

Valence-bond quantum Monte Carlo algorithms defined on trees

Andreas Deschner* and Erik S. Sørensen†

Department of Physics and Astronomy, McMaster University, Hamilton, Canada L8S 4M1

(Received 30 January 2014; revised manuscript received 5 July 2014; published 10 September 2014)

We present a class of algorithms for performing valence-bond quantum Monte Carlo of quantum spin models. Valence-bond quantum Monte Carlo is a projective $T = 0$ Monte Carlo method based on sampling of a set of operator strings that can be viewed as forming a treelike structure. The algorithms presented here utilize the notion of a worm that moves up and down this tree and changes the associated operator string. In quite general terms, we derive a set of equations whose solutions correspond to a whole class of algorithms. As specific examples of this class of algorithms, we focus on two cases. The *bouncing worm* algorithm, for which updates are *always accepted* by allowing the worm to bounce up and down the tree, and the *driven worm* algorithm, where a single parameter controls how far up the tree the worm reaches before turning around. The latter algorithm involves only a *single bounce* where the worm turns from going up the tree to going down. The presence of the control parameter necessitates the introduction of an acceptance probability for the update.

DOI: [10.1103/PhysRevE.90.033304](https://doi.org/10.1103/PhysRevE.90.033304)

PACS number(s): 05.10.Ln, 75.40.Mg

I. INTRODUCTION

Projective techniques are often used for determining the ground-state properties of strongly correlated models defined on a lattice. They were initially developed for nonlattice models [1] and then used for the study of fermionic lattice models [2]. They were subsequently applied to quantum spin models [3–8] as well as other models. The underlying idea is easy to describe. For a lattice Hamiltonian H , it is possible to choose a constant K such that the dominant eigenvalue E of $K\mathbb{1} - H$ corresponds to the ground-state wave function of H , $|\Psi_0\rangle$. We can then use $P = K\mathbb{1} - H$ as a projective operator in the sense that the repeated application of P to a trial wave function $P^n|\Psi_T\rangle$ will approach $E^n|\Psi_0\rangle$ for large n . Hence, if n can be taken large enough, $|\Psi_0\rangle$ can be projected out in this manner provided that $\langle\Psi_0|\Psi_T\rangle \neq 0$. Some variants of this approach are often referred to as Green's functions Monte Carlo (GFMC) [2,5–8]. Other projective operators such as $\exp(-\tau H)$ can be used depending on the model and its spectrum. For a review, see Refs. [9,10]. The convergence of such projective techniques may be nontrivial as can be shown by analyzing simple models [11]. If $P|\Psi_T\rangle$ can be evaluated exactly, this projective scheme is equivalent to the power method as used in exact diagonalization studies. As the number of sites in the lattice model is increased, exact evaluation quickly becomes impossible and Monte Carlo methods (projector Monte Carlo) have to be used.

The efficiency of the Monte Carlo sampling is crucial for the performance of implementations of the projective method and detailed knowledge of such Monte Carlo methods is of considerable importance. Here, we have investigated a class of Monte Carlo algorithms for projective methods for lattice models. We discuss these algorithms within the context of quantum Monte Carlo where the projection is

performed in the valence-bond basis [3,4,12–15], so-called valence-bond quantum Monte Carlo (VBQMC). The algorithms are, however, applicable to projective techniques in any basis.

VBQMC was first developed by Liang [3,4] and then, starting 15 years later, significantly further developed by Sandvik and collaborators [12–15] and it is now widely used. Since its inception, VBQMC has been improved and generalized in several ways: it can be used on systems with spins with $S \neq \frac{1}{2}$ [4] and states with total $S_z = \frac{1}{2}$ [16]. An efficient sampling algorithm with loop updates is known for systems with $S = \frac{1}{2}$ [15].

As outlined above, VBQMC works by projecting onto the ground state by repeatedly acting on a trial state $|\Psi_T\rangle$ with $P = K\mathbb{1} - H$, where the constant K is chosen such that the ground state has the biggest eigenvalue. For Hamiltonians with bounded spectrum, such a K can always be found. For a simple quantum spin model defined on a lattice, we have

$$H = J \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_j = \sum_{\langle i,j \rangle} h_{ij} \quad (1)$$

and we can write $P = K\mathbb{1} - H = \sum O_{ij}$ as a sum over N_B bond operators O_{ij} . Taking P to the n th power then results in a sum over products of these bond operators O_{ij} :

$$P^n = \sum_a \underbrace{O_{i(a,1)j(a,1)} \cdots O_{i(a,n)j(a,n)}}_{n \text{ operators}} := \sum_a S_a. \quad (2)$$

Each instance of this product then forms a string S_a of bond operators of length n . When selecting such a string of length n , one has to make a choice between the N_B bond operators at each position in the string. It is possible to view the construction of such a string as a specific path in a decision tree (see Fig. 1).

Although the algorithms we present can be extended to higher spin models, we shall restrict the discussion to quantum spin models with $S = \frac{1}{2}$ where one usually takes $K = JN_B/4$. The action of the bond operators then takes an attractively simple form.

*deschna@mcmaster.ca

†sorensen@mcmaster.ca




FIG. 1. (Color online) The branching tree of length 5 for the selection of an operator string in a system with a Hamiltonian of N_B (here 6) terms. The operator that acts on the state first is chosen at the first node on the left. This node is called the root and the direction towards the root we define to be up. The operator that acts on the state last is at the end of the string. The two colored paths differ in the choice of the last three operators. The last three branches, thus, contribute different operators and weights (s_i, t_i). The resulting strings S and T are different.

In a valence-bond basis state spins are paired into singlets. A specific pairing of all spins is usually referred to as a covering. All such coverings form an overcomplete basis for the singlet subspace of the model. We shall only be concerned with models defined on a bipartite lattice in which case a given valence bond covering C for a lattice with N spins can be denoted by listing all $N/2$ pairs of $[i, j]$ with i on sublattice A and j on sublattice B . Here, $[i, j] = (|\uparrow_{i_A} \downarrow_{j_B}\rangle - |\downarrow_{i_A} \uparrow_{j_B}\rangle) / \sqrt{2}$. We label the initial covering (trial state) as C_0 . The action of an operator O_{ij} can take two forms [3,4,12]:

(1) The sites i and j are in a singlet before the action of the operator. Then, the action of the operator does not change the state and we can associate a weight of $w = 1$:

$$O_{ij}[i, j] = 1[i, j]. \quad (3)$$

(2) The sites i and j are not in a singlet before the action of the operator. Then, after the action of the operator, the sites i and j form a singlet. The sites they were originally connected to are also returned in a singlet state. Furthermore, the state is multiplied by a weight equal to $w = \frac{1}{2}$:

$$O_{ij}[i, k][l, j] = \frac{1}{2}[i, j][l, k]. \quad (4)$$

A particularly nice feature is that the application of any of the O_{ij} to any given covering yields a unique other covering and *not* a linear combination of coverings. Although convenient, this feature of projections in the valence-bond basis is not strictly necessary for the algorithms we discuss here as they can be adapted to the case where a linear combination of states are generated [17]. For a given operator string $S_a = \prod_k O_k$, we can associate a weight given by $W_a = \prod w_k$. The state $S_a C_0$ will contribute to the final projected estimate of the ground state with this weight. One can then sample the ground state by performing a random walk in the space of all possible strings S_a according to the weight W_a .

This way of sampling is used in all VBQMC implementations [3,4,12] and quite different from GFMC's sampling even though VBQMC and GFMC are closely related projective techniques. GFMC, as it is used in, for instance, Ref. [6], is usually performed in the S^z basis but can also be done in terms of the valence-bond basis [18]. In GFMC, the projection is done by stochastically evaluating the action of the whole projection operator on a trial state. This is done by introducing probabilistic "walkers." In contrast, as mentioned, in VBQMC a single state results and the strings S_a are sampled according to their weight. A different approach for a projective method on a lattice system based on reptation quantum Monte Carlo was proposed [19].

Clearly, the efficient sampling of states resulting from the stochastic projection of the trial state is a difficult problem. Here, we propose to use worm (cluster) algorithms for this purpose.

In Monte Carlo calculations, one averages over many configurations of the system which are generated with appropriate probabilities. Usually, this is done in a Markov chain, where one configuration is chosen as a variation of the last. One important feature of an efficient algorithm is that these consecutive configurations are as uncorrelated as possible. This led to the introduction of algorithms where whole clusters and not just single elements are changed going from one configuration to the next [20,21] or where all elements in the path of a *worm* are changed [22].

Here, we show how it is possible to adapt such worm algorithms for projections in the context of VBQMC. The algorithms we have studied are based on the notion of a worm moving around in the decision tree described above. As in earlier worm algorithms, the change of many elements is achieved by moving the worm based on local conditions [22–26] and one might refer to the algorithms as tree-worm algorithms. In general, the algorithms can be viewed as *directed* [25] algorithms.

When we update the string, we start with a worm at the end of the tree and move it up the tree (see Fig. 1). The worm then moves around in the tree and where it goes the operator string is changed. When the worm finds its way back to the bottom of the tree, the update is complete. We derive a set of simple equations governing the movement of the worm. The solution of these equations leads to parameters defining a class of algorithms. Quite generally, many solutions are possible, leaving significant room for choosing parameters that will lead to the most optimal algorithm.

We focus on two specific choices of parameters corresponding to two different algorithms. The *bouncing worm* algorithm, for which every update is accepted and the *driven worm* algorithm, for which the update is accepted with some probability. With the driven worm algorithm, one can choose at will how much of the operator string is on average changed in a successful update.

In order to test the algorithms, we calculate the ground-state energy of the isotropic Heisenberg chain. This quantity is easy to calculate with VBQMC and can be exactly computed using the Bethe ansatz. It is thus a very convenient quantity to test the algorithms with. The algorithms presented in this paper can, however, be used for the same calculations as other VBQMC implementations (see, e.g., [13]).

In Sec. II, we derive the general equations governing the movement of the worm. Section III contains a description of the specific implementation corresponding to the two choices of parameter solutions we have studied. The *bouncing worm* is detailed in Sec. III A while the *driven worm* algorithm is described in Sec. III B. The algorithms are then compared in Sec. IV. We present our conclusions in Sec. V.

II. TREE ALGORITHMS

We now turn to a discussion of the general framework for the algorithms we have investigated. We begin by deriving the equations governing their behavior in a general way. Let us take the Hamiltonian to have N_B terms. We now imagine a tree where each node indicates the decision to choose one of the N_B bond operators composing the string (see Fig. 1). Each branch of the tree corresponds to one of the N_B bond operators. A given operator string then corresponds to selecting a path in the tree. Consider two such paths S and T that are identical for the part of the operator string first applied to the trial state. The last three operators, however, differ. This leads to different weights, which we denote with s_i and t_i .

As it is done in most Monte Carlo methods, we set out to construct a Markov chain. Here, it is a chain of different strings. If the probabilities to go from one string to the next have detailed balance, the Markov chain contains the strings with the desired probability. For detailed balance, the probabilities for starting from operator string S and going to operator string T and reverse have to satisfy

$$\frac{P(S \rightarrow T)}{P(T \rightarrow S)} = \frac{t_3 t_4 t_5}{s_3 s_4 s_5}. \quad (5)$$

We can achieve this ratio of probabilities by imagining a worm (tree worm) working its way up the tree to the point p where it turns around and then working its way down again.

Let us call the valence-bond covering of the trial state C_0 . Up to numerical factors, the application of an operator string S of length n will yield a new valence-bond covering $SC_0 \propto C_n$. The worm is started by removing the last applied bond operator and considering the resulting covering C_{n-1} . A decision now has to be made if the worm is to continue “up” the tree by removing more bond operators from the string or if it should instead go “down” the tree by adding a new bond operator to the string. At each node in the tree, the decision to continue up or turn around is made according to a set of *conditional probabilities* $P(\text{up}|s)$ and $P(t|s)$. Here, $P(\text{up}|s)$ denotes the probability for going up after coming from a bond operator that carried weight s and $P(t|s)$ is the probability for turning around by applying a bond operator of weight t coming from an operator with weight s . Likewise, $P(s|\text{up})$ denotes the probability of choosing an operator with weight s given that the worm is coming from further up the tree. With these conditional probabilities, the left-hand side of Eq. (5) can be written as

$$\frac{P(S \rightarrow T)}{P(T \rightarrow S)} = \frac{P(t_5|\text{up})P(t_4|\text{up})P(t_3|s_3)P(\text{up}|s_4)P(\text{up}|s_5)}{P(\text{up}|t_5)P(\text{up}|t_4)P(s_3|t_3)P(s_4|\text{up})P(s_5|\text{up})}. \quad (6)$$

Clearly, Eq. (6) is satisfied if we choose

$$\frac{P(\text{up}|s)}{P(s|\text{up})} = \frac{c}{s} \quad \text{and} \quad \frac{P(t|s)}{P(s|t)} = \frac{t}{s}, \quad (7)$$

where c is an additional free parameter included for later optimization of the probabilities. If we can choose conditional probabilities with these properties, we can go between different operator strings always accepting the new string. The rejection probability is then zero. This is a very desirable property for any Monte Carlo algorithm. We mostly focus on so-called *zero bounce* algorithms for which if the worm turns around the probability for replacing a bond operator with the same operator is zero. Then, the two operator strings S and T are always different. This means that

$$P(s|s) = 0. \quad (8)$$

Quite generally, it is easy to find *many* solutions to the equations (7) leading to many Monte Carlo algorithms which can be tuned for efficiency.

We now focus on $S = \frac{1}{2}$ Heisenberg models defined on bipartite lattices. As has been described above, for these models only two weights can occur: $1, \frac{1}{2}$. The two weights correspond to the two different actions the bond operators can have on the state. It is 1 if the operator acts on two sites that are in a valence bond. The state is not altered under the action of such an operator. We call such operators diagonal. The weight is $\frac{1}{2}$ if the operator acts on two sites that are *not* in a valence bond. After the action of the operator, the two sites are connected by a bond as well as the sites they were connected to. We call such operators nondiagonal.

Because it depends on the state acted on, whether an operator is diagonal or nondiagonal is not an inherent feature of an operator. For two different states, the same operator may be diagonal for one state and nondiagonal for the other.

If a decision has to be made at the node at position m , the conditional probabilities depend on how many of the N_B bond operators will yield a weight of 1 (are diagonal) or $\frac{1}{2}$ (are nondiagonal) when applied to the present covering C_{m-1} . We shall denote these numbers by N_1 and $N_{1/2}$, respectively. When the worm is started, N_1 and $N_{1/2}$ therefore have to be calculated for C_{n-1} , if they are not already known from an earlier update. It is thus sensible to store N_1 or $N_{1/2}$ at all nodes. N_1 can only be zero at the node furthest up the tree (the root) and only if the trial state is chosen such that the Hamiltonian contains no diagonal operators. $N_{1/2}$ cannot be smaller than $N_B/2$.

We can now write an $(N_B + 1) \times (N_B + 1)$ matrix M of conditional probabilities for each node of the tree. The j th column of the matrix describes the probability for going in any of the $N_B + 1$ directions when coming from the direction j . For clarity, we order the rows and columns such that the first $N_{1/2}$ correspond to the nondiagonal operators and the next N_1 to the diagonal operators. The last column contains the probabilities for going down the tree when coming from above and the last row the probabilities for going up the tree when coming from below. The remaining part of the matrix describes the probabilities for replacing one operator with another when the worm turns from going up to going down. The matrix M

has the form

$$M = \left(\begin{array}{c|ccc|ccc|c} & \overbrace{\hspace{4cm}}^{N_{\frac{1}{2}}} & & & \overbrace{\hspace{4cm}}^{N_1} & & & \overbrace{\hspace{1cm}}^{\text{up}} \\ \hline P(\frac{1}{2}|\frac{1}{2}) & P(\frac{1}{2}'|\frac{1}{2}) & P(\frac{1}{2}'|\frac{1}{2}) & \dots & P(\frac{1}{2}'|\frac{1}{2}) & P(\frac{1}{2}|1) & P(\frac{1}{2}|1) & \dots & P(\frac{1}{2}|1) & P(\frac{1}{2}|\text{up}) \\ P(\frac{1}{2}'|\frac{1}{2}) & P(\frac{1}{2}|\frac{1}{2}) & P(\frac{1}{2}'|\frac{1}{2}) & \dots & \cdot & \cdot & \cdot & \dots & \cdot & \cdot \\ \cdot & \cdot & \cdot & & \cdot & \cdot & \cdot & & \cdot & \cdot \\ \cdot & \cdot & \cdot & & \cdot & \cdot & \cdot & & \cdot & \cdot \\ \hline P(\frac{1}{2}'|\frac{1}{2}) & P(\frac{1}{2}'|\frac{1}{2}) & P(\frac{1}{2}'|\frac{1}{2}) & \dots & P(\frac{1}{2}|\frac{1}{2}) & P(\frac{1}{2}|1) & P(\frac{1}{2}|1) & \dots & P(\frac{1}{2}|1) & P(\frac{1}{2}|\text{up}) \\ P(1|\frac{1}{2}) & P(1|\frac{1}{2}) & P(1|\frac{1}{2}) & \dots & P(1|\frac{1}{2}) & P(1|1) & P(1'|1) & \dots & P(1'|1) & P(1|\text{up}) \\ \cdot & \cdot & \cdot & & P(1|\frac{1}{2}) & P(1'|1) & P(1|1) & \dots & \cdot & \cdot \\ \cdot & \cdot & \cdot & & \cdot & \cdot & \cdot & & \cdot & \cdot \\ \cdot & \cdot & \cdot & & \cdot & \cdot & \cdot & & \cdot & \cdot \\ \hline P(1|\frac{1}{2}) & P(1|\frac{1}{2}) & P(1|\frac{1}{2}) & \dots & P(1|\frac{1}{2}) & P(1'|1) & P(1'|1) & \dots & P(1|1) & P(1|\text{up}) \\ P(\text{up}|\frac{1}{2}) & P(\text{up}|\frac{1}{2}) & P(\text{up}|\frac{1}{2}) & \dots & P(\text{up}|\frac{1}{2}) & P(\text{up}|1) & P(\text{up}|1) & \dots & P(\text{up}|1) & P(\text{up}|\text{up}) \end{array} \right),$$

where $P(s'|s)$ refers to the conditional probability of coming from an operator with weight s and going to a *different* operator with the same weight. As mentioned above, $P(\text{up}|s)$ denotes the probability for going up coming from an operator with weight s and $P(t|s)$ is the probability for turning around by choosing a bond operator of weight t coming from an operator with weight s . Likewise, $P(s|\text{up})$ denotes the probability of choosing an operator with weight s coming from further up the tree.

To shorten the notation, we introduce the shorthand

$$x = P(1/2'|1/2), \quad y = P(1/2|1), \quad z = P(1'|1). \tag{9}$$

Furthermore, we define the bounce probabilities

$$b_{1/2} = P(1/2|1/2), \quad b_1 = P(1|1), \quad b_u = P(\text{up}|\text{up}). \tag{10}$$

Here, it is implied that the probabilities are for going from one operator to the *same* operator. Finally, we also need to define the branching probabilities

$$u = P(1/2|\text{up}), \quad w = P(1|\text{up}), \tag{11}$$

from which it follows [using Eq. (7)] that

$$2cu = P(\text{up}|1/2), \quad cw = P(\text{up}|1). \tag{12}$$

The matrix M is then given by

$$M = \left(\begin{array}{cccc|cccc|c} b_{\frac{1}{2}} & x & x & \dots & x & y & y & \dots & y & u \\ x & b_{\frac{1}{2}} & x & \dots & x & y & y & \dots & y & u \\ \cdot & \cdot & \cdot & & \cdot & \cdot & \cdot & & \cdot & \cdot \\ \cdot & \cdot & \cdot & & \cdot & \cdot & \cdot & & \cdot & \cdot \\ \hline x & \cdot & \cdot & \dots & b_{\frac{1}{2}} & y & y & \dots & y & u \\ 2y & 2y & 2y & \dots & 2y & b_1 & z & \dots & z & w \\ \cdot & \cdot & \cdot & & 2y & z & b_1 & \dots & \cdot & \cdot \\ \cdot & \cdot & \cdot & & \cdot & \cdot & \cdot & & \cdot & \cdot \\ \cdot & \cdot & \cdot & & \cdot & \cdot & \cdot & & \cdot & \cdot \\ \hline 2y & 2y & 2y & \dots & 2y & z & z & \dots & b_1 & w \\ 2cu & 2cu & 2cu & \dots & 2cu & cw & cw & \dots & cw & b_u \end{array} \right). \tag{13}$$

There are only few distinct entries in the matrix because there are only two different bond operators. For less special Hamiltonians (for example with higher spin), there may be more distinct entries. For most common Hamiltonians, the

matrix will contain very many redundant entries. Storing its information at every node can therefore easily be done.

The requirement that this matrix be stochastic (i.e., some branch is chosen with probability one) means that the entries in

each column have to sum to 1. This leads to the set of equations

$$\begin{aligned} 1 &= N_{1/2}u + N_1w + b_u, \\ 1 &= N_{1/2}y + (N_1 - 1)z + cw + b_1, \\ 1 &= (N_{1/2} - 1)x + N_12y + 2cu + b_{1/2}. \end{aligned} \quad (14)$$

These simple equations are the central equations governing the behavior of the algorithms. To find an algorithm, we need to solve these three equations with the constraints that $0 \leq x, y, z, b_{1/2}, b_1, b_u, u, w \leq 1$, a straightforward problem.

At the root, the equations are modified slightly: since it is not possible to go further up the tree, $2cu, cw, b_u$ are not meaningful and can be set to zero. For convenience, we set $b_1 = x$ and $b_2 = z$ at the root. This allows one to just choose diagonal operators twice as often as nondiagonal operators. Since the number of diagonal operators does not change at the root, a table generated at the beginning of the calculation suffices to perform this task.

It can be very useful to choose different c 's at different nodes. Then, calculating the probabilities to choose operators according to the rules introduced in this section will not lead to an algorithm with detailed balance because c_i from different strings will not cancel in Eq. (6). It is necessary to work with an acceptance probability. We find

$$\frac{P(S \rightarrow T)}{P(T \rightarrow S)} = \frac{t_3 t_4 t_5}{s_3 s_4 s_5} \frac{c_4^S c_5^S}{c_4^T c_5^T} \frac{P_{\text{acc}}(S \rightarrow T)}{P_{\text{acc}}(T \rightarrow S)}. \quad (15)$$

Here, c_i^S and c_i^T denote c at the different nodes in the strings S and T , respectively. To validate the algorithm, we must therefore introduce an acceptance probability that must cancel the factor $(c_4^S c_5^S)/(c_4^T c_5^T)$. This can be achieved by choosing

$$P_{\text{acc}}(S \rightarrow T) = \min \left(1, \frac{c_4^T c_5^T}{c_4^S c_5^S} \right), \quad (16)$$

meaning that when a new string is generated through a worm move, it is accepted with this probability.

The very first string is filled just as one fills the string any time the worm reaches the root. This first string is always accepted. Since we always start from the bottom of the tree (the last operator applied), the worm algorithms presented in this paper always change a block of consecutive branches at the end of the string. This is favorable to changes across the whole string because changes far up the string might be partly undone by changes closer to the end of the string [14]. In this way, the most important part of the string is updated most substantially.

It is also important to note that the algorithm will conserve certain topological numbers. For instance, for a two-dimensional system $S = \frac{1}{2}$ Heisenberg model the number of valence bonds crossing a cut in the x or y direction is either odd or even. Hence, the initial covering, C_0 is characterized by these two parities. It is easy to see that the application of P to any covering can not change these parities and they are therefore preserved under the projection.

III. IMPLEMENTATIONS OF TREE-WORM ALGORITHMS

As is explained in the last section, many different algorithms can be found because many different solutions to the equations (14) exist. In this section, we present two different algorithms: one pure worm algorithm where every update is accepted (the *bouncing worm* algorithm), and an algorithm that allows for control over how far in the tree updates are attempted (the *driven worm* algorithm). To test and compare the different algorithms, we calculate the ground-state energy of the antiferromagnetic Heisenberg chain.

The same quantities as with other VBQMC implementations can be calculated with the worm algorithms presented here. How this can be done is explained in earlier publications about VBQMC [12–14]. Since we only use the ground-state energy E_0 , we restrict ourselves to an explanation of its calculation.

The Néel state $|\text{Néel}\rangle$ has equal overlap with all valence-bond states. This can be used to very directly estimate the ground-state energy E_0 :

$$\begin{aligned} E_0 &= \frac{\langle \text{Néel} | H | \Psi_0 \rangle}{\langle \text{Néel} | \Psi_0 \rangle} \\ &= \lim_{n \rightarrow \infty} \frac{\langle \text{Néel} | H P^n | C_0 \rangle}{\langle \text{Néel} | P^n | C_0 \rangle} \\ &= \lim_{n \rightarrow \infty} \sum_{a=1}^{N_B^n} \frac{\langle \text{Néel} | H S_a | C_0 \rangle}{\sum_{a=1}^{N_B^n} \langle \text{Néel} | S_a | C_0 \rangle} \\ &= \lim_{n \rightarrow \infty} \sum_{a=1}^{N_B^n} \frac{W_a}{\sum_{a=1}^{N_B^n} W_a} \frac{\langle \text{Néel} | H | C_a \rangle}{\langle \text{Néel} | C_a \rangle}. \end{aligned} \quad (17)$$

If we take $E_a C_b = H C_a$ and assume that the Monte Carlo sampling will visit strings according to their weight W_a , then for a Monte Carlo sequence of length N of independent strings we find

$$E_0 = \frac{1}{N} \sum_{a=1}^N E_a, \quad (18)$$

where again we have used the fact that $\langle \text{Néel} | C \rangle$ is independent of the covering C .

To analyze the correlation properties of the worm algorithms, we use the energy autocorrelation time, which we take to be the number of updates it takes the energy autocorrelation function

$$A_E(t) = \frac{\langle E_i E_{i+t} \rangle - \langle E \rangle^2}{\langle E^2 \rangle - \langle E \rangle^2} \quad (19)$$

to decay to 0.1. The results of all update attempts enter the calculation of the expectation values. The autocorrelation time is measured in units of one update attempt. The shorter the autocorrelation time is, the more substantially do Monte Carlo steps decorrelate samples.

This way of estimating the autocorrelation time does not take into account the possibility of subleading modes with large autocorrelation times. It is intuitive, however, and sufficient in the context of this first exhibition of the worm algorithms.

The reader should be advised that other observables may have different autocorrelation times. We expect this to affect all algorithms used in this paper in the same way. If not stated otherwise, an operator string of 20 000 operators was used for calculations with worm algorithms. To calculate the errors given for estimates of E_0 , we binned measurements taken one autocorrelation time apart. The number of measurements in one bin was chosen such that the size of the error did not change appreciably when multiplying or dividing it by two.

A. Bouncing worm algorithm

The first algorithm we discuss is the bouncing worm algorithm. Only a few of the variables that appear in the equations (14) are chosen to be nonzero. We choose to set

$$x = y = 0, \quad b_1 = b_{1/2} = 0, \quad (20)$$

while $z = P(I'|1) \neq 0$ as is u, w . We leave b_u as a parameter that can be zero or nonzero allowing for tuning of the algorithm. With this choice, when the worm is moving up the tree, the only possibility for it to turn around is by opting to replace one diagonal operator with another diagonal operator. The c_i are chosen to be the same at all nodes: $c_i = c$.

The equations for the nonzero parameters are then

$$\begin{aligned} u &= \frac{1}{2c}, \\ w &= \frac{1}{N_1} \left(1 - \frac{N_{1/2}}{2c} - b_u \right), \\ z &= \frac{1 - cw}{N_1 - 1}. \end{aligned} \quad (21)$$

The requirements that $z, w > 0$ imply that

$$\frac{N_{1/2}}{2(1 - b_u)} \leq c \leq \frac{N_B + N_1}{2(1 - b_u)}. \quad (22)$$

To satisfy Eq. (22) with node independent c , we set

$$c = \frac{N_B}{2(1 - b_u)}. \quad (23)$$

With this choice of parameters, we find the probability to go up the tree if the worm is at a node with a nondiagonal operator to be

$$P(\text{up}|1/2) = 2cu = 1, \quad (24)$$

for any b_u . Likewise, if the worm is at a node with a diagonal operator, the probability to go up is given by

$$P(\text{up}|1) = cw = 1/2, \quad (25)$$

independent of b_u . The probability for going up the tree is therefore independent of b_u .

We define the penetration depth (p. depth), which we denote by r , as the maximal height that the worm reaches. The mean p. depth $\langle r \rangle$ is a measure of how much of the string is on average changed per update. The actual length of the worm is denoted by l and with $b_u = 0$ we find $l = 2r$. The penetration depth r will determine how much the operator string is changed. Obviously, it is desirable to have the worm reach as far up the tree as possible. It is possible to force the worm farther up the tree by having it bounce back to going up after it has turned to go down (see Fig. 2). In that case, the actual length of the




FIG. 2. (Color online) A possible path that contains one bounce and connects the string S and the string T . The worm first goes up to the node 2 where it turns to go down to node 3. The worm bounces back and goes all the way to node 1. Then the worm turns around and does not bounce again.

worm, l , will then be substantially different from twice the penetration depth since the worm can turn many times, a point we shall return to later. Such bounces occur with a likelihood of b_u which was left as a free parameter and can now be used as a tuning parameter.

The algorithm is straightforward to implement and the acceptance probability for a worm update is 1. The move is *always* accepted. Specific details of an implementation of the bouncing worm update can be found in the Appendix, Sec. A 1.

We begin by discussing the case of $b_u = 0$. In this case, the worm first moves up the operator string, turns around once, and then proceeds down to the bottom of the tree. It does not go back up the operator string since $b_u = 0$. In order to measure the performance of the algorithm, we did calculations on an antiferromagnetic Heisenberg chain with 50 sites using an operator string of length 100 000. As can be seen in Table I, this leads to a rather small mean penetration depth (p. depth) of about 5. The maximal penetration depth of 50 is substantially larger. Both these numbers are, however, substantially smaller than the length of the operator string (100 000) and it appears that the algorithm with b_u is not very effective.

We now turn to the case $b_u \neq 0$. In this case, the worm can now switch directions many times during construction (see Fig. 2). The results for the mean and maximal penetration

TABLE I. Data for several runs at different b_u . At $b_u \approx 0.25$ increasing the bounce probability starts to significantly slow down the algorithm. The last column contains the runtimes divided by the runtime for $b_u = 0$. The data were generated with an operator string of 100 000 operators. The maximal penetration and the expected slowdown could thus not be resolved for $b_u = 0.2790$. We used 10^6 measurements and a chain with 50 sites.

b_u	Mean p. depth	Max p. depth	Slowdown
0.0000	4.561(4)	50	1
0.2500	7.38(1)	305	1.7
0.2750	10.44(3)	1,465	9.7
0.2789	15.64(9)	41,010	316.7
0.2790	19.7(5)	>100,000	<5,535.7




FIG. 3. (Color online) The autocorrelation time of the energy as a function of system size N . The lines indicate tentative power-law fits to the data at large N with a power of 2.21 for $b_u = 0.26$ and 2.16 for $b_u = 0.275$. Operator strings of 20 000 operators were used.

depths are also listed in Table I. As b_u is increased from zero, the maximal penetration depth first increases very slowly until about $b_u = 0.25$. It then grows dramatically and, not surprisingly, reaches the length of the operator string. This occurs at $b_u \approx 0.2790$. At the same time, the mean penetration depth only increases by a factor of roughly 4, from 5 to about 20. For bounce probabilities bigger than $b_u \approx \frac{1}{4}$, the program is slowed down significantly compared to the algorithm with $b_u = 0$ as indicated in the last column in Table I. Thus, even though a large penetration depth is desirable, the computational cost can become so big that increasing b_u might not be worthwhile.

In contrast to the maximal penetration depth, the mean penetration depth grows very slowly for the values of b_u we have been able to study. For computations of reasonable computational cost, it never reaches the size of the system and thus also not the length of the operator string which has to be chosen to be several times the size of the system. That only a small part of the string is updated regularly is directly reflected in the energy autocorrelation time (see Fig. 3). The number of bonds that can change in one update of the worm calculations is twice the penetration depth. Typical updates never reach far into the operator string. Thus, the bigger the system is, the less it is perturbed by the update and the more correlated are the energies measured after consecutive updates.

As shown in Fig. 3, increasing b_u decreases the autocorrelation time. However, for large system sizes, the overall scaling of the autocorrelation time with the system size appears independent of b_u . At $b_u = 0.26$ we find a power law with an exponent of 2.21 while a slightly larger $b_u = 0.275$ yields a power of 2.16.

The shorter autocorrelation time for larger b_u comes with more work that has to be done per update since the worm is much longer and more operators have to be changed per update. It is thus instructive to look at a scaled autocorrelation time that takes this into account. Such a scaled autocorrelation time can be obtained by multiplying the autocorrelation time with the length of the worm, $\langle l \rangle$, divided by two. The scaled autocorrelation time indicates how many operators have to be changed in order to substantially change the operator string.




FIG. 4. (Color online) The scaled autocorrelation time of the energy as a function of system size N . The data from Fig. 3 were multiplied with $\langle l \rangle / 2$ which was estimated in independent calculations. Operator strings of 20 000 operators were used.

Interestingly, for large N , the amount of operators that have to be changed because of the meandering worm for large b_u and deep penetration compensates for the shorter autocorrelation time measured in terms of updates of the operator string, as is shown in Fig. 4.

Even though the mean penetration depth is small, one can still obtain high quality results. In particular, it is *not* necessary for the mean penetration depth to reach a value close to the length of the operator string (the projection power) in order to get reliable results. Since the maximal penetration depth is substantially larger than the mean, the operator string is often updated deeper than the mean penetration depth. Hence, the mean penetration depth can be much smaller than the length of the operator string has to be for otherwise equivalent calculations with conventional VBQMC. We discuss this effect in more detail in Sec. IV.

As one increases the probability to bounce up the tree, a longer part of the string actually partakes in the projection. Thus, also the quality of the projection is better (see Fig. 5). If operators are never updated, they do not contribute to the projection; they do, however, modify the trial state used in the projection. This leads to the irregularly scattered pattern the data show for small b_u . In this sense, one can think of rare updates that go high up the tree as effectively changing the trial state and the whole calculation as an averaging over these trial states. Equivalently, one can also view this irregular pattern as the result of nonergodic sampling of all operator strings because on average only a small part of the string is changed during one update.

As is evident from Fig. 5, the bouncing worm algorithm yields good results for $b_u > 0.25$. It is an attractively simple algorithm with zero bounce probability and a probability of 1 for accepting a new string. The autocorrelation time can be reduced by increasing b_u , whereas the overall power of the growth at large system sizes appears independent of b_u . However, the increased computational cost associated with increasing b_u is considerable and we have therefore investigated another parameter choice leading to a different algorithm, the driven worm algorithm. We now turn to a discussion of this algorithm.




FIG. 5. (Color online) With a bigger probability to bounce b_u , the bouncing worm algorithm yields a better approximation of the ground-state energy. The ground-state energy was calculated using the Bethe ansatz and is indicated by a dotted line. Since some operators in the string are never updated, calculations with different b_u were effectively done with different trial states. The calculation was done for a chain with 50 sites. Operator strings of 20 000 operators were used and for each point 10^7 measurements taken.

B. Driven worm algorithm

Clearly, it is desirable to have all updates result in a substantial change of the operator string. Then, fewer updates have to be performed. For the problem at hand, this means that we need the worm to go far up the tree as often as possible without increasing the computational cost too much. Direct control over the associated probability would be very convenient. We achieve this by setting the probabilities to go up the tree to be

$$2cu = cw = \alpha. \quad (26)$$

The value of α is the probability to, at each node, decide to go up the tree. Since u and w depend on N_1 and N_2 , this is only possible by allowing c to vary with the node. As explained at the end of Sec. II, the acceptance step of Eq. (16) thus has to be introduced. Updated strings may be rejected.

We set all bounce probabilities to be zero, $b_1 = b_{1/2} = b_u = 0$. Hence, the worm will move up the tree and then turn around once. To get a working algorithm, we have to find solutions to the equations (14) which will determine the transition probabilities [see Eq. (13)]. If $N_{1/2}, N_1 > 1$ the solutions to equations (14) are given by

$$\begin{aligned} x &= \frac{1}{N_{1/2}(N_{1/2} - 1)} [2N_1(N_1 - 1)z + (1 - \alpha)(N_{1/2} - 2N_1)], \\ y &= \frac{1}{N_{1/2}} [(1 - N_1)z + 1 - \alpha], \\ u &= 1/(N_{1/2} + 2N_1), \\ w &= 2u, \end{aligned} \quad (27)$$

where

$$\frac{1 - \alpha}{N_1 - 1} \left[1 - \frac{N_{1/2}}{2N_1} \right] \leq z \leq \frac{1 - \alpha}{N_1 - 1}. \quad (28)$$

If $N_1 \neq 1$, we set

$$z = \frac{1 - \alpha}{N_1 - 1} \left[1 - \frac{1}{2} \frac{N_{1/2}}{2N_1} \right] \quad (29)$$

if it results in $z > 0$ or

$$z = \frac{1}{2} \frac{1 - \alpha}{N_1 - 1} \quad (30)$$

otherwise. In this way, Eq. (28) is always satisfied and $z \geq 0$. If $N_1 = 1$ we set $z = 0$. Finally, we note that the worm update in this case has to be accepted or rejected according to the probability (16). Specific details of an implementation of this driven worm algorithm can be found in the Appendix, Sec. A 2.

How far up the tree updates are attempted can in this case easily be calculated. The probability for the worm to have length l and turn around after going up $r = l/2$ nodes is given by $P(r) = \alpha^r (1 - \alpha)$. The expectation value of r is given by

$$\langle r \rangle = \frac{1}{1 - \alpha}. \quad (31)$$

The probability distribution for the worm to penetrate the tree r nodes deep during a computation of m updates is given by

$$\begin{aligned} P_{m,\alpha}(r_{\max}) &= \underbrace{\left(1 - (1 - \alpha) \sum_{q=r}^{\infty} \alpha^q \right)^m}_{\text{probability that in } m \text{ attempts no worm turns at a node with } r > r_{\max}} \\ &\quad - \underbrace{\left((1 - \alpha) \sum_{q=1}^{r-1} \alpha^q \right)^m}_{\text{probability that in } m \text{ attempts all worms turn at a node with } r < r_{\max}} \\ &= (1 - \alpha^r)^m - (1 - \alpha^{r-1})^m. \end{aligned} \quad (32)$$

How far up the tree is updated is *not* given by how far the worm goes up the tree since the update might be rejected. The mean penetration depth is therefore not equal to $\langle r \rangle$. In Fig. 6, we show results for the mean and maximal penetration depths for two different system sizes ($N = 50, 1000$) as a function of $1/(1 - \alpha)$. As expected, both the mean and maximal penetration depths increase monotonically with $1/(1 - \alpha)$.

Another measure of the performance of the algorithm can be established by simply looking at the calculated ground-state energy and its error. This is done in Fig. 7 where the ground-state energy is shown as a function of $1/(1 - \alpha)$. Operators that are never updated only change the effective trial state the ground state is projected out of. The sampling in turn is nonergodic. By forcing the worm further up the tree, one can have a bigger part of the operator string partake in the projection (see Fig. 6). This leads to a better approximation of the ground-state energy as can be seen in Fig. 7.

For the driven worm algorithm, we have also studied the behavior of the scaled autocorrelation time of the energy. Here, the scaled autocorrelation time can very easily be calculated from the autocorrelation time because the number of attempted operator updates is given by $1/(1 - \alpha)$. The scaled autocorrelation time is therefore calculated by multiplying the




FIG. 6. (Color online) The mean and the maximal penetration depths grow as α approaches one. The mean penetration depth is always smaller than $\langle r \rangle = 1/(1 - \alpha)$. This behavior is independent of the system size. For the chain with 50 sites 3×10^8 and for the chain with 1000 sites 10^7 updates were performed. Operator strings of 20 000 operators were used.

autocorrelation time by $1/(1 - \alpha)$. Our results are shown in Fig. 8 as a function of $1/(1 - \alpha)$. The behavior is in this case not monotonic. At first, the scaled autocorrelation time decreases, but then it starts to grow at larger $1/(1 - \alpha)$.

This can be understood in the following way: As long as $\langle r \rangle$ is much smaller than the size of the system, the scaled autocorrelation time *decreases* with increasing α . This follows naturally from the fact that increasing $1/(1 - \alpha)$ will increase $\langle r \rangle$ and therefore lead to larger and more effective updates. This decreases the correlations between operator strings. The farther the worm travels up the string, the smaller is the probability that an update is accepted (see Table II). For bigger α , and thus also $\langle r \rangle$, this effect dominates and the scaled autocorrelation time grows. A characteristic minimum can be




FIG. 7. (Color online) The bigger the penetration probability α , the better the approximation of the ground-state energy given by the driven worm algorithm. The ground-state energy was calculated using the Bethe ansatz. It is indicated by a dotted line. Since some operators in the string are never updated, calculations with different α were effectively done with different trial states. The calculation was done for a chain with 50 sites. Operator strings of 20 000 operators were used and for each point 10^7 measurements taken.




FIG. 8. (Color online) The scaled autocorrelation time decreases with $\langle r \rangle$, if typical updates are smaller than the system size. Since bigger $\langle r \rangle$ means better projection, this implies that the autocorrelation time *decreases* as the quality of the projection is improved. Operator strings of 20 000 operators were used.

identified as is clearly evident in Fig. 8. The acceptance rate grows with the size of the system (see Table III)

IV. COMPARISON OF ALGORITHMS

In the following, we compare worm updates to simple conventional VBQMC updates as described, for example, in Refs. [12–14]. This means that for VBQMC we attempt to change four randomly selected operators during one update. We do not compare to loop updates as introduced in Ref. [15] since we anticipate the worm algorithms to be of most use with problems for which loop updates are not known although our current implementations of them are similar to conventional VBQMC.

We first consider the convergence of the energy with the projection power (the length of the operator strings). Our results for a 50 site Heisenberg chain are shown in Fig. 9. It turns out that if the worm algorithms penetrate the tree sufficiently deeply, the results do not depend on the type of algorithm in use. In particular, the dependence of the results on the length of the string is the same for *all three* algorithms (see Fig. 9), just as one might have expected since the power method underlies all three algorithms.

When using the worm algorithms, the operator string is usually chosen so long that the worm never or very rarely reaches the root of the tree. This means that there are almost always nodes close to the root with operators that are never updated and thus act on the trial state after every update. In this way, we are effectively using an optimized trial state. The effect

TABLE II. The acceptance rate drops when the string is updated more substantially. The calculations were done for a chain with 50 sites.

$1/(1 - \alpha)$	Acceptance rate
200	0.41
1000	0.13
5000	0.03

TABLE III. The acceptance rate is higher for bigger systems.

$1/(1 - \alpha)$	N	Acceptance rate
1000	100	0.15
1000	1000	0.31
100	100	0.63
100	1000	0.88

is similar to generating the trial state by performing several updates on a randomly chosen trial state and taking the resulting state for the actual calculation. We used such a trial state for the conventional VBQMC calculations shown in this section.

A useful measure of the effectiveness of an algorithm can be obtained from the autocorrelation function. If simply measured as a function of the number of updates, it decreases dramatically faster for the worm algorithms when compared to conventional VBQMC. However, just using one update as the temporal unit puts conventional VBQMC at an unfair disadvantage. The reason is that in calculations with conventional VBQMC, one attempts to change four operators per update while for the worm algorithms it could be many more. We take this into account by, just as above, scaling the temporal axis of the autocorrelation function. Then, one can easily see which algorithm has the smallest correlations when on average the same number of changes has been attempted.

In Fig. 10, we therefore show results for the energy autocorrelation function for the two worm algorithms as well as for conventional VBQMC with the temporal axis rescaled by the number of operators one attempts to change in a single update. During one update with worm algorithms, one tries to update $l/2$ operators. The scaled number of updates is simply $\#\text{updates} \times \langle l/2 \rangle$ with $\langle l/2 \rangle = 4$ for conventional VBQMC




FIG. 9. (Color online) Upon increasing the quality of the projection by using longer operator strings, the estimate of the ground-state energy converges to the correct value in the same way for worm and conventional VBQMC algorithms as long as the string is penetrated sufficiently deeply. The driven worm algorithm was run with the probability to go up the tree $\alpha = 0.995$, which corresponds to a mean penetration depth of about 90 and full penetration of the string. The bouncing worm algorithm was run with a bounce probability $b_u = 0.2675$, which corresponds to a mean penetration depth of roughly 8 and full penetration of the string. The calculation was done for a chain with 50 sites. For the worm algorithms 10^7 and for the conventional VBQMC 10^6 measurements were taken at each point.




FIG. 10. (Color online) The autocorrelation function versus scaled number of updates for the two worm algorithms and conventional VBQMC. The number of updates was scaled by the number of operators attempted to be changed in an update $\langle l/2 \rangle$. Hence, data for conventional VBQMC updates, the driven worm algorithm, and the bouncing worm algorithm were multiplied by 4, 200, and 63.758, respectively. For the worm algorithms, the same parameters as in Fig. 9 were used. This means that $\alpha = 0.995$ and $b_u = 0.2675$. An operator string of length 1000 was used for all three algorithms. The horizontal line at 0.1 was added to allow for easy visual estimation of the scaled autocorrelation time.

and $\langle l/2 \rangle = \langle r \rangle = 1/(1 - \alpha)$ for the driven worm algorithm. For the bouncing worm algorithm, $\langle l \rangle$ has to be measured during the simulation since the bouncing worm can go up and down the tree many times. Thus, $\langle l/2 \rangle$ can be orders of magnitudes bigger than the mean penetration depth. For instance, for the calculations shown in Fig. 10, the mean penetration depth was approximately 7.8 whereas $\langle l/2 \rangle = 63.758$. Even including such a rescaling of the temporal axis, it is clear that the autocorrelation times are much shorter for the worm algorithms, as shown in Fig. 10.

The two worm algorithms change operators of the string starting from one end while the conventional VBQMC selects four operators at random to be changed. As mentioned in Sec. III A, the mean and the maximum penetration depths are usually much smaller than the length of the operator string (the projection power). It is therefore natural to ask if one can reach a similar quality of results using worm algorithms and conventional VBQMC.

That this is so can be seen by plotting the absolute deviation from the ground-state energy $|\Delta E|$ versus the mean penetration depth. As shown in Fig. 11, the mean penetration depth can, in fact, be much smaller than the projection power of a conventional VBQMC calculation and still yield results of the same accuracy.

Finally, we look at how the scaled autocorrelation times of the worm algorithms depend on the size of the system studied. Our results are shown in Fig. 12 for a fixed length operator string of 20 000. For the simulations shown in Fig. 12, the mean penetration depths for the driven worm algorithm were about 100. The scaled autocorrelation time for the driven worm algorithm starts to increase appreciably at this system size, while it is initially almost flat. We conclude that a significant increase in the autocorrelation time appears once




FIG. 11. (Color online) The deviation $|\Delta E|$ from the exact Bethe-ansatz results for a chain with 100 sites. The results are shown for the driven and bouncing worm algorithms versus the mean penetration depth and for conventional VBQMC updates versus the projective power (length of operator string). The worm algorithms reach the same small value of $|\Delta E|$ with a mean penetration depth an order of magnitude smaller than the projective power used for the calculation with conventional VBQMC updates. The bars on the markers indicated the statistical uncertainty. The colored (dark) surfaces are due to overlapping error bars. Operator strings of 20 000 operators were used for the calculations with the worm algorithms. For each point, 10^5 measurements were taken.

the system size significantly exceeds the mean penetration depth. A similar effect can be observed for the bouncing worm algorithm. The mean penetration depths for the bouncing worm algorithm are, however, much smaller (around 9; see Fig. 11). Results for N smaller than the mean penetration depth are therefore not shown in Fig. 12. The autocorrelation times remain manageable for the system sizes studied, even though they are consistently increasing.




FIG. 12. (Color online) The scaled autocorrelation time for the driven and bouncing worm algorithms as a function of system size. A fixed operator string of length 20 000 was used in the calculations. In the calculations shown, the bouncing worm algorithm was run with $b_u = 0.275$ and the driven worm algorithm was run with $\alpha = 0.995$. For the scaling we use $\langle l/2 \rangle = 200$ for the driven worm algorithm and an $\langle l/2 \rangle$ between 60 for $N = 10$ and 221 for $N = 1000$ for the bouncing worm algorithm.

Compared to simple implementations of VBQMC, the worm algorithms have significant overhead. This is largely compensated by the large number of operators that can be changed in an update and resulting shorter autocorrelation times, as we found in all computations. Given the somewhat different properties of the two worm algorithms, a realistic implementation could combine the two by performing updates with the driven worm algorithm mixed with updates using the bouncing worm algorithm (and perhaps conventional VBQMC updates).

V. CONCLUSION

We have shown that valence-bond quantum Monte Carlo can be implemented with an update build around the notion of a worm propagating through a tree. Many different such algorithms are possible. We studied the validity and efficiency of two of them: one for which no update is rejected (the bouncing worm algorithm) and one for which big parts of the operator string are updated (the driven worm algorithm). Both algorithms are attractively simple and straightforward to implement and produce high quality results.

While they may not be computationally competitive with state of the art loop update algorithms [15] for VBQMC, the algorithms presented here are intrinsically interesting since they represent a class of algorithms that should be generally applicable to projective methods. These algorithms are not restricted to the valence-bond basis and preliminary results show that they can be quite efficient in the S^z basis [17] method and might spark further development of it. We also note that many other algorithms can easily be found with the results contained in this paper and that it is possible that the parameter space allows for much more efficient algorithms than the two we have studied here.

In terms of further optimizing the algorithms, several directions may be interesting to pursue. Not updating some of the operators the worm visits might boost the acceptance ratio of the driven worm algorithms and thereby reduce the autocorrelation times. This could be combined with attempting to reduce the overhead of the driven worm calculations by always forcing the worm all the way down to the root. This would eliminate the need to keep track of the state at each node. With the current practice of updating all operators after turning around, going all the way to the node during every update leads to very small acceptance ratios.

ACKNOWLEDGMENT

We acknowledge computing time at the Shared Hierarchical Academic Research Computing Network (SHARCNET:www.sharcnet.ca) and research support from NSERC.

APPENDIX: PSEUDOCODE FOR IMPLEMENTATIONS OF THE TREE-WORM ALGORITHM

This appendix contains pseudocode that shall serve to clarify the algorithms proposed in this paper. To simplify notation, we refer to diagonal operator as NDOP and nondiagonal operators as NDOP.

1. Bouncing worm algorithm

In this section, we give detailed information on a straightforward (albeit not optimized) implementation of the bounce algorithm (see Sec. III A). Shown is an outline of the central part of the algorithm: the update of the operator string and the state.

The algorithm works its way up the tree. It starts at the last branch which is assigned the n th position. At each position it is decided if the worm goes up the tree or down, in which case a new operator is chosen for the branch at this position. The necessary probabilities are calculated according to the expressions given in Eqs. (21) and (23). If an operator is chosen for the n th branch, the update is complete.

It is assumed that the tree is so high (the operator string so long) that the root is never reached. If the root is reached, one has to choose an operator for the first branch according to the probabilities outlined in Sec. II after Eq. (14). The first string is initialized by filling it as if the worm had just reached the root.

Schematically, using pseudocode, a bouncing worm update of a tree with n nodes can be outlined as follows:

```

pos = n                                     ▷ start at last branch
going_up = TRUE
while pos != n + 1 do
  ran = uniform(0,1)
  if going_up then
    if operator at pos is NDOP then
      pos = pos - 1
    else if ran < cw then
      pos = pos - 1
    else
      going_up = FALSE
      choose new DOP at pos
      update state,  $w$  and  $N_1$  at pos
      pos = pos + 1
    end if
  else
    if ran <  $b_u$  then
      going_up = TRUE
    else
      if ran -  $b_u$  <  $wN_1$  then
        choose DOP at pos
      else
        choose NDOP at pos
      end if
    end if
    update state,  $w$  and  $N_1$  at pos
    pos = pos + 1
  end if
end if
end while

```

The weights w, u and b_u, N_1 as well as the state are stored at each node.

2. Driven worm algorithm

We now turn to a description of a (not optimized) implementation of the driven worm algorithm (see Sec. III B). As above, we show an outline of the central part of the algorithm: the update of the operator string and the state.

The worm works its way up the tree. It starts at the last branch which is assigned the position n . While going up the tree, the worm, at each node, goes further up the tree with probability α or turns around with probability $1-\alpha$. After turning around, the worm keeps going down until it reaches the end. At the nodes the worm visits, new operators are chosen. When the worm reaches the end, it has to be decided whether or not the update should be accepted. The associated probabilities are calculated according to the expressions given in the main text [see Eqs. (26), (27) and (16)].

As above, we assume that the tree is so high (the operator strings so long) that the root is never reached. If the root is reached, one has to choose an operator for the first branch according to the probabilities outlined in Sec. II after Eq. (14). The first string is initialized by filling it as if the worm had just reached the root. This first string is always accepted.

Shown is the driven worm update of a tree with n nodes. Using pseudocode language, a driven worm update then takes the following form for a tree with n nodes:

```

pos = n                                     ▷ start at last branch
going_up = TRUE
while pos != n + 1 do
  ran = uniform(0,1)
  while going_up do
    if ran <  $\alpha$ 
      pos = pos - 1
    else
      going_up = FALSE
      if operator at pos is DOP then
        if ran -  $\alpha$  <  $yN_{1/2}$  then
          choose new NDOP at pos
        else
          choose DOP at pos
        end if
      else
        if ran -  $\alpha$  <  $zN_1$  then
          choose new DOP at pos
        else
          choose NDOP at pos
        end if
      end if
      update state, weights,  $c$ , and  $N_1$  at pos
      pos = pos + 1
    end if
  end if
end while
if ran <  $wN_1$  then
  choose DOP at pos
else
  choose NDOP at pos
end if
update state, weights,  $c$ , and  $N_1$  at pos
pos = pos + 1
end while
Accept or reject using old and new  $c$ 's

```

The weights c, N_1 as well as the state are stored at each node. A new string is not always accepted.

- [1] D. Ceperley and M. Kalos, in *Monte Carlo Methods in Statistical Physics*, Topics in Current Physics, Vol. 7, edited by K. Binder (Springer, Berlin, 1986), pp. 145–194.
- [2] R. Blankenbecler and R. L. Sugar, *Phys. Rev. D* **27**, 1304 (1983).
- [3] S. Liang, B. Douçot, and P. W. Anderson, *Phys. Rev. Lett.* **61**, 365 (1988); S. Liang, *Phys. Rev. B* **42**, 6555 (1990).
- [4] S. Liang, *Phys. Rev. Lett.* **64**, 1597 (1990).
- [5] N. Trivedi and D. M. Ceperley, *Phys. Rev. B* **41**, 4552 (1990).
- [6] K. J. Runge, *Phys. Rev. B* **45**, 7229 (1992).
- [7] M. Calandra, F. Becca, and S. Sorella, *Phys. Rev. Lett.* **81**, 5185 (1998).
- [8] L. Capriotti, A. E. Trumper, and S. Sorella, *Phys. Rev. Lett.* **82**, 3899 (1999).
- [9] W. von der Linden, *Phys. Rep.* **220**, 53 (1992).
- [10] N. J. Cerf and O. C. Martin, *Int. J. Mod. Phys. C* **06**, 693 (1995).
- [11] J. H. Hetherington, *Phys. Rev. A* **30**, 2713 (1984).
- [12] A. W. Sandvik, *Phys. Rev. Lett.* **95**, 207203 (2005).
- [13] K. Beach and A. W. Sandvik, *Nucl. Phys. B* **750**, 142 (2006).
- [14] A. W. Sandvik and K. S. D. Beach, [arXiv:0704.1469](https://arxiv.org/abs/0704.1469).
- [15] A. W. Sandvik and H. G. Evertz, *Phys. Rev. B* **82**, 024407 (2010).
- [16] A. Banerjee and K. Damle, *J. Stat. Mech.: Theor. Exper.* (2010) P08017.
- [17] A. Deschner and E. S. Sørensen (unpublished).
- [18] G. Santoro, S. Sorella, L. Guidoni, A. Parola, and E. Tosatti, *Phys. Rev. Lett.* **83**, 3065 (1999).
- [19] G. Carleo, F. Becca, S. Moroni, and S. Baroni, *Phys. Rev. E* **82**, 046710 (2010).
- [20] R. H. Swendsen and J.-S. Wang, *Phys. Rev. Lett.* **58**, 86 (1987).
- [21] H. G. Evertz, G. Lana, and M. Marcu, *Phys. Rev. Lett.* **70**, 875 (1993).
- [22] N. Prokof'ev and B. Svistunov, *Phys. Rev. Lett.* **87**, 160601 (2001).
- [23] N. Prokof'ev and B. Svistunov, [arXiv:0910.1393](https://arxiv.org/abs/0910.1393).
- [24] F. Alet and E. S. Sørensen, *Phys. Rev. E* **67**, 015701 (2003).
- [25] F. Alet and E. S. Sørensen, *Phys. Rev. E* **68**, 026702 (2003); P. Hitchcock, E. S. Sørensen, and F. Alet, *ibid.* **70**, 016702 (2004).
- [26] J.-S. Wang, *Phys. Rev. E* **72**, 036706 (2005).