Stochastic series expansion method for quantum Ising models with arbitrary interactions

Anders W. Sandvik
Department of Physics, Åbo Akademi University, Porthansgatan 3, FIN-20500 Turku, Finland
(Received 27 March 2003; published 11 November 2003)

A quantum Monte Carlo algorithm for the transverse Ising model with arbitrary short- or long-range interactions is presented. The algorithm is based on sampling the diagonal matrix elements of the power-series expansion of the density matrix (stochastic series expansion), and avoids the interaction summations necessary in conventional methods. In the case of long-range interactions, the scaling of the computation time with the system size N is therefore reduced from N^2 to $N\ln(N)$. The method is tested on a one-dimensional ferromagnet in a transverse field, with interactions decaying as $1/r^2$.

DOI: 10.1103/PhysRevE.68.056701 PACS number(s): 02.70.Ss, 05.30.−d, 75.40.Mg, 75.10.Jm

I. INTRODUCTION

Monte Carlo studies of classical and quantum many-body systems with long-range interactions are limited by time-consuming summations over the interacting particle pairs, the number of which grows quadratically with the system size. Many important problems in both basic and applied sciences can be mapped onto long-range interacting spin models, and hence it would be desirable to develop more efficient numerical techniques for tackling them. For classical Ising models, considerable progress has indeed been made on algorithms scaling almost linearly with the system size [1].

In the context of simulated annealing [2], where the ground state of a classical system (typically with complicated interactions) is obtained through a simulation where the temperature is slowly lowered to zero, it has been suggested [3] that a more rapid convergence could be achieved by using a quantum model, e.g., the Ising model in a transverse (spin-flipping) field. Even in an imaginary-time path-integral formulation, the quantum fluctuations can, at least in some cases [4], relax the system towards its classical ground state more rapidly than thermal fluctuations. This is a strong motivation for developing more efficient simulation methods for quantum Ising models. Another important reason is the continued prominence of the transverse Ising model in the theory of magnetism, particularly in the context of quantum phase transitions [5–7]. Whereas transverse Ising models with short-range interactions have recently been actively studied using quantum Monte Carlo methods [6,7], numerical work on long-range models has so far been limited to special cases [8]. In some of the best experimental realizations of the transverse Ising model the interactions are in fact long ranged [9].

Here a stochastic series expansion (SSE) [10] algorithm for transverse Ising models with long-range interactions is introduced in which the direct summation over the interacting spins is avoided. The computation time scales with the system size N as $N\ln(N)$ times the spatial integral of the absolute value of the interaction [which normally converges as $N\to\infty$, or diverges only as $\ln(N)$]. Both local and cluster-type updates are developed for the transverse Ising model with arbitrary interactions. The cluster update is a generalization of the classical Swendsen-Wang cluster method [11] to the transverse Ising model, and shares some features with a scheme previously used within the continuous-time world-line algorithm [7]. The way to treat the long-range interactions generalizes the scheme developed for the classical Ising model by Luijten and Blöte [1,12]. The integration of these features in the SSE formalism should open new opportunities for detailed numerical studies of a wide range of important models. The algorithm is here tested on a ferromagnetic chain with interactions decaying as $1/r^2$, for which many results in the classical limit [13–17] are available for comparison. Some analytical studies have also recently been carried out on the quantum model with the transverse field [18].

In Sec. II the application of the SSE method to the transverse Ising model is described in detail. Local updates as well as classical and quantum-cluster updates are discussed. Results for the model with $1/r^2$ interactions are presented in Sec. III. Section IV concludes with a brief discussion.

II. STOCHASTIC SERIES EXPANSION

The SSE method [10] is an efficient alternative to world-line quantum Monte Carlo [19]. It is based on a generalization of the power-series scheme for the Heisenberg ferromagnet that was developed by Handscomb in the early 1960s [20]. Handscomb's method was later extended to some other models [21], but the requirement of analytically calculable traces of the terms of the expansion inhibited further progress. In the SSE method a basis is instead chosen, and the traces are also evaluated stochastically, in combination with the sampling of the operator products in the series expansion of $\exp(-\beta H)$. This starting point for quantum Monte Carlo is as generally applicable as the world-line (imaginary-time path-integral) approach. Recently, loop-type cluster updates [22] have been developed and generalized for efficient SSE simulations of a wide range of models [23,24]. However, since the loop updates rely heavily on the presence of off-diagonal pair (or multiparticle) interactions, they cannot be directly adapted to the transverse Ising model in the standard basis where the Ising term is diagonal. In the basis where the field is diagonal, loop updates can be easily implemented [23,24] but then sign problems [25] appear when the interaction is frustrated. Here the SSE method is applied to an arbitrary transverse Ising model, i.e., with no limitations on the sign and range of the spin-spin interaction. Several types of local and cluster-type updates will be described.

A. Configuration space

Consider the general Hamiltonian for the Ising model in a transverse field of strength h,

$$H=\sum_{i,j} J_{ij}\sigma_i^z\sigma_j^z - h\sum_i \sigma_i^x, \tag{1}$$

where $\boldsymbol{\sigma}_i$ is a Pauli spin operator ($\sigma_i^z=\pm 1$) and J_{ij} is the strength of the interaction between spins i and j, which can be random or uniform and of any sign. The dimensionality is arbitrary. Define the operators

$$H_{0,0}=1, \tag{2a}$$

$$H_{i,0}=h(\sigma_i^+ + \sigma_i^-), \qquad i>0, \tag{2b}$$

$$H_{i,i}=h, \qquad i>0, \tag{2c}$$

$$H_{i,j}=|J_{ij}|-J_{ij}\sigma_i^z\sigma_j^z, \qquad i,j>0,\ i\neq j. \tag{2d}$$

Up to a constant, the Hamiltonian can be written as

$$H=-\sum_{i=1}^{N}\sum_{j=0}^{N} H_{i,j}. \tag{3}$$

The constants $H_{i,i}$ are introduced for purposes that will become clear below. Note that $H_{0,0}$ is not included as a term in the Hamiltonian (3) but will be important in the simulation scheme.

In the SSE approach [10] to finite-temperature quantum Monte Carlo, the partition function $Z=\mathrm{Tr}\{\exp(-\beta H)\}$ is written as a power-series expansion, with the trace expressed as a sum over diagonal matrix elements in a suitably chosen basis. Using Eq. (3) then gives

$$Z=\sum_\alpha \sum_{n=0}^{\infty}\sum_{S_n} \frac{\beta^n}{n!}\langle\alpha|\prod_{l=1}^{n} H_{i(l),j(l)}|\alpha\rangle, \tag{4}$$

where S_n denotes a sequence of n operator-index pairs (hereafter referred to as operators):

$$S_n=[i(1),j(1)],\ \ldots,[i(n),j(n)], \tag{5}$$

with $i(l)\in\{1,\ldots,N\}$ and $j(l)\in\{0,\ldots,N\}$. The standard basis $\{|\alpha\rangle\}=\{|\sigma_1^z,\ldots,\sigma_N^z\rangle\}$ is used.

Because of the constants added to $H_{i,j}$ in Eq. (2d), the eigenvalues of these operators are $2|J_{ij}|$ and 0. All nonzero terms in Eq. (4) are therefore positive and can be used as relative probabilities in an importance sampling scheme. A term is specified by a state $|\alpha\rangle$ and an operator sequence S_n. One can show that the total internal energy (including the constants added to H) is given by [10,20] $E=-\langle n\rangle/\beta$. Hence, the size of the operator sequence to be stored in computer memory scales as $\beta N I_N(J)$, where

$$I_N(J)=\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{N}|J_{ij}|, \tag{6}$$

which converges or grows much slower than N for most cases of interest.

In order to construct an efficient sampling scheme, it is useful to cut expansion (4) at some power $n=L$, sufficiently high for the remaining truncation error to be exponentially small and completely negligible [L clearly has to be $\sim\beta N I_N(J)$]. One can then obtain an expansion for which the length of the operator sequence is constant, by considering random insertions of $L-n$ unit operators $H_{0,0}$ in the product in Eq. (4). Adjusting for the $\binom{L}{n}$ possible insertions gives

$$Z=\frac{1}{L!}\sum_\alpha\sum_{S_L}\beta^n(L-n)!\langle\alpha|\prod_{l=1}^{L} H_{i(l),j(l)}|\alpha\rangle, \tag{7}$$

where $[i(l),j(l)]=[0,0]$ is now also an allowed operator in the sequence S_L, and n denotes the number of non-$[0,0]$ operators. Note again that $H_{0,0}$ is not part of the Hamiltonian, but is introduced only for the purpose of constructing a computationally simpler updating scheme where the operator list has a fixed length.

It is useful to define states $|\alpha(p)\rangle=|\sigma_1^z(p),\ldots,\sigma_N^z(p)\rangle$ obtained by propagating $|\alpha\rangle=|\alpha(0)\rangle$ by the first p operators in S_L:

$$|\alpha(p)\rangle=r\prod_{l=1}^{p} H_{i(l),j(l)}|\alpha\rangle, \tag{8}$$

where r is a normalization factor. A nonvanishing matrix element in Eq. (7) then corresponds to the periodicity condition $|\alpha(L)\rangle=|\alpha(0)\rangle$, which requires that for each site i there is an even number (or zero) of spin-flipping operators $[i,0]$ in S_L. Definition (2d) implies that the Ising operators $[i,j]$ may act only on states with $\sigma_i^z=\sigma_j^z$ if $J_{ij}<0$ (ferromagnetic), or $\sigma_i^z=-\sigma_j^z$ if $J_{ij}>0$ (antiferromagnetic). There are no other constraints.

An SSE configuration is illustrated in Fig. 1. The vertical direction in this representation will be referred to as the SSE *propagation direction*. It can be related to the imaginary-time direction in standard path-integral representations [26]. Note that this full configuration, including all the states $|\alpha(p)\rangle$ explicitly, does not have to be stored in the simulation. A single state and the operator sequence suffice for reproducing all the states, and such a representation is used in some stages of the simulation. For some updates it is convenient to generate other representations, as will be discussed below.

B. Local updates

The sampling of Eq. (7) can be carried out using simple operator substitutions of the types

$$[0,0]_p \leftrightarrow [i,j]_p, \qquad i,j\neq 0, \tag{9a}$$

$$[i,i]_{p_1}[i,i]_{p_2} \leftrightarrow [i,0]_{p_1}[i,0]_{p_2}, \qquad i\neq 0, \tag{9b}$$

where the subscript p indicates the position ($p=1,\ldots,L$) of the operator in the sequence S_L. The power n is changed by ± 1 in the *diagonal update* (9a) and is unchanged in the

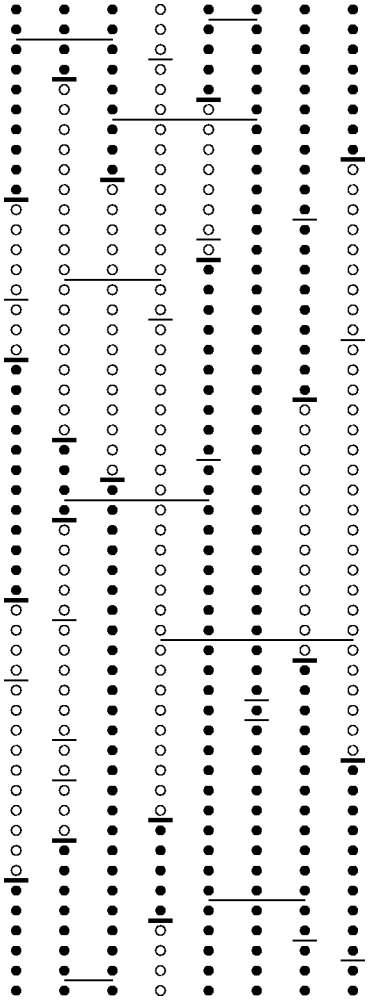

FIG. 1. An SSE configuration for an eight-site one-dimensional system. Here the truncation $L=49$, and the expansion order of the term (i.e., the number of Hamiltonian operators present) $n=40$. The solid and open circles represent the spins $\sigma_i^z(p)=\pm 1$, with the propagation index $p=0,\ldots,L$ corresponding to the different eight-spin rows. The thick and thin short horizontal bars represent spin-flip operators $H_{i,0}$ and constants $H_{i,i}$, respectively. The longer lines represent Ising operators $H_{i,j}$ $(i\neq j)$ acting on the spins at the line ends.

off-diagonal update (9b). In the diagonal update the Ising terms $[i,j]$ and the constants $[i,i]$ are sampled. The constants are used in the off-diagonal update as a means of achieving easy insertions and removals of two spin-flipping operators $[i,0]$. With the value h chosen for the constant in Eq. (2c), the operator replacements do not change the weight of the SSE configuration. However, the off-diagonal update also leads to spin flips in the propagated states between p_1 and p_2; $\sigma_i^z(p_1),\ldots,\sigma_i^z(p_2-1)\rightarrow-\sigma_i^z(p_1),\ldots,-\sigma_i^z(p_2-1)$. $[p_1>p_2$ also has to be considered, leading to flipped $\sigma_i^z(p_1),\ldots,\sigma_i^z(L-1)\sigma_i^z(0),\ldots,\sigma_i^z(p_2-1)]$, which is allowed if (and only if) no Ising operators acting on site i are present in S_L between positions p_1 and p_2. *Note that this constraint is completely local, regardless of the range of the interaction, and that the update requires no knowledge of the spin state*. This is the reason for the advantage of this simulation scheme over world-line methods [19,7], where calculating the acceptance probability for every update requires a summation over all the spins interacting with those flipped. Here an allowed off-diagonal update (9b) leaves the weight unchanged and can be carried out with probability 1.

If $h\neq 0$, the above updates of the operator sequence suffice for achieving ergodicity. If there are no Ising operators acting on a site i, $\sigma_i^z(0),\ldots,\sigma_i^z(L-1)$ can also be flipped without changes in S_L. This update in principle makes simulations using the present scheme possible also for $h=0$, but in practice unconstrained spins occur frequently only at high temperatures, when $\langle n\rangle$ is small. Other types of ''classical'' spin flips—flips of clusters—are also possible, and will be discussed in Sec. II C.

The simulation can be started with a random state $|\alpha(0)\rangle$ and a sequence S_L containing only $[0,0]$ operators. The truncation L can be chosen arbitrarily (small); it is adjusted during the equilibration part of the simulation, e.g., by requiring $L>(4/3)n$ after each update. This ensures than n never reaches L during the remainder of the simulation, and hence that there will be no detectable systematic errors arising from the truncation of the expansion [10]. In the beginning of an updating cycle, the operator sequence S_L and the state $|\alpha(0)\rangle$ are stored.

The diagonal update (9a) is attempted successively for all $p=1,\ldots,L$. In the course of this process, the spin state is propagated by flipping spins σ_i^z as off-diagonal operators $[i,0]$ are encountered in S_L, so that the states $|\alpha(p)\rangle$ are generated successively. For an $[i,j]\rightarrow[0,0]$ update, i.e., removing a Hamiltonian operator, there are no constraints and the update should always be accepted with some nonzero probability. In the case of $[0,0]\rightarrow[i,j]$, i.e., inserting an operator from the Hamiltonian, there are constraints, and the update may not be allowed for all i,j. However, initially the indices i,j are left undetermined and it is assumed that any $[i,j]$ would be allowed. Under this assumption, the acceptance probabilities for the diagonal update are given by

$$P([0,0]\rightarrow[i,j])=\frac{\beta\left(Nh+2\sum_{ij}|J_{ij}|\right)}{L-n+\beta\left(Nh+2\sum_{ij}|J_{ij}|\right)}, \quad (10a)$$

$$P([0,0]\rightarrow[i,j])=\frac{L-n+1}{L-n+1+\beta\left(Nh+2\sum_{ij}|J_{ij}|\right)}, \quad (10b)$$

where Σ_{ij} does not include $i=j$ and $P>1$ should be interpreted as probability 1, as usual. These heat-bath probabilities are simply obtained from the ratio of the new and old prefactors in Eq. (7) when $n\rightarrow n\pm 1$:

$$\beta^{\pm 1}\frac{[L-(n\pm 1)]!}{(L-n)!}, \quad (11)$$

and the ratio between the matrix element 1 of the $[0,0]$ operator and the sum $Nh+2\Sigma_{ij}|J_{ij}|$ of the nonzero matrix elements of all $[i,j]$ operators. Staying with the assumption that any $[i,j]$ is allowed in the update $[0,0]\rightarrow[i,j]$, the rela-

tive probability of an operator with the first index i is $P(i) = \Sigma_j M_{ij}$, where M_{ij} is the nonzero matrix element corresponding to H_{ij} (i.e., h for $i=j$ and and $2|J_{ij}|$ else). The normalized cumulative probabilities $P_c(k=1,\ldots,N)$ are stored in a pregenerated table:

$$P_c(k) = \frac{\sum_{i=1}^{k} P(i)}{\sum_{i=1}^{N} P(i)}. \tag{12}$$

In order to select the first index i of the operator $[i,j]$ to be inserted, a random number $0 \leqslant R < 1$ is generated. The table P_c is searched (using, e.g., a simple binary search) for the smallest k for which $P_c(k) \geqslant R$; the first index of the operator $[i,j]$ is then $i=k$. The second index can be chosen in a completely analogous way, with the relative probability for j, given i, being M_{ij}. For a random system with long-range interactions, a pregenerated table with N^2 elements is hence needed for storing all the cumulative probabilities for the second index. For nonrandom interactions in a translationally invariant system, the first index can be selected at random with equal probabilities without searching a table, and the size of the second table is reduced to N. For a short-range or truncated interaction the table size is smaller, corresponding to the number of spins within the range of the interaction; clearly, the whole selection process should then be reduced to a single step for obtaining both i and j (e.g., selecting one out of a total number $\sim N$ of operators and reading the corresponding i,j from a table). The two-step procedure is advantageous for nonrandom long-range interactions, where it allows for the reduction of the size of the probability table from N^2 to N. For random models, the storage requirement is always N^2, and it may then again be better to combine the first and second index searches, using a single size-N^2 table for all the cumulative probabilities of $[i,j]$. For short-range random interactions the size of the table is N times the number of spins within the interaction range.

The operator $[i,j]$ generated as above may or may not be allowed in the current spin configuration $|\alpha(p)\rangle$. If $\sigma_i^z(p)$ and $\sigma_j^z(p)$ indeed are in an allowed state, $[i,j]$ is inserted at position p. Otherwise, the operator [0,0] is left unchanged. This accept/reject step leads to the correct probabilities for selecting among all the allowed diagonal operators $[i,j]$.

The off-diagonal update (9b) can be efficiently carried out if S_L is first partitioned into separate subsequences for each site i. Subsequence i contains only spin-flipping operators $[i,0]$ and constants $[i,i]$. Their positions in S_L are also stored, to be used for recombining the subsequences after the update. The constraints on modifications at site i imposed by Ising operators $[i,j]$ or $[j,i]$ (for any j) can be stored as flags indicating the presence of one or several of these operators between neighboring subsequence operators. Updating a subsequence amounts to selecting two nonconstrained neighboring operators at random from the subsequence, and carrying out substitution (9b) if the two operators are identical. If they are different, they can be permuted. A number proportional to the subsequence length of such pair updates is carried out for each subsequence, after which they are recombined into a new S_L.

The diagonal update (9a) at all positions in S_L requires $\sim L\ln(N) \sim \beta N \ln(N) I_N(J)$ operations, where the factor $\ln(N)$ is the scaling of the average number of operations needed to search the cumulative probability table(s) in the case of long-range interactions. Partitioning S_L into subsequences and updating all of them according to Eq. (9b) requires on the order of L operations. Hence, the number of operations for a full updating cycle of the degrees of freedom of the system (one Monte Carlo step) scales as $\beta N \ln(N) I_N(J)$. This should be compared to the βN^2 scaling in world-line methods [7,19], where one power of N is due to the summation required to calculate the weight change when flipping a spin interacting with N other spins. Here this summation has been circumvented by writing the interactions in the SSE formalism as fluctuating constraints that are purely local.

C. Classical cluster update

In the Swendsen-Wang cluster algorithm [11] for the classical Ising model, i.e., with $h=0$ and a uniform nearest-neighbor interaction of strength J, auxiliary bond variables b_{ij} are introduced in order to construct clusters of spins that can be flipped independently of each other. Given a spin configuration, and with initially all bond variables $b_{ij}=0$, for every interacting spin pair for which $\sigma_i\sigma_j = -J/|J|$ (i.e., the orientation energetically favored) the bond variable is set, $b_{ij}=1$, with probability $P = 1 - e^{-2|J|\beta}$. When all bonds have been visited, clusters of spins connected by $b_{ij}=1$ bonds are formed, and each of these clusters is flipped with probability 1/2. Single spins not connected to any $b_{ij}=1$ bond are single-spin clusters. After the clusters have been flipped, all the bond variables are again set to zero and the process is repeated. This scheme can in fact be constructed using the SSE formalism, as an alternative to the Fortuin-Kasteleyn mapping [27], on which the Swendsen-Wang algorithm is based.

The relation to the Swendsen-Wang algorithm is shown as follows, by applying the SSE method to the classical Ising model, now again considering a general form of the interaction J_{ij} and with the bond operator $H_{ij} = |J_{ij}| - J_{ij}\sigma_i^z\sigma_j^z$ as in Eq. (2d). Since all operators H_{ij} commute, the operator $e^{-\beta H}$ can be written as a product of operators $e^{\beta H_{ij}} = 1 + \beta H_{ij} + \cdots$. The uniqueness of the power-series expansion then implies that in the SSE, where $e^{-\beta H}$ is expanded directly, the probability of having one or more operators H_{ij} on a bond i,j when $\sigma_i\sigma_j = -J_{ij}/|J_{ij}|$ is $1 - e^{-2|J_{ij}|\beta}$, i.e., exactly the probability of having the bond variable $b_{ij}=1$ in the Swendsen-Wang scheme. In a configuration $\sigma_i\sigma_j = J_{ij}/|J_{ij}|$ there can be no operators on the bond in the SSE, and the Swendsen-Wang $b_{ij}=1$ probability is also zero per construction. One can hence make the connection that one or more operators acting on a spin pair in the SSE scheme correspond to a filled bond ($b_{ij}=1$) in the Swendsen-Wang algorithm. The definition of a cluster is then exactly the same in the two algorithms. Clearly, such a cluster in the SSE can also always be flipped, since the Ising operators only impose constraints

on the relative orientations of connected spins, which is maintained when the cluster is flipped. Since the weight does not change, the flip should be done with probability 1/2. The scheme is hence identical to the Swendsen-Wang algorithm, except that the filled bonds $b_{ij}=1$ in SSE are generated in a different way, using the diagonal update (9a). Note that for a classical model, all the propagated SSE states (8) are identical, i.e., $\sigma_i^z(p)=\sigma_i^z(0)$ for all $p=0,\ldots,L-1$, and hence no state propagations have to be considered as the diagonal update is carried out.

It is interesting to note that the SSE scheme for the classical Ising model should in fact be more efficient than the standard Swendsen-Wang algorithm at high temperatures. This is because the number of operators in the SSE operator list scales as $E(T)/T$, where $E(T)$ is the total internal energy at temperature T ($E\sim N$) and for large T the construction of the clusters based on the operator list should then be faster than visiting all the bonds, as is done in the Swendsen-Wang algorithm. However, in practice the interesting physics occurs when the number of SSE operators per interacting spin pair is of the order of 1 or larger, and then there are no advantages of the SSE classical cluster algorithm relative to Swendsen-Wang.

The classical SSE cluster update can also be used in the presence of the transverse field ($h>0$). The clusters are defined in terms of bonds signifying the presence of one or more Ising operator, as above, without regard for the single-spin-flipping operators $H_{i,0}$ and constants $H_{i,i}$. These operators can be neglected because when a cluster is flipped, all spins σ_i^z belonging to the cluster are implicitly flipped in all propagated states (8), i.e., $\sigma_i^z(p)\rightarrow-\sigma_i^z(p)$ for all $p=0,\ldots,L-1$ (this is the reason for the term "classical cluster" even when $h>0$) and hence all operations with the single-spin operators remain valid and produce the same factors in the weight before and after the cluster flips. Note again that only the first state, i.e., $\sigma_i^z(0)$, $i=1,\ldots,N$, has to be stored when constructing the classical clusters.

In the case of long-range interactions, a cluster can consist of several intertwined pieces on the lattice, as illustrated for a two-dimensional case in Fig. 2. Regardless of the range of the interaction, the construction of the clusters, given a SSE operator list, can be easily carried out using a number of operations scaling as the number of operators in the list.

Since the classical SSE cluster update is equivalent to the Swendsen-Wang algorithm in the classical limit and only takes the Ising terms into account also in the quantum case, it cannot be expected to be efficient much beyond the classical limit $h=0$. For a nonrandom system that undergoes a phase transition at $T_c(0)$ when $h=0$, the critical temperature is reduced by the transverse field; $T_c(h)<T_c(0)$. Hence, the classical clusters will percolate for $T>T_c$ and this update will not be efficient close to T_c. The primary reason to introduce the classical cluster update here was to demonstrate the relationship between SSE and the Swendsen-Wang algorithm. In the case of long-range interactions, the scheme becomes very similar to the Luijten-Blöte algorithm [1], again just differing in the way the bonds are generated.

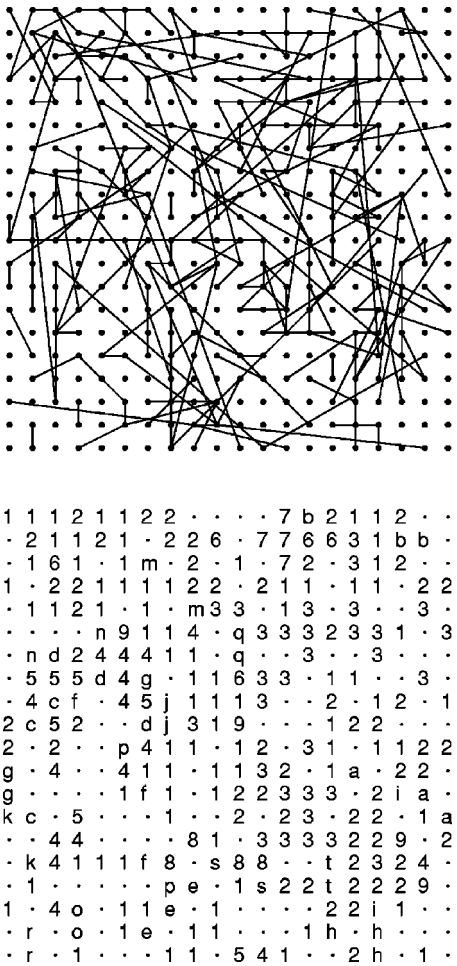

FIG. 2. Upper panel: Interaction bonds in a configuration for a 2D system with long-range interactions. Lower panel: The clusters constructed from the bonds. Sites with equal symbols belong to the same cluster. Dots indicate spins not acted on by any Ising operator and constitute single-spin clusters.

D. Quantum-cluster update

The purpose of the quantum-cluster update is to effect flips of spins $\sigma_i^z(p)$ only in a limited number of propagated states p, in different states for different sites i. In other words, these clusters will be finite and irregularly shaped both in the space and SSE propagation (imaginary-time) direction. In the process, operator substitutions $H_{i,i}\leftrightarrow H_{i,0}$ (constant to spin flip, and vice versa) will also be accomplished. This update hence replaces the local off-diagonal update (9b).

To discuss the quantum-cluster update, it is useful to introduce the notion of *vertices* [23,24]. Looking at the graphical representation of a configuration in Fig. 1, it can be noted that the vertical "lines" of same spins between two operators acting on a given site constitute redundant information. The full configuration can be represented by a list of positions (on the lattice) of the operators, and the spin states (on one or two sites for the model considered here) before and after the operators act. These relevant spins are called *legs* of the two-spin vertices (corresponding to constant and spin-flip operators) or four-spin vertices (corresponding to Ising bond op-

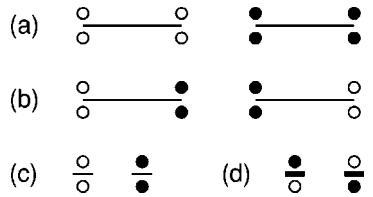

FIG. 3. All the possible four-leg and two-leg vertices. (a) Ferromagnetic Ising vertices, (b) antiferromagnetic Ising vertices, (c) constant vertices, and (d) spin-flip vertices.

erators). All possible vertices for the transverse Ising model are shown in Fig. 3. Note that only those Ising vertices that are compatible with the sign of the interaction between a given pair of spins are allowed for those spins; again, this is due to the choice of constant in the bond operator (2d). In the computer, the vertices are linked to each other by pointers, so that from a given vertex leg one can reach the next or previous vertex that has a leg on the same site (i.e., there are links that replace the segments of vertical lines of same spins in Fig. 1). A detailed discussion of the practical implementation of a linked vertex list has been given in Ref. [24].

To construct and flip a quantum cluster, one of the legs of one of the n vertices is picked at random, and the corresponding spin is flipped. Depending on the type of the vertex, different actions are taken, examples of which are given in Fig. 4. The arrow pointing into the vertex indicates the *entrance leg*. In the case of an Ising vertex, all the four spins are flipped and the cluster building process branches out from all the legs, as indicated by the arrows pointing out from the vertex. Using the pointers of the linked vertex list, the arrows point to legs of other vertices; these become new entrance legs which are put on a stack and subsequently processed one by one. If the entrance leg is on a constant or spin-flip vertex, only the entrance spin is flipped. The vertex type then also changes, in terms of operators from $H_{i,0}$ to $H_{i,i}$, and vice versa. In these cases there is no branching-out and no new legs are put on the stack, i.e., this particular branch of the cluster terminates. If a link points to a spin that has already been flipped (i.e., two arrows point toward each other), that leg should not be used again as an entrance and is hence not put on the stack. Therefore, each vertex leg can be visited at most once (each spin can be flipped at most once) and the cluster is completed when there are no more entrance legs on the stack. The reason that the cluster can always be flipped is again that the SSE weight is not affected; the matrix element of the Ising bond operator is not affected when

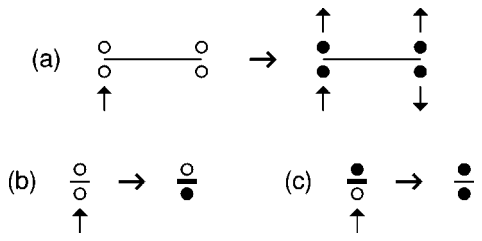

FIG. 4. Examples of vertex processes: (a) reversal of a ferromagnetic Ising vertex, (b) constant to spin flip, and (c) spin flip to constant.

both spins are flipped (in the absence of an external field in the z direction, which would necessitate a modified approach), and the matrix elements for the constant and spin-flip operators are both equal to h.

The construction of a single cluster, which is flipped with probability 1, is a quantum-mechanical analog of the classical Wolff algorithm [28]; in the absence of the transverse field the clusters are identical to those of the Wolff algorithm. Note, however, that there is a difference when constructing more than one cluster: The number of operators in the SSE operator list and their positions on the lattice do not change in the quantum-cluster update. The clusters are therefore completely deterministic once the operator list is given. Hence, when constructing several clusters using the same SSE operator list, it is quite likely that the same cluster is constructed and flipped multiple times. This is clearly not desirable. However, one can also construct all clusters, as in the Swendsen-Wang scheme, and only flip them with probability 1/2. This is done by always starting a new cluster from a vertex leg which has not yet been visited. Every vertex leg belongs uniquely to one cluster, and clearly the number of operations required to complete this update then scales as L, i.e., typically as βN.

A natural definition of a Monte Carlo step including the quantum-cluster update is a full sweep of diagonal updates, followed by the construction of the linked list of vertices, in which all clusters are constructed and flipped with probability 1/2. After that, the updated vertex list is mapped back into a state $|\alpha(0)\rangle$ and an operator sequence S_L. Free spins, i.e., those that are not acted on by any operators, can again be considered as single-spin clusters and should also be flipped with probability 1/2. No local off-diagonal updates (9b) are needed.

Since the quantum-cluster update explicitly includes the quantum-mechanical features of the configurations (i.e., the presence of spin-flip operators), it can be expected to work well also close to a quantum phase transition ($T_c=0$) driven by varying h. There are no problems in principle in taking the $T\to 0$ limit, although, as in all finite-T methods, very large inverse temperatures β have to be used to converge large lattices to the ground state (especially in the case of randomized interactions [29]).

III. ONE-DIMENSIONAL (1D) INVERSE-SQUARE FERROMAGNET

As a nontrivial demonstration of the method, a ferromagnetic chain with interactions decaying as $1/r^2$ is considered next. The interaction is summed over all i,j in Eq. (1), i.e., each pair is counted twice. Periodic boundary conditions are used. J_{ij} includes both distances in the periodic system, i.e.,

$$J_{ij}=J_{ji}=\frac{J}{2}\left(\frac{1}{|i-j|^2}+\frac{1}{(N-|i-j|)^2}\right), \tag{13}$$

where J sets the overall energy scale.

The classical $1/r^2$ Ising chain has been the subject of numerous studies [13–17]. The long-range interaction allows for a finite-T phase transition even in one dimension. The

transition is of an unusual kind, with the correlation length exponent $\nu=\infty$, and a discontinuous jump in the magnetization at T_c. It can be thought of as a one-dimensional analog of the Kosterlitz-Thouless (KT) transition, with the topological excitations being kink solitons [14]. The model is also important because it can be mapped onto the Kondo problem [13].

For small h/J, one can expect a behavior similar to the classical case, i.e., a finite-T phase transition to a ferromagnetic state. For $h\rightarrow\infty$ the system becomes disordered, and there should therefore be a finite h_c for which the system undergoes a quantum phase transition (i.e., $T_c=0$). For $h<h_c$, $T_c>0$ and one would then expect the same universality class as in the classical case, since the quantum fluctuations become irrelevant at T_c [18]. Here only a single field strength $h/J=0.5$ is considered; the simulations show that $T_c>0$ in this case. A more systematic study of the h dependence and the quantum phase transition are left for future studies.

The model is invariant with respect to flipping all spins, which means that for any finite system the average magnetization vanishes. The squared magnetization,

$$M^2=\left\langle\left(\frac{1}{N}\sum_i \sigma_i^z\right)^2\right\rangle, \tag{14}$$

is therefore calculated. Results for M^2 with statistical errors in the fifth decimal place can easily be obtained for systems with several hundred spins (and there are no problems in going to considerably larger systems). For small systems the results are in perfect agreement with exact diagonalization data.

A ''tempering'' scheme, where β is considered as an additional discretized dimension of the configuration space [30], was also implemented in the simulations. Transitions satisfying detailed balance are carried out between neighboring β values. This way, results can be obtained on a dense temperature grid with much less effort than by several fixed-β simulations. A temperature spacing $\Delta T/J=0.01$–0.02 was used.

Figure 5(a) shows results for systems with N up to 512. At high temperatures, M^2 decreases with increasing N, as expected, and there is a slight increase with N at low T. The curves intersect at $T/J\approx 1.4$. A discontinuous magnetization jump at T_c in the thermodynamic limit implies that M^2 should become size independent at T_c for sufficiently large N. A notable difference between the finite-size behavior of $M(T)$ seen in Fig. 5 and the magnetization curves for the classical system is that in the latter case the curves do not intersect, but the infinite-size value $M(T_c)$ is approached with a logarithmic correction [17]. Figure 5(b) shows in greater detail the behavior in the region where the curves intersect. The point of intersection moves slowly towards higher T as N increases, and larger N would be needed to determine whether a fixed crossing point is reached asymptotically. Under the assumption that the behavior for the largest systems reflects the asymptotic behavior, it would be natural to associate the crossing point with the critical temperature, and then based on the data, $T_c/J=1.42\pm 0.01$. This

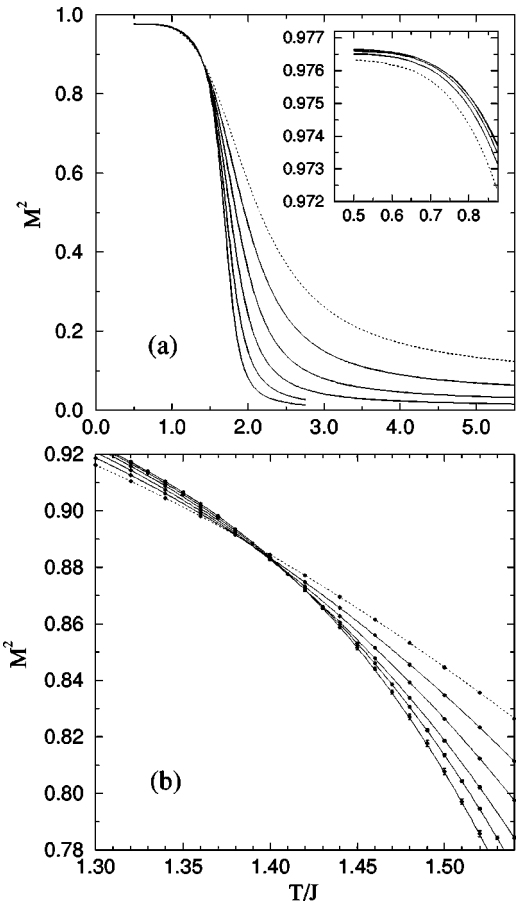

FIG. 5. (a) Magnetization squared vs temperature for system sizes $N=16$ (dotted curve), 32, 64, 128, 256, and 512 (solid curves). The statistical errors are smaller than the width of the curves. (b) The same quantity on a more detailed scale in the intersection region. The points with barely visible error bars are the simulation results. The curves are third-order polynomial fits.

can be compared with $T_c(h=0)\approx 1.53J$ for the classical model [17]. A reduction of T_c is expected on account of quantum fluctuations for $h>0$. The quite small reduction for $h/J=0.5$ is consistent with the $T\rightarrow 0$ magnetization being only slightly reduced from the classical value $M(0)=1$.

The reason for the different form of the finite-size scaling for $h>0$ should be clarified. Recently, a related model with a long-range interaction in the imaginary-time direction (dissipative transverse Ising chain) has been studied [32,31]. Some results for the quantum phase transition were obtained, which should also be relevant when $T_c\rightarrow 0$ in the model considered here (by switching the roles of the spatial and imaginary-time dimensions). The crossings observed in Fig. 5 could reflect a nonasymptotic behavior related to the quantum phase transition, i.e., for some larger L the curves would cease to cross as the classical $h=0$ scaling [17] sets in. The crossing behavior would, if this is the true asymptotic behavior, imply that the KT transition of the $h=0$ model is modified when $h>0$. A modified scaling would normally not be expected [18], but considering that the classical model with interactions decaying as $1/r^\alpha$ has continuously varying exponents for $3/2<\alpha<2$ and a KT transition only exactly at α

$=2$ [13–17], the system could potentially be very sensitive to the modified space-time interactions affected by the transverse field.

IV. DISCUSSION

An efficient approach to long-range interacting quantum models has been developed here within the framework of transverse Ising models. It is important to note that the technique can also be generalized to other types of systems, with the usual caveat of sign problems [25]. What is particular about the Ising interaction is that it can be written so that a spin-spin term either gives zero or a constant when acting on an arbitrary basis state. This is what is needed in order to reduce the interactions to local constraints in the SSE formalism. However, the algorithm can easily be modified to cases where the diagonal interaction can take several nonzero values. The first modification is in the diagonal update. For the Ising model, the probability of selecting a given bond (2d) is given by a matrix element corresponding to the spin pair being in a configuration energetically favored by the interaction. If the spins are in a nonfavored configuration (corresponding here to a vanishing matrix element) the update is simply rejected. In the general case, the probability to use in this update should correspond to the largest diagonal matrix element on a given bond, and if the actual configuration corresponds to a smaller matrix element the update should be accepted only with a probability reflecting this smaller value (i.e., the ratio between the actual value and the largest matrix element). The quantum-cluster update can be modified by using ideas developed within the ''directed-loop'' algorithm [24]. For example, there could be four-particle vertex processes where the whole vertex is not necessarily reversed as in Fig. 4(a). The process could instead either go straight through the vertex (modifying the vertex only at the entrance and exit legs) or ''bounce'' back without modifying the vertex at all. The details of how this is done in practice will of course depend on the types of diagonal and off-diagonal terms in the Hamiltonian. The main point to note is that in the SSE approach all the information needed to update the vertices is contained in the vertices themselves, which are always local and can be generated in the diagonal update based purely on local decisions.

The transverse Ising simulation algorithm has here been tested on a one-dimensional model with long-range interactions decaying as $1/r^2$. The program requires almost no modifications for higher-dimensional systems, and random interactions are also very easy to implement. The high accuracy of these simulations demonstrates that the algorithm indeed is very efficient. The computer resources used for this work were quite modest; on the order of 200 CPU hours on an SGI Origin2000. The scaling of the CPU time is close to linear in N for the $1/r^2$ interaction, for which the interaction sum (6) converges rapidly. Only the local updates discussed in Sec. II B were used in these simulations. The cluster updates have been tested as well and improve the performance. The quantum-cluster update should be particularly useful for studying the quantum phase transition, where there will be a broad distribution of the sizes of the clusters constructed in this update.

The initial study of $1/r^2$ model presented here was primarily intended as a demonstration of the algorithm. The results suggest the possibility of differences between the model with and without a transverse field and motivate further large-scale studies.

Future studies will also address how well the method works in practice for a variety of other systems that are more challenging because of frustrated interactions, long-range frustrated interactions, or even randomly frustrated long-range interactions. For short-range interactions, it would also be interesting to see how the SSE quantum-cluster method constructed here compares to the transverse Ising cluster method previously developed for continuous-time world-line simulations [7].

ACKNOWLEDGMENTS

I would like to thank P. Henelius and O. Syluuåsen for discussions of the SSE algorithm, and M. Troyer, K. Völker, and P. Werner for communicating their unpublished results on the dissipative Ising model. This work was supported by the Academy of Finland (Project No. 26175).

[1] E. Luijten and H.W.J. Blöte, Int. J. Mod. Phys. C **6**, 359 (1995).
[2] S. Kirkpatrick, S. Gelatt, and M.P. Vecchi, Science **220**, 671 (1983); S. Kirkpatrick, J. Stat. Phys. **34**, 975 (1984).
[3] A.B. Finnila, M.A. Gomez, C. Sebenik, C. Stenson, and J.D. Doll, Chem. Phys. Lett. **219**, 343 (1994); T. Kadowaki and H. Nishimori, Phys. Rev. E **58**, 5355 (1998).
[4] G.E. Santoro, R. Martonak, E. Tosatti, and R. Car, Science **295**, 2427 (2002).
[5] S. Sachdev and A.P. Young, Phys. Rev. Lett. **78**, 2220 (1997); S. Sachdev, Nucl. Phys. B **464**, 576 (1996).
[6] M. Guo, R.N. Bhatt, and D.A. Huse, Phys. Rev. Lett. **72**, 4137 (1994); H. Rieger and A.P. Young, *ibid.* **72**, 4141 (1994); C. Pich, A.P. Young, H. Rieger, and N. Kawashima, *ibid.* **81**, 5916 (1998).
[7] H. Rieger and N. Kawashima, Eur. Phys. J. B **9**, 233 (1999).
[8] M.J. Rozenberg and D.R. Grempel, Phys. Rev. Lett. **81**, 2550 (1998).
[9] W. Wu, D. Bitko, T.F. Rosenbaum, and G. Aeppli, Phys. Rev. Lett. **71**, 1919 (1993); D. Bitko, T.F. Rosenbaum, and G. Aeppli, *ibid.* **77**, 940 (1996).
[10] A.W. Sandvik and J. Kurkijärvi, Phys. Rev. B **43**, 5950 (1991); A.W. Sandvik, J. Phys. A **25**, 3667 (1992); Phys. Rev. B **56**, 11 678 (1997).
[11] R.H. Swendsen and J.S. Wang, Phys. Rev. Lett. **58**, 86 (1987).
[12] Schemes for treating long-range interactions have also been previously investigated within Handscomb's method: I.V. Rozhdestvensky and I.A. Favorsky, Mol. Simul. **9**, 213 (1992); O.N. Vassiliev, I.V. Rojdestvenski, and M.G. Cottam, Physica A **294**, 139 (2001).

[13] P.W. Anderson and G. Yuval, J. Phys. C **4**, 607 (1971).
[14] J.M. Kosterlitz, Phys. Rev. Lett. **37**, 1577 (1976).
[15] Z. Glumac and K. Uzelac, J. Phys. A **22**, 4439 (1989).
[16] S.A. Cannas, Phys. Rev. B **52**, 3034 (1995).
[17] E. Luijten and H. Meßingfeld, Phys. Rev. Lett. **86**, 5305 (2001).
[18] A. Dutta and J.K. Bhattacharjee, Phys. Rev. B **64**, 184106 (2001).
[19] M. Suzuki, S. Miyashita, and A. Kuroda, Prog. Theor. Phys. **58**, 1377 (1977); J.E. Hirsch, R.L. Sugar, D.J. Scalapino, and R. Blankenbecler, Phys. Rev. B **26**, 5033 (1982).
[20] D.C. Handscomb, Proc. Cambridge Philos. Soc. **58**, 594 (1962); **60**, 115 (1964).
[21] J.W. Lyklema, Phys. Rev. Lett. **49**, 88 (1982); D.H. Lee, J.D. Joannopoulos, and J.W. Negele, Phys. Rev. B **30**, 1599 (1984); S. Chakravarty and D.B. Stein, Phys. Rev. Lett. **49**, 582 (1982).
[22] H.G. Evertz, Adv. Phys. **52**, 1 (2003).
[23] A.W. Sandvik, Phys. Rev. B **59**, R14 157 (1999).
[24] O.F. Syljuåsen and A.W. Sandvik, Phys. Rev. E **66**, 046701 (2002).
[25] P. Henelius and A.W. Sandvik, Phys. Rev. B **62**, 1102 (2000).
[26] A.W. Sandvik, R.R.P. Singh, and D.K. Campbell, Phys. Rev. B **56**, 14 510 (1997).
[27] P.W. Kasteleyn and C.M. Fortuin, J. Phys. Soc. Jpn. **26**, 11 (1969); C.M. Fortuin and P.W. Kasteleyn, Physica (Amsterdam) **57**, 536 (1972).
[28] U. Wolff, Phys. Rev. Lett. **62**, 361 (1989).
[29] A.W. Sandvik, Phys. Rev. B **66**, 024418 (2002).
[30] E. Marinari and G. Parisi, Europhys. Lett. **19**, 451 (1992).
[31] P. Werner, K. Völker, and M. Troyer (private communication).
[32] S. Pankov, S. Florens, A. Georges, G. Kotliar, and S. Sachdev, e-print cond-mat/0304415.