

Density-matrix quantum Monte Carlo methodN. S. Blunt,¹ T. W. Rogers,¹ J. S. Spencer,^{1,2} and W. M. C. Foulkes¹¹*Department of Physics, Imperial College London, Exhibition Road, London, SW7 2AZ, United Kingdom*²*Department of Materials, Imperial College London, Exhibition Road, London, SW7 2AZ, United Kingdom*

(Received 20 March 2013; revised manuscript received 26 April 2014; published 18 June 2014)

We present a quantum Monte Carlo method capable of sampling the full density matrix of a many-particle system at finite temperature. This allows arbitrary reduced density matrix elements and expectation values of complicated nonlocal observables to be evaluated easily. The method resembles full configuration interaction quantum Monte Carlo but works in the space of many-particle operators instead of the space of many-particle wave functions. One simulation provides the density matrix at all temperatures simultaneously, from $T = \infty$ to $T = 0$, allowing the temperature dependence of expectation values to be studied. The direct sampling of the density matrix also allows the calculation of some previously inaccessible entanglement measures. We explain the theory underlying the method, describe the algorithm, and introduce an importance-sampling procedure to improve the stochastic efficiency. To demonstrate the potential of our approach, the energy and staggered magnetization of the isotropic antiferromagnetic Heisenberg model on small lattices, the concurrence of one-dimensional spin rings, and the Renyi S_2 entanglement entropy of various sublattices of the 6×6 Heisenberg model are calculated. The nature of the sign problem in the method is also investigated.

DOI: [10.1103/PhysRevB.89.245124](https://doi.org/10.1103/PhysRevB.89.245124)

PACS number(s): 02.70.Ss, 31.15.V-, 71.15.-m, 03.67.Ac

I. INTRODUCTION

Quantum Monte Carlo (QMC) methods are well established as vital tools in the study of complex many-body quantum systems, often providing highly accurate results. Projector methods such as diffusion Monte Carlo [1,2] (DMC) and Green's function Monte Carlo [3,4] (GFMC) grant access to zero-temperature properties by stochastically applying a projection operator to a starting wave function to obtain a statistical sampling of the ground state. The fixed-node approximation [5–8] allows projector QMC to be applied to systems with sign problems and often produces very good results, but its accuracy depends on the quality of the trial nodal surface and is difficult to assess. Another drawback of projector QMC is that expectation values of quantum mechanical observables that do not commute with the Hamiltonian are difficult to calculate [1,2,9,10].

Finite-temperature QMC methods take a different approach. Path-integral Monte Carlo (PIMC) calculations express the partition function, $Z = \text{Tr}(e^{-\beta H})$, as a sum of contributions from paths through Hilbert space [11]. With an appropriate update procedure, the paths can be sampled with the correct probabilities, thus allowing finite-temperature expectation values to be evaluated. The stochastic series expansion (SSE) method [12] has much in common with PIMC. These algorithms also allow access to ground-state properties in principle, but in practice the sign problem is often insurmountable at low temperatures.

The full configuration interaction quantum Monte Carlo (FCIQMC) method recently introduced by Booth, Thom, and Alavi [13,14] is a projector method for studying zero-temperature properties, and, as such, has much in common with DMC and GFMC. However, unlike DMC and GFMC, where the sampling of the ground-state wave function is performed in real space, FCIQMC samples the components of the wave function in a discrete basis. Crucially, no prior knowledge of the nodal structure of the ground-state wave function is required to reach the exact ground state. Rather, the

sign problem manifests itself in the large but system-specific population of quantum Monte Carlo walkers required in order for the ground state of the Hamiltonian to emerge [15] from the background noise. The system sizes accessible to FCIQMC are limited by the amount of memory available to store these walkers. However, the method has proven highly successful in many chemical systems, reducing the memory needed to achieve FCI-quality results by several orders of magnitude [16–20]. This has led to much interest in this direction and research into fundamental improvements and new applications of the algorithm continues [17,21].

This article presents a closely related QMC method, which we call density-matrix quantum Monte Carlo (DMQMC). Like the path-integral and SSE methods, DMQMC allows finite-temperature results to be calculated. However, it uses a projection approach to achieve this and thus has more in common with zero-temperature QMC methods. DMQMC was inspired by FCIQMC and shares many of its features, but samples the elements of the density matrix instead of the components of the wave function in a discrete basis. This enables expectation values of arbitrary quantum mechanical observables to be calculated easily, even when the operator corresponding to the observable does not commute with the Hamiltonian. Such expectation values are difficult to calculate using other QMC methods [1,2,9,10]. Moreover, the ability to directly sample the density matrix means that many quantum information measures are accessible. These advantages cannot be expected to come without drawbacks, and, indeed, the systems to which we have successfully applied DMQMC to date are small by the standards of other finite-temperature methods. However, the potential uses of a direct stochastic sampling of the density matrix are such that DMQMC deserves investigation. This paper demonstrates the use of DMQMC by studying the antiferromagnetic Heisenberg model, but DMQMC is a general method, applicable to bosonic and fermionic systems and to real molecules and solids.

Section II summarizes the FCIQMC method, setting the stage for the description of DMQMC in Sec. III. In Sec. IV an importance-sampling procedure is introduced. The DMQMC method is then applied to the isotropic antiferromagnetic Heisenberg model in Sec. V and used to calculate the energy and staggered magnetization of small square lattices, the concurrence of one-dimensional rings, and the Renyi S_2 entanglement entropy of various sublattices of the 6×6 lattice. The sign problem in DMQMC is investigated in Sec. VI. We discuss our results and offer some concluding remarks in Sec. VII. Hartree atomic units are used throughout.

II. FULL CONFIGURATION INTERACTION QUANTUM MONTE CARLO

The DMQMC algorithm was formulated in analogy with the FCIQMC method. The two methods share many of the same features and it is often useful to compare them. We therefore begin by providing a brief summary of the FCIQMC method. For more detailed discussions readers are referred to Refs. [13,15].

Consider the imaginary-time Schrödinger equation

$$\frac{d|\Psi\rangle}{d\tau} = -\hat{H}|\Psi\rangle. \quad (1)$$

The general solution to this equation is

$$|\Psi(\tau)\rangle = e^{-\tau\hat{H}}|\Psi(\tau=0)\rangle, \quad (2)$$

for some initial wave function $|\psi(\tau=0)\rangle$. If the initial wave function has a nonzero ground-state component, $c_0(0)$, it is easy to see that $|\Psi(\tau)\rangle$ will become proportional to the ground state in the limit of large τ ,

$$|\Psi(\tau \rightarrow \infty)\rangle = c_0(0)e^{-\tau E_0}|E_0\rangle, \quad (3)$$

where $|E_0\rangle$ is the ground state and E_0 the ground-state energy. The factor of $e^{-E_0\tau}$ can be removed by choosing the zero of energy such that $E_0 = 0$. In practice, since E_0 is usually unknown, we solve

$$\frac{d|\Psi\rangle}{d\tau} = -(\hat{H} - S\hat{1})|\Psi\rangle = \hat{T}|\Psi\rangle, \quad (4)$$

where we have defined $\hat{T} = -(\hat{H} - S\hat{1})$. The energy shift S is adjusted slowly during the simulation to keep the normalization approximately constant. The long-time average of S provides a measure of the ground-state energy.

The above theory is common to all projector methods; the difference is how they achieve the evolution to the ground state. FCIQMC works in a discrete basis of kets $|X_i\rangle$, which are normally Slater determinants for fermions or permanents for bosons. The components of the wave function in this basis are represented stochastically by a collection of markers. Following Anderson [5,6], we refer to these markers as psips. Each psip has an associated sign, $q = \pm 1$, which we refer to as its charge, and resides on a particular basis state $|X_i\rangle$ (or on site i). The expected value of the net psip charge on a basis state is proportional to the amplitude of that state in the expansion of the wave function. At any point in a simulation, the distribution of psip charges does not need to provide

an accurate representation of the wave function. Rather, the FCIQMC method only requires that the *expectation values* of the site charges represent the ground state [18]. Thus, at any point in the simulation, the memory required to sample the wave function may be many orders of magnitude smaller than the memory required to store the whole state.

Booth and coworkers [13] introduced an algorithm to evolve the population of psips according to the imaginary-time Schrödinger equation. This can be summarized as follows. For each time step $\Delta\tau$ we loop over the entire population of psips and perform the following steps.

(i) Spawning: Allow a psip with charge q_i on site i to spawn onto connected sites j , where $T_{ij} \neq 0$ and $i \neq j$, with probability $|T_{ji}|\Delta\tau$. If the spawning attempt is successful, a psip is born at site j with charge $q_j = \text{sign}(T_{ji})q_i$.

(ii) Diagonal death/cloning: Each psip has the chance to either clone or die with probability $|T_{ii}|\Delta\tau$. The consequence of a successful death/cloning event depends on the sign of the diagonal matrix element: if $T_{ii} > 0$ the psip is cloned; otherwise the psip is removed from the simulation.

(iii) Annihilation: Pairs of psips on the same site with opposite charges cancel out (annihilate) and are removed from the simulation, leaving a population of only a single charge type on each site.

The FCIQMC algorithm samples the solution of a first-order Euler finite-difference approximation to Eq. (4). Hence, the distribution of psips gives a stochastic representation of the wave function and, as $\tau \rightarrow \infty$, the psips settle down to sample the ground-state wave function.

The psip annihilation step does not alter the total charge on a site. However, it has been shown to be vital in order for the true ground-state wave function to emerge in systems with sign problems [15]. Similar and more complex walker cancellation mechanisms have been attempted in continuum DMC and GFMC calculations with less success [22–27]. Walker cancellation in FCIQMC works better because psips are more likely to encounter each other in a discrete and finite Hilbert space.

The ability of FCIQMC to tackle the sign problem through the annihilation step is perhaps its most significant advantage. Annihilation leads to the characteristic population dynamics and allows the true ground state to emerge without any knowledge of the nodal structure of the ground-state wave function [15]. At the start of an FCIQMC simulation, the shift is held constant at a value large enough to ensure that the psip population grows exponentially. Eventually, when a system-specific population of psips is reached, the annihilation rate becomes equal to the spawning rate and the population spontaneously plateaus. During this plateau period the ground state of the Hamiltonian emerges, after which the population begins to grow again. The population can then be controlled by varying the shift. The psip population at the plateau must be a small fraction of the number of basis states in the Hilbert space in order for FCIQMC to be more (memory) efficient than an exact diagonalization.

The ground-state energy in FCIQMC can be calculated using a projected estimator based on the expression

$$E_0 = \frac{\langle\Psi^T|\hat{H}|E_0\rangle}{\langle\Psi^T|E_0\rangle} = \frac{\sum_{i,j}\psi_i^T H_{ij}\chi_j}{\sum_i\psi_i^T\chi_i}. \quad (5)$$

Here $|\Psi^T\rangle$ is an appropriate trial function with components ψ_i^T in the many-particle basis, $|X_i\rangle$, and χ_i is a component of the exact ground state in this basis. FCIQMC was initially performed using only the Hartree-Fock determinant as the trial function; multireference trial functions have subsequently been shown to substantially decrease statistical fluctuations [21]. The ground-state energy is obtained by averaging the numerator and denominator separately, with the total psip charge on each site, q_i^{tot} , used in the place of corresponding exact amplitude, χ_i .

Calculating the ground-state expectation value of an operator \hat{O} that does not commute with the Hamiltonian is more difficult because a projected estimator cannot be used. Instead, assuming that $|E_0\rangle$ is real, it is necessary to evaluate

$$\langle \hat{O} \rangle = \frac{\langle E_0 | \hat{O} | E_0 \rangle}{\langle E_0 | E_0 \rangle} = \frac{\sum_{i,j} O_{ij} \chi_i \chi_j}{\sum_i \chi_i^2}, \quad (6)$$

where $O_{ij} = \langle X_i | \hat{O} | X_j \rangle$. Although $\langle q_i^{\text{tot}} \rangle = \chi_i$, the expectation value of a product is not equal to the product of the expectation values: $\langle q_i^{\text{tot}} q_j^{\text{tot}} \rangle \neq \chi_i \chi_j$. This means that $\chi_i \chi_j$ *cannot* be obtained by averaging the products of the instantaneous psip weights. One could in principle average q_i^{tot} over many iterations to obtain χ_i before multiplying χ_i and χ_j , but this would involve storing a number for every basis function, which is impractical due to memory limitations.

This problem is not easy to overcome and there is currently no way to evaluate general expectation values exactly and efficiently within the FCIQMC framework. Indeed, the calculation of general expectation values in other Monte Carlo methods is often a difficult task.

III. DENSITY-MATRIX QUANTUM MONTE CARLO

We now show how an FCIQMC-like dynamics can be used to sample both finite-temperature and ground-state density matrices. We first consider the thermal density matrix and how it can be evolved as a function of inverse temperature by solving the symmetrized Bloch equation. We then draw upon analogies with FCIQMC to formulate the DMQMC algorithm before discussing the calculation of estimators for a general quantum mechanical observable. This section ends with an explanation of how to sample a reduced density matrix in order to calculate estimators of entanglement measures.

A. Theory

Since the psip population (and hence the normalization) varies during a quantum Monte Carlo simulation, it is convenient to work with the unnormalized thermal density matrix

$$\hat{\rho}(\beta) = e^{-\beta \hat{H}}, \quad (7)$$

where \hat{H} is the Hamiltonian operator and $\beta = 1/k_B T$ is the inverse temperature. The canonical partition function $Z(\beta)$ is given by

$$Z(\beta) = \text{Tr}[\hat{\rho}(\beta)]. \quad (8)$$

Differentiating $\hat{\rho}(\beta)$ with respect to β shows that it obeys both the Bloch equation,

$$\frac{d\hat{\rho}}{d\beta} = -\hat{H}\hat{\rho}, \quad (9)$$

and the symmetrized Bloch equation,¹

$$\frac{d\hat{\rho}}{d\beta} = -\frac{1}{2}(\hat{H}\hat{\rho} + \hat{\rho}\hat{H}) = -\frac{1}{2}\{\hat{H}, \hat{\rho}\}. \quad (10)$$

The symmetrized version turns out to be more useful for our purposes. Consider the general solution to Eq. (10),

$$\hat{\rho}(\beta) = e^{-\frac{\beta}{2}\hat{H}} \hat{\rho}(\beta=0) e^{-\frac{\beta}{2}\hat{H}}, \quad (11)$$

and let $|E_i\rangle$ denote an energy eigenstate with energy E_i . Without loss of generality, we assume that $E_0 = 0$. By expanding $\hat{\rho}(\beta=0)$ as

$$\hat{\rho}(\beta=0) = \sum_{i,j} \rho_{ij}(0) |E_i\rangle \langle E_j| \quad (12)$$

and applying Eq. (11) it is seen that $\hat{\rho}(\beta)$ becomes proportional to the ground-state density matrix as $\beta \rightarrow \infty$:

$$\hat{\rho}(\beta \rightarrow \infty) = \rho_{00}(0) |E_0\rangle \langle E_0|. \quad (13)$$

A similar analysis shows that the ground-state density matrix is *not* reached if a general initial density matrix is propagated using the unsymmetrized Bloch equation. Hence, in cases where ground-state properties are desired, Eq. (10) is more useful. If we wish to sample thermal properties at any finite temperature T , we must also impose the correct initial condition on Eq. (10). The most convenient is at $\beta = 0$, where

$$\hat{\rho}(\beta=0) = \hat{1}. \quad (14)$$

Equations (10) and (14) provide everything required to stochastically sample the thermal density matrix. As in FCIQMC we introduce a collection of psips, which in DMQMC occupy not basis states $|X_i\rangle$ but basis operators $|X_i\rangle \langle X_j|$ of the expansion of the full many-particle density operator $\hat{\rho}(\beta)$ in some convenient basis set. At the start of a simulation the psips are distributed randomly along the diagonal of the infinite-temperature density matrix $\rho_{ij}(\beta=0) = \delta_{ij}$. The psips then evolve under a set of rules so that they sample the matrix elements $\rho_{ij}(\beta)$ in the chosen basis at each iteration. In order to prevent the population of psips

¹There is a close relationship between the symmetrized Bloch equation in Eq. (10) and the von Neumann and quantum Liouville equations, $i\hbar \frac{\partial \rho}{\partial t} = [H, \rho]$. The von Neumann equation may be derived by using the time-dependent Schrödinger equation to work out how the arbitrary initial density matrix of Eq. (12) evolves. Analogously, the Bloch equation may be derived by using the imaginary-time Schrödinger equation to work out how Eq. (12) evolves in imaginary time. The switch from real to imaginary time replaces the commutator of \hat{H} and ρ appearing in the von Neumann equation with the anticommutator appearing in the Bloch equation. The additional factor of $\frac{1}{2}$ is introduced for convenience: it ensures that evolving Eq. (10) for β units of imaginary time, starting from the unit operator, $\rho(\beta=0) = \hat{1}$, yields the canonical density matrix at inverse temperature β , not 2β .

from exploding, we introduce an energy shift, S , and let $\hat{H} \rightarrow \hat{H} - S\hat{1}$ as in FCIQMC. Eq. (10) thus becomes

$$\frac{d\hat{\rho}}{d\beta} = \frac{1}{2}(\hat{T}\hat{\rho} + \hat{\rho}\hat{T}), \quad (15)$$

where $\hat{T} = -(\hat{H} - S\hat{1})$ is the update matrix.

B. Algorithm

We use the following algorithm to evolve a collection of psips according to Eq. (15). For a single step in inverse temperature of $\Delta\beta$, we loop over the entire population of psips and perform the following steps.

(i) Spawning along columns of the density matrix: Allow a psip with charge q_{ij} on site (i, j) to spawn onto connected sites (k, j) , where $T_{ik} \neq 0$ and $i \neq k$, with probability $\frac{1}{2}|T_{ik}|\Delta\beta$. If the spawning attempt is successful, a psip is born at (k, j) with charge $q_{kj} = \text{sign}(T_{ik})q_{ij}$.

(ii) Spawning along rows of the density matrix: Similarly, allow a psip with charge q_{ij} on site (i, j) to spawn onto connected sites (i, k) with probability $\frac{1}{2}|T_{jk}|\Delta\beta$. If the spawning attempt is successful, a psip is born at site (i, k) with charge $q_{ik} = \text{sign}(T_{jk})q_{ij}$.

(iii) Diagonal death/cloning: If there is a psip on site (i, j) and $T_{ii} + T_{jj} < 0$, then with probability $p_d = \frac{1}{2}|T_{ii} + T_{jj}|\Delta\beta$ that psip is killed and removed from the simulation; if $T_{ii} + T_{jj} > 0$, the psip is cloned (i.e., a new psip is created on the same site and with the same charge) with probability p_d .

(iv) Annihilation: Pairs of psips inhabiting the same site (i, j) with opposite charges annihilate and are removed from the simulation, leaving a population of only a single charge type on each site.

The first three steps describe a stochastic algorithm to sample the solution of a first-order Euler finite-difference approximation to Eq. (10). The distribution of psip charges at $\beta + \Delta\beta$ is thus proportional to the density matrix at this inverse temperature, provided that the distribution of charges at β was correct.

The DMQMC method shares many similarities with FCIQMC, namely,

(i) Annihilation does not alter the expected (normalized) psip distribution but serves to overcome the sign problem [13,15].

(ii) The underlying finite-difference approximation is stable if $0 < \Delta\beta < 2/(E_{\max} - E_0)$, where E_{\max} is the largest eigenvalue of the Hamiltonian matrix [15]. This is a sufficient condition to ensure correct projection onto the exact ground state, but the finite value of $\Delta\beta$ leads to an error of $\mathcal{O}(\Delta\beta)$ in the density matrix at temperatures greater than zero. It is therefore necessary to check that finite-temperature results have converged with respect to $\Delta\beta$.

(iii) The familiar shift-update algorithm already used in DMC [28] and FCIQMC [13] simulations is employed to modify S and thus to control the population. The shift, S , is adjusted according to

$$S(\beta + A\Delta\beta) = S(\beta) - \frac{\zeta}{A\Delta\beta} \ln \left(\frac{N_p(\beta + A\Delta\beta)}{N_p(\beta)} \right), \quad (16)$$

where A is the number of β steps between shift updates, ζ is a shift damping parameter, and $N_p(\beta)$ is the total number of

psips at the inverse temperature β . During simulations, ζ is chosen carefully to prevent large fluctuations in S .

(iv) Rather than attempting all possible spawning events from a psip on a given site, it is computationally efficient to attempt just one (or a small number) of the many possible spawning events and reweight the acceptance probabilities accordingly [13].

(v) The algorithm is highly parallelizable—only the annihilation step requires communication between CPU cores and so using a large psip population is a viable option.

As with all projection methods, the ground state is approached as $\beta \rightarrow \infty$. Once convergence to the ground state has been attained the density matrix no longer depends on β (remember that S is chosen to keep the normalization fixed) and the statistical errors of measured ground-state expectation values can easily be reduced by averaging over many iterations. A single simulation is therefore sufficient to obtain accurate ground-state expectation values. The estimation of finite-temperature properties is more difficult because the inverse temperature β changes continuously as the simulation progresses; it is not possible to hold the simulation at a specific temperature while statistics are accumulated.

With access to a stochastic sampling of the unnormalized density matrix at a given temperature, the expectation value of any quantum mechanical observable, \hat{O} , at that temperature can be calculated using

$$\langle \hat{O} \rangle = \frac{\text{Tr}(\hat{\rho} \hat{O})}{\text{Tr}(\hat{\rho})} = \frac{\sum_{i,j} \bar{q}_{ij}^{\text{tot}} O_{ji}}{\sum_i \bar{q}_{ii}^{\text{tot}}}. \quad (17)$$

The numerator and denominator must be sampled and averaged separately at each temperature to achieve the desired statistical accuracy. Hence a simulation involves repeatedly evolving the density matrix from $\beta = 0$ to some chosen maximum value of β , a process we call a “ β loop”. This also gives us the freedom to allow Eq. (14) to be satisfied only *on average*, which greatly reduces the memory demands. Each β loop is initialized with a different random number seed and with psips randomly distributed with uniform probability along the diagonal of the density matrix. Statistics are accumulated over a sufficiently large number of β loops in order to obtain the desired statistical accuracy. As each β loop is completely independent, performing multiple β loops is an embarrassingly parallel task and each β loop gives statistically independent data points.

A primary concern is that averages over an impractically large number of β loops may be required. Indeed, it is not uncommon to have to average over $\mathcal{O}(10^6)$ iterations in ground-state calculations. In practice, however, reaching a sufficiently large number of β loops does not appear to be a problem. The crucial difference from the ground-state case is that each β loop provides statistically independent data whereas ground-state calculations need to overcome inherent correlations between nearby iterations. Furthermore, as will be shown, the quality of estimates obtained using a given number of β loops is better at high temperature than at low temperature. Another advantage is that a single simulation can provide data across the entire temperature range being studied.

The energy shift S is varied in order to control the normalization. This, however, introduces a bias: if the shift is varied by ΔS , then the psip population is effectively altered by

a factor $e^{\Delta\beta\Delta S}$ over the next β step. As a result, configurations in which the average energy of the psip population happens to be less negative than usual are effectively given too large a weight. This problem is particularly severe in DMQMC because the results at each temperature are obtained by averaging over separate β loops, and thus, for each quantity contributing to a given estimator, the corresponding shift (and hence population) profiles can be very different.

The population bias can be greatly reduced using a method suggested by Umrigar *et al.* in the context of DMC [28], in which each sampled quantity proportional to the psip population is multiplied by the factor

$$\Pi(\beta, B) = \prod_{m=0}^{\tilde{B}-1} e^{-S(\beta-m\Delta\beta)}, \quad (18)$$

where B is some chosen number of factors and $\tilde{B} = \min(\frac{\beta}{\Delta\beta}, B)$. By multiplying by this factor, we remove the last \tilde{B} factors of $e^{\Delta\beta\Delta S}$ introduced by varying the shift. As $B \rightarrow \infty$ the population control bias should be completely removed. The population control bias can also be reduced by using a larger population of psips.

Another concern is that the number of elements in the density matrix is the square of the dimension of the Hilbert space of many-particle states. At first glance it might seem that, without dramatically increasing the number of psips, the density matrix would be very poorly sampled. However, the dimension of the Hilbert space, D , rises exponentially with the number of particles or sites, N , so that $D \propto e^{\alpha N}$ and thus $D^2 \propto e^{\alpha(2N)}$. The doubling of the exponent implies that a DMQMC simulation for an $N/2$ -site lattice model requires approximately the same number of psips as an N -site FCIQMC simulation to achieve the same sampling quality. Moreover, DMQMC estimators for operators that do not commute with the Hamiltonian often have significantly smaller variance than the forward-walking or other estimators required to evaluate the ground-state expectation values of such operators in FCIQMC.

C. Entanglement measures

Entanglement measures are well established as an important concept in quantum information theory and have recently become a subject of active research in the condensed matter community. For example, changes in entanglement are observed at quantum phase transitions [29] and used to classify properties of Fermi liquids [30,31] and bonding in small molecules [32]. The Lanczos method can be used to calculate entanglement measures in small systems and the density-matrix renormalization group method is applicable in one-dimensional systems [33], but the study of entanglement in large systems in more than one dimension is less straightforward. Relationships between reduced density matrices (RDMs) and spin correlation functions [34] allow QMC methods to access some entanglement measures in certain situations [35], and Sandvik recently introduced a QMC method formulated in the valence-bond basis [36] that allowed certain new entanglement entropy measures to be evaluated [37,38]. Hastings *et al.* showed that Renyi S_2 entanglement entropy can also be calculated [33]. However, in general, the inability of QMC methods to directly access the

density matrix and reduced density matrices of systems has hindered their use in this area.

Within DMQMC it is straightforward to obtain a stochastic representation of any reduced density matrix element from a stochastic representation of the full density matrix. A RDM can be sampled by tracing out unwanted psips. Consider a composite quantum system C , which can be partitioned into two subsystems A and B so that $\mathcal{H}_C = \mathcal{H}_A \otimes \mathcal{H}_B$, where \mathcal{H}_A (\mathcal{H}_B) denotes the Hilbert space of system A (B). The RDM ρ_A that describes sublattice A is defined by taking the partial trace of the full density matrix, ρ_C , over all the sites on sublattice B :

$$\rho_A = \text{Tr}_B(\rho_C). \quad (19)$$

Our implementation of DMQMC represents the many-particle basis functions as bit strings [39], where each bit refers to the state of a single spin. To evaluate ρ_A we construct a mask, I_B , which only has bits set that correspond to spins in subsystem B . The (i, j) density matrix element of ρ_C contributes to ρ_A if the result of the logical AND operation of the i string with the I_B mask is identical to that of the j string. The corresponding element in the RDM can be found by taking AND with an analogous I_A mask.

Evaluating the von Neumann entanglement entropy, $S_1 = -\text{Tr}(\rho \log \rho)$, and other nonlinear functions of a RDM ρ is challenging because $\langle f(\rho) \rangle \neq f(\langle \rho \rangle)$. Thus, although it is easy to accumulate contributions to $\langle f(\rho) \rangle$ on the fly, this does not yield the correct result $f(\langle \rho \rangle)$. A well-averaged estimate of the full RDM is required before the nonlinear function can be calculated accurately. Accumulating this estimate requires storing the full RDM, limiting studies to relatively small subsystems. It would be reasonable to store the reduced density matrix of a subsystem of 15 spins, but much larger subsystems are out of reach. The von Neumann entanglement entropy is particularly challenging because its derivative with respect to any RDM eigenvalue tends to infinity as the eigenvalue tends to zero. Thus, small errors in small RDM eigenvalues can lead to large errors in the final estimate of S_1 .

Two-qubit entanglement measures are more straightforward to calculate as the small size of the RDM allows it to be estimated accurately with relative ease. In this paper a concurrence estimate was taken from each β loop by first averaging the relevant RDM in the ground state. These concurrence estimates were themselves averaged to provide a final value and associated error.

The replica trick approach, also used by Hastings *et al.* [33], allows Renyi entropies to be calculated in a rigorously unbiased manner without storing and averaging the full RDM. Two DMQMC calculations performed simultaneously, each starting from a different random number seed and thus following a different Markov chain, are statistically independent. The Renyi-2 entropy can therefore be evaluated in an unbiased fashion:

$$S_2 = -\log_2 \left(\left\langle \left(\sum_{i,j} q_{ij}^{\text{tot}} w_{ij}^{\text{tot}} \right)^2 \right\rangle \right), \quad (20)$$

where q_{ij}^{tot} (w_{ij}^{tot}) is the total psip amplitude on the ij th element of the reduced density matrix in the first (second) simulation. Thus, by averaging $\sum_{ij} q_{ij}^{\text{tot}} w_{ij}^{\text{tot}}$ at each temperature over

many β loops, it is straightforward to calculate an accurate temperature-dependent estimate of S_2 .

In general, performing N replica simulations enables the unbiased sampling of ρ^N and thus S_N . S_2 is particularly simple computationally because, as can be seen from Eq. (20), only amplitudes from psips on the same sites need to be multiplied. Note that, unlike in Ref. [33], the replica trick is used here only to ensure the statistical independence of the two reduced density matrices.

IV. IMPORTANCE SAMPLING

This article considers the application of DMQMC to the $S = 1/2$ antiferromagnetic Heisenberg model,

$$\hat{H} = J \sum_{\langle i,j \rangle} \hat{S}_i \cdot \hat{S}_j, \quad (21)$$

where $J > 0$ and the $\langle i,j \rangle$ implies that the summation is over nearest-neighbor pairs of spins only. Periodic boundary conditions are applied. We work in the standard basis set where the many-spin states are tensor products of the one-spin eigenstates, $|\uparrow\rangle$ and $|\downarrow\rangle$, of the \hat{S}_z operator.

The algorithm described in Sec. II allows a stochastic sampling of the exact finite-temperature density matrix, assuming $\Delta\beta$ is sufficiently small. However, for the antiferromagnetic Heisenberg model, the sampling method described is found to be insufficient for all but the smallest systems. The ground-state wave function is highly delocalized over the states in the basis set, and thus the ground-state density matrix, $\rho = |E_0\rangle\langle E_0|$, has many off-diagonal elements with magnitudes comparable to the diagonal elements. As a result, when the density matrix is sampled via the DMQMC algorithm, only a small fraction of psips reside on or near diagonal elements. Estimators for the expectation values of most operators of interest only receive contributions from psips on or near the diagonal, so very few psips contribute and these estimators suffer from large statistical errors at low temperatures. Importance sampling can greatly improve the sampling quality.

We start by defining the excitation level between basis states $|X_i\rangle$ and $|X_j\rangle$ to be the smallest number of pairs of opposite spins that must be flipped in order to reach $|X_i\rangle$ from $|X_j\rangle$. For the Heisenberg model, Eq. (21), the excitation level can change by at most ± 1 in a single application of the Hamiltonian.

One straightforward way to improve the quality of sampling is to reduce the probability of psips spawning far from the diagonal of the density matrix. Psips that *do* reside on higher excitation levels are given a correspondingly larger weight, so that expectation values of operators are unchanged. However, the increase in the population of low-weight psips near the diagonal of the density matrix reduces the stochastic error in near-diagonal expectation values.

With this motivation in mind, we define the following importance-sampling procedure; a more rigorous formulation is given in Appendix B. Every time a psip on excitation level γ attempts to spawn a new psip on excitation level δ , the probability of successful spawning is altered by a factor $P_{\gamma\delta}$. Thus, if a psip on a diagonal element attempts to spawn a new psip onto the first excitation level, the probability of successful spawning is altered by a factor P_{01} , where $P_{01} < 1$. The first excitation level will thus be

occupied by P_{01} times as many psips as it would have been with unaltered spawning probabilities. The reduced (relative) population must be accounted for by giving all psips in the first excitation level a weight of $W_1 = 1/P_{01}$ when evaluating estimators. The number of spawning attempts from the first to the second excitation level is also altered by a factor P_{01} (due to the reduced number of psips at level 1) and the chance of successful spawning for each attempt is multiplied by the factor P_{12} . Hence, the weight given to the second excitation level must be $W_2 = 1/P_{01}P_{12}$. Furthermore, since there are P_{01} times as many psips on the first excitation level, the probability of spawning from the first excitation level to the diagonal elements must be enhanced by a factor $1/P_{01}$ in order to achieve consistent spawning dynamics.

In general $P_{\gamma\delta} = 1/P_{\delta\gamma}$ and the weight given to a psip on excitation level γ is

$$W_\gamma = \prod_{\delta=1}^{\delta=\gamma} \frac{1}{P_{\delta-1,\delta}}. \quad (22)$$

The estimator for an expectation value, $\langle \hat{O} \rangle$, is thus

$$\langle \hat{O} \rangle = \frac{\sum_{i,j} \tilde{\rho}_{ij} W_{E(i,j)} O_{ji}}{\sum_i \tilde{\rho}_{ii}}, \quad (23)$$

where $E(i,j)$ is the excitation level between $|X_i\rangle$ and $|X_j\rangle$, and $\tilde{\rho}_{ij}$ is the importance-sampled density matrix. The expected value of the importance-sampled psip charge on site (i,j) is proportional to $\tilde{\rho}_{ij}$.

There is clearly some freedom in the numerical values chosen for the factors $P_{\gamma\delta}$. In this study we adjust $P_{\gamma\delta}$ such that all excitation levels have similar psip populations in the ground state. While this choice is not necessarily optimal, it successfully increases the quality of ground-state estimates by many orders of magnitude while still allowing the entire density matrix to be sampled. Figures 1(a) and 1(b) show the fractions of psips on the first four excitation levels for the 4×4 square Heisenberg lattice as obtained without and with importance sampling. Both simulations used an initial population of 10^5 psips and a single β loop. It is clear that importance sampling greatly assists in keeping a non-negligible population on the lower excitation levels.

In practice, the values of $P_{\gamma\delta}$ are very small for small γ and δ and decrease in magnitude as the lattice size increases. In order to avoid an unnecessarily large suppression of psip spawning from the diagonal at high temperatures, we introduce the weights W_γ and probabilities $P_{\gamma\delta}$ gradually from $\beta = 0$ until they reach their desired final values at β_{target} . For $\beta < \beta_{\text{target}}$, the weight at excitation level γ is set equal to $(W_\gamma)^{\beta/\beta_{\text{target}}}$; for $\beta > \beta_{\text{target}}$, the weight is held constant. The value of β_{target} is chosen to be the inverse temperature at which the ground state is deemed to have been reached. However, the value of this parameter is not critical for the quality of the sampling.

V. RESULTS

Results are presented for the $S = 1/2$ antiferromagnetic Heisenberg model. In order to calculate finite-temperature properties, it is necessary to include contributions from all M_S (total spin) subspaces. Because the Heisenberg Hamiltonian conserves M_S , different subspaces can be studied using




FIG. 1. (Color online) The inverse-temperature dependence of the fractions of psips on the first four excitation levels during a DMQMC simulation of a 4×4 square antiferromagnetic Heisenberg lattice [40]. A single β loop was carried out using an initial population of 10^5 psips on the diagonal elements. The results shown in (a) were obtained without importance sampling. The results shown in (b) were obtained using importance sampling to ensure that every excitation level had a roughly equal number of psips at zero temperature.

separate simulations for estimators of the form in Eq. (17), which is an embarrassingly parallel computational task. Combining such results for different values of M_S is straightforward but requires additional calculations and reveals nothing of interest about DMQMC. The energy and staggered magnetisation results presented in this paper are for the $M_S = 0$ subspace only. For more general quantities such as S_2 it is necessary to perform a simulation with psips spanning all M_S subspaces simultaneously.

A. Temperature-dependent energy on the 4×4 lattice

The 4×4 lattice is small enough that an exact diagonalization can be performed, thus allowing a direct check of our DMQMC results. Figure 2(a) shows the energy as a function of temperature using an initial population of 100 psips at the start of each β loop and accumulating statistics over 1000 such loops. The shift was allowed to vary throughout the simulation and there were typically 700–800 psips in the simulation by the end of each β loop. No importance sampling was applied, and so statistical fluctuations increase with inverse temperature as explained in Sec. IV. Increasing the initial population to 10^5 psips [Fig. 2(b)] reduces the statistical errors such that the agreement with the exact results is essentially perfect.

B. Temperature-dependent staggered magnetization on the 8×8 lattice

The square of the staggered magnetization is represented by the operator

$$\hat{M}^2 = \hat{M} \cdot \hat{M}, \quad \text{with } \hat{M} = \frac{1}{N} \sum_i (-1)^{x_i+y_i} \hat{S}_i, \quad (24)$$




FIG. 2. (Color online) The exact and DMQMC energies for the 4×4 antiferromagnetic Heisenberg lattice (Hilbert space dimension $D \approx 1.29 \times 10^4$) with periodic boundary conditions [40]. The β step size, $\Delta\beta$, obeys $JN\Delta\beta = 0.1$, where $N = 16$, and results are shown every 30 iterations. Each simulation consisted of 1000 β loops. For the simulation reported in (a), 100 psips were introduced at the start of each β loop; for (b), 10^5 psips were introduced. The error bars in (b) are smaller than the size of the markers.

where x_i and x_j denote the coordinates of the square lattice. This operator does not commute with the Hamiltonian. Its expectation value, $\langle \hat{M}^2 \rangle$, is plotted as a function of β in Fig. 3 for an 8×8 lattice; a ground-state value obtained using the SSE method [41] is also shown. The dimension of the $M_S = 0$ Hilbert space of an 8×8 Heisenberg Hamiltonian is approximately 1.83×10^{18} , so the density matrix sampled in this simulation has approximately 3.36×10^{36} elements.

C. Ground-state energy and staggered magnetization on the 6×6 lattice

For ground-state calculations statistics are accumulated over a single β loop after the ground state has been reached. A relatively modest simulation of a 6×6 lattice using 6.5×10^6




FIG. 3. (Color online) The square of the staggered magnetization for an 8×8 antiferromagnetic Heisenberg lattice ($D \approx 1.83 \times 10^{18}$) using 1.4×10^7 psips and 143 β loops [40]. The importance-sampling procedure described in the main text was applied. The DMQMC value of $\langle \hat{M}^2 \rangle$ is plotted every 50th iteration. Error bars, where not visible, are smaller than the size of the marker. The ground-state value of $\langle \hat{M}^2 \rangle$ obtained from an SSE simulation [41] is plotted for comparison.

TABLE I. DMQMC estimates of the ground-state concurrence, C_{gs} , for antiferromagnetic spin rings containing N sites in the absence of an external magnetic field. Each ring corresponds to a Hilbert space of D basis functions. The density matrix was sampled using N_p psips, and statistics accumulated over N_l β loops. Exact results for the energy and concurrence for rings up to $N = 10$ are taken from Ref. [45]. An FCIQMC calculation was used to find the energy of the $N = 36$ chain, from which an exact value of C_{gs} was determined using Eq. (25). The full density matrix of the $N = 36$ chain has approximately $D^2 = 8.24 \times 10^{19}$ elements.

| N | E_0/N | C_{gs} | | N_p | N_l | D |
|-----|-------------|------------|-----------|-------------------|-----------------|--------------------|
| | | Exact [45] | DMQMC | | | |
| 4 | -0.5000 | 0.5000 | 0.5005(4) | 2.5×10^3 | 250 | 6 |
| 6 | -0.4671 | 0.4343 | 0.4342(5) | 2.5×10^3 | 1×10^3 | 20 |
| 8 | -0.4564 | 0.4128 | 0.4129(5) | 1×10^4 | 1×10^3 | 70 |
| 10 | -0.4515 | 0.4031 | 0.4031(4) | 4×10^3 | 1×10^3 | 252 |
| 36 | -0.44374(2) | 0.38748(4) | 0.3873(8) | 1×10^6 | 12 | 9.08×10^9 |

psips gave a ground-state energy of $-0.67888(24)JN$ and a value of $\langle \hat{M}^2 \rangle$ equal to $0.20985(7)$, where errors were obtained via a blocking analysis [42]. These values agree with exact results [43] of $-0.678872JN$ and 0.20983 for the ground-state energy and staggered magnetization, respectively, to within statistical errors [44]. Remarkably, despite the fact that $[\hat{M}^2, \hat{H}] \neq 0$, the statistical error in the value of $\langle \hat{M}^2 \rangle$ was smaller than that of the energy. It was not necessary to reweight these results using Eq. (18) as the psip population was large enough to render such biases negligible.

D. Ground-state concurrence on spin rings

The ground-state concurrence, C_{gs} , for neighboring spins (qubits) on an antiferromagnetic Heisenberg ring was studied by Wootters and O'Connor in 2001 [45]. They showed that, for an even number of spins, C_{gs} has a simple relationship with the ground-state energy,

$$C_{gs} = -\frac{1}{2}(4E_0/N + 1). \quad (25)$$

The exact results calculated from this formula provide a useful test of our DMQMC estimates of C_{gs} , which are obtained from the sampled reduced density matrix (see Appendix A).

The DMQMC estimates of C_{gs} are presented in Table I, along with Wootters and O'Connor's exact analytic values for up to $N = 10$ sites. A ring with $N = 36$ sites was also studied, using FCIQMC to calculate E_0 and then Eq. (25) to obtain a value of C_{gs} for comparison with the DMQMC results. For lengths up to $N = 10$, the calculation of the concurrence can easily be carried out using other methods and provides a straightforward test of the DMQMC algorithm. The $N = 36$ chain is far from trivial and it is promising that such accurate results can be obtained.

We note that the DMQMC estimates of C_{gs} were obtained by sampling the reduced density matrix for a single pair of spins. Due to translational invariance, these estimates could have been improved by sampling the reduced density matrix for every neighboring pair and combining the results.

E. Temperature-dependent Renyi-2 entropy on the 6×6 lattice

Temperature-dependent Renyi-2 entropies were calculated for various sublattices of the 6×6 lattice using the replica approach described in Sec. III. Translational invariance was used to improve the quality of the statistics and importance

sampling was once again applied. Note that, unlike other calculations presented, the initial population spanned the entire Hilbert space rather than just the $M_S = 0$ subspace in order to capture the correct behavior of S_2 at nonzero temperatures. Results are presented in Fig. 4 for square sublattices and in Fig. 5 for strip sublattices. We were unable to find results in the literature for comparison.

VI. SIGN PROBLEM IN DMQMC

As discussed in Sec. II, the annihilation step in FCIQMC leads to the characteristic population dynamics, whereby the sign problem can only be overcome once a critical psip population (the plateau) has been exceeded. Since the annihilation steps in FCIQMC and DMQMC are identical, we would expect DMQMC to possess similar population dynamics. The annihilation rate for a given psip population in DMQMC will be significantly smaller than in FCIQMC because the number of density matrix elements is the square of the number of basis functions. As such, it is to be expected that the plateau height for a given Hamiltonian will be higher in DMQMC than in FCIQMC.

To investigate this issue, we considered the antiferromagnetic Heisenberg model on a 4×4 triangular lattice, for which an exact diagonalization is easily performed. The triangular




FIG. 4. (Color online) Renyi-2 entropy for square sublattices of a 6×6 antiferromagnetic Heisenberg lattice. Error bars, where not shown, are smaller than the line thickness. The inset zooms in on the 4×4 sublattice, which has the largest error bars. S_2 values were taken from every 25th iteration and averaged over 16–180 loops over the temperature range.




FIG. 5. (Color online) Renyi-2 entropy for strip sublattices of a 6×6 lattice. Error bars are shown in the inset (color scheme as in the main plot) for large β , where they are largest, and are smaller than the linewidth where not visible. S_2 values were taken from every 25th iteration and averaged over 32–240 loops over the temperature range. As expected, S_2 for the 2×6 and 4×6 sublattices tend to the same value in the zero-temperature limit.

lattice is the archetypal example of a frustrated lattice and has a severe sign problem. We carried out a single β loop and allowed the population to grow with a fixed shift while simultaneously investigating the accuracy of the DMQMC energy estimate. Figure 6(a) demonstrates that the population plateaus as expected. Figure 6(b) shows the accuracy of the energy estimate throughout the plateau period.

At high temperatures accurate results are obtained. This is also the case in other finite-temperature methods, where it is generally found that the sign problem is less severe for small β ; indeed, there is no sign problem at all at infinite temperature. However, at lower temperatures the energy estimates suffer large fluctuations, and it is not until the population exceeds the plateau that a good agreement with exact results is once




FIG. 6. (Color online) Population dynamics and energy estimate for the antiferromagnetic Heisenberg model on a 4×4 triangular lattice with periodic boundary conditions. A fixed shift was used and a single β loop performed. (a) shows the emergence of a plateau in the psip population and (b) shows the exact energy and the DMQMC estimate of the energy as functions of β . Once the plateau has been exited, the DMQMC energy is in good agreement with the exact results. The severity of the sign problem increases with inverse temperature.

again obtained. The plateau height occurs at $\sim 1.75 \times 10^8$ psips. For comparison, the plateau population in an FCIQMC calculation of the same system is at $\sim 2.0 \times 10^4$ psips. We find (in this case) that the DMQMC plateau height is approximately the square of the plateau height in FCIQMC, matching the increase in the size of the space being sampled. It is possible to obtain accurate results for the entire temperature range simply by starting the simulation with an initial population greater than that of the plateau. Moreover, even if the plateau cannot be reached due to memory restrictions, one can nevertheless systematically reach lower temperatures by increasing the population of psips.

The sign problem in the antiferromagnetic Heisenberg model on a triangular lattice is severe in both FCIQMC and DMQMC [15]. However, the efficiency of the annihilation procedure in FCIQMC varies substantially with the system studied and with the basis set used [15]. Furthermore, the initiator approximation (i-FCIQMC) proposed by Cleland *et al.* [17] has been shown to reduce the memory requirements for FCIQMC calculations by several orders of magnitude in many cases. In i-FCIQMC, spawning events onto previously unoccupied sites are forbidden unless the psip population of the parent site exceeds a threshold. This increases the annihilation rate relative to FCIQMC and ameliorates the sign problem at the expense of introducing a systematically improvable approximation. We have yet to investigate the DMQMC equivalent of the initiator method. It is expected that ground-state properties will be available, but it is not yet clear to what extent the modified spawning will affect finite-temperature results.

VII. DISCUSSION

This article has described DMQMC, a quantum Monte Carlo method that allows direct sampling of the finite-temperature and ground-state density matrices in a discrete basis. The validity of the method has been verified by reproducing exact and well-established results for some small systems, including the calculation of the concurrence of one-dimensional spin rings by directly sampling reduced density matrices. In all cases investigated, DMQMC has proved capable and accurate.

The introduction of an importance-sampling procedure allows larger lattices to be investigated—the largest system that we have simulated successfully to date is a 10×10 antiferromagnetic Heisenberg model. Larger systems could easily be tackled at high temperatures, but it is unlikely that our very simple approach to importance sampling will allow simulations of lattices of more than 10×10 sites over the entire range of temperatures from infinity to zero. If a better importance-sampling procedure can be devised, as we believe likely, many more avenues will be opened.

Like FCIQMC, DMQMC uses an annihilation procedure that allows it to overcome the sign problem if a system-specific population of psips is reached. In FCIQMC the sign problem can often be ameliorated, for example by changing to a more appropriate basis set [15,46] or by applying the initiator approximation [17]. Given the similarities between DMQMC and FCIQMC, it is likely that these ideas and future developments in FCIQMC will also apply to DMQMC. Due

to the relative youth of FCIQMC, the rate of theoretical and algorithmic improvements is rapid [17,21,47–49].

Obtaining expectation values of operators that are nonlinear in the full density matrix or functions of a reduced density matrix is challenging. We used the replica approach [33], in which two DMQMC simulations are performed side by side, to obtain unbiased estimates of the Renyi-2 entropy for several subsystems of a Heisenberg lattice across a range of temperatures.

The replica approach might also be useful in FCIQMC for unbiased sampling of expectation values of operators that do not commute with the Hamiltonian. However, a general calculation of this type would require running over every pair of psips in both calculations, which would be computationally demanding and have a major impact on the efficient parallelization of the FCIQMC algorithm.

The lattices studied with DMQMC so far are small by the standards of path-integral quantum Monte Carlo methods. However, the unique and defining feature of DMQMC is that it samples the full density matrix. This allows the calculation of arbitrary expectation values and quantum information measures that are inaccessible using other quantum Monte Carlo methods. We wish to finish by emphasizing that, while the applications presented here have been to lattice models, the formulation of DMQMC is general. For example, DMQMC could be used to calculate the two-electron reduced density matrix for small molecules, which could then be used to construct accurate exchange-correlation functionals for density-functional theory [50].

ACKNOWLEDGMENTS

We are grateful for enlightening discussions with Professor Terry Rudolph, Dr. Alex Thom, and Tom Poole. This work made use of the Imperial College London High Performance Computing facilities. J.S.S. and W.M.C.F. acknowledge the stimulating research environment provided by the Thomas Young Centre under Grant No. TYC-101. N.S.B. was supported by an Undergraduate Research Opportunities Scholarship in the Centre for Doctoral Training on Theory and Simulation of Materials at Imperial College funded by EPSRC under Grant No. EP/G036888/1.

APPENDIX A: CONCURRENCE

For the special case of a two-qubit state, the entanglement of formation can be calculated from a quantity known as the concurrence. Given a reduced density matrix ρ_A , where in this case A refers to a subsystem of two qubits, the concurrence is defined as

$$\mathcal{C}(\rho_A) \equiv \max(0, \gamma_1 - \gamma_2 - \gamma_3 - \gamma_4), \quad (\text{A1})$$

where $\gamma_1 > \gamma_2 > \gamma_3 > \gamma_4$ are the eigenvalues of the matrix

$$\mathbf{R} = \sqrt{\sqrt{\rho_A} \tilde{\rho}_A \sqrt{\rho_A}} \quad (\text{A2})$$

and

$$\tilde{\rho}_A = (\sigma_y \otimes \sigma_y) \rho_A^* (\sigma_y \otimes \sigma_y), \quad (\text{A3})$$

with σ_y the Pauli spin matrix for the y direction. This expression is only valid in the standard basis set.

The value of the concurrence \mathcal{C} ranges from zero to one and is monotonically related to the entanglement of formation [51]; the concurrence can therefore be regarded as a measure of entanglement. The entanglement of formation for two qubits [52] is given by

$$\mathcal{E}(\mathcal{C}) = h\left(\frac{1 + \sqrt{1 - \mathcal{C}^2}}{2}\right), \quad (\text{A4})$$

where

$$h(x) = -x \log_2 x - (1 - x) \log_2 (1 - x). \quad (\text{A5})$$

If the Hamiltonian and thus the reduced density matrix are real (as is the case in the Heisenberg model), the calculation of \mathcal{C} reduces to the calculation of the moduli of the eigenvalues of

$$\mathbf{R} = \rho_A(\sigma_y \otimes \sigma_y). \quad (\text{A6})$$

Since \mathbf{R} is a 4×4 matrix, it is trivial to compute the concurrence by direct diagonalization once ρ_A is known.

Using DMQMC, it is possible to sample the unnormalized reduced density matrix as described in Sec. III and thus to estimate the concurrence using

$$\langle \mathcal{C} \rangle = \frac{\max(0, \gamma_1 - \gamma_2 - \gamma_3 - \gamma_4)}{\text{Tr}(\rho_A)}, \quad (\text{A7})$$

where ρ_A is now the unnormalized reduced density matrix and $\{\gamma_i\}$ are the eigenvalues of the unnormalized matrix \mathbf{R} .

APPENDIX B: IMPORTANCE SAMPLING

The DMQMC evolution equation is

$$\frac{d\rho_{ij}}{d\beta} = \frac{1}{2} \sum_k (T_{ik} \rho_{kj} + \rho_{ik} T_{kj}), \quad (\text{B1})$$

where ρ_{ij} is an element of $\hat{\rho}$ in the many-particle basis chosen for the simulation. Instead of ρ_{ij} , we would like to sample the importance-sampled density matrix

$$\tilde{\rho}_{ij} = \frac{\rho_{ij}}{W_{E(i,j)}}, \quad (\text{B2})$$

where $E(i, j)$ is the excitation level of the pair (i, j) and W_α is defined in Eq. (22). Following the standard procedure of importance sampling, we introduce a trial function

$$\rho_{ij}^T = \frac{1}{W_{E(i,j)}}. \quad (\text{B3})$$

This matrix is symmetric, $\rho_{ij}^T = \rho_{ji}^T$, as $E(i, j) = E(j, i)$. The importance-sampled density matrix then has components $\tilde{\rho}_{ij} = \rho_{ij}^T \rho_{ij}$, where *no* summation is performed over indices. Multiplying Eq. (B1) by ρ_{ij}^T yields the following evolution equation for $\tilde{\rho}_{ij}$:

$$\frac{d(\rho_{ij}^T \rho_{ij})}{d\beta} = \frac{1}{2} \sum_k (\rho_{ij}^T T_{ik} \rho_{kj} + \rho_{ik} \rho_{ij}^T T_{kj}) \quad (\text{B4})$$

$$= \frac{1}{2} \sum_k \left[\left(\rho_{ij}^T T_{ik} \frac{1}{\rho_{kj}^T} \right) (\rho_{kj}^T \rho_{kj}) + (\rho_{ik}^T \rho_{ik}) \left(\rho_{ij}^T T_{kj} \frac{1}{\rho_{ik}^T} \right) \right] \quad (\text{B5})$$

$$\Rightarrow \frac{d\tilde{\rho}_{ij}}{d\beta} = \frac{1}{2} \sum_k \left[\left(\rho_{ij}^T T_{ik} \frac{1}{\rho_{kj}^T} \right) \tilde{\rho}_{kj} + \tilde{\rho}_{ik} \left(\rho_{ij}^T T_{kj} \frac{1}{\rho_{ik}^T} \right) \right]. \quad (\text{B6})$$

The above differential equation is entirely analogous to Eq. (15). As such, the finite-difference version of Eq. (B6) can be simulated in an almost identical manner to the standard DMQMC algorithm: the extra factors of ρ^T simply act to alter the spawning probabilities. Consider the case where the excitation level of (i, j) is γ and excitation level of (k, j) is

$\gamma - 1$. The probability that a spawning attempt from (k, j) to (i, j) is successful is altered by the factor

$$\rho_{ij}^T \frac{1}{\rho_{kj}^T} = \frac{W_{E(k,j)}}{W_{E(i,j)}} = \frac{W_{\gamma-1}}{W_{\gamma}} = P_{\gamma-1,\gamma}, \quad (\text{B7})$$

where Eq. (22) has been used to simplify the expression. Similarly, when spawning in the opposite direction, the probability is altered by the reciprocal of this factor. Spawning events that do not alter the excitation level are unaffected.

-
- [1] B. L. Hammond, W. A. Lester, Jr., and P. J. Reynolds, *Monte Carlo Methods in Ab Initio Quantum Chemistry* (World Scientific, Singapore, 1994).
- [2] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, *Rev. Mod. Phys.* **73**, 33 (2001).
- [3] M. H. Kalos, *Phys. Rev.* **128**, 1791 (1962).
- [4] D. M. Ceperley and M. H. Kalos, in *Monte Carlo Methods in Statistical Physics*, edited by K. Binder (Springer-Verlag, Berlin, 1979), p. 145.
- [5] J. B. Anderson, *J. Chem. Phys.* **63**, 1499 (1975).
- [6] J. B. Anderson, *J. Chem. Phys.* **65**, 4121 (1976).
- [7] J. W. Moskowitz, K. E. Schmidt, M. A. Lee, and M. H. Kalos, *J. Chem. Phys.* **77**, 349 (1982).
- [8] P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester, Jr., *J. Chem. Phys.* **77**, 5593 (1982).
- [9] R. Gaudoin and J. M. Pitarke, *Phys. Rev. Lett.* **99**, 126406 (2007).
- [10] J. Casulleras and J. Boronat, *Phys. Rev. B* **52**, 3654 (1995).
- [11] R. P. Feynman, *Phys. Rev.* **91**, 1291 (1953).
- [12] A. W. Sandvik, *AIP Conf. Proc.* **1297**, 135 (2010).
- [13] G. H. Booth, A. J. W. Thom, and A. Alavi, *J. Chem. Phys.* **131**, 054106 (2009).
- [14] G. H. Booth, A. Gruneis, G. Kresse, and A. Alavi, *Nature (London)* **493**, 365 (2012).
- [15] J. S. Spencer, N. S. Blunt, and W. M. C. Foulkes, *J. Chem. Phys.* **136**, 054110 (2012).
- [16] G. H. Booth and A. Alavi, *J. Chem. Phys.* **132**, 174104 (2010).
- [17] D. M. Cleland, G. H. Booth, and A. Alavi, *J. Chem. Phys.* **132**, 041103 (2010).
- [18] D. Cleland, G. H. Booth, C. Overy, and A. Alavi, *J. Chem. Theor. Comput.* **8**, 4138 (2012).
- [19] J. J. Shepherd, G. Booth, A. Gruneis, and A. Alavi, *Phys. Rev. B* **85**, 081103 (2012).
- [20] C. Daday, S. Smart, G. H. Booth, A. Alavi, and C. Filippi, *J. Chem. Theor. Comput.* **8**, 4441 (2012).
- [21] F. R. Petruzielo, A. A. Holmes, H. J. Changlani, M. P. Nightingale, and C. J. Umrigar, *Phys. Rev. Lett.* **109**, 230201 (2012).
- [22] M. H. Kalos and F. Pederiva, *Phys. Rev. Lett.* **85**, 3547 (2000).
- [23] J. Carlson and M. H. Kalos, *Phys. Rev. C* **32**, 1735 (1985).
- [24] J. B. Anderson, C. A. Traynor, and B. M. Boghosian, *J. Chem. Phys.* **95**, 7418 (1991).
- [25] S. Zhang and M. H. Kalos, *Phys. Rev. Lett.* **67**, 3074 (1991).
- [26] Z. Liu, S. Zhang, and M. H. Kalos, *Phys. Rev. E* **50**, 3220 (1994).
- [27] M. H. Kalos, *Phys. Rev. E* **53**, 5420 (1996).
- [28] C. J. Umrigar, M. P. Nightingale, and K. J. Runge, *J. Chem. Phys.* **99**, 2865 (1993).
- [29] L. Amico, R. Fazio, A. Osterloh, and V. Vedral, *Rev. Mod. Phys.* **80**, 517 (2008).
- [30] J. McMinis and N. M. Tubman, *Phys. Rev. B* **87**, 081108 (2013).
- [31] B. Swingle, J. McMinis, and N. M. Tubman, *Phys. Rev. B* **87**, 235112 (2013).
- [32] N. M. Tubman and J. McMinis, [arXiv:1204.4731](https://arxiv.org/abs/1204.4731) [cond-mat.str-el].
- [33] M. B. Hastings, I. González, A. B. Kallin, and R. G. Melko, *Phys. Rev. Lett.* **104**, 157201 (2010).
- [34] U. Glaser, H. Büttner, and H. Fehske, *Phys. Rev. A* **68**, 032318 (2003).
- [35] T. Roscilde, P. Verrucchi, A. Fubini, S. Haas, and V. Tognetti, *Phys. Rev. Lett.* **94**, 147208 (2005).
- [36] A. W. Sandvik, *Phys. Rev. Lett.* **95**, 207203 (2005).
- [37] F. Alet, S. Capponi, N. Laflorencie, and M. Mambrini, *Phys. Rev. Lett.* **99**, 117204 (2007).
- [38] Y.-C. Lin and A. W. Sandvik, *Phys. Rev. B* **82**, 224414 (2010).
- [39] P. Knowles and N. Handy, *Chem. Phys. Lett.* **111**, 315 (1984).
- [40] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevB.89.245124> and doi:10.6084/m9.figshare.655964 for raw data.
- [41] A. W. Sandvik, *Phys. Rev. B* **56**, 11678 (1997).
- [42] H. Flyvbjerg and H. Petersen, *J. Chem. Phys.* **91**, 461 (1989).
- [43] The definitions of M^2 given in Eq. (24) and in Ref. [44] have slightly different normalization factors. The result quoted from Ref. [44] has been renormalized accordingly.
- [44] H. J. Schulz, T. A. L. Ziman, and D. Poilblanc, *J. Phys. I France* **6**, 675 (1996).
- [45] K. M. O'Connor and W. K. Wootters, *Phys. Rev. A* **63**, 052302 (2001).
- [46] M. H. Kolodrubetz, J. S. Spencer, B. K. Clark, and W. M. Foulkes, *J. Chem. Phys.* **138**, 024110 (2013).
- [47] A. J. W. Thom, *Phys. Rev. Lett.* **105**, 263004 (2010).
- [48] A. Mukherjee and Y. Alhassid, *Phys. Rev. A* **88**, 053622 (2013).
- [49] A. Roggero, A. Mukherjee, and F. Pederiva, *Phys. Rev. B* **88**, 115138 (2013).
- [50] R. Q. Hood, M. Y. Chou, A. J. Williamson, G. Rajagopal, R. J. Needs, and W. M. C. Foulkes, *Phys. Rev. Lett.* **78**, 3350 (1997).
- [51] S. Hill and W. K. Wootters, *Phys. Rev. Lett.* **78**, 5022 (1997).
- [52] W. K. Wootters, *Quant. Inf. Comp.* **1**, 27 (2001).