error bar is only shown at $\tau = 1.21$ and 7.61 for clarity, but in the lower panel, it is plotted for each data point. The cumulatively averaged results are shown in the inset. The horizontal black solid line in the upper panel is from exact diagonalization. The system is $L_x = L_y = 4$, $U/t = 4$, at half-filling with periodic boundary conditions. D is computed using finite difference as in Eq. (14), with $\delta U = 0.01$. Note that the vertical axis of the lower panel is on log scale.

results are shown in Fig. 5. Here, each random walk step corresponds to an increment of τ of $\Delta\tau = 0.01$. Results at the same imaginary time relative to the reset points are averaged to estimate the statistical error. We can see that results are all consistent with the exact answer. In the uncorrelated run, the statistical fluctuation is large throughout the imaginary time propagation. While correlated sampling without population control shows a significantly reduced error bar at small τ , this reduction deteriorates as τ increases, as expected. Population controlled correlated sampling is seen to significantly extend the correlation, reducing the statistical error by more than a factor of 10 throughout the entire convergence window in τ .

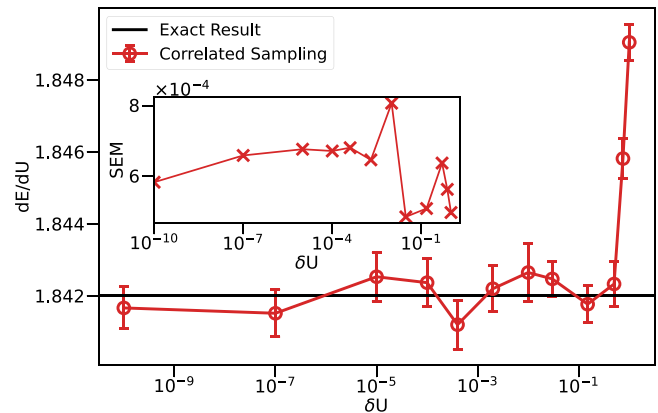


FIG. 6. Accuracy of correlated sampling calculations of the energy derivative vs finite-difference size. Results are shown for a range of δU across many decades, for the same system as in Fig. 5. The main figure is the computed result, and the inset shows the estimated error bar. Computational cost is the same for each calculation. Note the log scale on the x axes. The exact result is obtained from exact diagonalization.

We next examine the behavior of correlated sampling vs the proximity of the correlated systems, which is specified by δU in the present system. In uncorrelated runs, the statistical error grows in inverse proportion to δU , as given by Eq. (14). The computed energy derivative dE/dU with correlated sampling is shown in Fig. 6, for a range of δU values spanning many decades. The results are surprisingly robust. Only at very large δU values are systematic biases seen, indicating the breakdown of the finite difference formula in approximating the derivative. The statistical error is seen to be essentially independent of δU . This is highly advantageous, as very small values of δU can be used to ensure that the finite difference yields an accurate estimate of the derivative, without increase in computational cost.

The algorithm works equally well in larger system sizes, and major computational efficiency gain is seen in realistic system sizes in state-of-the-art calculations [38,40]. In Table II, we show a simple comparison of correlated sampling with population control vs uncorrelated calculations of the double occupancy. For both types of calculations, δU was fixed at 0.01. This is a reasonably safe choice to ensure that the finite difference approximation in Eq. (14) remains reliable. The systems have $U/t = 8$, which are more strongly

TABLE II. Double occupancy computed from calculations using correlated sampling with population control (Corr. samp.) and uncorrelated independent runs (Uncorr. samp.). Results are for the Hubbard model at half-filling with $U/t = 8$, in periodic supercells of size 4×4 , 8×8 and 16×16 . Note that the listed values are D as defined in Eq. (12), not per site. Statistical errors are in the last digit(s) as indicated in parentheses.

L	Corr. samp.	Uncorr. samp.	Efficiency gain
4×4	0.857(2)	0.92(9)	5.5×10^3
8×8	3.472(5)	2.95(32)	1.4×10^4
16×16	13.93(4)	13.0(14)	5.7×10^3




FIG. 7. Monitoring the level of correlation in correlated sampling. The efficiency gain over uncorrelated calculations, indicated by the variance ratio $\sigma/\sigma_{\text{ref}}$, is plotted in the main figure vs the metric Q of Eq. (9). Two examples with $\delta U = 1$ and 2 are shown, for the same system as in Fig. 5. The horizontal dashed line marks $\sigma/\sigma_{\text{ref}} = 1$. The inset plots Q vs projection time τ . All the curves were smoothed by averaging many neighboring points.

correlated than the previous example. In all the runs here, the error bars are estimated from 40 resets (or 40 repeat runs in uncorrelated calculations) after equilibration. We quantify the computational efficiency gain as

$$\left(\frac{\sigma_{\text{uncorrelated}}}{\sigma_{\text{correlated}}} \right)^2 \frac{T_{\text{uncorrelated}}}{T_{\text{correlated}}}, \quad (15)$$

where T is the total computational cost of a calculation (correlated or uncorrelated), and σ is the corresponding statistical error bar. A smaller choice of δU would make the computational efficiency gain grow inversely proportional to δU^2 .

In Fig. 7, we illustrate how the index Q defined in Eq. (9) can help monitor the quality of the correlation. We quantify the superiority of the correlated sampling calculation over the uncorrelated calculation through the variance ratio $\sigma/\sigma_{\text{ref}}$, where σ_{ref} denotes the variance of the uncorrelated sampling calculation, which works as a reference. When this ratio exceeds 1, there is no longer any efficiency gain with correlated sampling. To see the crossover more clearly, we deliberately choose unnaturally large δU values. As is shown in the main figure, $\sigma/\sigma_{\text{ref}}$ increases with Q . In the inset, we see that Q increases with projection time, and at the same τ value, Q is larger for larger δU . These observations are all as expected from the nature of correlated sampling. The similar value of Q where $\sigma/\sigma_{\text{ref}}$ crosses 1 in the main figure indicates that Q gives a reasonably generic metric for measuring the proximity of the correlated systems.

As a cross-check between the static and dynamic realization, we tested the static realization in the system computed in Fig. 5 with the same number of measurement blocks and thus the same CPU time. It turns out the dynamic variant

gives slightly smaller statistical error [$D = 1.8407(12)$] than the static variant [$D = 1.8441(15)$].

B. Real materials: Solid Si

Here, we perform a complimentary set of tests in *ab initio* calculations of real materials in bulk silicon. Our computations use plane-wave basis AFQMC [41] with multiprojector norm-conserving pseudopotentials [42,43]. A single-determinant trial wave function from density functional theory is used. We focus on the diamond-structured Si, with 32 atoms in a body-centered cubic (BCC) supercell. [We apply a twist boundary condition to the supercell using as the \mathbf{k} point the BCC Baldereschi point: $(\frac{1}{4}, \frac{1}{4}, -\frac{1}{12})$ in fractional coordinates.] As such, the system is an interacting many-body system with over 100 electrons and in excess of 9000 plane-waves (after the use of pseudopotentials), presenting a stringent test of any correlated sampling approach. We measure the energy difference ΔE between two systems: one at the minimum-energy or equilibrium geometry and the other with one of the 32 Si atoms displaced by 0.01 Å. The energy

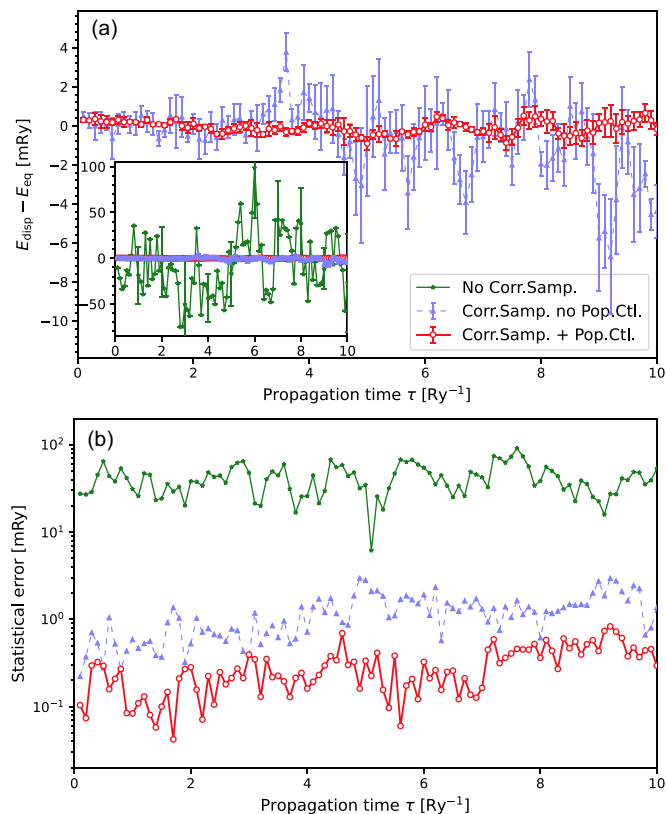


FIG. 8. Performance improvement by correlated sampling with population control in *ab initio* computations of the real material silicon. The difference in ground-state energies between two structures in bulk Si is computed [E_{eq} at equilibrium (diamond structure), and E_{disp} with one of the atoms displaced by 0.01 Å] by Ph-AFQMC, with and without correlated sampling or population control. (a) shows this energy difference, with a zoom-out view in the inset to demonstrate the scale of the no-correlated-sampling run. (b) shows how the statistical error evolves over τ . The system is a $2 \times 2 \times 2$ supercell with 32 Si atoms. The calculation consists of five pairs of runs per method, each run with a total population of 800 random walkers.

difference between the two systems can be used to compute the force exerted on the displaced atom. Moreover, such energy differences are crucial for structural optimizations or reaction pathway studies. Here, the static version of correlated sampling population control is tested.

Figure 8 shows a comparison of the computed energy difference of the two systems, using Ph-AFQMC without correlated sampling, with correlated sampling but without population control, and finally with our full algorithm of correlated sampling with population control. The total population size (reflecting the total computational cost) is the same for all three. The computed energy difference is around zero since the structure is at equilibrium and the atomic force on the displaced atom vanishes. Without correlated sampling, the calculations show fluctuations which are 1–2 orders of magnitude larger than with correlated sampling, as seen in the inset of panel (a) or in panel (b). The main part of panel (a) omits these results and shows only a magnified view of the two correlated sampling results. Without population control, correlated sampling exhibits a clear increase in statistical fluctuations with projection time. With population control, the fluctuations are reduced. Furthermore, the growth with random walk steps is much suppressed; in fact, it is barely discernible in the range of projection time studied, which is far greater than the time needed for the targeted difference to reach convergence.

V. CONCLUSIONS AND OUTLOOK

In this paper, we proposed a branching and population control algorithm for correlated sampling. The algorithm

is generally applicable to MC calculations that involve branching random walks. We outlined several variants for implementing the algorithm, which can be adapted based on the calculational setup and the behavior of the underlying correlated sampling method. We also discussed the quantification and validation of the effectiveness of correlated sampling. We illustrated and tested our algorithm in the Hubbard model and in *ab initio* solid calculations, using the ground-state AFQMC method. The population control algorithm was shown to significantly increase the efficiency of correlated sampling and extend the duration of correlation as a function of random walk time.

We expect that the algorithm will expand the range of applicability of correlated sampling to larger systems and more diverse and challenging problems. Many applications in different areas can be pursued along these lines. In addition, there is considerable room to improve the algorithm itself. For example, in defining the function f or in implementing branching/population control after the reference weights are produced, we have only taken a first step and expect considerable efficiency gain can be available with further study.

ACKNOWLEDGMENTS

S.C. acknowledges support from the U.S. Department of Energy under Grant No. DE-SC0001303. Y.Y. acknowledges support from the Simons Foundation. S.C. and Y.Y. thanks the Flatiron Institute for hospitality, computational resources, and technical support. Flatiron Institute is a division of Simons Foundation.

-
- [1] M. H. Kalos and P. A. Whitlock, *Monte Carlo Methods* (WILEY-VCH Verlag GmbH & Co. KGaA, Weinheim, 2008), pp. 7–34.
 - [2] R. Jain, *The Art of Computer Systems Performance Analysis Techniques for Experimental Design Measurement Simulation and Modeling* (John Wiley & Sons, Inc., New York, 1991).
 - [3] S. Zhang, *Auxiliary-Field Quantum Monte Carlo at Zero- and Finite-Temperature* (Verlag des Forschungszentrum Jülich, Jülich, 2019), Vol. 9.
 - [4] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Quantum Monte Carlo simulations of solids, *Rev. Mod. Phys.* **73**, 33 (2001).
 - [5] J. Spanier and E. M. Gelbard, *Monte Carlo Principles and Neutron Transport Problems* (Addison-Wesley Publishing Company, Reading, 1969).
 - [6] J. R. Howell, Application of Monte Carlo to heat transfer problems, *Adv. Heat Transf.* **5**, 1 (1969).
 - [7] H. W. Bertini, Spallation reactions: Calculations, in *Spallation Nuclear Reactions and their Applications*, edited by B. S. P. Shen and M. Merker (Springer, Dordrecht, 1976).
 - [8] B. M. Ayyub and R. H. McCuen, *Probability, Statistics, and Reliability for Engineers and Scientists* (CRC Press, Boca Raton, 2003).
 - [9] R. Y. Rubinstein, *Monte Carlo Optimization, Simulation, and Sensitivity of Queueing Networks* (Wiley, New York, 1986).
 - [10] D. Mcleish, *Monte Carlo Simulation and Finance* (Wiley, Hoboken, 2005).
 - [11] J. H. Hetherington, Observations on the statistical iteration of matrices, *Phys. Rev. A* **30**, 2713 (1984).
 - [12] N. Trivedi and D. M. Ceperley, Ground-state correlations of quantum antiferromagnets: A Green-function Monte Carlo study, *Phys. Rev. B* **41**, 4552 (1990).
 - [13] M. Calandra Buonaura and S. Sorella, Numerical study of the two-dimensional Heisenberg model using a Green function Monte Carlo technique with a fixed number of walkers, *Phys. Rev. B* **57**, 11446 (1998).
 - [14] H. Nguyen, H. Shi, J. Xu, and S. Zhang, CPMC-lab: A MATLAB package for constrained path Monte Carlo calculations, *Comput. Phys. Commun.* **185**, 3344 (2014).
 - [15] S. Zhang, J. Carlson, and J. E. Gubernatis, Constrained path Monte Carlo method for fermion ground states, *Phys. Rev. B* **55**, 7464 (1997).
 - [16] S. Zhang and H. Krakauer, Quantum Monte Carlo method using phase-free random walks with Slater determinants, *Phys. Rev. Lett.* **90**, 136401 (2003).
 - [17] S. Zhang, Finite-temperature Monte Carlo calculations for systems with fermions, *Phys. Rev. Lett.* **83**, 2777 (1999).
 - [18] M. Motta and S. Zhang, Communication: Calculation of interatomic forces and optimization of molecular geometry with auxiliary-field quantum Monte Carlo, *J. Chem. Phys.* **148**, 181101 (2018).

- [19] S. Chen and S. Zhang, Computation of forces and stresses in solids: Towards accurate structural optimization with auxiliary-field quantum Monte Carlo, *Phys. Rev. B* **107**, 195150 (2023).
- [20] M. Qin, C.-M. Chung, H. Shi, E. Vitali, C. Hubig, U. Schollwöck, S. R. White, and S. Zhang, Absence of superconductivity in the pure two-dimensional Hubbard model, *Phys. Rev. X* **10**, 031016 (2020).
- [21] W. A. Lester, Jr. and B. L. Hammond, Quantum Monte Carlo for the electronic structure of atoms and molecules, *Annu. Rev. Phys. Chem.* **41**, 283 (1990).
- [22] C. J. Umrigar, Two aspects of quantum Monte Carlo: Determination of accurate wavefunctions and determination of potential energy surfaces of molecules, *Int. J. Quantum Chem.* **36**, 217 (1989).
- [23] Y. Kwon, D. M. Ceperley, and R. M. Martin, Quantum Monte Carlo calculation of the Gerni-liquid parameters in the two-dimensional electron gas, *Phys. Rev. B* **50**, 1684 (1994).
- [24] R. L. Coldwell, Zero Monte Carlo error or quantum mechanics is easier, *Int. J. Quantum Chem.* **12**, 215 (1977).
- [25] Z. Sun, Jr., R. N. Barnett, and W. A. Lester, Quantum and variational Monte Carlo interaction potentials for $\text{Li}_2(X^1\Sigma_g^+)$, *Chem. Phys. Lett.* **195**, 365 (1992).
- [26] C. Filippi and C. J. Umrigar, Correlated sampling in quantum Monte Carlo: A route to forces, *Phys. Rev. B* **61**, R16291 (2000).
- [27] C. A. Traynor and J. B. Anderson, Parallel Monte Carlo calculations to determine energy differences among similar molecular structures, *Chem. Phys. Lett.* **147**, 389 (1988).
- [28] B. H. Wells, The differential Green's function Monte Carlo method, the dipole moment of LiH, *Chem. Phys. Lett.* **115**, 89 (1985).
- [29] J. Vrbik and S. M. Rothstein, Infinitesimal differential diffusion quantum Monte Carlo study of diatomic vibrational frequencies, *J. Chem. Phys.* **96**, 2071 (1992).
- [30] J. Shee, S. Zhang, D. R. Reichman, and R. A. Friesner, Chemical transformations approaching chemical accuracy via correlated sampling in auxiliary-field quantum Monte Carlo, *J. Chem. Theory Comput.* **13**, 2667 (2017).
- [31] T. Poole, Calculating derivatives within quantum Monte Carlo, Ph.D. thesis, Imperial College London, 2014.
- [32] A. Mahajan, J. S. Kurian, J. Lee, D. R. Reichman, and S. Sharma, Response properties in phaseless auxiliary field quantum Monte Carlo, *J. Chem. Phys.* **159**, 184101 (2023).
- [33] W. Krauth, *Statistical Mechanics: Algorithms and Computations* (Oxford University Press, Oxford, 2006).
- [34] M. Motta and S. Zhang, *Ab initio* computations of molecular systems by the auxiliary-field quantum Monte Carlo method, *WIREs Comput. Mol. Sci.* **8**, e1364 (2018).
- [35] J. Shee, J. L. Weber, D. R. Reichman, R. A. Friesner, and S. Zhang, On the potentially transformative role of auxiliary-field quantum Monte Carlo in quantum chemistry: A highly accurate method for transition metals and beyond, *J. Chem. Phys.* **158**, 140901 (2023).
- [36] H. Shi and S. Zhang, Some recent developments in auxiliary-field quantum Monte Carlo for real materials, *J. Chem. Phys.* **154**, 024107 (2021).
- [37] J. Hubbard and B. H. Flowers, Electron correlations in narrow energy bands, *Proc. R. Soc. Lond. A* **276**, 238 (1963).
- [38] M. Qin, H. Shi, and S. Zhang, Benchmark study of the two-dimensional Hubbard model with auxiliary-field quantum Monte Carlo method, *Phys. Rev. B* **94**, 085103 (2016).
- [39] M. Qin, H. Shi, and S. Zhang, Coupling quantum Monte Carlo and independent-particle calculations: Self-consistent constraint for the sign problem based on the density or the density matrix, *Phys. Rev. B* **94**, 235119 (2016).
- [40] H. Xu, H. Shi, E. Vitali, M. Qin, and S. Zhang, Stripes and spin-density waves in the doped two-dimensional Hubbard model: Ground state phase diagram, *Phys. Rev. Res.* **4**, 013239 (2022).
- [41] M. Suewattana, W. Purwanto, S. Zhang, H. Krakauer, and E. J. Walter, Phaseless auxiliary-field quantum Monte Carlo calculations with plane waves and pseudopotentials: Applications to atoms and molecules, *Phys. Rev. B* **75**, 245123 (2007).
- [42] D. R. Hamann, Optimized norm-conserving Vanderbilt pseudopotentials, *Phys. Rev. B* **88**, 085117 (2013).
- [43] F. Ma, S. Zhang, and H. Krakauer, Auxiliary-field quantum Monte Carlo calculations with multiple-projector pseudopotentials, *Phys. Rev. B* **95**, 165103 (2017).