

Engineering local optimality in quantum Monte Carlo algorithms

Lode Pollet^{a,*}, Kris Van Houcke^b, Stefan M.A. Rombouts^b

^a *Theoretische Physik, ETH Zürich, CH - 8093 Zürich, Switzerland*

^b *Vakgroep subatomaire en stralingsfysica, Proeftuinstraat 86, Universiteit Gent, Belgium*

Received 19 August 2005; received in revised form 22 December 2006; accepted 20 March 2007

Available online 28 March 2007

Abstract

Quantum Monte Carlo algorithms based on a world-line representation such as the worm algorithm and the directed loop algorithm are among the most powerful numerical techniques for the simulation of non-frustrated spin models and of bosonic models. Both algorithms work in the grand-canonical ensemble and can have a winding number larger than zero. However, they retain a lot of intrinsic degrees of freedom which can be used to optimize the algorithm. We let us guide by the rigorous statements on the globally optimal form of Markov chain Monte Carlo simulations in order to devise a locally optimal formulation of the worm algorithm while incorporating ideas from the directed loop algorithm. We provide numerical examples for the soft-core Bose–Hubbard model and various spin- S models.

© 2007 Elsevier Inc. All rights reserved.

PACS: 02.70.Tt; 05.50.+q; 82.20.Wt

1. Introduction

Monte Carlo methods have become a standard numerical tool in many branches of science, offering exact results in a statistical sense [1]. In physics, some Monte Carlo algorithms still resemble the original description by Metropolis [2] in the fifties. For instance, Ising and other classical lattice models were for a very long time simulated using random spin flips. Such algorithms suffer from a critical slowing down in the neighborhood of the second order phase transition. However, some twenty years ago Swendsen and Wang found a solution using cluster updates [3,4], completely overcoming the critical slowing down for the classical Ising model. Monte Carlo methods have also been applied to quantum many-body systems, where one tries to sample either the wavefunction or the partition function [5]. The quantum analog of cluster updates, namely the loop algorithm [6], triggered the development of the worm algorithm [7] and the operator loop algorithm [8] (and later the directed loop algorithm [9,10]) in the stochastic series expansion representation. These algorithms share the properties that they sample the one-body Green function, that they are formulated in continuous

* Corresponding author. Tel.: +41 1 633 2592.

E-mail address: pollet@itp.phys.ethz.ch (L. Pollet).

time [7] and are based on a world-line representation. Thus both algorithms are similar [11]. These algorithms have successfully been applied to spin systems and to the Bose–Hubbard model [12]. Recently, the worm algorithm has been formulated in the canonical ensemble, allowing to study more systems including superconducting grains and the nuclear pairing Hamiltonian [13,14].

The efficiency of a numerical simulation method is primordial: efficient algorithms lead to more accurate results at the same computational cost and allow for the study of larger systems. The algorithms mentioned above are based on a Markov process that results in a random walk in a specific configuration space. Configurations are visited by the random walker proportionally to their respective weights. By the Markov process, subsequent measurements are trivially correlated. The transition matrix specifies the probability of going from one configuration to another, and has to be defined in advance. In practice, one requires that transition matrices satisfy the principle of detailed balance [1]. This, however, still leaves some freedom in the choice of the transition matrices, which can be used to optimize the efficiency of the algorithm. A convenient updating scheme is the Metropolis-Hastings algorithm [2,15]: a limited number of configurations can be reached from the current one by defining a proposal distribution, and the transition to the new configuration is accepted or rejected according to detailed balance [1].

In previous work, we introduced the notion of ‘locally optimal Monte Carlo’ [16], when the degrees of freedom of every transition matrix are chosen in such a way as to mimic the globally optimal transition matrix, which, at least in principle, can be written down exactly. This approach is a best guess for obtaining optimal efficiency, and was found successful in practice for the directed loop algorithm in the stochastic series expansion representation [16]. Here we investigate the consequences of the locally optimal Monte Carlo idea for the worm algorithm, and try to combine the advantages of the directed loop algorithm with the advantages of the worm algorithm. This results in a new formulation of the worm algorithm, hereafter called the locally optimal worm algorithm (LOWA) [17]. We show results for various spin models and the Bose–Hubbard model, and compare the efficiency of the LOWA with that of the directed loop algorithm in the stochastic series expansion framework.

2. The algorithm

We consider a two-body Hamiltonian H defined on a discrete lattice of system size L . The Hamiltonian can be written as $H = H_0 - V$, where the terms H_0 and V are to be specified later on. We also assume a single particle basis $|i_v\rangle = |i_{v_1}, \dots, i_{v_L}\rangle$ of H_0 such that the action of any term in the Hamiltonian on a basis state yields a single basis state. The models that we typically have in mind are the Bose–Hubbard model,

$$H = -t \sum_{\langle i,j \rangle} b_i^\dagger b_j + \frac{U}{2} \sum_i n_i(n_i - 1) + V_{mn} \sum_{\langle i,j \rangle} n_i n_j - \sum_i \mu n_i, \quad (1)$$

and a general spin- S model,

$$H = J \sum_{\langle i,j \rangle} \frac{1}{2} (S_i^+ S_j^- + S_i^- S_j^+) + \Delta S_i^z S_j^z - h \sum_i S_i^z. \quad (2)$$

In both expressions the notation $\langle i, j \rangle$ refers to the sum over nearest-neighbor sites only. In the Bose–Hubbard model, bosons are created on site j by the operator b_j^\dagger and the number of bosons on site j is counted by the number operator n_j . The kinetic term describes hopping of the bosons with tunneling amplitude t , while we consider two kinds of potential energy: on-site repulsion with strength U and nearest-neighbor repulsion with strength V_{mn} . For the spin anti-ferromagnet (spin exchange amplitude $J > 0$), we require that the lattice is bipartite. All matrix elements remain positive as long as the model does not exhibit any frustration which can for instance be induced by second nearest-neighbor hopping. As the calculations serve to demonstrate the ideas related to efficiency, we restrict the discussion to one dimension.

A worm algorithm [7] is a quantum Monte Carlo algorithm where the decomposition of the partition function,

$$Z = \text{Tr} \sum_{n=0}^{\infty} \int_0^\beta dt_n \int_0^{t_n} dt_{n-1} \cdots \int_0^{t_2} dt_1 e^{-t_1 H_0} V e^{-(t_2 - t_1) H_0} \cdots e^{-(t_n - t_{n-1}) H_0} V e^{-(\beta - t_n) H_0}, \quad (3)$$

is sampled indirectly by performing local moves in the extended configuration space of open world-line configurations in the Green function sector,

$$Z_e = \text{Tr}[\mathcal{T}((b_i(t_0)b_j^\dagger(\tau) + \text{h.c.}) \exp(-\beta H))]. \quad (4)$$

The symbol \mathcal{T} denotes time ordering, and we have introduced the Heisenberg operators $(b_i(t_0)b_j^\dagger(\tau) + b_i^\dagger(t_0)b_j(\tau))$ which we call the ‘worm’ operators. Summing over all possible worms requires the extra sum evaluation of

$$\sum_{i,j,t_0,\tau} \dots e^{-t_0 H_0} b_i^\dagger e^{t_0 H_0} \dots e^{-t_\tau H_0} b_j e^{t_\tau H_0} \dots, \quad (5)$$

where we have explicitly written out the Heisenberg worm operators, and where the ‘ \dots ’ mean that they should be inserted at the right time, as implied by the time ordering of Eq. (4). Insertion of complete sets of basis states allows to replace the operators H_0 and V by matrix elements. We will choose to work in the number occupation basis for the Bose–Hubbard model and in the spin S_z basis for the spin models. A natural choice is to consider the one-body tunneling operators as perturbations V , and collect the diagonal one-body and two-body operators in H_0 . This leads to the path-integral formulation. The stochastic series expansion representation is found back when all operators of the Hamiltonian are put into V and $H_0 = 0$. The integration over time is then immediate. In the path-integral formulation, the extended partition function can be written as

$$Z_e = \sum_{n=0}^{\infty} \sum_{i_1, \dots, i_n} \sum_{i_j} \int_0^\beta dt_n \int_0^{t_n} dt_{n-1} \dots \int_0^{t_2} dt_1 W, \quad (6)$$

where the terms $W(n, t_1, \dots, t_n, i, j, t_0, t_\tau, |i_v\rangle)$ can be interpreted as weights when positive. The weights depend on the order n and the integration times $t_j, j = 1, \dots, n$ of the perturbative expansion of the partition function Z in V and also on the position and the time of the worm operators, and all possible inserted basis sets $|i_v\rangle$.

$$W = \langle i_1 | V | i_2 \rangle e^{-(t_2 - t_1) E_{i_2}} \langle i_2 | V | i_3 \rangle e^{-(t_3 - t_2) E_{i_3}} \dots e^{-(t_k - t_{k-1}) E_{i_k}} \langle i_k | b_i^\dagger | i_{k+1} \rangle \dots e^{-(t_l - t_{l-1}) E_{i_l}} \langle i_l | b_j | i_{l+1} \rangle \dots e^{-(t_n - t_{n-1}) E_{i_n}} \langle i_n | V | i_1 \rangle e^{-(\beta + t_1 - t_n) E_{i_1}}. \quad (7)$$

The Monte Carlo algorithm has to sample over all expansion orders n , all interaction times t_i , all possible worm times and positions, and all basis sets $|i_v\rangle$. In a worm algorithm this is achieved by moving one of the worm operators through the configuration space. Such updates shift, annihilate and create interactions and sample indirectly over all possible basis states.

It is convenient to represent the weights in a graphical representation using world lines. An example is shown in Fig. 1.

Whenever the worm operators reach each other, the discontinuities in the graphical representation cancel and the resulting configuration belongs to the partition function Z , apart from the worm matrix element $\langle i_k | b_i(0) b_i^\dagger(0) | i_k \rangle$. At this point we are free to assign any value to this matrix element. The standard choice, which we follow, is to take it constant for all states $|i_k\rangle$ (see Eq. (24)). This is graphically represented by the absence of ‘circles’ (which denoted the discontinuities in Fig. 1) in a world-line picture.

In the literature there have been two different implementations of the worm idea. First, there was the worm algorithm by Prokof’ev et al. [7], which was formulated in the path-integral representation. Later the operator loop [8] and directed loop algorithm [9,10] in the stochastic series expansion representation were formulated.

Compared to the original formulation of the worm algorithm by Prokof’ev et al. [7], the efficiency of the directed loop algorithm in the stochastic series expansion representation was found superior for spin systems and even for a homogeneous Bose–Hubbard model in both phases (except for the extreme soft-core case) [18]. This is an expected result for spin systems where the diagonal energies are of the same order as the spin exchange amplitudes, but for the Bose–Hubbard model this result feels unsatisfactory. For a trapped Bose–Hubbard model however, the worm algorithm was found superior [18].

In the present paper we think of a new formulation of the worm idea, trying to combine to advantages of the directed loop algorithm with those of the worm algorithm. More specifically, we want an algorithm where the worm inserts and annihilates the interactions (as in the worm algorithm), but we also want to use the direc-




Fig. 1. Graphical representation of a typical configuration in the Green function sector. Time goes from left to right in the figure, there are five sites. World lines are denoted by single lines (site is once occupied), double lines (site has occupancy two) or dashed lines (site is not occupied). Interactions (hopping of a particle) are denoted by vertical lines. The two circles mark a discontinuity in the world lines and correspond to the worm operators. One of them creates an extra particle, the other one annihilates it. As a consequence of the $U(1)$ symmetry, total particle number is conserved at every interaction.

ted worms as in the directed loop algorithm. The modification of the diagonal factors and the hopping factors contributing to the weight, Eq. (7), should be done in a locally optimal way [16].

3. Updates and detailed balance

The easiest way to ensure that the Markov chain converges to the correct invariant probability distribution is by assuring detailed balance [2],

$$W(X)T(X \rightarrow Y)q = W(Y)T(Y \rightarrow X). \quad (8)$$

The current configuration is denoted by X , the new configuration by Y . q is the acceptance factor to be used in the Metropolis algorithm, where the transition rule for going from X to Y is given by $P(X \rightarrow Y) = T(X \rightarrow Y) \min(1, q)$, while the transition rule for the reverse update is given by $P(Y \rightarrow X) = T(Y \rightarrow X) \min(1, 1/q)$, such that $W(X)P(X \rightarrow Y) = W(Y)P(Y \rightarrow X)$. We will now show that detailed balance is fulfilled for every possible update occurring in the algorithm.

3.1. Move

The simplest possible update is to move the mobile worm to another (later/earlier) time. We assume that the mobile worm is on site j , and corresponds to an operator b_j . The configurations before and after the update are shown in Fig. 2.

The relevant factors contributing to the weights before ($W(X)$) and after ($W(Y)$) the update are

$$W(X) = e^{-t_k E_{i_k}} \langle i_k | b_j | i_{k+1} \rangle e^{t_k E_{i_{k+1}}}, \quad (9)$$

$$W(Y) = e^{-t_{k'} E_{i_k}} \langle i_k | b_j | i_{k+1} \rangle e^{t_{k'} E_{i_{k+1}}}. \quad (10)$$

The notation i_k means the complete state over all sites as in Eq. (7), but the subscript j means the particular site j . The acceptance factor reads

$$q = \frac{e^{-\Delta t E_{i_k}} T(Y \rightarrow X)}{e^{-\Delta t E_{i_{k+1}}} T(X \rightarrow Y)}, \quad (11)$$

with $\Delta t = t_{k'} - t_k$. The exponentials can be canceled by choosing the transition probability densities as




Fig. 2. Graphical illustration of the Move update. The worm, originally at time t_k tries to jump to a later time $t_{k'}$. The state to the left of the worm is $|i_k\rangle$, with energy E_{i_k} . The state to the right of the worm is $|i_{k+1}\rangle$, with energy $E_{i_{k+1}}$.

$$T(X \rightarrow Y)dt = E_{i_k} e^{-\Delta t E_{i_k}} dt, \tag{12}$$

$$T(Y \rightarrow X)dt = E_{i_{k+1}} e^{-\Delta t E_{i_{k+1}}} dt. \tag{13}$$

The normalization factors E_{i_k} and $E_{i_{k+1}}$ enter into q . These factors will be taken into account explicitly together with the interaction matrix elements at the moment that interactions can be inserted/annihilated, see Section 3.2.

Let u be an uniform random deviate, $u \in [0, 1[$. Then, $p = -\log u$ follows an exponential distribution and we can compute the time shift window Δt as $\Delta t = p/E_{i_k} = -\log(u)/E_{i_k}$, with $E_{i_k} \geq 0$. We come back to this issue in Section 3.4. The recipe for the move update is thus

1. Draw a random deviate $u \in [0, 1[$.
2. When moving to the right (left), compute the time shift window $\Delta t = -\log(u)/E$, where E is the diagonal energy to the left (right) of the worm.
3. If no interaction is encountered, update the current worm time from time t to time $t + \Delta t$.

This amounts to a random walk based on Poisson steps. The exponential factors contributing to the weight, Eq. (7), might fluctuate heavily. The strong point of the Poisson moves is that these exponential factors are canceled exactly.

3.2. Inserting, removing and relinking an interaction

So far we assumed that the worm did not encounter any interaction during its propagation. What happens if a worm does reach an interaction? At that point a decision must be taken: can the worm pass the interaction or not? Or shall we delete the interaction? In that case we must also define updates to insert interactions that are in balance with the removal of interactions. In this section, we will first discuss two cases where the worm can pass the interaction, leaving it unchanged. Second, we discuss the insertion and removal of interactions which, as we will see, also incorporates the modification of interactions.

When the worm b_i^\dagger encounters an interaction $b_k^\dagger b_l$ whose sites are different from the site on which the worm head resides ($i \neq k, l$), the worm can pass the interaction with probability one. This is a consequence of the commutator of the worm operator and the interaction being zero, $[b_i^\dagger, b_k^\dagger b_l] = 0$ [17]. Second, when a worm operator b_i^\dagger encounters an interaction $b_i^\dagger b_j$, it can pass the interaction with probability one. This is also the result of the commutator being zero, $[b_i^\dagger, b_i^\dagger b_j] = 0$ [17].

Suppose we move to the right and want to insert an interaction. The update consists of a move from imaginary time t_w to the time t_k , inserting an interaction and then continue the move to the imaginary time t'_k , as shown in Fig. 3. Let's suppose that the worm pair was created at t_w , and that we want to evaluate estimators at time t'_k , meaning that the worm has to halt at time t'_k . We will later discuss generalizations. The relevant contributions to the weights of the configurations before and after the update are




Fig. 3. Graphical illustration of the insertion of an interaction by the worm. The worm, originally at the time t_w , before the time t_k , jumps to the time t_k and inserts an interaction there from site j to site $j - 1$. Afterwards the worm jumps to a later time t'_k , where we assume the worm always has to halt. Later we will relax this condition. Between times t_k and t'_k a new state i'_k is created. Note that this update is only possible if the occupation on site $j - 1$ is larger than zero.

$$W(X) = e^{-t_w E_{i_k}} \langle i_k | b_j | i_{k+1} \rangle e^{t_w E_{i_{k+1}}}, \tag{14}$$

$$W(Y) = e^{-t_k E_{i_k}} \langle i_k | b_j b_{j-1}^\dagger | i_{k'} \rangle e^{-(t_{k'} - t_k) E_{i_{k'}}} \langle i_{k'} | b_{j-1} | i_{k+1} \rangle e^{t_{k'} E_{i_{k+1}}}. \tag{15}$$

The transition to move to the new configuration Y is given by the probability density

$$T_M(X \rightarrow Y) dt = E_{i_k} e^{-\Delta t E_{i_k}} e^{-\Delta t' E_{i_{k'}}} dt, \tag{16}$$

with $\Delta t = t_k - t_w$ and $\Delta t' = t_{k'} - t_k$. The second exponential does not have an E prefactor with it, since all time shifts larger than $\Delta t'$ lead to the same final configuration Y . Similarly, for the reverse update

$$T_M(Y \rightarrow X) = \int_{\Delta t}^{\infty} E_{i_{k+1}} e^{-\tau E_{i_{k+1}}} d\tau \times \int_{\Delta t'}^{\infty} e^{-\tau E_{i_{k+1}}} E_{i_{k+1}} d\tau = e^{-\Delta t E_{i_{k+1}}} e^{-\Delta t' E_{i_{k+1}}} = e^{-(t_{k'} - t_w) E_{i_{k+1}}}. \tag{17}$$

Once again, all the exponential factors cancel in the acceptance factor q . There remains a factor $\langle i_k | V | i_{k'} \rangle / E_{i_k}$ (with V the interaction), which is taken into account in the equations of detailed balance for the actual insertion of a new interaction. If the worm was not forced to halt at the time t'_k , the update would continue with the insertion of another interaction at t'_k (which would take the extra $E_{i_{k'}}$ (i.e., the normalization factor of the second exponential in the right hand side of Eq. (16)) into account).

When inserting a new interaction, we can either make the current site j interact with site $j + 1$ or site $j - 1$ in one dimension. In higher dimensions d , interactions can be inserted to all $2d$ neighboring sites. We have to define the (conditional) probability distribution function that samples the three configurations of Fig. 4. When we are in configuration a , we have to define $P(a \rightarrow b)$ and $P(a \rightarrow c)$, both corresponding to the insertion of a new interaction. Similarly, when we are in configuration b we have to define $P(b \rightarrow a)$ and $P(b \rightarrow c)$, corresponding to the removal and the modification (relinking) of an interaction, respectively. Updating between the configurations a , b and c should be done proportional to the following factors contributing to the weights:

$$w(a) = \langle i_k | b_j | i_{k+1} \rangle E_{i_k}, \tag{18}$$

$$w(b) = \langle i_k | b_j b_{j-1}^\dagger | i' \rangle \langle i' | b_{j-1} | i_{k+1} \rangle, \tag{19}$$

$$w(c) = \langle i_k | b_j b_{j+1}^\dagger | i'' \rangle \langle i'' | b_{j+1} | i_{k+1} \rangle. \tag{20}$$

The energy term in Eq. (18) will be explained in Section 3.5. In Fig. 4, the worm is on site j in configuration (a), on site $j - 1$ in configuration (b) and on site $j + 1$ in configuration (c). We will discuss three possibilities to sample configurations (a)–(c).




Fig. 4. When the worm moves to the right and tries to insert an interaction in configuration (a), the two possible new configurations are configurations (b) and (c). The third possibility is bouncing back and changing direction in configuration (a). The transition matrix is thus a 3×3 matrix in one dimension, or a $(2d + 1) \times (2d + 1)$ matrix in d dimensions.

- The ‘Metropolis-like’ way. Assume that we are in configuration (a). We choose with equal probability between (b) and (c). Say (b) was chosen. The transition to (b) is then accepted with probability $\min[1, w(b)/w(a)]$. If the update is not accepted, we stay in configuration (a). This approach has the advantage that only the matrix elements for configuration (b) and (a) have to be calculated or known, the ones for (c) are not needed. This approach is thus recommendable for high dimensions, and for long-range interactions. The (normalized) 3×3 transition probability matrix reads

$$T^{\text{Met}} = \begin{bmatrix} 1 - q_{ab} - q_{ac} & q_{ab} & q_{ac} \\ q_{ba} & 1 - \dots & q_{bc} \\ q_{ca} & q_{cb} & 1 - \dots \end{bmatrix}, \tag{21}$$

with $q_{kl} = (1/2) \min[1, w(l)/w(k)]$. The diagonal element corresponding to the lowest weight is thus zero.

- The heat-bath way. We choose between (a), (b) and (c) according to their relative weights, $\pi(j) = w(j)/(\sum_k w(k))$, irrespective of the current configuration. The weights of all configurations are needed in order to evaluate their sum. The transition probability matrix is

$$T^{\text{Hb}} = \begin{bmatrix} \pi(a) & \pi(b) & \pi(c) \\ \pi(a) & \pi(b) & \pi(c) \\ \pi(a) & \pi(b) & \pi(c) \end{bmatrix}. \tag{22}$$

- The locally optimal way. Heat-bath updates can have relatively large diagonal elements, which are associated with large rejection ratios or large bounce probabilities. The principle of locally optimal Monte Carlo updating suggests an optimal transition probability matrix,

$$T^{\text{Lo}} = \begin{bmatrix} 0 & \frac{\pi_2}{1-\pi_1} & \frac{\pi_3}{1-\pi_1} \\ \frac{\pi_1}{1-\pi_1} & 0 & \frac{1-2\pi_1}{1-\pi_1} \\ \frac{\pi_1}{1-\pi_1} & \frac{\pi_2}{\pi_3} \frac{1-2\pi_1}{1-\pi_1} & 1 - \dots \end{bmatrix}, \tag{23}$$

where the normalized weights $\pi(j)$ are now ordered in ascending order, $\pi_1 \leq \pi_2 \leq \pi_3$. The ordering of the weights makes sense only if they can be tabulated before the start of the actual simulation. The locally optimal matrix is the stochastic matrix with the lowest possible second largest eigenvalue, which is negative.

The non-zero diagonal elements correspond to bounces, meaning that the worm head changes its direction and undoes its changes until it reaches another point where an interaction can be changed, removed or inserted.

3.3. Insertion and removal of a worm pair

The only remaining update to discuss is the insertion/removal of a worm pair, which is the connection between the partition function sector Z and the Green function sector, Z_e . The update is depicted in Fig. 5.

In case of Fig. 5, the weights before and after the update are

$$W(X) = C_Z, \tag{24}$$

$$W(Y) = \langle i_k | b_j | i' \rangle \langle i' | b_j^\dagger | i_k \rangle = (n_j + 1), \tag{25}$$

where C_Z is a constant specifying the relative weights of the partition function and the Green function sectors, and the number n_j is the occupation on site j before a worm pair is inserted. This is the usual convention in a worm algorithm. However, we are free to choose the weight of the worm matrix elements, and in SSE they are usually taken as unity. One of the worm ends will be moved through configuration space and is called the mobile worm, while the other end remains stationary. We choose to always annihilate a worm pair when the mobile worm ‘bites’ into the stationary worm,

$$P(Y \rightarrow X) = 1. \tag{26}$$

A worm pair is inserted by choosing a random time and a random site. Thus (in case of Fig. 5),

$$P(X \rightarrow Y) = \frac{1}{\beta} \frac{1}{L} P(\text{an}) P(\text{dir}). \tag{27}$$

We have to define the probabilities $P(\text{cr})$ and $P(\text{an})$ that the occupation between the two worm operators is higher (‘an’) or lower (‘cr’) than outside this infinitesimal interval. We choose with equal probability among those, except when the occupation is zero or equal to the maximum occupation allowed. In case of zero occupancy, we choose with probability 50% not to insert a worm pair, and analogous to the case of maximum occupation number. Thus,

$$P(\text{an}) = \frac{1 - \delta_{n_j, N_{\text{Max}}}}{2}, \tag{28}$$

$$P(\text{cr}) = \frac{1 - \delta_{n_j, 0}}{2}. \tag{29}$$

In case of hard-core bosons or spin $-1/2$ systems, one can modify this relation so as to always insert a worm pair.

When a worm pair is inserted, there are two possibilities: of we move forward in time creating a particle, or we move backward in time creating a particle. Physically, moving forward in time and annihilating a particle is the same as the latter, while moving backward in time and annihilating a particle is the same as the former.




Fig. 5. Graphical illustration of the insertion of a worm pair. An arbitrary site j and an arbitrary time t_k is chosen. We have shown here the case that the occupation between the worm ends is increased by one.

Yet, the decorrelation benefits if both a direction (forward/backward) and an operation (creation/annihilation) are chosen. At this point, we explicitly choose a direction, and we take

$$P(\text{dir} \Rightarrow) = P(\text{dir} \Leftarrow) = 1/2. \quad (30)$$

This equation is correct in the directed loop algorithm since there are on average as many moves to the left as there are to the right.

Finally, we have to fix the constant C_Z . Apart from the cases $n_j = 0$ and $n_j = N_{\text{Max}}$, which we have already discussed, the update is always accepted when we choose (in case of Fig. 5)

$$C_Z = 4\beta L. \quad (31)$$

With this choice, the Green function $G(x, t)$ is automatically correctly normalized. The procedure to measure will be discussed in Section 3.5, while the value of the measurements is the same as the ones discussed in the generalized directed loop algorithm [19] and in the canonical worm algorithm [13].

3.4. Refinements

The perturbative expansion of the partition function is a Poisson process: the interactions (events) are distributed according to a Poisson distribution, with the intervals between them following an exponential distribution. We sampled these intervals exactly using exponential deviates divided by an energy. According to Eq. (7), these energies are the diagonal energies of the system. Taking into account that the direction of propagation is fixed in the directed loop algorithm, we only consider positive time shift windows Δt . Because the exponential distribution $\exp(-x)$ is only defined for positive x , the energies have to be positive, which is not a priori guaranteed. Let $E_{L(R)}$ be the energy to the left (right) of the mobile worm. We can proceed in the following ways:

1. We shift all energies with a large, positive constant, $\epsilon_{LR} = E_0 + E_{LR}$. This will result in small time shift windows. Since these energies also enter in the equations of detailed balance for inserting/removing/relinking an interaction, they might lead to large bounce ratios.
2. We try to make use of the fact that only energy differences are physically relevant. We have the algorithmic freedom to choose any pair of ϵ_L and ϵ_R such that $\epsilon_L - \epsilon_R = E_L - E_R$ and $\epsilon_L, \epsilon_R \geq 0$. A good choice is $\epsilon_{L,R} = E_{L,R} - \min[E_L, E_R]$. This quantity can be zero, meaning that the jump in imaginary time is infinite, i.e., that the next interaction is always reached. Compared to the previous approach, we make thus larger jumps in imaginary time, and the energies that enter the detailed balance equations of inserting/removing/relinking an interaction are of the same order of magnitude as the hopping matrix elements. However, we found that this parameter choice resulted in some anomalously long (non-closing) loops and in problems with ergodicity when these long loops are discarded.
3. We suggest to use the shifted energies

$$\epsilon_{L,R} = E_{L,R} - \min[E_L, E_R] + E_{\text{off}}, \quad (32)$$

where the energy offset E_{off} overcomes the aforementioned problems with ergodicity. It makes sense to choose $E_{\text{off}} = \langle V \rangle$ with $\langle V \rangle$ a typical matrix element of the interaction.

After the worm has passed an interaction at time t_k , we have to calculate the new energy parameters $\epsilon_{L,R}$ and consider a new time shift. The latter can be accomplished by either drawing a new exponential deviate (as in step 3), or by adjusting the time shift window $\Delta\tau \rightarrow \Delta\tau - (t_k - t_{k-1})\epsilon_L$ (when moving to the right), using the property that the exponential distribution has no ‘memory’.

3.5. Detailed balance and the directed loop algorithm

To prove detailed balance one has to consider the global updates. Such a global update starts and ends with a jump in imaginary time (thus only a forced halting at a chosen time can end a global update). In between, there can be any number of insertions, annihilations or relinks. For example, the following sequence satisfies detailed balance: insertion of a worm pair-jump-insertion of an interaction-jump-insertion of an interaction-jump-removal of a worm pair. We have already shown that the insertion and removal of a worm pair are bal-

anced, so we just have to look at the sequence starting and ending with the jump. For every such sequence there exists also the exact opposite sequence. Writing down the acceptance factor for the global move

$$q = \frac{W(Y)T(Y \rightarrow X)}{W(X)T(X \rightarrow Y)}, \tag{33}$$

one finds that all exponential factors in the weights are cancelled out by the probabilities for the Poisson jumps. For each insertion there enters an interaction matrix element $\langle i_L | V | i_R \rangle$ in the numerator due to the ratio $W(Y)/W(X)$, and a factor E (E_L for moves to the right, E_R for moves to the left) in the denominator due to the normalization of the preceding Poisson jump. Both these factors are balanced through the locally optimal transition matrix of Section 3.2. As a result one finds that q is exactly equal to one, hence all moves can be accepted, which greatly simplifies the computer code. The resulting algorithm samples the configurations that appear in the decomposition of Eq. (3). Note that the above reasoning still holds if we assume that the worm continues to move in the same direction after each Poisson move and only allow bounces at the moment when one decides whether or not to insert or remove interactions. This leads to a version of the worm algorithm that is similar to the directed loop algorithm for SSE [9,10,17].

The intermediate steps of the algorithm correspond to configurations in the decomposition of the extended partition function Z_e of Eq. (6). To understand this, consider a point in imaginary time at a distance τ of the time where the worm was inserted. Now suppose we *halt* (this is crucial for measuring the Green function) the worm head at the moment it passes the point τ , and suppose we choose with probability 1/2 whether to continue the move in the same direction or to move in the opposite direction (we will see further on that this latter condition is not necessary). Following exactly the same reasoning as above, we see that this algorithm leads to detailed balance between configurations in the decompositions of Z and $Z_e(\tau)$. Now from here we can derive that this is true even if we continue the worm move in the same direction at the moment of passing the point τ . To see this it is instructive to consider both directions as two branches of a normalized transition kernel Ψ that depends on two variables, namely the time Δt and the direction D , and fulfills detailed balance:

$$W(X)\Psi(\Delta t, \rightarrow) = W(Y)\Psi(\Delta t, \leftarrow), \tag{34}$$

for $\Delta t = |t_Y - t_X|$ and $t_X < t_Y$. We have that

$$P(X \rightarrow Y) = \Psi(\Delta T, \rightarrow), \quad X < Y = \Psi(\Delta T, \leftarrow), \quad X > Y. \tag{35}$$

Thanks to time reversal symmetry, the statement that a worm creates a particle and propagates forward in time is completely equivalent to the statement that the worm annihilates a particle and propagates backward in time. Therefore both directions will occur with equal probability. Suppose that at a given moment a configuration X is the actual one with a probability proportional to its weight $W(X)$. Then the probability for a configuration Y to occur at the next step is proportional to

$$\begin{aligned} \int W(X)P(X \rightarrow Y)dX &= \sum_D \int W(X)\Psi(\Delta t = |t_Y - t_X|, D)dX \\ &= \int_{t_Y > t_X} W(X)\Psi(\Delta t, \rightarrow)dX + \int_{t_Y < t_X} W(X)\Psi(\Delta t, \leftarrow)dX \\ &= \int_{t_Y > t_X} W(Y)\Psi(\Delta t, \leftarrow)dX + \int_{t_Y < t_X} W(Y)\Psi(\Delta t, \rightarrow)dX \\ &= \sum_D \int W(Y)\Psi(\Delta t, D)dX = W(Y) \int P(Y \rightarrow X)dX = W(Y). \end{aligned}$$

This equation holds for any algorithm where the direction is fixed, and also forms the basis of the *bounce algorithm* of Ref. [20].

So, even when we do not change the direction of the worm at the moment that it passes the point τ , the probability to pass a point τ is still proportional to the weight of the corresponding configuration in the extended partition function Z_e . This is quite subtle. Suppose we do a jump of the worm operator, without encountering an interaction. Then it is not possible to immediately come back to the original configuration in the next step since the direction is preserved. This observation lies at the heart of the proof of convergence

of the algorithm given by the authors of Refs. [9,10,19]. They prove that detailed balance is satisfied between any two diagonal configurations, from which it follows that every local step is balanced (in the SSE representation). Here, we see that detailed balance is fulfilled every time the worm is forced to halt at a chosen time, but we emphasize that precisely at the moment of inserting or annihilating or relinking an interaction (without further jump of the worm) detailed balance is not fulfilled.

Because the probability distribution for two consecutive Poisson steps in the same direction is identical to the probability distribution for a single Poisson step, one finds that in this case the dynamics of the algorithm is completely equivalent to the dynamics without considering a special point τ . Therefore one can state in general that the probability for the worm head to pass a point at a distance τ from the worm tail is given by the weight of the corresponding configuration in the extended partition function Z_e .

This observation allows for an efficient and unbiased evaluation of the equal-time and unequal time Green function $G_{ij}(\tau) = Z_{e,ij}(\tau)/Z$: each time the worm head at site i passes the worm tail on a different site j , one has a measurement for the equal time Green function $G_{ij}(0)$, i.e., the one-body density matrix. Counting the times that the worm passes at a distance τ , one obtains a measurement of the unequal time Green function $G_{ij}(\tau)$.

3.6. Stochastic series expansion representation

As we have already mentioned, we end up in the stochastic series expansion (SSE) representation if we treat all terms in the Hamiltonian as interactions. In the present formulation this means that the diagonal energies are always zero, leading to infinite time shift windows all the time. The mobile worm will jump from interaction to interaction, either deleting or relinking it or bouncing back. There is a serious problem with this algorithm, because it will never insert a new interaction. Therefore, in the SSE representation one needs two updates: one update is to scan over all (discrete) times in order to insert and remove interactions, the other one consists of modifying interactions with a fixed graph. The proof of convergence with directed loops in the second update proceeds in the same way as for the LOWA algorithm.

3.7. Summary of the LOWA algorithm

We recapitulate and write down the full algorithm for the soft-core Bose–Hubbard model.

1. Pick an arbitrary site and an arbitrary time and call it (i_0, τ_0) . Find the occupation on all sites at that particular time τ_0 . Calculate the corresponding diagonal energy. A direction (left or right) is chosen with equal probability. Assume propagation to the right was chosen.
2. At (i_0, τ_0) a worm-pair (tail-head) is inserted. If the occupation is higher than zero, the occupation between the worm ends can either be increased or decreased. If the occupation at (i_0, τ_0) is zero, then with probability 50% a worm is inserted with increased occupation between the worm ends. With probability 50% no worm is inserted.
3. When moving to the right (left), we denote by ϵ the shifted energy to the left (right) of the worm head, Eq. (32). Draw an exponential deviate, $p = -\ln(u)$ with u an uniform random number, $0 < u \leq 1$. Evaluate the imaginary time shift window $\Delta t = p/\epsilon$ and the new worm time $\tau' = \tau + \Delta t$.
4. If the worm head encounters the worm tail (at the same site) during its propagation, the update ends with probability one and we arrive at a new diagonal configuration.
5. If the new worm time is larger than the time to the next interaction, the new worm time only equals the time of the next interaction. The worm can either bounce back, pass, annihilate or relink the interaction, according to the locally optimal transition matrix.
6. If no interaction is encountered in the imaginary time shift window, the worm shifts to its new time where an interaction is inserted or a bounce occurs, according to the locally optimal transition matrix.
7. Go back to step 3.

Every local single step respects the invariant distribution in the Green function sector. When the mobile worm reaches the stationary worm, one can measure diagonal observables such as the energy, winding number, density, etc. They can be updated in the same way as in the directed loop [9] and worm [7] algorithms.

The algorithm described above is valid when the diagonal energies involve a single site, but also when the diagonal energies contain nearest-neighbor repulsion terms, and even longer range interactions.

The worm algorithm in path-integral representation is in essence the same algorithm as the LOWA algorithm. The only difference is the way ergodicity and convergence to the correct invariant distribution are implemented with respect to the direction of worm propagation. In the worm algorithm, one chooses at every step between forward and backward propagation in time with equal probability, while in the LOWA the direction is maintained until an interaction forces the worm to alter its direction of propagation. In the context of the canonical worm algorithm, even other choices have been implemented [13]. All these algorithms are the same in spirit, they are slightly different implementations of the same idea of performing local updates in the Green function sector. The directed loop algorithm has previously been formulated in the path-integral formulation [9]. Although there are some resemblances, the LOWA is different, more general, and the principles that lie at heart of the derivation of the algorithm are different.

4. Is the locally optimal worm algorithm efficient?

Although we have argued why we believe the proposed LOWA is efficient, we can only verify by doing numerics in order to get a definite answer on its efficiency. The results are compared with data obtained by the directed loop algorithm in the stochastic series expansion representation of Refs. [9,19,16], which will be abbreviated as DLSSE. We used the method of Ref. [16] for the actual computation of the DLSSE. The error and autocorrelation estimation was done using a binning analysis.

A direct comparison is complicated because present implementations use different data structures for both algorithms. In the DLSSE, a doubly linked list is constructed before the loop update. Since the graph is fixed, the number of elements in this list cannot change. This allows to allocate memory statically. In the worm algorithm on the other hand, the number of interactions can change at any time. In our Fortran code this was implemented using two arrays of a predetermined fixed length, corresponding to the interactions before and after the current mobile worm time.

4.1. Bose–Hubbard model

We have calculated the standard deviations on the kinetic energy and on the squared density for a one-dimensional Bose–Hubbard model of size $L = 32$ sites at an inverse temperature of $\beta = 32$ and with a fixed chemical potential $\mu = 2$ in the absence of nearest-neighbor repulsion, $V_{nn} = 0$. We work in units $t = 1$. Simulations consisted of 40 bins that each ran 300 s on a Pentium III processor. We imposed a particle number cutoff of ten particles per site for $U = 1–3$, while a cutoff of five particles per site was taken for the other values of U , ranging from $U = 4$ to 10. Imposing a cut-off is a necessity for the DLSSE, but not for the LOWA. The number of loops per update was optimized along the guidelines of Ref. [19]. The Mott phase is reached for $U = 6$.

We have calculated the standard deviation on the kinetic energy and on the density squared (i) for the DLSSE. (ii) for the LOWA where the diagonal energy parameters were chosen according to approach (a) (iii) for the LOWA where the diagonal energy parameters were chosen according to approach (c). In (ii) and (iii) the 3×3 transition matrices were taken as the locally optimal ones as in Eq. (23). We shall discuss this in Section 4.3. In Figs. 6 and 7 the results for algorithms (i) and (iii) are shown.

Among the different LOWA optimization parameters, approach (iii) is almost always the most efficient one. The DLSSE seems to be the preferred model for very low values of U . However, approach (ii) performs lots better than approach (iii) in this regime, and its efficiency is comparable to that of the DLSSE. When the diagonal matrix elements are much larger than the off-diagonal ones, as in the Mott phase $U > 6$, the present algorithm is superior. Admittedly, we recognize that it is not unambiguous how the directed loop simulations should be performed in the Mott phase because of the very short loop sizes.

4.2. Spin systems

For spin systems, the magnitude of the diagonal and the off-diagonal matrix elements is of the same order. One can thus expect that the DLSSE is more efficient for spin models than for soft-core bosonic models, since

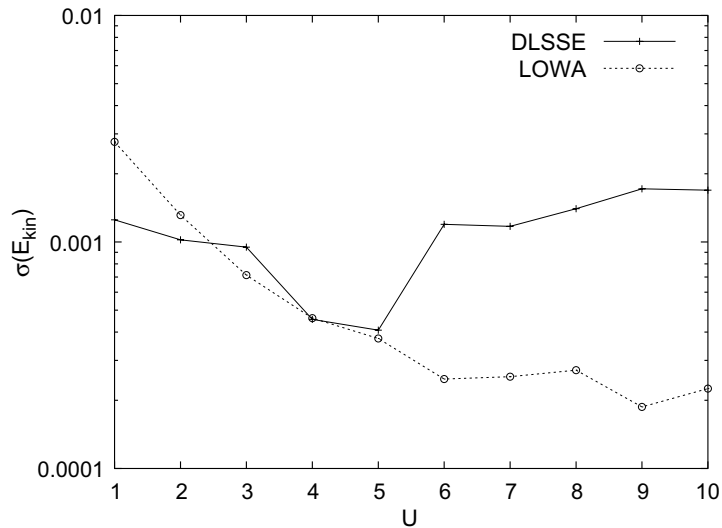


Fig. 6. Standard deviation (statistical error) on the kinetic energy for the directed loop algorithm in the stochastic series expansion framework (DLSSE), and the locally optimal worm (LOWA) using an energy offset added to the energy difference for the diagonal energy parameters and using locally optimal updates (Eq. (23)). Simulations consisted of 40 bins that each ran 300 s on a Pentium III processor for a one-dimensional Bose–Hubbard model of $L = 32$ sites at an inverse temperature of $\beta = 32$ and with a fixed chemical potential $\mu = 2$. The accuracy of the data points is about 10%.

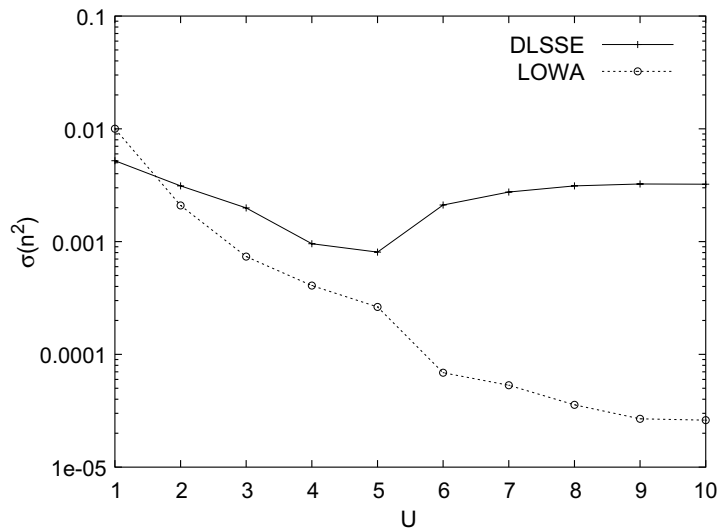


Fig. 7. Idem as in Fig. 6, but now for the standard deviation on the average square density, $\langle n^2 \rangle = \langle (\sum_i n_i)^2 \rangle / L$.

diagonal and off-diagonal operators are treated on equal footing in the stochastic series expansion representation. It is thus interesting to compare the efficiency of the locally optimal worm with the efficiency of the DLSSE for a spin-1/2 chain. The LOWA was most efficient when the energy offset parameter was set to $E_{\text{off}} = 0.5$. In Fig. 8 the standard deviations of the total energy and of the kinetic energy are shown, which are obtained by applying the LOWA and the DLSSE to a spin-1/2 chain subject to a magnetic field H . We find that the locally optimal worm is superior to the DLSSE in our implementations, but it is more meaningful to say that both algorithms behave similarly as the magnetic field is increased, while the algorithms are most




Fig. 8. Standard deviations on the total energy ($\sigma(E)$) and on the kinetic energy ($\sigma(E_{kin})$) obtained after a simulation of a spin-1/2 chain with $J = \Delta = 1$ consisting of $L = 64$ sites and with inverse temperature $\beta = 16$. Simulations have been performed with the directed loop algorithm in the stochastic series expansion framework (DLSSE) and with the locally optimal worm algorithm (LOWA). Computations consisted of +40 bins that each ran for +60 sec on a Pentium-III processor.

efficient near $H = 0$. Analogous conclusions were found for a spin +1 chain. For a spin +2 chain however, the DLSSE was found to be superior.

4.3. Heat-bath updates and the locally optimal matrix

A comparison between the heat-bath and locally optimal approach is made in Fig. 9 for a one-dimensional Bose–Hubbard model with parameters $U = 1, \mu = 0.5, \beta = 32, L = 32, V_m = 0.4$ and varying tunneling amplitude t . The ratio of the standard deviation obtained by the locally optimal approach to the




Fig. 9. Simulation of a one-dimensional Bose–Hubbard model with parameters $U = 1, \mu = 0.5, \beta = 32, L = 32, V_m = 0.4$ and varying tunneling amplitude t . Plotted is the ratio r for the condensate fraction ($r(cf)$) and for the total energy ($r(E)$), being the ratio of the standard deviation σ_{Lo} obtained by the locally optimal approach to the standard deviation σ_{Hb} obtained by the heat-bath approach ($r = \frac{\sigma_{Hb}}{\sigma_{Lo}}$). Simulations consisted of 40 bins of 300 s on a Pentium III processor per data point.

standard deviation obtained by the heat-bath approach is shown for the condensate fraction and the total energy. We see that the locally optimal approach is on average ten to fifteen percent better, but the effect is less pronounced than in the DLSSE [9,10,19,16]. Even smaller differences were found for some other parameter regimes.

4.4. Scaling of the worm with system size

Both the worm and the DLSSE are $\mathcal{O}(N)$ methods. In the absence of correlations between subsequent measurements, the needed computation time for a desired accuracy scales linearly with system size and inverse temperature. The scaling efficiency is further determined by the dynamical exponent z , which describes how the integrated autocorrelation time scales with system size and inverse temperature. The worm and directed loop algorithm have very low dynamical exponents; z is even zero in some high-dimensional cases. This beneficial scaling is the cornerstone for the study of very large system sizes at very low temperatures. Since the present algorithm is based on the same principles as the directed loop algorithm and the worm algorithm, one expects that the dynamical exponent is similar (at least of the same order) but the prefactor of the scaling behavior might be different.

We studied the scaling behavior for the critical system of an isotropic spin-1/2 Heisenberg chain ($\Delta = J = 1$) in zero magnetic field ($H = 0$), for which the worm updates are fast. We investigated the effects of increasing system size L at fixed inverse temperature β on the one hand and of increasing the inverse temperature β at fixed system size L on the other hand. This allows us to see whether the algorithm scales symmetrically with system size and temperature or not. All calculations ran for a fixed time of 40 bins of 300 s each on a Pentium III processor. We used optimal parameters for the simulation: we use the locally optimal transition matrix of Eq. (23) and we set $E_{\text{off}} = 0.5$. We focused on the standard deviation and the integrated autocorrelation time of the kinetic energy, since the modes of this observable couple to the slowest mode of the simulation while the measurements can be calculated at low computational cost. In Fig. 10 we study the standard deviation (statistical error) on the average kinetic energy per site. When the system size and the inverse temperature are sufficiently large, we see a gradual increase in the statistical error with an exponent of $\sigma \sim (L\beta)^{0.5}$, irrespective of whether L , β or both (note that the quantum imaginary time direction scales as




Fig. 10. Standard deviation on the average kinetic energy per site as a function of the system size L multiplied by the inverse temperature β for an isotropic spin-1/2 chain in zero magnetic field. Plots are shown when the system size and the inverse temperature are increased simultaneously ($L = \beta$), when the system size is held constant at $L = 128$ and only the temperature β is varied ($L = 129$), and when the temperature is held constant at $\beta = 16$ and the system size L varies ($\beta = 16$).




Fig. 11. Idem as in Fig. 10, but now the loop size (S_{loop}) is shown. The loop size S_{loop} is defined as the total number of interactions passed, inserted, annihilated and modified by the mobile worm in a single update.




Fig. 12. Idem as in Fig. 10, but now the integrated autocorrelation time of the kinetic energy, $\tau(E_{kin})$ is shown.

one classical space direction) are increased. Since, for larger lattices, the worm will visit each site less often, this result is intuitively understandable. From the data in Fig. 10 we can already see that the dynamic exponent z obeys $0 < z \leq 1$. Strangely, when β is taken too low (see the data points at $L\beta = 16$ in Fig. 10), the present algorithm loses its efficiency. Below we will relate this to higher integrated autocorrelation times. A similar picture results when we look at the loop size in Fig. 11. We define the loop size S_{loop} as the total number of interactions passed, inserted, annihilated and modified by the mobile worm in a single update. We see that the loop size S_{loop} scales as $S_{loop} \approx (L\beta)^1$. The loop size increases thus linearly with the increase in area in space-time. However, if we are very close to the ground state at a fixed system size, then increasing β does not result in longer loops and the algorithm loses its scaling properties.

The integrated autocorrelation time $\tau(E_{\text{kin}})$ of the kinetic energy in Fig. 12 shows a similar pattern: for large enough space-time areas, the integrated autocorrelation time scales as $\tau(E_{\text{kin}}) \sim (L\beta)^{0.42 \pm 2}$, which was fitted to the curve $L = \beta$ in Fig. 12. When the inverse temperature is too low, we unexpectedly find high autocorrelation times which explains the high standard deviations for the same data points in Fig. 10. This behavior could be due to the fact that the mobile worm is always forced to make relatively large jumps in the time direction.

With the DLSSE, the integrated autocorrelation time scales as $\tau(E_{\text{kin}}) \sim (L\beta)^{0.38 \pm 2}$. This result is in agreement with Ref. [9]. The dynamic exponent of the DLSSE is thus lower, yet in our calculations the standard deviations of the DLSSE increased more rapidly with increasing system size. This is due to the increase in computational cost for a single update, which scales worse for the DLSSE. The computational cost of a single update depends strongly on the way of implementation, and the scaling of the standard deviations with system size should be interpreted accordingly. In addition, when looking at the magnetization on every site, the worm algorithm performs much better than the DLSSE for all system sizes. We conclude that efficiency largely depends on the implementation of the algorithm and on the observables of interest [19].

5. Conclusion

In conclusion, we have presented a new formulation of the worm algorithm. The present algorithm has been derived using the concept of locally optimal Monte Carlo [16] and incorporates ideas both from the worm algorithm [7] and the directed loop algorithm [9] in the stochastic series expansion representation [8]. We have compared the efficiency of the present algorithm with that of the directed loop algorithm for spin chains and for the Bose–Hubbard chain. Especially when there are large diagonal matrix elements, the present worm algorithm is very successful. We have shown that choosing the locally optimal matrix for the transition matrices occurring in the stochastic subprocesses yields an efficient algorithm. We found that the loop size increases linearly with the increase in area in space-time, and that the dynamic exponent equals $z = 0.84 \pm 4$ for an isotropic Heisenberg chain without magnetic field. Seen the efficiency of the method and its advantageous scaling properties, the algorithm is suitable for large scale calculations of spin systems and soft-core bosonic models.

Acknowledgments

The authors thank the Research Board of the Universiteit Gent and the Fund for Scientific Research, Flanders for financial support. The authors acknowledge D. Ceperley, E. Gull, N.V. Prokof'ev, B.V. Svistunov, M. Troyer, S. Wessel and F.F. Assaad for valuable discussions.

References

- [1] J.S. Liu, Monte Carlo Strategies in Scientific Computing, (Springer Verlag, New York, USA, 2001).
- [2] N. Metropolis, A.W. Rosenbluth, M.N. Metropolis, A.H. Teller, E. Teller, J. Chem. Phys. 21 (1953) 1087.
- [3] R.H. Swendsen, J.S. Wang, Phys. Rev. Lett. 58 (1987) 86.
- [4] U. Wolff, Phys. Rev. Lett. 62 (1989) 361.
- [5] D. Ceperley, Rev. Mod. Phys. 71 (1999) S438.
- [6] H.G. Evertz, G. Lana, M. Marcu, Phys. Rev. Lett. 70 (1993) 875.
- [7] N.V. Prokof'ev, B.V. Svistunov, I.S. Tupitsyn, Zh. Eksp. Theor. Fiz. 114 (1998) 570 (Sov. Phys. JETP 87, 310); N.V. Prokof'ev, B.V. Svistunov, I.S. Tupitsyn, Phys. Lett. A 238 (1998) 253.
- [8] A.W. Sandvik, Phys. Rev. B 59 (1999) 14157(R).
- [9] O.F. Syljuåsen, A.W. Sandvik, Phys. Rev. E 66 (2002) 046701.
- [10] O.F. Syljuåsen, Phys. Rev. E 67 (2003) 046701.
- [11] M. Troyer, F. Alet, S. Trebst, S. Wessel, AIP Conf. Proc. 690, 2003, see cond-mat/0306128.
- [12] N. Kawashima, K. Harada, J. Phys. Soc. Jpn. 73 (2004) 1379.
- [13] S.M.A. Rombouts, K. Van Houcke, L. Pollet, Phys. Rev. Lett. 96 (2006) 180603.
- [14] K. Van Houcke, S.M.A. Rombouts, L. Pollet, Phys. Rev. E 73 (2006) 056703.
- [15] W.K. Hastings, Biometrika 57 (1970) 97.

- [16] L. Pollet, S.M.A. Rombouts, K. Van Houcke, K. Heyde, *Phys. Rev. E* 70 (2004) 056705.
- [17] L. Pollet, PhD thesis, Ghent University, Belgium, unpublished, see <http://inwpent4.ugent.be/PhDthesis.html>, 2005.
- [18] S. Wessel, F. Alet, S. Trebst, D. Leumann, M. Troyer, G.G. Batrouni, *J. Phys. Soc. Jpn.* 74 (Suppl.) (2005) 10.
- [19] F. Alet, S. Wessel, M. Troyer, *Phys. Rev. E* 71 (2005) 036706.
- [20] C. Pierleoni, D.M. Ceperely, *Chem. Phys. Chem.* 6 (9) (2005) 1872–1878.