

7 Monte Carlo Simulations of Quantum Spin Models

Stefan Wessel

Institute for Theoretical Solid State Physics

RWTH Aachen University

Contents

1	Introduction	2
2	World lines and local updates	2
2.1	Suzuki-Trotter decomposition	3
2.2	World lines	6
2.3	Local updates	7
2.4	The continuous-time limit	8
3	Stochastic series expansion	10
3.1	Configuration space	11
3.2	Observables	13
3.3	Local updates	15
4	The loop algorithm	16
5	Worms, operator loops, and directed loops	21
6	The sign problem	23
A	Quantum Heisenberg Model 101	26

1 Introduction

These lecture notes aim to provide a basic introduction to the central concepts behind modern quantum Monte Carlo (QMC) techniques for the simulation of quantum spin systems in thermal equilibrium. Acquaintance with the Monte Carlo approach to simulate classical statistical physics models is assumed, and can be obtained for example from the lecture of W. Krauth in this volume. In devising QMC simulation methods, two major steps need to be taken: (i) The quantum partition function must be reformulated in a way as to allow stochastic sampling over a space of effective configurations, such that each contributing configuration has a positive statistical weight. (ii) An efficient update-scheme needs to be devised in order to sample this configuration space through a Markov process. The first part of these lecture notes (Secs. 2 and 3) concerns the derivation of different representations of the quantum partition function, and thus relates to the first of these steps. Similarly as in the case of classical Monte Carlo schemes, a major breakthrough concerning the second step was the invention of efficient global update schemes, which will be introduced in the second part of these notes (Secs. 4 and 5). Finally, in section 6, we will discuss the QMC sign problem, which still poses the most severe restriction in applying the presented QMC methods to (frustrated) quantum spin systems. There exist already various excellent recent reviews on the methods that we discuss below, and we will mention some of these at the appropriate places. Thus, in these notes, we do not attempt to provide a comprehensive account on the subject, but instead we will highlight the key ideas behind these techniques.

2 World lines and local updates

We will discuss QMC methods for the simulation of (finite) quantum spin systems at finite temperatures. Upon tuning the simulation temperature sufficiently low (sufficiently below the finite-size spin gap), however, ground state properties of a finite system can usually be explored as well. As will be seen below, these QMC methods are feasible for non-frustrated spin models. In particular, this restricts for example the simulation of the antiferromagnetic Heisenberg model to bipartite lattices. But, there are no restrictions regarding e.g., the dimensionality of the system, and also very large systems with several thousands or even millions of spins can be treated, depending on the temperature range of interest; the computational effort of these methods scales linearly in both the system size and in the inverse temperature $\beta = 1/T$ (we fix $k_B = 1$ here). For the sake of clarity, we consider in the following the spin-1/2 Heisenberg model on an open chain with N_s lattice sites,

$$H = J \sum_{i=1}^{N_s-1} \vec{S}_i \cdot \vec{S}_{i+1} \quad (1)$$

as an example system when deriving the QMC algorithms. We refer to appendix A concerning the notation employed here. In the following, we are interested in calculating thermal expectation values, such as for example the internal energy $E = \langle H \rangle$ of the system at an inverse

temperature β . From statistical physics, we know that

$$\langle H \rangle = \frac{1}{Z} \text{Tr} (H e^{-\beta H}) = \frac{1}{Z} \sum_n E_n e^{-\beta E_n}, \quad (2)$$

where in the last step we expressed the trace in the basis of eigenstates of the Hamiltonian, and the partition function $Z = \text{Tr}(e^{-\beta H}) = \sum_n e^{-\beta E_n}$. While the above expressions in the basis of eigenstates of the Hamiltonian closely resemble the corresponding formulas for a classical system, the problem in the quantum case is that usually one does not know the full spectrum of the Hamiltonian, and thus the above sums cannot be directly calculated. Knowing the spectrum of H in the quantum case, one would often have essentially already solved the problem of interest.

QMC methods circumvent the full diagonalization of H by mapping the quantum partition function Z to a partition function of an effective classical model, and then performing a Monte Carlo sampling of the states contributing to the effective classical partition function. In order to see how such a mapping can be realized, we consider first the Suzuki-Trotter decomposition of the partition function, as pioneered by Suzuki in 1976 [1]. This approach will set the stage for more advanced QMC approaches.

2.1 Suzuki-Trotter decomposition

Our first task will be to map the quantum partition function to that of an effective classical model. In the following, we consider in particular the quantum Heisenberg model on an open chain, given by the Hamiltonian Eq. (1), which can be decomposed into contributions from each bond,

$$H = \sum_i H_i, \quad H_i = J \vec{S}_i \cdot \vec{S}_{i+1}, \quad (3)$$

where H_i is a bond Hamiltonian, that corresponds to the bond between site i and site $i + 1$ on the chain. We can furthermore separate H into two parts,

$$H = H_A + H_B, \quad (4)$$

where H_A (H_B) consists of only the even (odd) bond contributions, i.e.,

$$H_A = \sum_{i \text{ even}} H_i, \quad H_B = \sum_{i \text{ odd}} H_i. \quad (5)$$

While $[H_A, H_B] \neq 0$, the bond Hamiltonians within H_A or H_B commute among themselves, since $[H_i, H_j] = 0$ for i and j both even (or both odd). We have thus separated H into two parts, each of which consists of commuting terms. Now, we rewrite the statistical operator as a product of many terms, each with a small prefactor β/M in front of H ,

$$e^{-\beta H} = \left(e^{-\frac{\beta}{M} H} \right)^M = \left(e^{-\Delta\tau H} \right)^M, \quad (6)$$

with a (large) integer number (called the Trotter number) M , and $\Delta\tau = \beta/M$. The Suzuki-Trotter (or split-operator) approximation [1, 2] now consists in approximating the exponential

of H , expressed in terms of the two non-commuting pieces $H = H_A + H_B$, by a product of exponentials. While many such decompositions are possible, the most commonly used approximations (also employed beyond QMC methods) are

$$e^{-\Delta\tau H} = e^{-\Delta\tau(H_A+H_B)} = \begin{cases} e^{-\Delta\tau H_A} e^{-\Delta\tau H_B} & +O(\Delta\tau^2) \\ e^{-\Delta\tau H_B/2} e^{-\Delta\tau H_A} e^{-\Delta\tau H_B/2} & +O(\Delta\tau^3) \end{cases}, \quad (7)$$

where the errors are also proportional to the commutator $[H_A, H_B]$. Using either of the two approximations in Eq. (6), we obtain

$$\begin{aligned} Z &= \text{Tr} \left(e^{-\beta H} \right) \\ &= \text{Tr} \left(\underbrace{e^{-\Delta\tau H_A} e^{-\Delta\tau H_B}}_M \underbrace{e^{-\Delta\tau H_A} e^{-\Delta\tau H_B}}_{M-1} \dots \underbrace{e^{-\Delta\tau H_A} e^{-\Delta\tau H_B}}_1 \right) + O(\Delta\tau^2). \end{aligned} \quad (8)$$

When using the second-order approximation from Eq. (7), the final expression is obvious. However, when using the third-order approximation, one actually obtains the same final expression, due to the cyclic invariance of the trace, which can be used to move the most left exponential to the very right, and then coalesce every other two consecutive terms. This also shows, that the systematic error in Z actually scales as $M\Delta\tau^3 \propto \Delta\tau^2$, even when using the second-order approximation, where one might have expected an $O(\Delta\tau)$ error in the final expression. Now, consider a basis of the Hilbert space, e.g., in terms of the local eigenstates of S_i^z , which we write

$$\begin{aligned} |\sigma_1\rangle &= |\uparrow\uparrow \dots \uparrow\uparrow\rangle \\ |\sigma_2\rangle &= |\uparrow\uparrow \dots \uparrow\downarrow\rangle \\ &\vdots \\ |\sigma_{2N}\rangle &= |\downarrow\downarrow \dots \downarrow\downarrow\rangle. \end{aligned} \quad (9)$$

Since this set forms a basis, we obtain a completeness relation:

$$\sum_{\sigma} |\sigma\rangle\langle\sigma| = 1, \quad (10)$$

where $|\sigma\rangle\langle\sigma|$ is a projection operator onto the basis state $|\sigma\rangle$. Within this basis of the Hilbert space, we can thus express the partition function as

$$\begin{aligned} Z &= \sum_{\sigma^0} \langle\sigma^0| e^{-\beta H} |\sigma^0\rangle \\ &\approx \sum_{\sigma^0} \langle\sigma^0| \underbrace{e^{-\Delta\tau H_A} e^{-\Delta\tau H_B}}_M \underbrace{e^{-\Delta\tau H_A} e^{-\Delta\tau H_B}}_{M-1} \dots \underbrace{e^{-\Delta\tau H_A} e^{-\Delta\tau H_B}}_1 |\sigma^0\rangle \\ &= \sum_{\sigma^0} \langle\sigma^0| e^{-\Delta\tau H_A} e^{-\Delta\tau H_B} \dots e^{-\Delta\tau H_A} \left(\sum_{\sigma^1} |\sigma^1\rangle\langle\sigma^1| \right) e^{-\Delta\tau H_B} |\sigma^0\rangle \\ &= \sum_{\sigma^0} \sum_{\sigma^1} \langle\sigma^0| e^{-\Delta\tau H_A} e^{-\Delta\tau H_B} \dots \left(\sum_{\sigma^2} |\sigma^2\rangle\langle\sigma^2| \right) e^{-\Delta\tau H_A} |\sigma^1\rangle\langle\sigma^1| e^{-\Delta\tau H_B} |\sigma^0\rangle \\ &= \sum_{\sigma^0, \sigma^1, \sigma^2} \langle\sigma^0| e^{-\Delta\tau H_A} e^{-\Delta\tau H_B} \dots \left(\sum_{\sigma^3} |\sigma^3\rangle\langle\sigma^3| \right) e^{-\Delta\tau H_B} |\sigma^2\rangle\langle\sigma^2| e^{-\Delta\tau H_A} |\sigma^1\rangle\langle\sigma^1| e^{-\Delta\tau H_B} |\sigma^0\rangle. \end{aligned}$$

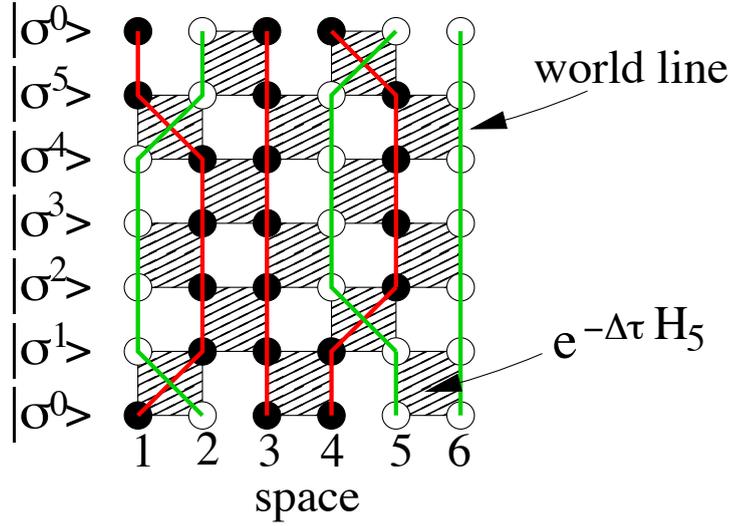


Fig. 1: A configuration contributing to the checkerboard decomposition of the partition function of an open spin-1/2 Heisenberg chain with 6 lattice sites, for a Trotter number $M = 3$. Black (white) circles denote spin up (down), and the world lines of the up (down) spins are denoted by red (green) lines. Each dashed square represents the exponential of a bond Hamiltonian.

In the above summations, we added a (superscript) label to distinguish the basis sets that arise from the trace ($|\sigma^0\rangle$) and from the various inserted partitions of unity ($|\sigma^1\rangle$, $|\sigma^2\rangle$, and $|\sigma^3\rangle$). Continuing this way, we eventually arrive at the following expression:

$$\begin{aligned}
 Z = \sum_{\{\sigma^i\}} & \underbrace{\langle \sigma^0 | e^{-\Delta\tau H_A} | \sigma^{2M-1} \rangle}_{2M} \underbrace{\langle \sigma^{2M-1} | e^{-\Delta\tau H_B} | \sigma^{2M-2} \rangle}_{2M-1} \underbrace{\langle \sigma^{2M-2} | e^{-\Delta\tau H_A} | \sigma^{2M-3} \rangle}_{2M-2} \dots \\
 & \dots \underbrace{\langle \sigma^3 | e^{-\Delta\tau H_B} | \sigma^2 \rangle}_3 \underbrace{\langle \sigma^2 | e^{-\Delta\tau H_A} | \sigma^1 \rangle}_2 \underbrace{\langle \sigma^1 | e^{-\Delta\tau H_B} | \sigma^0 \rangle}_1 + \mathcal{O}(\Delta\tau^2).
 \end{aligned} \tag{11}$$

This representation of Z is referred to as the ‘‘Suzuki-Trotter decomposition’’. Furthermore, since H_A and H_B each consists of commuting parts, we find that

$$e^{-\Delta\tau H_A} = \prod_{i \text{ even}} e^{-\Delta\tau H_i} = e^{-\Delta\tau H_2} e^{-\Delta\tau H_4} \dots \tag{12}$$

$$e^{-\Delta\tau H_B} = \prod_{i \text{ odd}} e^{-\Delta\tau H_i} = e^{-\Delta\tau H_1} e^{-\Delta\tau H_3} \dots \tag{13}$$

and thus in the Suzuki-Trotter decomposition each exponential expression factorizes into exponentials for the bond Hamiltonians on either all even or all odd bonds. One can represent a given contribution to Z graphically, as is shown for a specific example with $N_s = 6$, and $M = 3$ in Fig. 1. This two-dimensional picture appears like a space-time picture of spins propagating in discrete steps from the initial configuration $|\sigma^0\rangle$ to $|\sigma^1\rangle, \dots$, and finally from $|\sigma^{2M-1}\rangle$ back to $|\sigma^0\rangle$, since the configurations on the first and the last step are equal. This pictorial illustration also explains why the underlying partitioning of the Hamiltonian is often referred to as the ‘‘checkerboard decomposition’’.

2.2 World lines

Each exponential $e^{-\Delta\tau H_i}$ couples two spin sites on a shaded plaquette, and the statistical weight of a configuration is the product of all the matrix elements of the exponentials on all shaded plaquettes. Note, that $e^{-\Delta\tau H_i}$ appears as an imaginary-time evolution operator, that propagates the two spins at sites i and $i+1$ from one “time slice” to the next. To appreciate this suggestive picture of a discrete time propagation, we introduced the corresponding “Trotter time-step” $\Delta\tau = \beta/M$ above. However, not all of the $2^4 = 16$ possible spin configurations along a shaded plaquette are allowed to appear within the checkerboard decomposition. In order to assess the allowed configurations and the corresponding matrix elements, we explicitly calculate the matrix elements of the operators

$$e^{-\Delta\tau H_i} = \sum_{k=0}^{\infty} \frac{1}{k!} (\Delta\tau)^k (-H_i)^k. \quad (14)$$

In the local two-sites basis of the two spins at site i and $i+1$, the Hamiltonian matrix is

$$H_i = J \begin{pmatrix} \frac{1}{4} & 0 & 0 & 0 \\ 0 & -\frac{1}{4} & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & -\frac{1}{4} & 0 \\ 0 & 0 & 0 & \frac{1}{4} \end{pmatrix}, \quad \text{in the local basis} \begin{matrix} |\uparrow\uparrow\rangle \\ |\uparrow\downarrow\rangle \\ |\downarrow\uparrow\rangle \\ |\downarrow\downarrow\rangle \end{matrix}. \quad (15)$$

Upon performing the Taylor expansion, and grouping back the resulting terms, one finds (this would make up a nice exercise)

$$e^{-\Delta\tau H_i} = e^{+\Delta\tau J/4} \begin{pmatrix} e^{-\frac{\Delta\tau J}{2}} & 0 & 0 & 0 \\ 0 & \cosh(\frac{\Delta\tau J}{2}) & -\sinh(\frac{\Delta\tau J}{2}) & 0 \\ 0 & -\sinh(\frac{\Delta\tau J}{2}) & \cosh(\frac{\Delta\tau J}{2}) & 0 \\ 0 & 0 & 0 & e^{-\frac{\Delta\tau J}{2}} \end{pmatrix}. \quad (16)$$

There are thus 6 possible allowed plaquettes with a finite weight. They are shown in Fig. 2. All other combinations would lead to a vanishing matrix element and are thus not allowed to occur as part of any allowed configuration. The allowed configurations are those that exhibit the same total local magnetization $S_i^z + S_{i+1}^z$ on the lower and the upper edge of each shaded plaquette. This property derives from the fact that each bond Hamiltonian H_i conserves the total magnetization of the two spins connected by the bond.

When we connect the positions of the up and down spins as they propagate through the shaded plaquettes, we obtain continuous lines. These are the “world lines” of the spins, and each such world line denotes the evolution of one spin up or down from $|\sigma^0\rangle$ back to $|\sigma^0\rangle$. The world lines for both the spin-up and the spin-down case are illustrated in Fig. 1. Note, that it would be sufficient to only show, say, the up spins, which is what will be done later on in these notes. If one would glue together the upper and lower boundary of the space-time configuration to make the periodicity constraint by the trace more explicit, these world lines are thus continuous, and are not broken anywhere. This set of unbroken world-line configurations defines a classical



Fig. 2: Allowed shaded plaquette configurations for the spin-1/2 Heisenberg model. Black (white) circles denote spin up (down), and the world lines of the up (down) spins are denoted by red (green) lines.

statistical model, which has the same partition function as the quantum partition function, if the weight of a given world-line configuration C equals the product of the matrix elements from all shaded plaquettes P

$$W(C) = \prod_P W_P(C|_P), \quad (17)$$

where $W_P(C|_P)$ is the corresponding matrix element of $e^{-\Delta\tau H_i}$ for plaquette P in the configuration C . At first sight, the effective model might appear to be just the two-dimensional Ising model on a square lattice of size $N_s \times 2M$, with periodic boundary conditions in the y -direction. However, the effective classical model is in fact more complex than the Ising model, which has an unconstrained configuration space. Only a subset of configurations of the two-dimensional Ising model are allowed to occur also in the effective classical model for Z . These configurations are those that correspond to continuous, unbroken world lines. In higher spatial dimensions, one can use the Suzuki-Trotter approach with decompositions very similar to the one discussed above. One thus finds that the quantum partition function of a d -dimensional system (here $d = 1$) corresponds to that of an effective classical model in $(d + 1)$ dimensions, with the $(d + 1)$ -th direction corresponding to an imaginary-time evolution of the original quantum model. This well known quantum-to-classical-mapping holds in fact much more generally, and exposes a deep connection between classical and quantum statistical physics.

Note however the minus sign in front of the off-diagonal matrix elements: for the allowed world-line configurations, all the accumulated signs actually cancel out, due to the periodicity constraint in imaginary time, in case of the open chain or for a closed chain with an even number of sites. This holds also true for higher-dimensional generalizations of the world-line approach whenever the underlying lattice structure is bipartite. However, for frustrated systems, e.g., for a closed 3-site chain, there appears a “sign problem”, since both configuration with positive and negative weights are allowed. We will discuss the consequences of this issue in Sec. 6.

2.3 Local updates

The world-line representation introduced above can be taken as a starting point to set up a QMC algorithm [3]. But, how does one sample configurations in the effective world-line model in a Monte Carlo algorithm, i.e., how do we generate new valid world-line configurations from a given one? Since the spin S^z -conservation prohibits the breaking of world lines, the updates need to move world-line segments instead of just changing local spin states like in the classical Ising model. In the early days, so-called local updates were performed, i.e., local manipulations

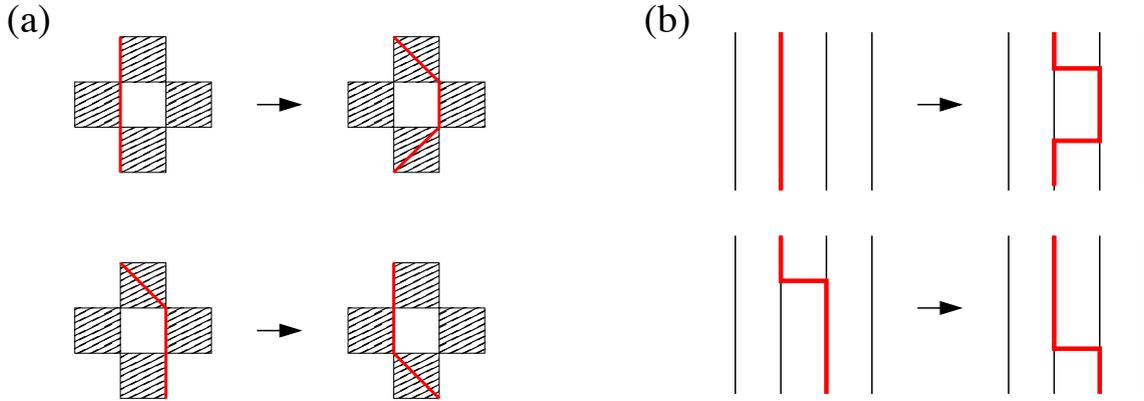


Fig. 3: Two different local updates of world-line configurations in discrete time (a) and in the continuous-time limit (b). For clarity, only the spin-up world lines are shown explicitly.

of the world-line configuration, such as those shown in Fig. 3 (a). These local updates are quite simple and either generate a pair of kinks across a white plaquette within a straight section of the world line, or they move a kink to a neighboring plaquette [3, 4], as shown in Fig. 3 (a). Slightly more complicated local moves are required for higher-dimensional spin systems [5], or for other quantum lattice models [6]. In any case, such local updates are accepted by considering the related change in the statistical weight according to Eq. (17), using e.g., the Metropolis [7] acceptance scheme.

However, such local updates are not efficient in two respects: First, global properties of the configurations cannot be changed by such local updates; the spatial winding number of the world lines and their total number, i.e. the total magnetization of the system, cannot be changed using local moves. Thus, they had to be complemented by special global updates as well [5]. Second, such updates lead to severe critical slowing down upon approaching critical regions, similarly as those familiar from the local Metropolis algorithm for the Ising model. A major breakthrough in overcoming this problem was the work by Evertz, Lana and Marcu [8], in which they presented an extension of the Swendsen-Wang cluster update idea known from classical Monte Carlo studies to world-line QMC methods [56]. We will present this “loop algorithm” in Sec. 4. Before doing so, we want to introduce two other QMC representations, which are relevant for modern QMC algorithms.

2.4 The continuous-time limit

Due to the finite number of time slices, M , the approach described above suffers from a systematic error – the discretization or Trotter error. It has been shown that in most cases, one can keep $\Delta\tau$ independent of N_s and β in order to ensure a constant error level [9] (for some observables, however, care has to be taken to avoid divergent errors in the zero temperature limit $\beta \rightarrow \infty$ [9, 10]). The Trotter error was controlled originally by extrapolation to the continuous-time limit $\Delta\tau \rightarrow 0$, from simulations with different values of the time step $\Delta\tau$. It was realized later [11] that the continuous-time limit can be taken already in the construction of

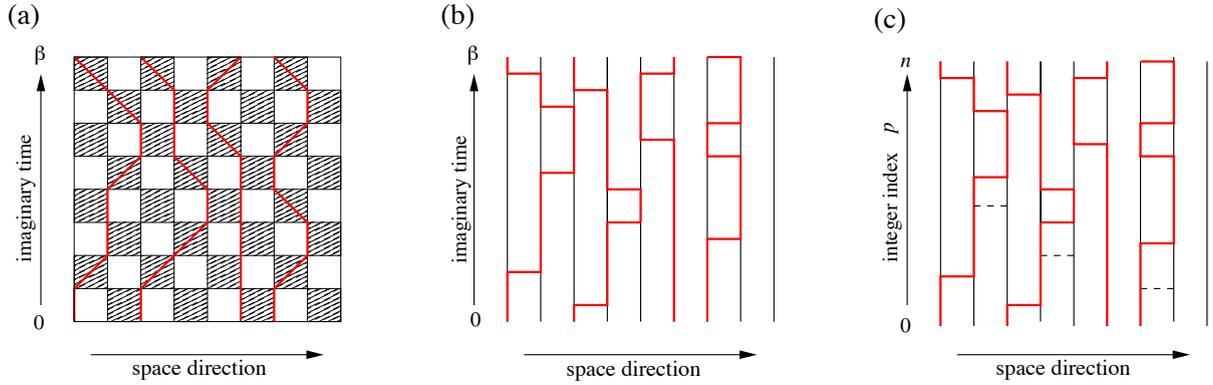


Fig. 4: Comparison of world-line QMC configurations in a) discrete time, b) continuous time, and c) in SSE representation. In the latter, the continuous-time index is replaced by an ordered integer index of the operators, and additional diagonal terms are indicated by dashed lines.

the algorithm, so that simulations can be performed directly in the limit $\Delta\tau \rightarrow 0$, i.e., $M \rightarrow \infty$, without the need to perform any final extrapolations. This appears feasible, once one realizes that a given world-line configuration can be represented by keeping only a list of times at which the configuration *changes*, instead of storing the configuration at each of the $2M$ time slices. Indeed, the probability for a jump of a world line (i.e., a kink in the world line) from one lattice site to a neighboring site across a given plaquette, and thus a change in the local configuration is proportional to $\sinh(\Delta\tau J/2) \propto \Delta\tau \propto 1/M$. Hence, the mean total number of such kinks remains finite in the limit $M \rightarrow \infty$. The continuous world-line representation, i.e., the limit $\Delta\tau \rightarrow 0$, is thus well defined, and the relevant configurational information can be efficiently represented on a computer. This is illustrated in Fig. 4 (a) and (b), by comparing the discrete and continuous-time world-line approach.

In addition to the configurational information, the local updates need to be revisited in the continuous-time limit as well. In particular, the probability P_{pk} for the insertion of a pair of kinks in the world line (the upper move in Fig. 3 (a)), vanishes in the continuous-time limit as

$$P_{pk} = \sinh^2(\Delta\tau J/2) / \cosh^2(\Delta\tau J/2) \propto \Delta\tau^2 \propto 1/M^2 \rightarrow 0. \quad (18)$$

To circumvent such a vanishing probability, one now proposes to insert a pair of jumps not at specific locations, but *anywhere* inside a finite time interval [11], as illustrated in Fig. 3 (b). The integrated probability for such a move then remains finite in the continuous-time limit. Similarly, instead of shifting a kink by $\Delta\tau$ (the lower move in Fig. 3 (a)), one now moves it within a finite time interval in the continuous-time algorithm. In addition to local updates, also the loop algorithm has been shown to allow for an efficient realization in the continuous-time limit [12], as will be discussed below.

Finally, we remark that the continuous-time limit of the Suzuki-Trotter formula, Eq. (11), in the above interpretation is in fact equivalent to a time-dependent perturbation theory in imaginary

time,

$$\begin{aligned} Z &= \text{Tr} \exp(-\beta H) = \text{Tr} \left[\exp(-\beta H_0) \mathcal{T} \exp \int_0^\beta d\tau V(\tau) \right] \\ &= \text{Tr} \left[\exp(-\beta H_0) \left(1 - \int_0^\beta d\tau V(\tau) + \frac{1}{2} \int_0^\beta d\tau_1 \int_{\tau_1}^\beta d\tau_2 V(\tau_1) V(\tau_2) + \dots \right) \right], \end{aligned} \quad (19)$$

where the symbol \mathcal{T} denotes the (imaginary-) time ordering of the exponential [11, 13]. In this representation, the Hamiltonian $H = H_0 + V$ is split up into a diagonal term H_0 and an off-diagonal perturbation V , which in the interaction picture is $V(\tau) = \exp(\tau H_0) V \exp(-\tau H_0)$. For the Heisenberg model, the diagonal term H_0 is given by the longitudinal $S_i^z S_j^z$ spin-spin interactions, while the off-diagonal perturbation V relates to the transverse, spin exchange terms, $\frac{1}{2}(S_i^+ S_j^- + S_i^- S_j^+)$ in H . In more detail, we express the Heisenberg Hamiltonian as

$$H = H_0 + V = \sum_{\langle i,j \rangle} J S_i^z S_j^z + \sum_{\langle i,j \rangle} \frac{J}{2} (S_i^+ S_j^- + S_i^- S_j^+). \quad (20)$$

The number of kinks in a given continuous-time world-line configuration is then equal to the expansion order of a specific term in the perturbation expansion, wherein each kink's space-time position is set by a specific spin exchange term from V between two lattice sites i and j , as well as a corresponding imaginary time τ between 0 and β . Furthermore, the exponential factors $\exp(\tau H_0)$ in $V(\tau)$ describe the vertical, unperturbed evolution of the spins along the world lines between the kinks, thus establishing explicitly the above-mentioned equivalence between these two formulations of the continuous-time limit.

3 Stochastic series expansion

Already before Suzuki's approach to QMC in the mid 1970s [1, 3], an alternative approach to QMC simulations had been put forward by Handscomb in the early 1960s [14, 15] for the specific case of the ferromagnetic Heisenberg model. It is based on a Taylor expansion of the statistical operator inside the partition function:

$$Z = \text{Tr} \exp(-\beta H) = \sum_{n=0}^{\infty} \frac{\beta^n}{n!} \text{Tr} ((-H)^n). \quad (21)$$

In Handscomb's approach and within later extensions to other models [16–18], the traces of H^n were evaluated employing projection operator expressions. Later, in the early 1990s, the power-series approach to QMC simulations was revisited by Sandvik and Kurkijärvi within the more generally applicable stochastic series expansion (SSE) formulation [19, 20], in which these traces are also sampled stochastically. A recent review of the SSE method, combined with a general introduction to computational methods for quantum spin systems, can be found in Ref. [21]. It contains also many basic details for implementing the SSE QMC method.

3.1 Configuration space

To formulate the SSE method, it proves convenient to first express the Heisenberg Hamiltonian as a sum of bond operators, which are either diagonal or off-diagonal in the standard S^z -basis,

$$H = - \sum_{b=1}^{N_b} (H_{1,b} - H_{2,b}) + JN_b/4, \quad (22)$$

with

$$H_{1,b} = J \left(\frac{1}{4} - S_{i(b)}^z S_{j(b)}^z \right), \quad (23)$$

$$H_{2,b} = \frac{J}{2} \left(S_{i(b)}^+ S_{j(b)}^- + S_{i(b)}^- S_{j(b)}^+ \right). \quad (24)$$

Here, N_b equals the total number of bonds in the system, and $i(b)$ and $j(b)$ denote the two lattice sites that are connected by the bond with bond index b . The explicit introduction of the minus sign in H is convenient, while the constant in the diagonal operators will make the series expansion positive definite, as we will see shortly. Note that this constant is irrelevant for the physics of the system, but has to be accounted for when calculating the system's energy, as we will see in the next section. We now insert the above form of the Hamiltonian into the expression Eq. (21) for the partition function, thereby obtaining

$$Z = \sum_{\sigma} \sum_{n=0}^{\infty} \frac{\beta^n}{n!} \sum_{\mathcal{S}_n} (-1)^{n_2} \langle \sigma | H_{t(n),b(n)} \cdots H_{t(2),b(2)} H_{t(1),b(1)} | \sigma \rangle. \quad (25)$$

Here we used the local S^z -basis to express the trace, and \mathcal{S}_n denotes products (strings) of the bond operators $H_{1,b}$ or $H_{2,b}$ that originate from expanding H^n in these bond operator terms using Eq. (22), namely, each such operator string \mathcal{S}_n is an sequence of n bond operators, specified by the type labels $t(p) \in \{1, 2\}$ (i.e. diagonal or off-diagonal) and $b(p) \in \{1, \dots, N_b\}$, so that we can also write such a sequence as

$$\mathcal{S}_n = [t(1), b(1)], [t(2), b(2)], \dots, [t(n), b(n)]. \quad (26)$$

The summation in Eq. (25) extends over all expansion orders n , and for each expansion order, over all operator sequences \mathcal{S}_n of length n , containing n bond operators. Finally, n_2 denotes the number of off-diagonal operators in the sequence \mathcal{S}_n , i.e., the number of sequence elements with $t(p) = 2$. It is important to note that, in the local S^z -basis, the bond operators exhibit a non-branching property, i.e., if one applies a bond operator to one of these basis states, the resulting state is either proportional to the same basis state, or to another basis state, or it vanishes. But in no case it is a superposition of two or more such basis states generated when traversing the action of the sequence of bond operators starting from any initial basis state $|\sigma\rangle$.

The propagated basis state that appears after the action of the first p operators, i.e., at propagation level p , will be denoted (after proper normalization) by $|\sigma(p)\rangle$, i.e.,

$$|\sigma(p)\rangle \propto \prod_{q=1}^p H_{t(q),b(q)} |\sigma\rangle. \quad (27)$$

Note that for any operator-state configuration $(\mathcal{S}_n, |\sigma\rangle)$ that contributes to Z , the final state, resulting after the action of all the n operators has to fulfill $\langle\sigma|\sigma(n)\rangle \neq 0$, thus we obtain the following periodicity constraint: $|\sigma(n)\rangle = |\sigma(0)\rangle = |\sigma\rangle$.

If one examines in more detail the action of the different bond operators $H_{1,b}$ and $H_{2,b}$ on the two spins at the related lattice sites $i(b)$ and $j(b)$ of a basis state $|\sigma\rangle$, one finds, that (i) the state is destroyed if both spins are parallel,

$$H_{1,b}|\uparrow_{i(b)}\uparrow_{j(b)}\rangle = 0, \quad H_{1,b}|\downarrow_{i(b)}\downarrow_{j(b)}\rangle = 0, \quad (28)$$

$$H_{2,b}|\uparrow_{i(b)}\uparrow_{j(b)}\rangle = 0, \quad H_{2,b}|\downarrow_{i(b)}\downarrow_{j(b)}\rangle = 0, \quad (29)$$

and (ii) if the two spins are anti-parallel, the spin state is preserved by $H_{1,b}$, while being flipped by $H_{2,b}$, in both cases with a matrix element of $J/2$,

$$\langle\uparrow_{i(b)}\downarrow_{j(b)}|H_{1,b}|\uparrow_{i(b)}\downarrow_{i(b)}\rangle = J/2, \quad \langle\downarrow_{i(b)}\uparrow_{j(b)}|H_{1,b}|\downarrow_{i(b)}\uparrow_{i(b)}\rangle = J/2, \quad (30)$$

$$\langle\downarrow_{i(b)}\uparrow_{j(b)}|H_{2,b}|\uparrow_{i(b)}\downarrow_{i(b)}\rangle = J/2, \quad \langle\uparrow_{i(b)}\downarrow_{j(b)}|H_{2,b}|\downarrow_{i(b)}\uparrow_{i(b)}\rangle = J/2. \quad (31)$$

From (i) it follows that for any operator-state configuration $(\mathcal{S}_n, |\sigma\rangle)$ that contributes to Z , the bond operators $H_{t(p),b(p)}$ act only on propagated states with anti-parallel spins on bond $b(p)$. Property (ii) implies that all the matrix elements that arise from the bond operators within the operator string are equal and positive. The effective configuration space of the SSE method thus essentially consists of all allowed operator-state configurations, i.e., those with non-zero weight, with the weight of a given configuration $C = (\mathcal{S}_n, |\sigma\rangle)$ obtained from Eq. (25),

$$Z = \sum_C W(C), \quad W(C) = \frac{\beta^n}{n!} \langle\sigma| \prod_{p=1}^n H_{t(p),b(p)} |\sigma\rangle. \quad (32)$$

Here, we also anticipated that on a bipartite lattice the number of off-diagonal operators in an allowed operator sequence must be even, such that $(-1)^{n_2} = +1$, due to the periodicity constraint $|\sigma(n)\rangle = |\sigma(0)\rangle = |\sigma\rangle$. For a non-bipartite (frustrated) lattice this would not be the case, and a QMC sign-problem would result; cf. Sec. 6 for a discussion of this case. The SSE QMC configurations can be visualized very similarly as in the previous case, cf. Fig. 4 for a comparison between the different formulations. Note that in the SSE formulation a discrete time-like index p is introduced. Furthermore, in addition to the jump events that appear in the continuous-time world-line formulation, the SSE configurations exhibit the additional presence of diagonal bond-operators. Indeed, the SSE representation can be formally related to the continuous-time world-line representation by observing that Eq. (25) is obtained from Eq. (19) upon setting $H_0 = 0$, $V = H$ and integrating over all times (compare also Fig. 4) [22]. This mapping exposes the respective advantages and disadvantages of the two representations: The SSE representation corresponds to a perturbation expansion in all terms of the Hamiltonian, whereas continuous-time world-line algorithms treat the diagonal terms exactly and perturb only in the off-diagonal terms of the Hamiltonian. The continuous-time algorithm hence needs fewer terms in the expansion, but we pay for it by having to handle continuous imaginary-time variables.

For spin systems in which the contributions to the total energy arising from diagonal and off-diagonal terms are well balanced, the SSE is typically preferable in practice, due to its discrete representation of the imaginary-time continuum. It is important to note, that this discretization does *not* introduce any Trotter-error as in the discrete-time world-line QMC; instead, the SSE formulation should be considered an essentially exact QMC technique.

It is possible to formulate Monte Carlo sampling algorithms for the SSE configuration space introduced above, i.e., containing operator strings of fluctuating length [14, 20, 23]. However, it is also possible to work with operator strings of a fixed length, which is in most cases more convenient computationally [19, 20]. Indeed, as will be shown in the next section, the mean length of the operator string $\langle n \rangle$ is essentially related to the energy of the system and scales linear with the system size N_s and β . Furthermore, the specific heat is given by $C = \langle n^2 \rangle - \langle n \rangle^2 - \langle n \rangle$, so that in particular for $T \rightarrow 0$, where C vanishes, the variance of the distribution of expansion orders is seen to be equal to $\langle n \rangle$. Hence, for a given system size and temperature, the series expansion order n of the operator-state configurations $C = (\mathcal{S}_n, |\sigma\rangle)$ can in practice be truncated at a sufficiently large cutoff Λ , which is typically determined during the thermalization phase of an SSE simulation, without introducing any detectable error. It is then feasible to work with fixed-length operator strings upon augmenting all operator sequences of length $n < \Lambda$ by $\Lambda - n$ unit operators I , denoted in the following by $H_{0,0} = I$. Allowing for all possible positions of the unit operators in the original operator strings, one is then lead to the following representation of the partition function:

$$Z = \sum_C W_\Lambda(C), \quad W(C) = \frac{\beta^n (\Lambda - n)!}{\Lambda!} \langle \sigma | \prod_{p=1}^{\Lambda} H_{t(p), b(p)} | \sigma \rangle, \quad (33)$$

where each operator-state configuration $(\mathcal{S}_\Lambda, |\sigma\rangle)$ now consists of a sequence of operators specified by the type labels $t(p) \in \{0, 1, 2\}$ and $b(p) \in \{1, \dots, N_b\}$ in case $t(p) \in \{1, 2\}$, or $b(p) = 0$, if $t(p) = 0$, which would represent a unit operator. In the above formula, n denotes the number of non-unit operators in the operator sequence, i.e., the number of true bond operators from H . The new combinatorial factor in $W_\Lambda(C)$ accounts for the $\binom{\Lambda}{n}$ equivalent terms that are generated this way as compared to the expression without unit operators in Eq. (25). Again, we left out the factor $(-1)^{n_2} = +1$ on bipartite lattices, where all allowed configurations thus contribute with a positive weight $W_\Lambda(C) > 0$ to the partition function, which is hence feasible for Monte Carlo sampling.

3.2 Observables

Thus far in this exposition of QMC methods, we focused on different representations of the quantum partition function in terms of an effective set of configurations that prove amenable to Monte Carlo sampling. However, we did not mention yet how observables can be calculated. Here, we provide an overview how typical observables are measured within the SSE framework [19, 20, 24]. In many cases, the generalization from the formulas provided below to the other world-line methods is straightforward. A remarkably simple formula holds for the total

inner energy $E = \langle H \rangle$, which can be expressed in terms of the mean expansion order:

$$\langle H \rangle = \frac{1}{Z} \text{Tr}[H e^{-\beta H}] = \frac{1}{Z} \sum_{n=0}^{\infty} \frac{\beta^n}{n!} \text{Tr}[(-H)^{n+1}] = -\frac{1}{Z} \sum_{n=1}^{\infty} \frac{\beta^n}{n!} \frac{n}{\beta} \text{Tr}[(-H)^n] = -\frac{\langle n \rangle}{\beta}. \quad (34)$$

Accounting for the constant shift that we added to the Hamiltonian in Eq. (22) to perform the SSE expansion, we thus obtain the estimator for the inner energy before the shift as

$$E = -\frac{\langle n \rangle}{\beta} + JN_b/4. \quad (35)$$

By the same procedure, that leads to Eq. (34), one finds that $\langle H^2 \rangle = \langle n(n-1) \rangle / \beta^2$, from which the specific heat estimator

$$C = \beta^2 (\langle H^2 \rangle - \langle H \rangle^2) = \langle n^2 \rangle - \langle n \rangle^2 - \langle n \rangle \quad (36)$$

is readily obtained, which we already advocated above. Note, that the constant introduced in Eq. (22) cancels out in the estimator for C , as would be expected. In practice, the statistical error of C becomes large at low temperatures (where C itself becomes small), since it is given by the difference of two large numbers ($\sim (N_s \beta)^2$). We can also readily obtain the expectation value of any operator A that is diagonal in the local S^z -basis, such as for example a spin-spin correlation function $\langle S_i^z S_j^z \rangle$, in which case $A = S_i^z S_j^z$, by using averages over the propagated states:

$$\langle A \rangle = \left\langle \frac{1}{n} \sum_{p=0}^{n-1} A(p) \right\rangle, \quad (37)$$

where $A(p) = \langle \sigma(p) | A | \sigma(p) \rangle$ has been introduced for convenience. From the spin-spin correlation function, the equal-time structure factor is then obtained by Fourier transformation.

Finally, let us mention, how one can measure generalized susceptibilities related to two diagonal operators A, B within the SSE approach. Consider that we add to the Hamiltonian H a linear coupling term $h_B B$, i.e., $H \rightarrow H - h_B B$, then one obtains for the linear response function

$$\chi_{AB} = \left. \frac{\partial \langle A(h_B) \rangle}{\partial h_B} \right|_{h_B=0} \quad (38)$$

the Kubo formula

$$\chi_{AB} = \int_0^\beta d\tau \left(\langle A(\tau) B(0) \rangle - \langle A \rangle \langle B \rangle \right), \quad (39)$$

with $A(\tau) = e^{-\tau H} A e^{\tau H}$ the imaginary time-evolved operator. In case A and B are both diagonal in the local S^z -basis, this integral can be evaluated using

$$\int_0^\beta d\tau \langle A(\tau) B(0) \rangle = \left\langle \frac{\beta}{n(n+1)} \left(\sum_{p=0}^{n-1} A(p) \right) \left(\sum_{p=0}^{n-1} B(p) \right) + \frac{\beta}{(n+1)^2} \left(\sum_{p=0}^n A(p) B(p) \right) \right\rangle. \quad (40)$$

A special case concerns the uniform magnetic susceptibility, for which $A = B = \frac{1}{N_s} \sum_{i=1}^{N_s} S_i^z$, and thus we obtain the simple expression

$$\chi_u = \frac{\beta}{N_s} \langle M_z^2 \rangle, \quad M_z = \sum_{i=1}^{N_s} S_i^z. \quad (41)$$

For further observables, such as the spin stiffness, we refer to Sandvik's review [21]. Let us finally note, that within the world-line approach and the SSE representation it is also possible to evaluate the expectation values of imaginary-time-dependent correlation functions [19, 24]. For an efficient evaluation of the corresponding SSE expressions, based on the explicit mapping between the SSE configuration space and the continuous-time representation, we refer to the review by Assaad and Evertz [25].

3.3 Local updates

In order to sample the SSE configuration space, we need to generate operator-state configurations $(\mathcal{S}_A, |\sigma\rangle)$ according to the appropriate statistical weights. Updates of the operator sequence can in general not be carried out without affecting also the spin configuration. For example, if one changes the operator at a given propagation level p , say from a diagonal operator $H_{1,b}$ on the bond b to the off-diagonal operator $H_{2,b}$ on the same bond, then the spin configuration $|\sigma(p)\rangle$ has to change as well. This change in the world-line configuration must be healed at some other propagation level p' , by also exchanging diagonal and off-diagonal operators, so that the resulting state $|\sigma(p')\rangle$ is the same as before. Hence, one needs to perform Monte Carlo updates by attempting to change the bond operators at two appropriately picked propagation levels. Indeed, in most cases update attempts with randomly picked positions would not be possible, such that one has to specifically search for operator pairs (or even more operators), which can be updated this way. In the end, such a procedures correspond to the local update schemes that we introduced above for the world-line QMC approach, and they suffer from similar problems.

However, more efficient update schemes have been developed [26, 27]. Such updates in fact consist of two sub-steps: (i) the diagonal update, wherein the expansion order n is modified while keeping fixed the spin configuration $|\sigma\rangle$ of $(\mathcal{S}_A, |\sigma\rangle)$ as well as the off-diagonal operator content of the operator sequence \mathcal{S}_A , and (ii) a non-diagonal update, the operator-loop update, that updates the operator content as well as the spin configuration simultaneously – similar to, but in a much more efficient way than, the two-operator case that we discussed above.

We will outline the operator-loop update in the next section, and concentrate here on the diagonal update step, that is an essential local update procedure, which we formulate here for the fixed operator-string length representation [26]. Within this diagonal update, the whole operator sequence is sequentially traversed, and attempts are made to exchange diagonal operators and identity operators, so that in each such step, the expansion order, i.e., the number of non-identity operators, n can change by one to $n \pm 1$. While moving through the operator sequence, the spin configuration $|\sigma\rangle$ is updated whenever an off-diagonal operator is encountered, such that the propagated spin configurations $|\sigma(p)\rangle$ are readily available, as they will be required within this

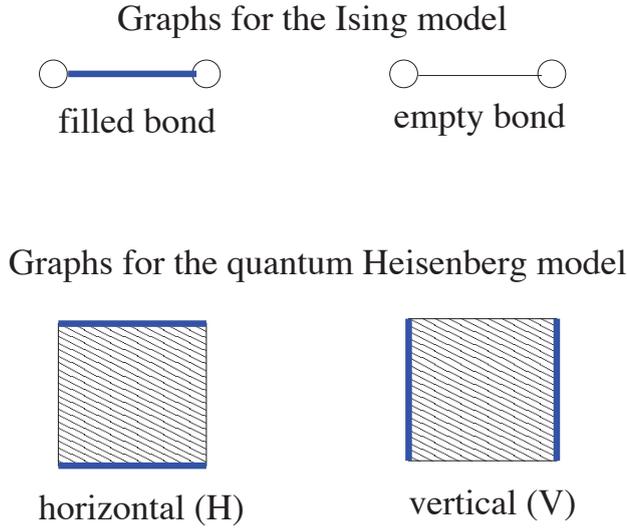


Fig. 5: Cluster components on a bond vs. graph components on a plaquette (breakups).

update scheme. In more detail, if at propagation level $p = 1, \dots, \Lambda$, an identity-operator is present, i.e., $t(p) = 0$, we attempt to replace it by a diagonal bond-operator on a bond b that is randomly chosen among all the N_b lattice bonds. We then use the Metropolis acceptance probability for this move, i.e.

$$P([0, 0] \rightarrow [1, b]) = \min \left[\frac{\beta N_b}{\Lambda - n} \langle \sigma(p-1) | H_{1,b} | \sigma(p-1) \rangle, 1 \right], \quad (42)$$

where n is the number of bond-operators in the sequence before the update. For the Heisenberg model considered here, this update can be immediately rejected, if the two spins at the chosen bond at the current propagation level are parallel, while in the other case, the matrix element simply equals $1/2$, cf. Eqs. (28)-(30). On the other hand, if the local operator at the current propagation level is a diagonal bond operator, $t(p) = 1$, an attempt is made to replace it by the identity operator with Metropolis acceptance probability

$$P([1, b] \rightarrow [0, 0]) = \min \left[\frac{\Lambda - n + 1}{\beta N_b} \langle \sigma(p-1) | H_{1,b} | \sigma(p-1) \rangle^{-1}, 1 \right]. \quad (43)$$

Once the full operator sequence has been traversed, one should have recovered the initial spin configuration, since $|\sigma(\Lambda)\rangle = |\sigma\rangle$. In case of long-ranged exchange interactions, the above update scheme can be easily generalized to ensure that the various interaction terms are efficiently sampled [28, 29].

4 The loop algorithm

We now present the improved global update schemes that have been developed during the 1990s and early 2000s for the world-line QMC approach. We will first discuss the idea behind the loop algorithm [8], based on the discrete-time formulation. An excellent, general account on the loop algorithm and related methods is provided by Evertz in his review [30].

To appreciate the idea behind the loop algorithm, it is useful to first recall the Swendsen-Wang [56] algorithm for the ferromagnetic Ising model. The goal there is to identify physically relevant clusters of parallel spins, and then collectively flip all spins within such a cluster in order to enhance the update dynamics near criticality. Essentially, one visits all bonds on the lattice and assigns a bond variable $\tau_b \in \{0, 1\}$ to each bond, by deciding if a given bond will be filled ($\tau_b = 1$) or kept empty ($\tau_b = 0$) according to certain probabilities, cf. the upper panel of Fig. 5. Namely, if the two spins connected by the bond are anti-parallel, the bond is kept empty for sure, while for parallel spins along a bond, the bond is filled with a finite probability $p = 1 - e^{-2\beta|J|}$, where $J < 0$ denotes the strength of the (ferromagnetic) Ising interaction on that bond. Then, one identifies clusters of spins connected by filled bonds, and attempts to flip each cluster individually with probability $1/2$. Within the Wolff-algorithm [57] one instead generates only one of these clusters by starting its construction from a randomly chosen lattice site, and then flips this cluster with probability 1.

For the case of the quantum model in the checkerboard decomposition, we cannot directly apply these schemes of assigning bond variables, since we need to ensure that after an update only valid, unbroken world-line configurations have been generated. Since we are given restrictions for the possible types of shaded plaquettes that can occur in valid configurations, we should thus instead of assigning bond variables assign plaquette variables. Therefore, in the loop algorithm, each plaquette P is assigned a plaquette variable (now called “graph” or “breakup”), with certain probabilities. For the antiferromagnetic Heisenberg model, there are two types of breakups that can be assigned to a plaquette; they are shown in the lower panel of Fig. 5 and we refer to them as the horizontal (H) and vertical (V) breakup. Each shaded plaquette with the state $C|_P$ is assigned the horizontal breakup (H) with probability

$$P(H) = \begin{cases} 0 & , \quad C|_P = \langle \uparrow\uparrow | e^{-\Delta\tau H_i} | \uparrow\uparrow \rangle, \langle \downarrow\downarrow | e^{-\Delta\tau H_i} | \downarrow\downarrow \rangle \\ \tanh(\frac{\Delta\tau J}{2}) & , \quad C|_P = \langle \uparrow\downarrow | e^{-\Delta\tau H_i} | \uparrow\downarrow \rangle, \langle \downarrow\uparrow | e^{-\Delta\tau H_i} | \downarrow\uparrow \rangle \\ 1 & , \quad C|_P = \langle \downarrow\uparrow | e^{-\Delta\tau H_i} | \uparrow\downarrow \rangle, \langle \uparrow\downarrow | e^{-\Delta\tau H_i} | \downarrow\uparrow \rangle \end{cases} \quad (44)$$

otherwise, it is assigned the vertical breakup (V). These graph assignment rules can be obtained upon considering an extended configuration space that combines spin and graph configurations, which we describe next. Such a general framework to describe cluster algorithms was presented by Kandel and Domany [31], generalizing the Fortuin-Kasteleyn [32] representation of the Ising model. Below we follow the formulation by Kawashima and Gubernatis [33], which makes this relation very transparent. We start from the original representation of the quantum partition function in terms of the spin (world-line) configurations,

$$Z = \sum_C W(C). \quad (45)$$

The phase space is now enlarged, by assigning a set of possible graphs G to the original configurations, such that

$$Z = \sum_C \sum_G W(C, G), \quad (46)$$

where the new weights $W(C, G) \geq 0$ are chosen as to ensure

$$\sum_G W(C, G) = W(C). \quad (47)$$

The algorithm then proceeds as follows: Given a configuration C (which implies $W(C) \neq 0$), we assign first a graph G to the configuration C , chosen with the correct probability,

$$P(G|C) = W(C, G)/W(C). \quad (48)$$

Then we choose a new configuration C' with probability $P[(C, G) \rightarrow (C', G)]$, keeping the graph G fixed. This completes a configurational update $C \rightarrow C'$ and the process repeats by choosing a new graph G' , etc. The first step, choosing graphs with probabilities $P(G|C)$, trivially obeys detailed balance. Detailed balance for the second step requires that

$$W(C, G)P[(C, G) \rightarrow (C', G)] = W(C', G)P[(C', G) \rightarrow (C, G)] \quad (49)$$

One possible solution of this detailed balance condition is provided by the heat-bath algorithm

$$P[(C, G) \rightarrow (C', G)] = \frac{W(C, G)}{W(C, G) + W(C', G)}. \quad (50)$$

The whole approach is apparently simplified a lot, if one can find an assignment of graph weights, such that $W(C, G)$ does not depend on the configuration C , whenever it is non-zero in that configuration, i.e., when $W(C, G)$ has the following form:

$$W(C, G) = \Delta(C, G)V(G), \quad (51)$$

where

$$\Delta(C, G) = \begin{cases} 1 & , \text{ if } W(C, G) \neq 0 \\ 0 & , \text{ otherwise} \end{cases}. \quad (52)$$

In this case, the heat-bath probability (50) simply becomes $P = 1/2$. Furthermore, the weight of the spin configuration $W(C)$ in fact decomposes into a product of plaquette weights, cf. Eq. (17). Further simplifications thus arise, if also the graph weight can be represented as a product of separate weights for each plaquette, i.e., if

$$W(C, G) = \prod_P W_P(C|_P, G|_P), \quad (53)$$

so that the whole graph assignment procedure can be performed locally on the level of the (shaded) plaquettes, i.e., in terms of the plaquette breakups.

While for a more general, anisotropic spin-1/2 model additional graph elements are required, the above conditions can be fulfilled for the Heisenberg antiferromagnet by employing only the two breakups (H and V) on each shaded plaquette and assigning the H breakup with the probability given in Eq. (44). An example of such an assignment of the plaquette breakups is shown in the central panel of Fig. 6. After having assigned to each plaquette a breakup, the

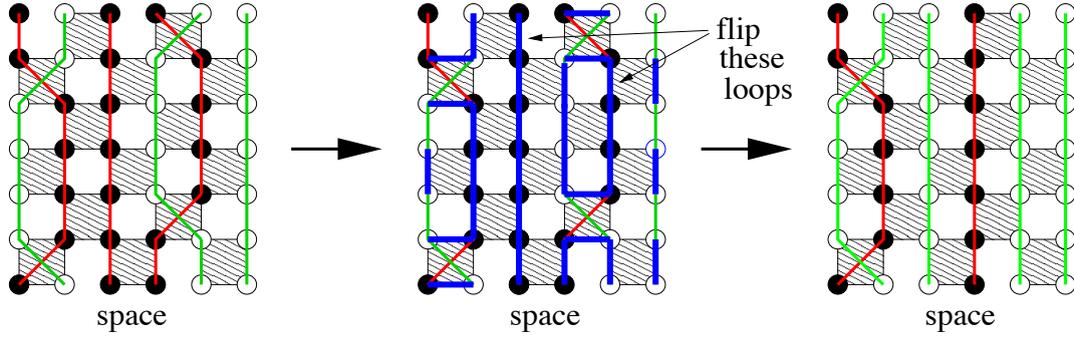


Fig. 6: An update step in the loop algorithm, where two clusters of spins are flipped.

graph edges are connected to form clusters of connected graph edges, and hence clusters (loops) of connected local spin states. Now one flips each such cluster independently with probability $P = 1/2$, thereby changing the local spins and thus the world-line configuration. An example of the whole update procedure is shown in Fig. 6. One sees from this example, that after a single loop update, one can realize large-scale changes of the world-line configuration. In particular, as seen in Fig. 6, it is possible to change the total magnetization of the system. The algorithm that we introduced above generates the whole graph, i.e., the complete set of all loops, G and flips each such cluster with probability $P = 1/2$. It is however also possible, to employ a single-loop version of the algorithm, in which, similar to the classical Wolff cluster algorithm, only a single cluster is constructed: One randomly picks a site from the world-line configuration and constructs only the cluster which includes that site. This can be done by determining the breakups (and thus the route of the loop) only on those plaquettes, that are indeed traversed during the loop construction. It is then feasible to employ the Metropolis acceptance for solving the detailed balance condition in Eq. (49), which results in the probability $P = 1$ to flip the just constructed loop. As in the classical case, the single-loop variant usually provides an even smaller dynamical critical scaling than the multi-loop variant.

Furthermore, it is also feasible to realize the loop algorithm directly in the continuous-time version of the world-line QMC formulation, in both multi- and single-loop versions [12]. Since the single-loop version paves a direct conceptual path to the worm, operator and directed loop algorithms to be introduced below, we focus here on the single-loop variant, even though it is technically a bit more involved than the multi-loop version. The key insight in realizing, that a continuous-time formulation of the loop algorithm is feasible, follows from considering the breakup probability per time, which has a continuous time limit. Indeed, for an imaginary time interval $\tau_1 < \tau_2$, during which the world-line configuration on a given bond does not change, the breakup probability is constant, and the probability density for a horizontal breakup within this imaginary time range becomes

$$\lim_{\Delta\tau \rightarrow 0} \frac{P(H)}{\Delta\tau} = \lim_{\Delta\tau \rightarrow 0} \frac{\tanh(\Delta\tau J/2)}{\Delta\tau} = \frac{J}{2}. \quad (54)$$

We are thus lead to the following procedure: To start the loop construction, a site i is picked randomly, as well as a random imaginary time τ_1 between 0 and β , from which the loop will be

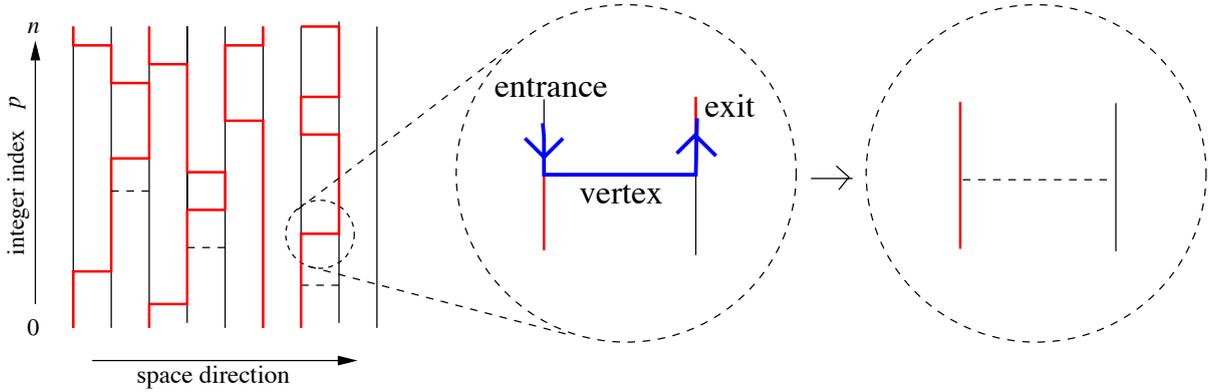


Fig. 7: A SSE configuration (left panel), the local construction of the operator loop at a 4-leg vertex corresponding to a bond operator from the SSE operator sequence (central panel), and the resulting final vertex (right panel). The path taken by the operator loop through the vertex is indicated in the central panel by the arrowed blue line (switch-and-reverse rule).

constructed, starting, e.g., initially moving upwards in time. In this forward time direction, one next identifies the time interval $\tau_1 < \tau < \tau_2$, within which the world-line configuration on all sites neighboring site i does not change. For each such neighbor j , draw a random number τ_{ij} from an exponential distribution based on the above probability density for a horizontal breakup, i.e., with $P(\tau_{ij}) \propto \exp(-J\tau_{ij}/2)$. Now, let \bar{j} be the neighbor with the smallest value of τ_{ij} , i.e., $\tau_{i\bar{j}} = \min_j(\tau_{ij})$. If $t_d = t_1 + \tau_{i\bar{j}} < t_2$, the loop-end on site i now moves up to the time t_d , and right there jumps to the site \bar{j} (a memorable analogy of this procedure is the radioactive decay with decay constant $J/2$ where the neighboring sites correspond to different decay channels). Since the constructed loops do not self-overlap, one has to exclude in the above procedure all temporal regions of neighboring sites, which have been visited by the current loop already. In case $t_d > t_2$, the loop end stays on site i and moves to time t_2 . In case the world-line jumps at time t_2 , the loop must jump as well. The whole procedure ensures the constant probability density Eq. (54) for a horizontal breakup. This process is iterated until the loop eventually closes and can be flipped as a whole.

Based on the loop algorithm, large scale simulations of up to a several million quantum spins can be performed even in (quantum) critical regions. After the original introduction of the loop algorithm for the spin-1/2 Heisenberg model in a six-vertex model formulation, it has seen various extensions to anisotropic- and higher-spin models, as well as to Hubbard and t - J models. For a detailed derivation of the breakup rules, also beyond the case of the isotropic Heisenberg model considered here, we refer to the review article by Evertz [30]. It is also possible to devise so-called “improved estimators” [34] for certain observables within the loop update, e.g., for correlation functions or magnetic susceptibilities, which provide a significant reduction of statistical errors compared to “naive” estimators for the same quantities. See Evertz’s article [30] for details.

Within the SSE formulation, which does not exhibit any systematic (Trotter) error, but still avoids the technical complications of the continuous-time formulation, it is feasible to imple-

ment a very simple version of a loop algorithm [26, 35]. This cluster update complements the local diagonal updates, that were introduced in Sec. 3.3. Namely, one considers to construct a discrete-time-like loop, based on a given SSE configuration: Consider a SSE configuration, such as the one shown in Fig. 4(c), redrawn in Fig. 7, as a collection of n vertexes (where n denotes the number of non-identity operators in the operator string), each coming along with 4 “legs”. In the central part of Fig. 7, one such 4-leg-vertex is highlighted. The loop construction starts from a randomly chosen vertex leg in the SSE configuration, called the “entrance leg”. Like in the discrete-time version of the loop algorithm, one then defines a local rule of how the loop continues through this vertex and leaves the vertex through an “exit leg”. One can interpret such rules as prescribing the scattering of the loop-head off the vertexes. The corresponding probabilities for choosing the exit leg among the four legs of the vertex need to satisfy again a local detailed balance condition. For the antiferromagnetic Heisenberg model it turns out that a very simple, deterministic scattering rule can be devised, namely, that the entrance and exit leg are always located on the same side of the vertex (the switch-and-reverse rule). This is illustrated in the central part of Fig. 7. After the exit leg has been assigned, the loop moves from this exit leg to the leg of another vertex, to which the exit leg is connected within the given SSE configuration, cf. Fig. 7. Now this leg becomes the new entrance leg on the new vertex and the whole process repeats, until the loop closes, i.e., when the last exit leg eventually equals the initial entrance leg on the starting vertex. In the example of Fig. 7, the loop closes when it reaches back to the initial entrance leg after it has visited the vertex atop the vertex from which it started. All spins along the path of the loops are then flipped. In addition to replacing all visited diagonal operators by off-diagonal operators and vice versa (cf. the right panel of Fig. 7), this operator loop update can also modify the initial spin configuration $|\sigma\rangle$, thus ensuring an ergodic sampling of the SSE configuration space if combined with the diagonal update. Of course, it is also possible in this SSE formulation to construct all the loops for a given SSE configuration and perform a multi-loop update, flipping each such loop with probability $1/2$.

5 Worms, operator loops, and directed loops

The cluster-updates that have been introduced above provide optimal performance for spin inversion symmetric Hamiltonians. However, terms in the Hamiltonian which break this symmetry, such as a magnetic field, are not taken into account during the loop construction. Instead, they would have to be included through the acceptance rate of the loop flips, which can become exponentially small at low temperatures, degrading the algorithm’s performance. Hence, it would be highly desirable to generalize the idea of a local construction of a global cluster of spins to cases, where spin-inversion symmetry is not present. The algorithms that we outline next provide such a general sampling scheme.

The worm algorithm [13] formally works in an extended configuration space, where in addition to closed world-line configurations one open world-line fragment (the “worm”) is allowed.

Formally, this is done by adding a source term to the Hamiltonian,

$$H_{\text{worm}} = H - \eta \sum_i (S_i^+ + S_i^-), \quad (55)$$

which allows world lines to be broken with a matrix element proportional to η . The worm algorithm then proceeds as follows: a worm, i.e., a world-line fragment, is created by randomly inserting a pair (S_i^+, S_i^-) of operators on a world line at nearby times in the world-line configuration. The ends of this worm are then moved randomly in space and time, using local Metropolis or heat bath updates, until the two ends of the worm meet again. Then an update which removes the worm is proposed, and if accepted we are back to a configuration with closed world lines only. This algorithm is straightforward, consisting just of local updates of the worm ends in the extended configuration space but it can perform non-local changes. Furthermore, a worm end can wind around the lattice in the temporal or spatial direction and that way change the magnetization and winding number. While not being as efficient as the loop algorithm in zero magnetic field (the worm movement follows a random walk while the loop algorithm can be interpreted as a self-avoiding random walk), the big advantage of the worm algorithm is that it remains efficient in the presence of a magnetic field. A similar algorithm was actually proposed more than a decade earlier [36]. Instead of a random walk fulfilling detailed balance at every move, the worm head in this earlier algorithm just performed a random walk. The *a posteriori* acceptance rates are then often very small and the algorithm is not efficient, just as the small acceptance rates for loop updates in magnetic fields make the loop algorithm inefficient. This highlights the importance of having the cluster-building rules of a non-local update algorithm closely tied to the physics of the problem.

Algorithms with a similar basic idea are the operator-loop update [26] in the SSE formulation and the directed-loop algorithms [27, 37] which can be formulated in both the SSE and the world-line representation. Like the worm algorithm, these algorithms create two world line discontinuities and move them around by local updates. The operator-loop algorithm for the SSE representation can be understood as a generalization of the loop algorithm for the SSE configuration, which we described at the end of Sec. 4 [26]. Again, a loop (“operator-loop”) is constructed for a given SSE configuration, starting from a random leg of a randomly chosen vertex as the entrance leg. However, for a general Hamiltonian, including, e.g., a magnetic field, the scattering rules at the vertices have to be chosen appropriately, and in general turn out to include a stochastic decision, instead of the deterministic switch-and-reverse-rule that holds for the isotropic Heisenberg model. One generic solution of the local detailed balance condition on the scattering rates is provided by a heat-bath choice among the 4 possible exit legs: the probability to choose one out of the four legs is taken to be proportional to the matrix element of the bond operator of the considered vertex in the final resulting spin configuration on this bond [26]. This choice already leads to a rather efficient algorithm. However, it is in general not excluded, that the chosen exit leg is equal to the leg on which the vertex was entered. Such a “bounce” move would result in the operator loop retracing its previous path, and thus in undoing a previously performed change to the spin configuration and the operator content of the operator string. It thus appears desirable to reduce the probability for such bounce

moves, or to even eliminate them completely. In fact, for many cases it can be shown that scattering rates can be optimized such as to eliminate completely the bounce move, while still ensuring detailed balance of the operator-loop construction [26, 27, 35, 37, 38]. An example is just the switch-and-reverse-rule for the isotropic Heisenberg antiferromagnet that was presented above. For more general models, a systematic approach to reduce the bounce probability has been formulated [27]. Such algorithms, which direct the loop away from the last changes, are called “directed loop” methods and can be formulated also for the continuous-time world-line approach [27]. Furthermore, it is possible to use linear programming techniques, in order to optimize the scattering rates with respect to the overall bounce probability for general quantum lattice models [38]. An alternative strategy to enhance the update dynamics of the operator loop algorithm is based on the concept of optimal Monte Carlo updates [39, 40]: Starting from the heat-bath solution, one iteratively improves the scattering rules by minimizing its higher eigenvalues in a matrix formulation. Both this “locally optimal solution” and the directed loops are superior to the heat-bath solution and typically perform equally well.

We furthermore want to point out that within the worm and operator-loop approaches, a natural implementation for evaluating transverse spin-spin correlation functions can be realized, based upon on-the-fly measurements of the corresponding matrix elements at the worm-ends, while constructing the worm’s path through the world-line configuration [11, 13, 24, 38]. Based on the close connection between the SSE and the continuous-time world-line approach, that we mentioned in Sec. 3.3, it is possible to implement such measurements in a rather efficient way [25]. Finally, we would like to point out that beyond the cluster update methods, which drastically reduce autocorrelation times in QMC simulations of quantum spin systems, it is also possible to adapt various extended ensemble methods such as parallel tempering [41], multicanonical methods [42, 43], and the Wang-Landau technique [44, 45], which are employed in classical Monte Carlo simulations for systems with a rough energy landscape or at first-order phase transitions, to further enhance the performance of QMC simulations in similar situations [23, 46–48].

6 The sign problem

Before closing this short review on world line QMC methods, we finally want to address a severe restriction of this approach: the sign problem. This is indeed its major limitation (note that also other unbiased QMC methods, in particular those for fermionic models, exhibit sign problems in many physically interesting situations) [49]. Currently, it seems rather unlikely that there is a general solution to this problem [50]. Let us explain what the sign problem is and where it comes from. For this purpose, consider again the weight of a plaquette in the Suzuki-Trotter decomposition that corresponds to a local spin exchange:

$$\langle \uparrow\downarrow | e^{-\Delta\tau H_i} | \downarrow\uparrow \rangle = -\sinh\left(\frac{\Delta\tau J}{2}\right). \quad (56)$$

In the antiferromagnetic case ($J > 0$), this matrix element is negative. Then we might actually worry that there will be allowed QMC world line configurations with an overall negative weight,

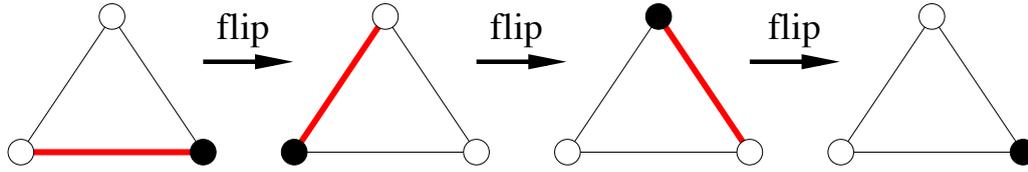


Fig. 8: Example of an odd number of spin flips on a triangle, returning to the initial state. The bond along which the spins are going to be flipped next is indicated by a red line.

which cannot be assigned as the probability for such a configuration, since probabilities ought to be positive. It is not too hard to convince yourself, that for an open chain or square lattice (in fact any bipartite lattice graph structure) the total number of spin exchange plaquettes in an allowed QMC configuration is even. In that case, the overall weight of the configuration will be positive. Hence, for the Heisenberg antiferromagnet, we are restricted to positive weights only, and can apply the Monte Carlo sampling method of the world-line configurations. However, if the lattice is not bipartite, this property is lost. To see why, consider a single triangle as an example of a non-bipartite lattice. In this case, we can start from one local spin configuration and return to the same configuration, after applying an odd number of (namely 3) spin exchanges, see Fig. 8 for an illustration. Hence, on the triangular lattice, and in fact on any frustrated, non-bipartite lattice, the weight $W(C)$ of a QMC configuration C is not necessarily positive, but can take on negative values as well. How can we then perform a MC sampling of such configurations? Well, at first sight, there seems to be an easy solution: Let us express the partition function of the quantum model in the effective classical representation as follows:

$$Z = \sum_C W(C) = \sum_C |W(C)| \cdot \text{Sgn}(C), \quad (57)$$

where we have written the weight as a product of its absolute value and its sign $\text{Sgn}(C) = W(C)/|W(C)| = \pm 1$. Now assume that we would perform a simulation of the antiferromagnet and ignore the sign, thus weighting each configuration according to the absolute values $|W(C)|$ of the weights. In fact, this precisely corresponds to simulating the ferromagnetic Heisenberg model on the same lattice, with the partition function

$$Z_F = \sum_C |W(C)|. \quad (58)$$

Using the generated configurations, we can then obtain the following expectation value of an observable A of the antiferromagnetic model,

$$\begin{aligned} \langle A \rangle &= \frac{1}{Z} \sum_C W(C) A(C) \\ &= \frac{1}{Z} \frac{Z_F}{Z_F} \sum_C |W(C)| \cdot \text{Sgn}(C) \cdot A(C) \\ &= \frac{Z_F}{Z} \frac{1}{Z_F} \underbrace{\sum_C |W(C)| \cdot \text{Sgn}(C) \cdot A(C)}_{\langle \text{Sgn} \cdot A \rangle_F}, \end{aligned} \quad (59)$$

from (i) the expectation value of $\langle \text{Sgn} \cdot A \rangle_F$ measured from the simulation of the ferromagnetic model (i.e., we would measure the expectation value of the product of the sign of a configuration times the value of A in that configuration), and (ii) the expectation value of the sign itself, since

$$\frac{Z}{Z_F} = \langle \text{Sgn} \rangle_F, \quad (60)$$

so that

$$\langle A \rangle = \frac{\langle \text{Sgn} \cdot A \rangle_F}{\langle \text{Sgn} \rangle_F}. \quad (61)$$

Hence, it looks like in order to obtain the expectation value of A for the antiferromagnetic model on a frustrated lattice, we could just simulate the ferromagnetic model (all weights positive), and perform the two measurements of $\langle \text{Sgn} \cdot A \rangle_F$ and $\langle \text{Sgn} \rangle_F$. In fact, such an approach works for large temperatures $T \gg J$, where the spins do not feel their exchange interactions too much, and behave weakly coupled in both the ferromagnetic and the antiferromagnetic case. But, once we are interested in the low temperature regime, then both models are characterized by very different important states. Can we really expect to learn something about the low temperature behavior of the antiferromagnet from the low temperature behavior of the ferromagnet? No, this should not be reasonable, and indeed, the above approach fails miserably at low temperatures. What happens is, that there will be almost the same number of configurations with positive weight as with negative weight, such that the average sign decreases exponentially to zero. Namely, at low temperatures we find

$$\langle \text{Sgn} \rangle_F \sim e^{-\beta N_s \Delta f}, \quad (62)$$

where Δf is the difference in the free energy per site between the antiferromagnetic and the ferromagnetic model – remember, that the free energy $F = -\ln(Z)/\beta$. Hence, in order to calculate the physical observable A from Eq. (61) we divide (in the denominator) by an exponentially small value with a finite statistical error, leading to an exponentially increasing statistical uncertainty for $\langle A \rangle$. Thus, in order to reach low temperatures and large system sizes, the computational time needs to grow exponentially and renders any useful simulation impossible. This behavior is called “the sign problem”, and makes QMC simulations of, e.g., frustrated quantum magnets practically impossible within the interesting (low-) temperature and large-system regime.

Up to date, no general solution of the QMC sign problem is known, although it can be overcome in certain special cases (cf., e.g., [35, 51, 52]). Moreover, it has been shown that a general solution to the sign problem essentially constitutes an NP-hard challenge [50]. It is however generally suspected, that no polynomial-time solutions to NP-hard problem exist. Hence, we urge the reader to contact us immediately in case she or he finds a serious path to a generic solution of the QMC sign problem! In the mean time, we hope to have stimulated your interest in performing some QMC simulations of your own for sign-problem free situations. An excellent opportunity to access application-ready codes for the methods that were introduced here is the ALPS (Algorithms and Libraries for Physics Simulations) library [53, 54]. Furthermore, basic SSE simulation codes are also available online [55]. Let the dice roll!

Appendix

A Quantum Heisenberg Model 101

For reference and to fix notation, we define the spin-1/2 quantum Heisenberg model on a finite lattice, such as a one-dimensional chain or a two-dimensional square lattice. A spin-1/2 quantum spin resides at each lattice site. Each spin \vec{S}_i , located at lattice site $i = 1, \dots, N_s$ is described by quantum mechanical spin operators

$$\vec{S}_i = \begin{pmatrix} S_i^x \\ S_i^y \\ S_i^z \end{pmatrix},$$

which fulfill the commutation relation $[S_i^\alpha, S_j^\beta] = i\hbar\varepsilon_{\alpha\beta\gamma}S_i^\gamma\delta_{ij}$, i.e., operators on different sites ($i \neq j$) commute, and locally ($i = j$) the above commutation relations reduce to the well-known algebra of spin operators. In order to set up a basis of the Hilbert space, we start from the eigenvectors of the local S_i^z operator at each lattice site i :

$$\begin{aligned} |S_i^z = +1/2\rangle &= |\uparrow\rangle_i = \begin{pmatrix} 1 \\ 0 \end{pmatrix}_i, \\ |S_i^z = -1/2\rangle &= |\downarrow\rangle_i = \begin{pmatrix} 0 \\ 1 \end{pmatrix}_i. \end{aligned}$$

In this local basis, the spin operators are given in matrix form as ($\hbar = 1$),

$$\begin{aligned} S_i^z = \frac{1}{2}\sigma_z &= \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & -\frac{1}{2} \end{pmatrix}, & \text{i.e.} & \begin{aligned} S_i^z|\uparrow\rangle_i &= +\frac{1}{2}|\uparrow\rangle_i \\ S_i^z|\downarrow\rangle_i &= -\frac{1}{2}|\downarrow\rangle_i \end{aligned} \\ S_i^+ = S_i^x + iS_i^y &= \frac{1}{2}\sigma_x + i\frac{1}{2}\sigma_y = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, & \text{i.e.} & \begin{aligned} S_i^+|\downarrow\rangle_i &= |\uparrow\rangle_i \\ S_i^+|\uparrow\rangle_i &= 0 \end{aligned} \\ S_i^- = S_i^x - iS_i^y &= \frac{1}{2}\sigma_x - i\frac{1}{2}\sigma_y = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, & \text{i.e.} & \begin{aligned} S_i^-|\uparrow\rangle_i &= |\downarrow\rangle_i \\ S_i^-|\downarrow\rangle_i &= 0 \end{aligned} \end{aligned}$$

with the spin raising and lowering operators S_i^+ and S_i^- , and the Pauli matrices

$$\begin{aligned} \sigma_x &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \\ \sigma_y &= \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \\ \sigma_z &= \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \end{aligned}$$

For a lattice of N_s spins, we have a total of 2^{N_s} basis states of the Hilbert space,

$$\{|S_1^z, \dots, S_{N_s}^z\rangle\} = \left\{ \begin{array}{l} |\uparrow, \uparrow, \dots, \uparrow, \uparrow\rangle \\ |\uparrow, \uparrow, \dots, \uparrow, \downarrow\rangle \\ |\uparrow, \uparrow, \dots, \downarrow, \uparrow\rangle \\ \vdots \\ |\uparrow, \downarrow, \dots, \downarrow, \downarrow\rangle \\ |\downarrow, \downarrow, \dots, \downarrow, \downarrow\rangle \end{array} \right\}.$$

Each spin operator on site i acts only on the corresponding local spin. Next, we introduce an exchange coupling between nearest neighbor spins, so that the quantum spin system is described by the Hamiltonian

$$H = J \sum_{\langle i, j \rangle} \vec{S}_i \cdot \vec{S}_j,$$

where $\langle i, j \rangle$ indicates a pair of neighboring lattice sites.

In case $J > 0$, the above model is the antiferromagnetic quantum Heisenberg model, which results, e.g., within second-order perturbation theory from the half-filled Hubbard model at large local repulsions. The Hamiltonian of the Heisenberg model can also be written using spin raising and lowering operators as

$$H = J \sum_{\langle i, j \rangle} \vec{S}_i \cdot \vec{S}_j = J \sum_{\langle i, j \rangle} \left(\frac{1}{2} (S_i^+ S_j^- + S_i^- S_j^+) + S_i^z S_j^z \right).$$

Consider for example a system of only two spins (a dimer). The Hamiltonian is then given by the matrix

$$H = J \begin{pmatrix} \frac{1}{4} & 0 & 0 & 0 \\ 0 & -\frac{1}{4} & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & -\frac{1}{4} & 0 \\ 0 & 0 & 0 & \frac{1}{4} \end{pmatrix} \text{ in the basis } \begin{array}{l} |\uparrow\uparrow\rangle \\ |\uparrow\downarrow\rangle \\ |\downarrow\uparrow\rangle \\ |\downarrow\downarrow\rangle \end{array}.$$

References

- [1] M. Suzuki, *Prog. of Theor. Phys.* **56**, 1454 (1976)
- [2] H.F. Trotter, *Proc. Am. Math. Soc.* **10**, 545 (1959)
- [3] M. Suzuki, S. Miyashita and A. Kuroda, *Prog. Theor. Phys.* **58**, 1377 (1977)
- [4] J.E. Hirsch, D.J. Scalapino, R.L. Sugar and R. Blankenbecler, *Phys. Rev. Lett.* **47**, 1628 (1981)
- [5] M.S. Makivić and H.-Q. Ding, *Phys. Rev. B* **43**, 3562 (1991)
- [6] F.F. Assaad and D. Würtz, *Phys. Rev. B* **44**, 2681 (1991)
- [7] N. Metropolis, A.R. Rosenbluth, M.N. Rosenbluth, A.H. Teller and E. Teller, *J. of Chem. Phys.* **21**, 1087 (1953)
- [8] H.G. Evertz, G. Lana, and M. Marcu, *Phys. Rev. Lett.* **70**, 875 (1993)
- [9] R.M. Fry, *Phys. Rev. B* **33**, 6271 (1986)
- [10] R.M. Fry and R.T. Scalettar, *Phys. Rev. B* **36**, 3833 (1987)
- [11] N.V. Prokof'ev, B.V. Svistunov, and I.S. Tupitsyn, *Pis'ma v Zh.Eks. Teor. Fiz.*, **64**, 853 (1996) [English translation: cond-mat/9612091]
- [12] B.B. Beard and U.-J. Wiese, *Phys. Rev. Lett.* **77**, 5130 (1996)
- [13] N.V. Prokof'ev, B.V. Svistunov, and I.S. Tupitsyn, *Sov. Phys. - JETP* **87**, 310 (1998)
- [14] D.C. Handscomb, *Proc. Cambridge Philos. Soc.* **58**, 594 (1962)
- [15] D.C. Handscomb, *Proc. Cambridge Philos. Soc.* **60**, 115 (1964)
- [16] J.W. Lyklema, *Phys. Rev. Lett.* **49**, 88 (1982)
- [17] D.H. Lee, J.D. Joannopoulos, and J. W. Negele, *Phys. Rev. B* **30**, 1599 (1984)
- [18] E. Manousakis and R. Salvador, *Phys. Rev. B* **39**, 575 (1989)
- [19] A.W. Sandvik and J. Kurkijärvi, *Phys. Rev. B* **43**, 5950 (1991)
- [20] A.W. Sandvik, *J. Phys. A* **25**, 3667 (1992)
- [21] A.W. Sandvik, *AIP Conf. Proc.* **1297**, 135 (2010)
- [22] A. W. Sandvik, R. R. P. Singh, and D. K. Campbell, *Phys. Rev. B* **56**, 14510 (1997)
- [23] S. Wessel, N. Stoop, E. Gull, S. Trebst, and M. Troyer, *J. Stat. Mech.*, P12005 (2007)

- [24] A. Dorneich and M. Troyer, *Phys. Rev. E* **64**, 066701 (2001)
- [25] F.F. Assaad and H.G. Evertz, in H. Fehske, R. Schneider and A. Weiße (eds.): *Computational Many-Particle Physics* (Springer, Heidelberg, 2008)
- [26] A.W. Sandvik, *Phys. Rev. B* **59**, R14157 (1999)
- [27] O.F. Syljuåsen and A.W. Sandvik, *Phys. Rev. E* **66**, 046701 (2002)
- [28] A.W. Sandvik, *Phys. Rev. E* **68**, 056701 (2003)
- [29] K. Fukui and S. Todo, *J. Comput. Phys.* **228**, 2629 (2009)
- [30] H.G. Evertz, *Adv. Phys.* **52**, 1 (2003)
- [31] D. Kandel and E. Domany, *Phys. Rev. B* **43**, 8539 (1991)
- [32] C. Fortuin and P. Kasteleyn, *Physica* **57**, 536 (1972)
- [33] N. Kawashima and J. Gubernatis, *J. Stat. Phys.* **80**, 169 (1995)
- [34] B. Ammon, H.G. Evertz, N. Kawashima, M. Troyer and B. Frischmuth, *Phys. Rev. B* **58**, 4304 (1998)
- [35] P. Henelius and A.W. Sandvik, *Phys. Rev. B* **63**, 1102 (2000)
- [36] J.J. Cullen and D.P. Landau, *Phys. Rev. B* **27**, 297 (1983)
- [37] O.F. Syljuåsen, *Phys. Rev. E* **67**, 046701 (2003)
- [38] F. Alet, S. Wessel, and M. Troyer *Phys. Rev. E* **71**, 036706 (2005)
- [39] L. Pollet, S.M.A. Rombouts, K. Van Houcke, and K. Heyde, *Phys. Rev. E* **70**, 056705 (2004)
- [40] L. Pollet, K. Van Houcke, and S.M.A. Rombouts, *J. Comp. Phys.* **225**, 2249 (2007)
- [41] E. Marinari and G. Parisi, *Europhys. Lett.* **19**, 451 (1992)
- [42] B.A. Berg and T. Neuhaus, *Phys. Lett. B.* **267**, 249 (1991)
- [43] B.A. Berg and T. Neuhaus, *Phys. Rev. Lett.* **68**, 9 (1992)
- [44] F. Wang and D. P. Landau, *Phys. Rev. Lett.* **86**, 2050 (2001)
- [45] F. Wang and D. P. Landau, *Phys. Rev. E* **64**, 056101 (2001)
- [46] P. Sengupta, A.W. Sandvik, and D.K. Campbell, *Phys. Rev. B* **65**, 155113 (2002)
- [47] M. Troyer, S. Wessel and F. Alet, *Phys. Rev. Lett.* **90**, 120201 (2003)

-
- [48] M. Troyer, F. Alet, and S. Wessel, *Braz. J. of Physics* **34**, 377 (2004)
- [49] E.Y. Loh Jr., J.E. Gubernatis, R.T. Scalettar, S.R. White, D.J. Scalapino, and R.L. Sugar, *Phys. Rev. B* **41**, 9301 (1990)
- [50] M. Troyer and U.-J. Wiese, *Phys. Rev. Lett.* **94**, 170201 (2005)
- [51] T. Nakamura, *Phys. Rev. B* **57**, R3197 (1998)
- [52] S. Chandrasekharan and U.-J. Wiese, *Phys. Rev. Lett.* **83**, 3116 (1999)
- [53] A. F. Albuquerque, F. Alet, P. Corboz, P. Dayal, A. Feiguin, S. Fuchs, L. Gamper, E. Gull, S. Gürtler, A. Honecker, R. Igarashi, M. Körner, A. Kozhevnikov, A. Läuchli, S. R. Manmana, M. Matsumoto, I. P. McCulloch, F. Michel, R. M. Noack, G. Pawłowski, L. Pollet, T. Pruschke, U. Schollwöck, S. Todo, S. Trebst, M. Troyer, P. Werner, and S. Wessel, *J. Magn. Mater.* **310**, 1187 (2007)
- [54] Codes available online at <http://alps.comp-phys.org>
- [55] Codes available online at <http://physics.bu.edu/~simonsandvik>
- [56] R.H. Swendsen and J.S. Wang, *Phys. Rev. Lett* **58**, 86 (1987)
- [57] U. Wolff, *Phys. Rev. Lett.* **62**, 361 (1989)