Fast, Hierarchical, and Adaptive Algorithm for Metropolis Monte Carlo Simulations of Long-Range Interacting Systems

Fabio Müller,[*] Henrik Christiansen,[†] Stefan Schnabel,[‡] and Wolfhard Janke[§]
Institut für Theoretische Physik, Universität Leipzig, IPF 231101, 04081 Leipzig, Germany

(Received 19 January 2022; revised 23 March 2023; accepted 12 May 2023; published 17 July 2023)

We present a fast, hierarchical, and adaptive algorithm for Metropolis Monte Carlo simulations of systems with long-range interactions that reproduces the dynamics of a standard implementation exactly, i.e., the generated configurations and consequently all measured observables are identical, allowing in particular for nonequilibrium studies. The method is demonstrated for the power-law interacting long-range Ising and XY spin models with nonconserved order parameter and a Lennard-Jones particle system, all in two dimensions. The measured run times support an average complexity $O(N \log N)$, where N is the number of spins or particles. Importantly, prefactors of this scaling behavior are small, which in practice manifests in speedup factors larger than 10^4. The method is general and will allow the treatment of large systems that were out of reach before, likely enabling a more detailed understanding of physical phenomena rooted in long-range interactions.

DOI: 10.1103/PhysRevX.13.031006 Subject Areas: Computational Physics

I. INTRODUCTION

The statistical physics of interacting N-body systems poses many important scientific problems that can be solved by analytic methods only approximately or in certain limits. Therefore, they are nowadays often investigated by means of computer simulations which can be categorized into two main groups: molecular dynamics (MD) simulations solve a system's equations of motion numerically and Monte Carlo (MC) simulations explore its phase space in a stochastic manner. In both cases the interaction among the constituents of the system has to be taken into account, for MD simulations as forces and for the MC method as energy changes associated with random moves of the components. With short-range interactions only a few other partners have to be considered while in the long-range case all the other constituents of the system are involved. This severely limits the accessible system size N, since updating all constituents once naively requires $\sim N^2$ operations, usually labeled as complexity $O(N^2)$. Since systems with long-range interactions are omnipresent in nature [1–8], fast algorithms for their investigation are highly desirable. Consequently, there has been a lot of research proposing several methods addressing this inherent computational challenge.

Two major classes of such methods are (i) methods based on splitting the evaluation of the potential into short- and long-range contributions, with one important example being the Ewald summation [9], and (ii) hierarchical methods where groups of components are treated collectively such as the Barnes-Hut algorithm [10]. Algorithms from these two classes reduce the computational complexity to $O(N^{3/2})$ and $O(N \log N)$, respectively. However, they have some disadvantages in certain situations (periodic versus free boundaries, very large prefactors, control of systematic errors) and cannot all equally well be employed in MD and MC simulations. Only in MD, where all components progress synchronously, advanced algorithms based on fast multipole methods, particle mesh Ewald, and multigrid techniques which calculate all forces at the same time lead to even further reduced computational complexity. Most of these studies focus on Coulomb interactions; for reviews, see Refs. [11,12]. In contrast, typically MC algorithms work asynchronously; i.e., they change only small parts of the system at a time. There, these advanced methods cannot be used successfully, since calculating all interactions after each local update is wasteful even if done efficiently. To achieve a similar improvement also for (asynchronous) Metropolis MC simulations of long-range systems, we present a hierarchical adaptive algorithm that for the here considered systems reduces the computational complexity while maintaining a small prefactor without introducing any systematic

[*]fabio.mueller@itp.uni-leipzig.de
[†]henrik.christiansen@itp.uni-leipzig.de
Present address: NEC Laboratories Europe GmbH, Kurfürsten-Anlage 36, 69115 Heidelberg, Germany.
[‡]stefan.schnabel@itp.uni-leipzig.de
[§]wolfhard.janke@itp.uni-leipzig.de

errors. The basic idea of our algorithm is the combination of the inverted Metropolis criterion with an adaptive treelike spatial decomposition of the interaction energy. Our algorithm reproduces exactly the same Markov chain as a traditional Metropolis implementation and can therefore be used as a one-to-one replacement, enabling in particular nonequilibrium studies as well.

This is a major conceptual difference to MC methods that deviate from conventional Metropolis dynamics. Prominent examples are nonlocal cluster algorithms [13–15] for the simulation of spin systems which can reduce complexity to $O(N)$, overcome critical slowing-down, and hence be more efficient for equilibrium studies close to criticality than any algorithm with local dynamics including the one presented here. Furthermore, there is the rejection-free event-chain MC method for systems with continuous degrees of freedom [16]. It was first applied with great success to hard-sphere systems and was later developed further to treat systems with general interactions [17,18]. It exploits that additive terms in the Hamiltonian transpose to factors in the Boltzmann weight and thus allow the application of a factorized Metropolis filter [17,19]. This idea has also been used in the recent development of the clock MC method [20] which in contrast to event-chain MC simulations is applicable to Ising systems as well and has a reported complexity of $O(N)$.

In this study, we demonstrate the power of our algorithm by applying it to nonconserved long-range O(n) vector spin models, the long-range Ising ($n = 1$) model (LRIM) and XY ($n = 2$) model (LRXYM), with power-law decaying potential, where we focus on integrable interactions, and an off-lattice Lennard-Jones (LJ) particle system, all in two dimensions. In equilibrium, we find that the algorithm's performance is excellent and run times are in agreement with a complexity of $O(N \log N)$ for all three systems. The LRIM undergoes a second-order phase transition at the critical temperature T_c while the LRXYM can have either a symmetry-breaking transition or a Berezinskii-Kosterlitz-Thouless (BKT) transition, or even both, depending on the decay of the long-range potential [21]. For both spin models we consider three nonequilibrium processes: quenches from a disordered configuration to a temperature $T < T_c$ into the ordered phase and to T_c itself and a Kibble-Zurek protocol. The first case is referred to as phase-ordering kinetics where coarsening and aging phenomena [22–24] occur. In the second setup critical aging [24,25] can be investigated. In the last protocol the central aspect concerns the loss of adiabaticity and the resulting formation of defects during the crossing of phase transitions [26–28]. Such nonequilibrium processes are governed by local dynamics and consequently need to be modeled by local algorithms like the one presented here. We show that in all cases a significant speedup is achieved and point out other systems where a similar performance can be expected.

The rest of the paper is structured as follows. In Sec. II we first describe the general procedure of our algorithm. We then apply it to the LRIM and the LRXYM in Sec. III and the LJ system in Sec. IV. In Sec. V we outline the applicability for a variety of other important (spin) models. Finally, we conclude and give an outlook in Sec. VI.

II. METHOD

In this section we present a general formulation of the algorithm, with no reference to any system-specific properties. We consider a system of N components $\mathbf{q} = (q_1, \ldots, q_i, \ldots, q_N)$ where q_i can, for example, stand for the spatial position $\mathbf{r}_i$, a binary Ising spin $s_i = \pm 1$, or other internal degrees of freedom (which may also be vector valued). The components q_i and q_j interact via a symmetric pairwise potential $V_{i,j} = V_{j,i}$ such that the total energy reads

$$E = \frac{1}{2} \sum_i \sum_{j \neq i} V_{i,j}. \tag{1}$$

For a traditional Metropolis MC simulation with local dynamics at each step the update of a single randomly chosen component q_i is proposed:

$$\begin{aligned} \mathbf{q}^{\text{old}} &= (q_1, \ldots, q_i^{\text{old}}, \ldots, q_N) \\ \rightarrow \mathbf{q}^{\text{new}} &= (q_1, \ldots, q_i^{\text{new}}, \ldots, q_N). \end{aligned} \tag{2}$$

The proposed update is accepted with the Metropolis probability,

$$P_{\text{acc}} = \min\left(1, e^{-\beta \Delta E}\right), \tag{3}$$

where $\Delta E = E^{\text{new}} - E^{\text{old}}$ is the energy difference resulting from the update and β is the inverse temperature. This acceptance probability is then compared to a (pseudo) random number $\rho \in [0, 1)$. If the calculated probability is larger than this random number, the update is accepted and otherwise rejected. As only one component q_i of the system is updated, we can write the change in energy of the whole system as

$$\Delta E = \sum_{j \neq i} \left(V_{i,j}^{\text{new}} - V_{i,j}^{\text{old}}\right). \tag{4}$$

Since in long-range interacting systems all constituents of the system interact with each other, the calculation of ΔE requires the evaluation of $2(N-1)$ interactions. A Metropolis MC sweep (MCS) consisting of N updates therefore has $O(N^2)$ complexity.

The traditional way of performing a Metropolis simulation is to calculate $\exp(-\beta\Delta E)$ first and then compare it to ρ. That is, an update is accepted only if

$$\rho \leq e^{-\beta\Delta E} \tag{5}$$

is fulfilled. Equivalently, one may write

$$\Delta E \leq -\frac{\ln\rho}{\beta} \equiv \Delta E_{\mathrm{th}} \tag{6}$$

and first draw the random number determining the threshold energy ΔE_{th}. This shifts the decision about a proposed update to the problem of determining whether the actual energy difference ΔE involved in the update lies below or above ΔE_{th}. In order to achieve that, ΔE does *not* need to be known exactly. Instead, it is enough to establish sufficiently narrow, rigorous bounds $\Delta E^{\mathrm{min}} \leq \Delta E \leq \Delta E^{\mathrm{max}}$. The update is either accepted if the upper bound ΔE^{max} is smaller than ΔE_{th} or rejected if the lower bound ΔE^{min} lies above ΔE_{th}. Avoiding the direct calculation of ΔE can reduce the complexity and, thus, result in considerable speedups. This procedure can easily be applied to other acceptance criteria such as, e.g., the Glauber acceptance rule, giving a different expression for ΔE_{th}.

To construct the bounds $\Delta E^{\mathrm{min/max}}$ we perform a spatial decomposition of the simulation domain which is based on an extrinsic treelike structure, in contrast to the intrinsic decomposition for self-avoiding walks [29,30] and polymers [31]. We note that any d-dimensional simulation box of linear size L can be split into 2^d boxes of size $L/2$. Of course, each of these boxes can again be split into 2^d boxes of size $L/4$ and so on. This is repeated until each box contains no more than one constituent. All theses boxes are thus automatically arranged hierarchically on a tree $\mathcal{T}$. Inner nodes contain only the collective information needed for the estimation of the interaction, whereas within each leaf the single contained constituent is stored. The construction of $\mathcal{T}$ has complexity $O(N)$ and rebuilding $\mathcal{T}$ completely at each update step would, therefore, be inefficient. Instead we update $\mathcal{T}$ locally after an accepted update. This requires $O(\log N)$ operations, since only the collective information of all the ancestor nodes of the leaf containing the updated component needs to be modified.

This spatial decomposition of the simulation domain allows us to split the energy difference which follows from an update,

$$\Delta E = \sum_{B\in\mathcal{D}} \Delta E_B, \tag{7}$$

where $\mathcal{D}$ is the set of currently selected, nonoverlapping boxes (which may be of different size) covering the simulation space and $\Delta E_B = E_B^{\mathrm{new}} - E_B^{\mathrm{old}}$ is the exact change in energy contributed by the interaction with the constituents of box B. Accordingly, if we can find strict lower and upper bounds $\Delta E_B^{\mathrm{min/max}}$ of ΔE_B, we can establish bounds for the total energy change as well:

$$\Delta E^{\mathrm{min/max}} = \sum_{B\in\mathcal{D}} \Delta E_B^{\mathrm{min/max}}. \tag{8}$$

General albeit not very tight bounds can be constructed by assuming that all constituents of a box are located at the points of minimal or maximal interaction, respectively.

We aim at finding bounds $\Delta E^{\mathrm{min/max}}$ that are *just* accurate enough to decide about the acceptance or rejection of a proposed update. The general strategy is illustrated with an example progression of the decomposition $\mathcal{D}$ in Fig. 1. For easy visualization it is shown in two dimensions and for a nonmoving component such as an Ising spin. We start with the initial decomposition which is just the box containing the whole system $\mathcal{D} = \{B_0\}$; see Fig. 1(a), where the position of the component q_i to be updated is marked by a red dot. The bounds $\Delta E^{\mathrm{min/max}} = \Delta E_{B_0}^{\mathrm{min/max}}$ of this initial decomposition are in most cases not accurate enough to take the decision about the proposed update.

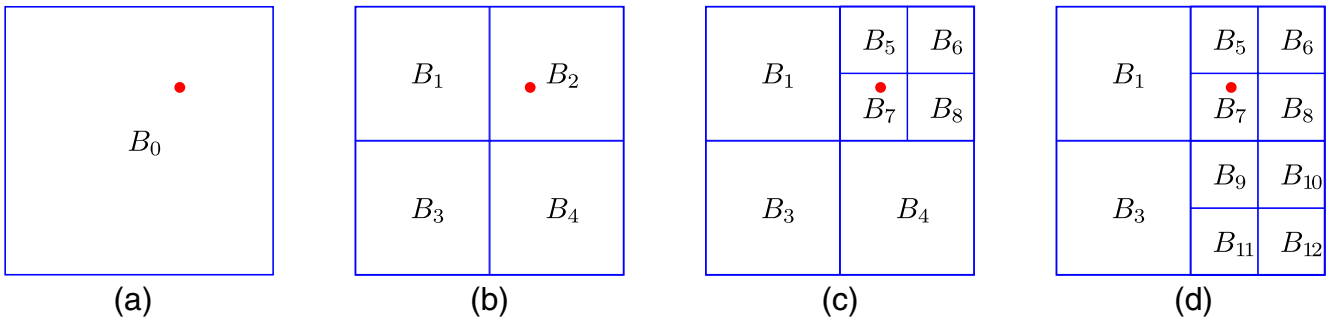

FIG. 1. A visual sketch of an example progression for the decomposition of the interaction. The red dot marks the position of the component for which the update is proposed. The interacting boxes are enumerated in ascending order in which they are placed into the decomposition. From left to right always one box with high uncertainty is split into smaller boxes leading to more accurate bounds for ΔE.

Thus we replace B_0 in $\mathcal{D}$ with the child boxes $B_1, \ldots, B_{2d}$ of B_0 in $\mathcal{T}$; see Fig. 1(b). The bounds $\Delta E^{\min/\max}$ are updated according to Eq. (8). With each pair of bounds $\Delta E_B^{\min/\max}$ there is an associated uncertainty $\Delta_B = \Delta E_B^{\max} - \Delta E_B^{\min}$ for each box in $\mathcal{D}$. The sum of the uncertainties of the newly inserted child boxes is by construction always smaller or equal to the uncertainty of the removed parent box. If the newly obtained bounds are not sufficiently narrow to reach a decision, further boxes are split until a decision can be made.

In order to do so, a strategy is needed to decide which box to split next. It is generally beneficial to split boxes which have a high uncertainty Δ_B, since there the potential for improving the bounds is greatest. A natural approach is, therefore, to always select the box with the greatest Δ_B for splitting. Since searching an unordered set is computationally very expensive, this would require us to keep all elements of $\mathcal{D}$ strictly ordered with respect to their uncertainties and new boxes could only be added to $\mathcal{D}$ in logarithmic time $O(\log|\mathcal{D}|)$. An overall faster way is to group boxes according to the integer part of their logarithmic uncertainties $\delta_B = \lfloor \log_2 \Delta_B \rfloor$ and always select some box from the nonempty set with the highest δ_B. The number of the necessary operations is now independent of the size of $\mathcal{D}$, and a significant computational overhead can be avoided this way.

The process of sequential decomposition of boxes is sketched out in Fig. 1. The box B_2 in Fig. 1(b) had a high uncertainty and was replaced by its four child boxes in $\mathcal{T}$; see Fig. 1(c). The next box to be split was B_4. Such an adaptive spatial decomposition can be performed in most cases: on lattices, on graphs with different geometries, and even in the case of continuous spatial degrees of freedom although the nodes of $\mathcal{T}$ might not always correspond to simple square or cubic boxes.

With this refinement protocol, we have effectively constructed a hierarchical, adaptive (spatial) decomposition of the interactions, which depends strongly on the current configuration, the energy difference due to the proposed update ΔE, and the threshold energy ΔE_{th}. After the decision has been made the next update starts again with $\mathcal{D} = \{B_0\}$, since if a different component is to be updated, the final decomposition will likely be completely different.

III. LONG-RANGE O(n) MODELS

A. Model

We first consider the broad class of long-range interacting O(n) vector spin models on two-dimensional ($d = 2$) $L \times L$ square lattices where the spins $\mathbf{s}_i$ are n-dimensional unit vectors and each spin interacts with every other spin of the system via a decaying power-law potential. For the LRIM ($n = 1$) there have been numerical studies regarding its equilibrium as well as its nonequilibrium properties. Results from simulations for other n, however, are very sparse in literature, which may partly be explained by the lack of a generally applicable and efficient algorithm for their simulation.

This model class is described by the Hamiltonian,

$$\mathcal{H} = -\frac{1}{2}\sum_i \sum_{j\neq i} J_{i,j}\mathbf{s}_i\mathbf{s}_j, \tag{9}$$

where for free boundary conditions the interaction couplings $J_{i,j}^{\text{FBC}}$ decay with distance like

$$J_{i,j}^{\text{FBC}} = r(i,j)^{-(d+\sigma)}. \tag{10}$$

Here, $r(i,j) = |\mathbf{r}(i,j)|$ is the Euclidean distance, d is the spatial dimension, and $\sigma > 0$ is a tunable parameter controlling the decay of the potential. The self-interaction of the individual spins is set to zero; i.e., $J_{i,i} = 0$. To reduce finite-size effects we employ periodic boundary conditions (PBC), implemented via Ewald summation [32], which takes all periodic images of the system into account. For spins on fixed lattice positions the Ewald summation can be incorporated directly into the couplings [33],

$$J_{i,j}^{\text{PBC}} = \sum_{\mu,\nu=-\infty}^{\infty} |\mathbf{r}(i,j) + \mu L\hat{e}_x + \nu L\hat{e}_y|^{-(d+\sigma)}, \tag{11}$$

which can be used in conjunction with the simple minimum image convention, avoiding a significant computational overhead.

For lattices with linear size $L = 2^m$ with m being a positive integer, it is intuitive to decompose the lattice into smaller squares, which for the purpose of hierarchical access are arranged on a quadtree. Starting with the full lattice as the original box B_0 (cf. Fig. 1) which encloses all the spins, we decompose the lattice into four boxes with half the linear size of the original box. This is repeated until reaching the single spin level, where each box contains only a single spin. We denote the level of decomposition as $\Gamma \in \{0, \ldots, m\}$, where $\Gamma = 0$ corresponds to the full lattice and $\Gamma = m$ is the single spin level. The linear size of a box is $L_\Gamma = 2^{m-\Gamma}$ and the number of spins inside this box is $N_\Gamma = L_\Gamma^2 = 4^{m-\Gamma}$.

B. Bounds $\Delta E_B^{\min/\max}$ for the spin-box interactions

Upon updating a spin $\mathbf{s}_i$, the contribution of the box B to the energy difference can be written as

$$\begin{aligned}\Delta E_B &= E_B^{\text{new}} - E_B^{\text{old}} \\ &= -\sum_{j\in B} J_{i,j}\mathbf{s}_i^{\text{new}}\mathbf{s}_j + \sum_{j\in B} J_{i,j}\mathbf{s}_i^{\text{old}}\mathbf{s}_j,\end{aligned} \tag{12}$$

where we identify the box B with the set of the indices of the contained spins. One can rewrite Eq. (12) as

$$\Delta E_B = -(\mathbf{s}_i^{\text{new}} - \mathbf{s}_i^{\text{old}}) \sum_{j \in B} J_{i,j} \mathbf{s}_j$$
$$= -\Delta \mathbf{s}_i \sum_{j \in B} J_{i,j} \mathbf{s}_j = \sum_{\alpha=1}^{n} \Delta E_{B,\alpha}. \tag{13}$$

The total energy difference is equal to the sum of the energy differences from each of the n components. For a lattice spin system with PBC, the set of couplings which are involved in a spin-box interaction is always uniquely determined by the vector $\mathbf{R}$ from the spin to the center of the box and the size of the box N_Γ. In order to construct bounds $\Delta E_B^{\min/\max}$, we need the minimal and maximal couplings in this set $J_\Gamma^{\min/\max}(\mathbf{R})$, which for monotonically decaying couplings can be calculated in constant complexity $O(1)$ and can, therefore, be determined on the fly. For the more advanced bounds we additionally need the integrated interaction,

$$J_\Gamma^{\text{int}}(\mathbf{R}) = \sum_{j \in B} J_{i,j}, \tag{14}$$

which corresponds to the total interaction strength with a fully magnetized box B.

In the following subsections we devise several bounds for the spin-box interaction. We start with the specific case of the LRIM, which is a good starting point because of the simpler form of the involved equations due to the scalar nature of the interaction. Following the same strategy, we then derive the corresponding expressions for general O(n) vector spin models.

1. Bounds for the LRIM

For the LRIM, the contribution of the box B to the energy difference Eq. (13) simplifies to

$$\Delta E_B = 2 s_i^{\text{old}} \sum_{j \in B} J_{i,j} s_j, \tag{15}$$

where $\Delta s_i = s_i^{\text{new}} - s_i^{\text{old}} = -2 s_i^{\text{old}}$ is used.

For bounds 1 we do not discriminate between spins pointing up or down and use only the number of spins N_Γ contained in the box. It is clear that for each summand of Eq. (15) $-J_\Gamma^{\max}(\mathbf{R}) \le J_{i,j} s_i^{\text{old}} s_j \le J_\Gamma^{\max}(\mathbf{R})$ holds. This allows us to formulate the following bounds for Eq. (15),

$$\Delta E_B^{\max} = -\Delta E_B^{\min} = 2 N_\Gamma J_\Gamma^{\max}(\mathbf{R}), \tag{16}$$

which are plotted in Fig. 2 as horizontal dash-dotted lines as a function of N_B^+/N_Γ for one example situation, where N_B^+ is the number of positive spins in the box which is related to the magnetization of the box

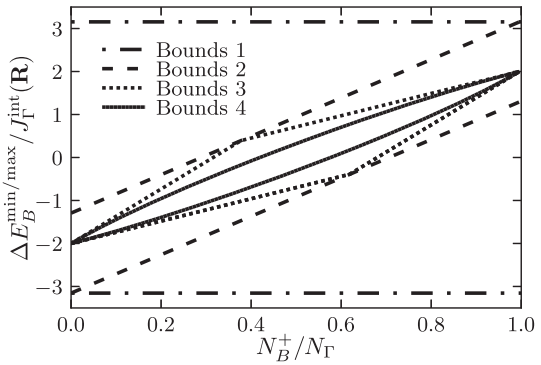

FIG. 2. Comparison of the different bounds $\Delta E_B^{\min/\max}$ for the interaction of a test spin pointing in positive direction with a box B, normalized by the total interaction strength J_Γ^{int} defined in Eq. (14). A pair of bounds always consists of an upper and a lower bound which are kept in the same style. N_B^+ is the number of spins pointing up and N_Γ the total number of spins in the box. Bounds 1 are given by $2N_\Gamma$ multiplied with the maximum interaction. Bounds 2 are refined by making use of the magnetization of the box. Bounds 3 additionally employ the value of the interaction with the fully magnetized box $J_\Gamma^{\text{int}}(\mathbf{R})$. Bounds 4 are the tightest bounds which can be obtained with the knowledge of the magnetization and the set of interactions involved in a spin-box interaction. For more information, see text.

$M_B = N_B^+ - N_B^- = 2N_B^+ - N_\Gamma$. The parameters for this example are linear lattice size $L = 128$, decay exponent of the potential $\sigma = 0.8$, distance $\mathbf{R} = (50.5, 44.5)$ from the spin to the center of the box B, and box size $N_\Gamma = L_\Gamma^2 = 256$. For other parameter values the curves in Fig. 2 would look different, but the main features would remain unchanged. Since these simple bounds do not make use of the box magnetization M_B, they are constant and much wider than the more refined bounds considered next.

In contrast, bounds 2 do depend on the magnetization M_B in a box B. To make use of M_B we split the box B into the sets of indices of spins pointing up B^+ or down B^-,

$$B = B^+ \cup B^-, \tag{17}$$

to rewrite Eq. (15) as

$$\Delta E_B = 2 s_i^{\text{old}} \left(\sum_{j \in B^+} J_{i,j} - \sum_{j \in B^-} J_{i,j} \right). \tag{18}$$

For each box of the spatial decomposition we keep track of the number of elements of $B^\pm$, i.e., the number of spins pointing up or down $N_B^\pm$. For the sums in Eq. (18) it follows that

$$N_B^\pm J_\Gamma^{\min}(\mathbf{R}) < \sum_{j \in B^\pm} J_{i,j} < N_B^\pm J_\Gamma^{\max}(\mathbf{R}). \tag{19}$$

If s_i^{old} points into the positive direction, the rhs of Eq. (18) becomes maximal if the first term is maximal and the second one is minimal and vice versa for its minimum. This yields the following lower and upper bounds for the spin-box interaction:

$$\begin{aligned}\Delta E_B^{\mathrm{min}} &= 2[N_B^+ J_\Gamma^{\mathrm{min}}(\mathbf{R}) - N_B^- J_\Gamma^{\mathrm{max}}(\mathbf{R})],\\ \Delta E_B^{\mathrm{max}} &= 2[N_B^+ J_\Gamma^{\mathrm{max}}(\mathbf{R}) - N_B^- J_\Gamma^{\mathrm{min}}(\mathbf{R})].\end{aligned} \tag{20}$$

If the test spin s_i^{old} points in the negative direction, N_B^+ and N_B^- have to be exchanged accordingly. In this picture the upper bound $\Delta E_B^{\mathrm{max}}$ corresponds to the situation where all spins pointing in the same direction as the test spin are placed at the position of the strongest interaction and the spins pointing in the opposite direction at the spot of the weakest interaction. For the lower bound $\Delta E_B^{\mathrm{min}}$ the positions are switched. Looking at the dashed lines in Fig. 2, one observes that these bounds depend on the box magnetization M_B and their uncertainty Δ_B is much smaller than for the rather loose bounds 1.

For bounds 3 we use $J_\Gamma^{\mathrm{int}}(\mathbf{R})$ from Eq. (14). Since the $J_\Gamma^{\mathrm{int}}(\mathbf{R})$ do not depend on the spin configuration they can be precalculated. The required computational effort is negligible compared to the typical simulation time, whereas the memory demands scale as $O(N \log N)$. This can become challenging for systems that are significantly larger than those we consider here. Using $J_\Gamma^{\mathrm{int}}(\mathbf{R})$, Eq. (15) can be rewritten as

$$\Delta E_B = 2 s_i^{\mathrm{old}} \left(J_\Gamma^{\mathrm{int}}(\mathbf{R}) - 2 \sum_{j \in B^-} J_{i,j} \right). \tag{21}$$

This equation shows that the interaction of the test spin with a box with some arbitrary configuration can be seen as the sum of the interaction with the fully magnetized box where all spins point in the same direction as the test spin and twice the interaction of the spins inside the box which point in the opposite direction. Alternatively, the interaction can also be calculated using the interaction with the fully magnetized box with all spins pointing in the opposite direction of the test spin and adding twice the interaction with the spins parallel to the test spin:

$$\Delta E_B = 2 s_i^{\mathrm{old}} \left(2 \sum_{j \in B^+} J_{i,j} - J_\Gamma^{\mathrm{int}}(\mathbf{R}) \right). \tag{22}$$

In order to derive two new pairs of bounds for the spin-box interaction we can again use Eq. (19), inserting it into Eq. (21) and Eq. (22). As both pairs of bounds are valid, we can combine the two bounds for the minimum and the two bounds for the maximum by taking the tighter of the two and obtain (again for $s_i^{\mathrm{old}} = 1$):

$$\begin{aligned}\Delta E_B^{\mathrm{min}} &= 2 \max[J_\Gamma^{\mathrm{int}}(\mathbf{R}) - 2N_B^- J_\Gamma^{\mathrm{max}}(\mathbf{R}),\\ &\qquad 2N_B^+ J_\Gamma^{\mathrm{min}}(\mathbf{R}) - J_\Gamma^{\mathrm{int}}(\mathbf{R})],\\ \Delta E_B^{\mathrm{max}} &= 2 \ \min[J_\Gamma^{\mathrm{int}}(\mathbf{R}) - 2N_B^- J_\Gamma^{\mathrm{min}}(\mathbf{R}),\\ &\qquad 2N_B^+ J_\Gamma^{\mathrm{max}}(\mathbf{R}) - J_\Gamma^{\mathrm{int}}(\mathbf{R})].\end{aligned} \tag{23}$$

These bounds are exact for fully magnetized boxes and still very accurate for almost fully magnetized ones (see the dotted lines in Fig. 2). This greatly enhances the performance of the algorithm in the presence of large magnetic domains, because the bounds of the boxes which fully lie inside a domain are very accurate and thus the decomposition can be coarser. At very large distances $|\mathbf{R}| \gg L_\Gamma$ we find $J_\Gamma^{\mathrm{int}} \approx N_\Gamma (J_\Gamma^{\mathrm{min}} + J_\Gamma^{\mathrm{max}})/2$, so that the crossing of the pairs of bounds for the minimum or the maximum would occur at $N_B^+/N_\Gamma \approx 0.5$.

Exploiting the fact that the set of $J_{i,j}$ can be sorted, we can obtain even narrower bounds 4. For the second term in the parentheses of Eq. (21) we replace the bounds from Eq. (19), for which we assume that all spins interact with $J_\Gamma^{\mathrm{min/max}}(\mathbf{R})$, with the sum of the first N_B^- of the couplings sorted in ascending or descending order. These are the tightest bounds which can be established using only the magnetization of the box and the set of couplings involved in the spin-box interaction. These bounds are plotted in Fig. 2 as solid lines and as the previous bounds can be calculated in $O(1)$ complexity if the abovementioned sums are all computed and stored before the simulation. The memory complexity of the algorithm using these bounds would scale as $O(N^2 \log N)$, however, which limits the applicability for large system sizes. Therefore, in the following we employ bounds 3 for the LRIM which embody the best compromise between performance and memory requirements. For the iterative refinement of the decomposition of the interaction, we proceed as described in the general outline of the algorithm, with one modification. It turns out to be beneficial to evaluate the interactions with small boxes via a direct summation of the spin-spin interactions since this is of comparable speed and has no uncertainty.

2. Bounds for general O(n) models

In this section we propose a more general formulation for bounds 1 and bounds 2 from the last section, valid for all O(n) models with $n > 1$.

For bounds 1 the generalization is straightforward since for each spin-spin interaction in Eq. (13) one has again $-2J_\Gamma^{\mathrm{max}}(\mathbf{R}) \leq J_{i,j} \Delta \mathbf{s}_i \mathbf{s}_j \leq 2J_\Gamma^{\mathrm{max}}(\mathbf{R})$. This results in exactly the same bounds as for the Ising model:

$$\Delta E_B^{\mathrm{max}} = -\Delta E_B^{\mathrm{min}} = 2 N_\Gamma J_\Gamma^{\mathrm{max}}(\mathbf{R}). \tag{24}$$

For bounds 2 the box is again split up, however, for each of the n components separately, into spins whose

component $s_{j,\alpha}$ points into positive or negative direction, respectively,

$$B = B_\alpha^+ \cup B_\alpha^-, \tag{25}$$

where for each α the decomposition may be different. This allows us to split the contribution $\Delta E_{B,\alpha}$ to the energy difference as well:

$$\Delta E_{B,\alpha} = -\Delta s_{i,\alpha}\left[\sum_{j\in B_\alpha^+} J_{i,j}|s_{j,\alpha}| - \sum_{j\in B_\alpha^-} J_{i,j}|s_{j,\alpha}|\right]. \tag{26}$$

During the simulation we keep track of the componentwise magnetization split into positive and negative contributions:

$$M_{B,\alpha}^\pm = \sum_{j\in B_\alpha^\pm} |s_{j,\alpha}|. \tag{27}$$

Since there is no such relation as $N_B^+ + N_B^- = N_\Gamma$, valid only for the LRIM, both directions need to be tracked for each component, resulting in a total of $2n$ values $M_{B,\alpha}^\pm$ which need to be stored for each box B. The two sums in Eq. (26) are bound by

$$M_{B,\alpha}^\pm J_\Gamma^{\min}(\mathbf{R}) < \sum_{j\in B_\alpha^\pm} J_{i,j}|s_{j,\alpha}| < M_{B,\alpha}^\pm J_\Gamma^{\max}(\mathbf{R}). \tag{28}$$

If $\Delta s_{i,\alpha} < 0$ (which corresponds to the case where s_i^{old} points into positive direction in the LRIM), Eq. (26) becomes minimal (maximal) if the first sum becomes minimal (maximal) and the second sum becomes maximal (minimal). This yields the following bounds:

$$\begin{aligned}\Delta E_{B,\alpha}^{\min} &= |\Delta s_{i,\alpha}|[M_{B,\alpha}^+ J_\Gamma^{\min}(\mathbf{R}) - M_{B,\alpha}^- J_\Gamma^{\max}(\mathbf{R})],\\ \Delta E_{B,\alpha}^{\max} &= |\Delta s_{i,\alpha}|[M_{B,\alpha}^+ J_\Gamma^{\max}(\mathbf{R}) - M_{B,\alpha}^- J_\Gamma^{\min}(\mathbf{R})].\end{aligned} \tag{29}$$

If $\Delta s_{i,\alpha}$ is positive, $M_{B,\alpha}^+$ and $M_{B,\alpha}^-$ have to be exchanged accordingly. This componentwise formulation of bounds 2 for general $\mathrm{O}(n)$ vector spin models is completely analogous to bounds 2 [cf. Eq. (20)] for the LRIM.

The complete bounds for the spin-box interaction follow as the sum of the bounds for the individual components. As in the case of the LRIM, these bounds are formulated in terms of the extremal values of the positive and the negative contributions to ΔE_B. In addition to this rather intuitive construction of the bounds, we present a more formalistic derivation in the Supplemental Material [34]. The framework developed there, including all the above presented bounds and, additionally, a formulation of bounds 3 for general $\mathrm{O}(n)$ models, may also be useful for the construction of bounds for completely different models.

C. Example decomposition

In Fig. 3 we demonstrate the basic principle of our algorithm in application to the LRIM using bounds 3 by showing a single example snapshot and the corresponding spatial decomposition of the interaction for a simulation with $\sigma = 0.8$ and $L = 256$ at $T = T_c$ in equilibrium (the system is chosen to be relatively small, so that details can still be observed). Here, the spin under consideration is positioned close to the center of the snapshot and is marked in red. In the vicinity of the test spin—the green shaded region—maximal resolution is reached: all the boxes are of size 1×1 and contain only a single spin each. To ensure the possibility of arbitrarily precise estimation of ΔE, interactions with these spins have to be considered exactly irrespective of which bounds are used, although we note that this is formally equivalent to the use of bounds 2–4.

From top left to bottom right, the estimates of $\Delta E_{\max}$ and $\Delta E_{\min}$ are more refined and approach ΔE. As one can see, this is achieved by reducing the box sizes. The decomposition adapts to the given configuration; i.e., regions which have a bigger influence on the decision are covered by smaller boxes. The scenario in Fig. 3 requires a rather fine-grained decomposition, but this is not representative and usually decisions can be made with a much coarser decomposition. It is nonetheless a very illustrative example, since it nicely demonstrates the progression of the algorithm. Looking at a decomposition for other $\mathrm{O}(n)$ would yield a similar picture with the direct neighborhood of the test spin resolved in detail and the decomposition of more distant regions composed of larger boxes.

D. Analysis of the run times

The speed of our algorithm strongly depends on the state of the simulation, i.e., the current configuration, the choice of the spin to be updated, the temperature, etc. The decision becomes harder the closer the involved change in energy ΔE is to the threshold energy ΔE_{th}. For the extreme case of $\Delta E = \Delta E_{\text{th}}$ the interaction of the updated spin with all other individual spins has to be evaluated exactly, implying a *worst case* complexity of the algorithm of $O(N^2)$ per sweep. However, ΔE and ΔE_{th} are independent random variables. Their respective distributions are different and typically have a nonzero spread, such that the influence of the *worst case* complexity of the algorithm on its *average* complexity will in most cases be negligible. In the case of high temperatures, one has $\Delta E_{\text{th}} \to \infty$, which means that the actual change in energy of the proposed update becomes irrelevant, so that the initial bounds are sufficiently narrow to accept it. Only the updating of the tree $\mathcal{T}$ has to be performed, which yields the *best case* complexity $O(N \log N)$.

It is not straightforward to predict the *average* complexity from the algorithm's design alone, so that, in the following, we record the resulting run times τ per sweep for different physical settings. For all measurements

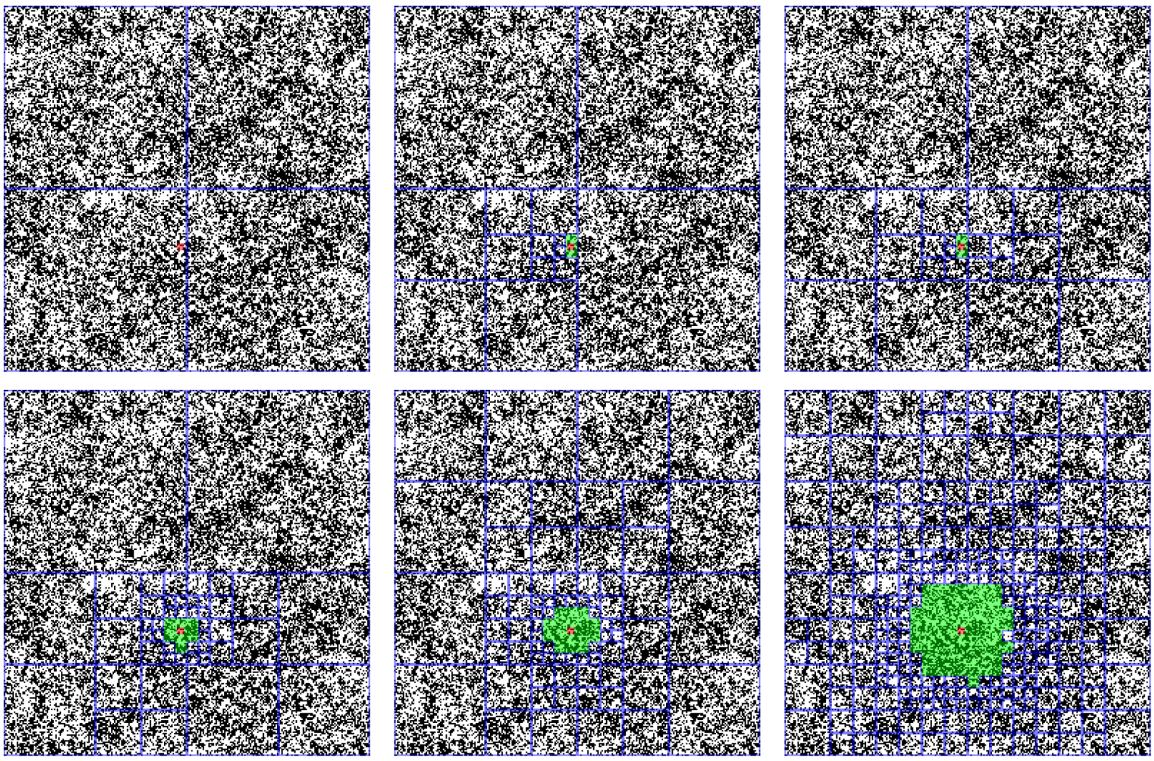

FIG. 3. Example decomposition of the LRIM lattice with $\sigma = 0.8$ and $L = 256$ simulated in equilibrium at $T = T_c$. The spin which is to be updated is marked by the red dot close to the center of the box. The boundaries of the boxes are represented as blue lines. Going from top left to bottom right, the accuracy of the bounds of $\Delta E^{\min}$ and $\Delta E^{\max}$ increases with shrinking size of the boxes in the decomposition. Boxes of size 8×8 are broken up directly into 1×1 boxes (green area); 2×2 and 4×4 boxes therefore do not occur.

we take care of minimizing possible hardware-related influences. To be on the safe side and to not draw any wrong conclusions in our analysis, we nonetheless assume an error of 10% on our estimated run times to account for remaining errors. For equilibrium simulations, the computational cost can relatively straightforwardly be extracted from rather short runs, which allows us to consider a broad temperature range. A further advantage is that the run times only weakly depend on the initial conditions and can be averaged over the full run after the equilibration. The benchmarking is performed on a single hardware configuration: single socket motherboard equipped with an Intel Core i5-8500T CPU and 16 GB DDR4-2667 dual-channel RAM. The algorithm is implemented in C++17 and compiled using GCC8.3. Since modern processors do not run using a constant frequency and use speculative execution, results of benchmarks can fluctuate. This is especially a problem if the compute nodes are occupied by other tasks. Therefore, we take care to exclusively run a single simulation at a time per compute node in order to minimize possible fluctuations.

1. LRIM

We first turn to the LRIM, which is the most prominent representative of the long-range interacting $\mathrm{O}(n)$ models, for which we use bounds 3 throughout, since their use yields the best performance. We have checked that using bounds 2 instead comes along with moderate, size-independent losses of factors 2–5 for all the scenarios except the equilibrium run times for very low temperatures, where the frequent occurrence of strongly magnetized boxes yields a more pronounced advantage for bounds 3, almost completely eliminating any system-size dependence of the run times.

Equilibrium.—In Fig. 4 the equilibrium run times per spin update τ/N (in units of μs) in dependence of the system size are presented for different fractions of T_c in a semilog plot for $\sigma = 0.8$ [Fig. 4(a)] and $\sigma = 1.5$ [Fig. 4(b)]. In both cases the growth of the run times crosses over to linear behavior on the semilog scale irrespective of the temperature T, in a manner compatible with $O(N \log N)$ complexity. This we deem very plausible considering the hierarchical progression of the algorithm through the use of a tree. Based on this data (and Fig. 5 where we plot the run times on a log-log scale), a power-law complexity $O(N^{1+\alpha})$ with a small exponent α cannot, however, be completely ruled out. In order to corroborate either of the two hypotheses, significantly larger systems need to be considered. While a dedicated investigation of moderately larger system sizes could still be feasible in principle, significantly larger sizes are out of reach with the hardware used in this study, due to the large memory requirements,

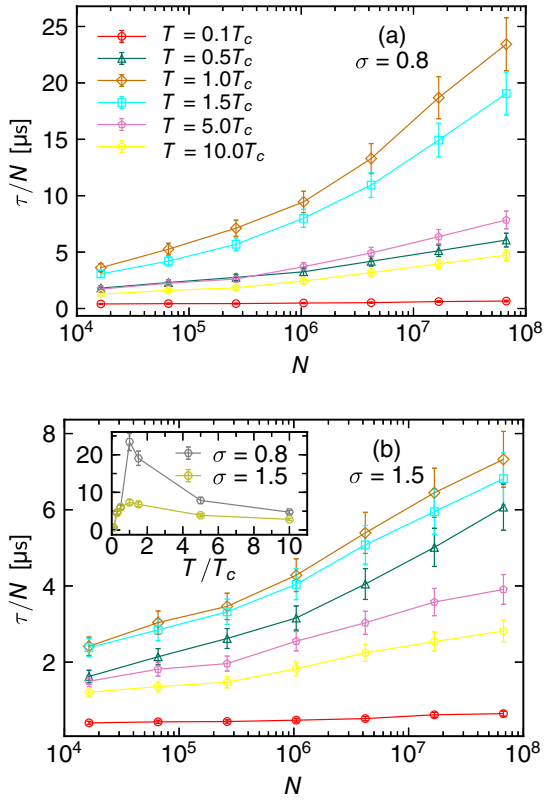

FIG. 4. Equilibrium run times per spin update τ/N for the LRIM versus system size N on a semilog scale for (a) $\sigma = 0.8$ and (b) $\sigma = 1.5$ and several temperatures T. In the inset of (b) we show the combined data for both values of σ and the biggest system size $L = 8192$ as a function of T, demonstrating the maximum around T_c.

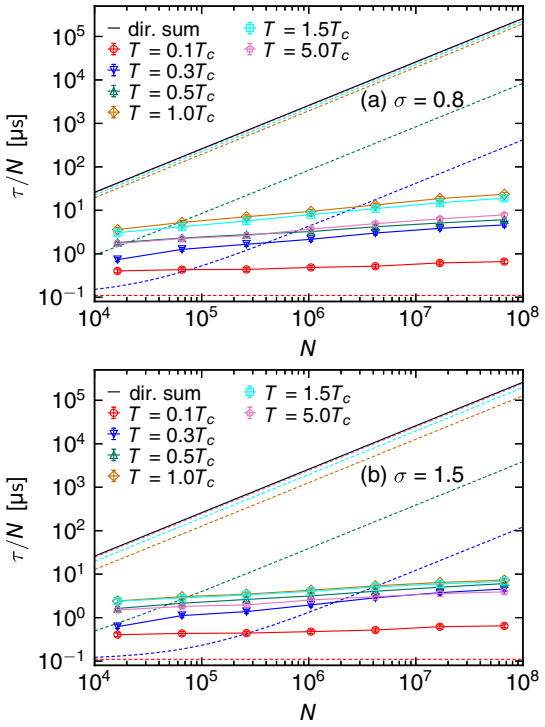

FIG. 5. Equilibrium run times per spin update τ/N for the LRIM versus system size N plotted on a log-log scale for (a) $\sigma = 0.8$ and (b) $\sigma = 1.5$ and several temperatures T. As dashed lines in the same color as the original data points, we have included the calculated run times for an effective field simulation. Also included is the run-time obtained from a naive Metropolis MC simulation using direct summation (dir. sum).

and a final assessment of the asymptotic scaling has to be left for future studies.

We find that there is a clear dependence of the equilibrium run times on the simulation temperature, which is also visualized in the inset of Fig. 4(b). This can be understood qualitatively from the considerations made before: The maximum close to the critical temperature stems from the fact that the average threshold energy E_{th} is on average close to the actual energy difference ΔE involved in the proposed spin flips, which makes a fine decomposition of the interaction necessary. In the case of very high temperature one has $\Delta E_{\mathrm{th}} \gg \Delta E$, and the actual change in energy of the proposed spin flip becomes irrelevant. At low temperatures an opposite mechanism is at work. Since $\Delta E_{\mathrm{th}} \to 0$, almost only spin flips are accepted that do not increase the energy, but a typical configuration at these temperatures is (nearly) ordered. Thus, a single spin flip on average has $\Delta E \gg \Delta E_{\mathrm{th}}$, so that also this decision can be made with loose bounds.

The comparison most relevant is likely the achieved speedup factor, which can be visually appreciated from Figs. 5(a) and 5(b), where we show the run times on a log-log scale in comparison to an effective field simulation [35] (dashed lines in the same color as the original data points) and a direct summation (solid black line) of all interactions (for the sake of better comparability, here we replace $T = 10T_c$ by $T = 0.3T_c$). The effective field approach uses a relatively simple storage trick to save many calculations whenever an update is *not* accepted. This yields a massive speedup whenever low acceptance rates are encountered, but requires the same number $O(N)$ of operations for an accepted update as in the direct summation, resulting in the same computational complexity $O(N^2)$ [36]. Thus, the resulting run times are strongly dependent on the acceptance rate of the simulation, which means that this approach is especially fast at low temperatures where the acceptance rates are low [37]. At high temperatures, many more proposed updates are accepted, which gives the effective field approach only a minor advantage over a direct summation. At the critical temperature (for which our algorithm has the highest run-time), we can report a speedup factor of ≈5500 for $\sigma = 0.8$ and ≈12 000 for $\sigma = 1.5$ compared to the effective field approach (≈11 000 and ≈35 000, respectively, compared to direct summation). For a temperature below T_c, e.g., for $T = 0.5T_c$, we observe a speedup of ≈1700 for $\sigma = 0.8$ and ≈500 for $\sigma = 1.5$ (≈40 000 and ≈43 000, respectively,

compared to direct summation). For increasing system size all these factors grow steadily. For each fixed temperature there will be a crossover size above which our new algorithm is faster than the effective field method. From Fig. 5 one reads off that the crossover for $T = 0.3T_c$ occurs at $N \approx 10^6$ for the considered values of σ and, respectively, $N \approx 10^5$ for $T = 0.5T_c$. For larger values of T the crossover happens already for much smaller systems while for the very low temperature $T = 0.1T_c$ the crossover cannot yet be observed for the considered system sizes of up to $N = 10^8$ due to the extremely low acceptance rates ($\approx 10^{-6}$) and hence a very small prefactor of the run times of the effective field approach. For this temperature the size dependence of the run times vanishes almost completely, which can be explained by the exactness of bounds 3 for fully magnetized boxes.

For investigations of the phase transition, i.e., for equilibrium simulations in the proximity of T_c, our method, like any other local algorithm, is not a serious contender, at least when nonlocal cluster algorithms are available as for the Ising model. In such cases, its main field of application is nonequilibrium scenarios.

Nonequilibrium.—We first focus on two cases: quenches from a disordered start configuration (i) to a temperature substantially below the critical temperature or (ii) to the critical temperature. To mimic a physical evolution, in these cases only local dynamics that preserves the dynamical properties of the system may be used. Nonlocal update schemes, including cluster algorithms, are not allowed, which makes these scenarios the prime field of application for our algorithm. In the first case, the system undergoes an ordering process and consequently the dynamics of growth of ordered structures in the system is of interest, both from coarsening and aging perspective involving single- and two-time quantities, respectively [22–24]. The physical properties of this model during this process have recently been investigated in Refs. [35,38–40] and are not part of the discussion here. In Fig. 6(a) we present the time dependence of the run-time of our algorithm per spin update τ/N (in units of μs) as a function of simulation time for a phase-ordering quench with $\sigma = 0.8$ to $T_q = 0.1T_c$ for large systems of linear sizes $L = 4096$ and 8192. As in a typical production run, the quoted times are averages over many different start configurations and time evolutions. Also shown, for sake of comparison, are the run times for the effective field method [35] and a direct summation of the interactions. We could quench to arbitrary temperatures with our new algorithm, but choose $T_q = 0.1T_c$ to have the same setting as in Ref. [35].

We observe that the time needed per update τ/N is strongly dependent on how far the system has proceeded in its ordering process for our algorithm and, even more, for the effective field approach. Since the temperature in our algorithm is set to a low temperature $T_q = 0.1T_c$, we typically draw threshold energies E_{th} comparatively close to zero; i.e., spin flips which significantly increase the energy are usually not accepted. At the start of our simulation the configuration is completely disordered, and for many proposed spin flips $\Delta E \approx 0 \approx \Delta E_{\text{th}}$, so that ΔE has to be known rather accurately. In the course of the simulation, when the configurations are locally ordered, proposed spin flips in domains have typically a large ΔE and can mostly be rejected with very loose bounds. The decision is more involved for spins positioned at domain boundaries, where ΔE is typically much closer to zero, the interactions often have to be resolved in more detail, and the probability of acceptance is higher. With growing domains fewer spins are situated at the domain boundaries so that the average acceptance rate decreases, and the effective field simulations become faster. Ultimately, the run times of both methods reach their respective equilibrium run times at T_q.

Over the full process, our new approach is significantly faster than the already very fast effective field method, resulting in a ≈100 times faster total run-time until finite-size effects are reached. A direct summation becomes in these cases prohibitively expensive (a factor of ≈40 000 slower than our new algorithm), and cannot be used to simulate systems of this size. This factor grows significantly for increasing system size since the run times of the two algorithms scale differently.

The second case of interest is critical aging, i.e., the behavior of two-time quantities during quenches from a disordered starting configuration to the critical temperature [24,25]. With some modifications, such as a small initial magnetization, it is also possible to investigate short-time dynamics during such processes [41], but we here focus on the completely disordered start. We present in Fig. 6(b) the obtained run times per spin update τ/N (again in units of μs) in our simulation with $\sigma = 0.8$, where we use the same notation as in Fig. 6(a) for the different methods. Here, the run times both for the new algorithm and the effective field approach remain more or less constant throughout the whole simulation, although the system has not yet reached equilibrium. In both cases only a small initial decay of the run times is visible. Here, the advantage of the effective field simulation over a direct summation is relatively small, since the acceptance rates are of the order of 1. Albeit the equilibrium simulation close to the critical temperature is also one of the most difficult situations for our algorithm, it nonetheless produces much smaller run times than the effective field approach and the direct summation. In this scenario the run times are close to those found in the equilibrium simulations. We find a speedup of ≈6000 compared to the effective field approach and ≈8000 to the direct summation, allowing for the investigation of this process for the presented system sizes, which was entirely out of reach before. A similar acceleration is also expected for other nonequilibrium simulations with comparably high acceptance rates.

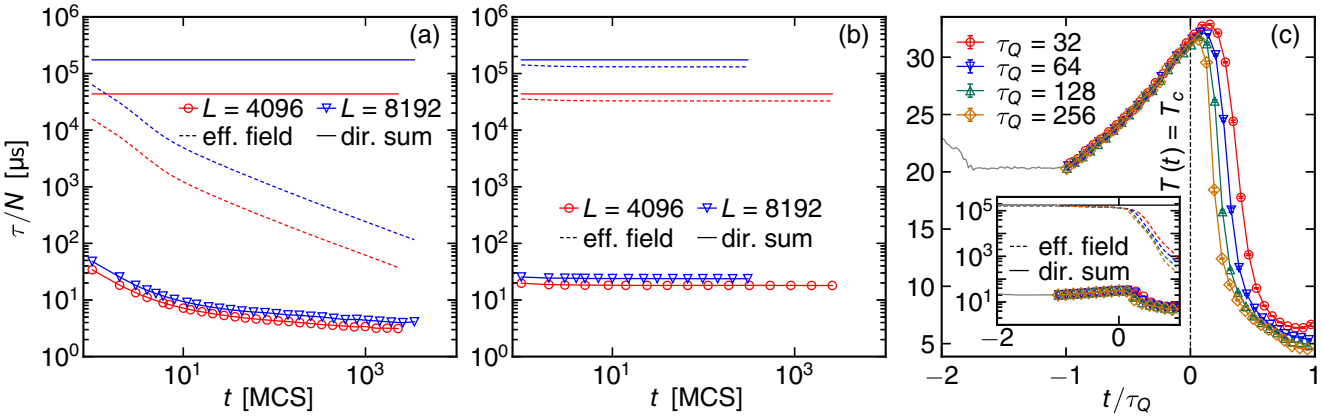

FIG. 6. Run times per spin update τ/N for the LRIM versus simulation time t for $\sigma = 0.8$ in quenches from a disordered start to (a) $T_q = 0.1T_c$ and (b) $T_q = T_c$ on large lattices with $L = 4096$ and 8192 and (c) run times per spin update for simulation with $L = 8192$ using a Kibble-Zurek protocol. The dashed lines in the same color show the run times calculated for the effective field (eff. field) algorithm and the solid lines without markers are the run times of the direct summation (dir. sum).

The third nonequilibrium process which we consider is the Kibble-Zurek mechanism (KZM). This is a phenomenon which is well accessible experimentally for the case of long-range interacting quantum systems. Here, also in the experiments the exponent σ of the long-range interaction can be tuned to a desired value [42–44]. Since it has deep relations to the physics of quantum annealers [45–47], with the most prominent example being the D-Wave quantum computer [48,49], it has recently attracted a lot of interest. Here, we adopt the protocol described in Ref. [50] where the KZM was investigated for the short-range XY model. The system is first equilibrated at $T = 2T_c$ and then annealed with different cooling rates τ_Q according to the schedule $T(t) = T_c(1 - t/\tau_Q)$ for $t \in [-\tau_Q, \tau_Q]$. The run times per update for the KZM for $L = 8192$ are shown in Fig. 6(c). The gray lines show the timings for the equilibration process which could in principle also be performed using nonlocal update schemes and which is therefore excluded from the comparison of the run times. After the equilibration period with nearly constant speed the run times ramp up, while staying very similar among the different cooling rates. Close to the time where the critical temperature is crossed, the run times reach a maximum (corresponding roughly to the equilibrium run times at T_c) After the crossing the broken adiabaticity becomes visible in the run times of the algorithm: the run times of the simulations with faster annealing schedules (small τ_Q) lag behind the run times of the simulations with the slower schedules. The same effect is visible for the run times of the effective field method which are shown in the inset, in comparison to the run times of the new algorithm, as dashed lines with the same color code for τ_Q. Here, again, the run times are calculated solely from measured acceptance rates during the annealing, illustrating that the acceptance rates also reflect the out-of-equilibrium nature of the KZM, which explains the annealing-rate-dependent run times of the effective field method. On average we find a speedup of more than 5000 compared to the effective field and more than 9000 compared to direct summation.

Another local algorithm for long-range interacting spin systems is the clock MC method [20] based on the factorized Metropolis filter [17,19]. It, too, is potentially applicable in the scenarios discussed above. Yet, so far it has only been applied to *disordered* long-range interacting spin systems in *equilibrium*. We have implemented this method for the ferromagnetic LRIM and tested it for quenches to low temperatures. In its basic form where spin-spin interactions are considered individually, we find that the times of crossover to the asymptotic scaling behavior become prohibitively large. This is due to drastically reduced acceptance rates of the factorized Metropolis filter, when compared to conventional Metropolis dynamics. In the framework of this method there is, however, the possibility to treat multiple factors collectively. This shifts the dynamics toward traditional Metropolis, increasing the acceptance rate, but also the computational effort. For each physical setting (i.e., combination of T, L, σ, etc.) a different grouping of spins may yield the best performance. A detailed analysis and comparison is beyond the scope of this study and will be presented elsewhere [51].

2. LRXYM

As a concrete example for O(n) vector spin models we consider the two-dimensional long-range XY model ($n = 2$) for $L \in [128, 8192]$. Recently this model has attracted a considerable interest in analytical work. It has a much richer phase diagram than its short-range counterpart with different types of phase transition depending on the exponent σ of the power-law interaction. There can be transitions between the paramagnetic, the ferromagnetic, and the Berezinskii-Kosterlitz-Thouless phase [21]. For intermediate values of $\frac{7}{4} < \sigma < 2$ the system is predicted to

undergo both a BKT transition at T_{BKT} as well as a symmetry-breaking transition at $T_c < T_{\mathrm{BKT}}$ [21]. For $\sigma > 2$ it is expected to fall into the same universality class as its short-range counterpart. Simulation results, however, are very sparse in literature, and to our knowledge the model has previously only been simulated on diluted graphs [52,53]. There exists a long-standing theoretical prediction for the growth of the characteristic length scale of the LRXYM [54] which has never been tested. Especially for the intermediate regime it is very interesting to investigate if this prediction holds true and if there is any influence of the quench temperature (and of the phase in which the system equilibrates) in this case.

The spins in the XY model are two-dimensional unit vectors whose state can be represented by an angle $\phi \in [0, 2\pi]$. Here, an update consists of the proposal of a random angle $\phi \in [0, 2\pi]$ for a randomly chosen spin and the decision about the acceptance which is performed by the new algorithm using bounds 2 as derived for the case of general O(n) models. For a meaningful benchmarking of the algorithm we need the σ-dependent transition temperatures. As we could not find any values in the literature, we roughly estimate the transition temperatures by means of short-time dynamics [55,56] for the different values of σ considered here for which we obtain $T_c(\sigma = 0.8) \approx 4.87$, $T_c(\sigma = 1.5) \approx 2.74$, and $T_{\mathrm{BKT}}(\sigma = 2.5) \approx 1.66$, respectively.

Equilibrium.—As for the LRIM we again look at the equilibrium run times in dependence of the system size for different fractions of T_c and several values of σ. In addition to $\sigma = 0.8$ and $\sigma = 1.5$, we include $\sigma = 2.5$ since there the system at low temperatures is in the BKT phase which potentially may have influence on the run times. The run times are plotted in Fig. 7 where, compared to the LRIM, we see an earlier and more clear-cut crossover to the expected logarithmic scaling at $N \approx 2.5 \times 10^5$ [57]. Another difference to the run times of the LRIM are the comparably higher run times for $T = 0.1T_c$, which can be blamed on the use of bounds 2 which, in contrast to bounds 3, do not become arbitrarily accurate for fully magnetized boxes. In general, the temperature dependence of the run times seems less pronounced. The speculation that the nature of the low-temperature phase (BKT or ferromagnetic) may manifest in the run times of the algorithm, however, is not confirmed. For $L = 8192$ and the same values of σ, the run times for temperatures close to T_c are less than a factor of 2 larger than in the LRIM, while for temperatures farther away (excluding $T = 0.1T_c$), this factor can reach up to 3.

Compared to the LRIM, the effective field method for the LRXYM is much less efficient at low temperatures since the acceptance rates even at $T = 0.1T_c$ stay rather high ≈ 0.1 (as compared to $\approx 10^{-6}$ for the LRIM) yielding only a speedup factor of ≈ 10 compared to a direct summation. In conclusion, for $L = 8192$ and the considered values of σ and T, we find speedups of the order of 2000–20 000, which are comparable to the speedups observed for the LRIM.

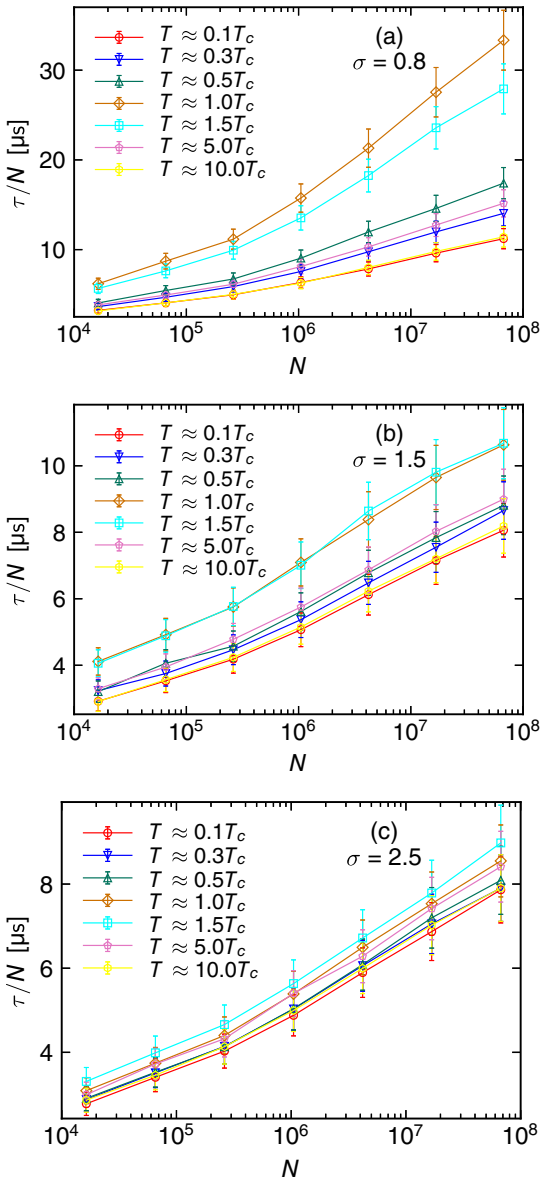

FIG. 7. Equilibrium run times per spin update τ/N for the LRXYM versus system size N plotted on a semilog scale for (a) $\sigma = 0.8$, (b) $\sigma = 1.5$, and (c) $\sigma = 2.5$ and several temperatures T.

Nonequilibrium.—As nonequilibrium settings we consider the same processes for which we have benchmarked the algorithm in the LRIM case. The measured run times per spin update are plotted in Fig. 8. Similarly to the LRIM case, also here we find a rather pronounced dependence of the run times on t for the quench to $T_q = 0.1T_c$, while for the critical quench and the KZM, this dependence is again much weaker. Overall, the situation in the latter two protocols is fairly similar, since they are both concerned with phenomena related to the critical behavior of the system. For the quench to $0.1T_c$ we find an overall speedup

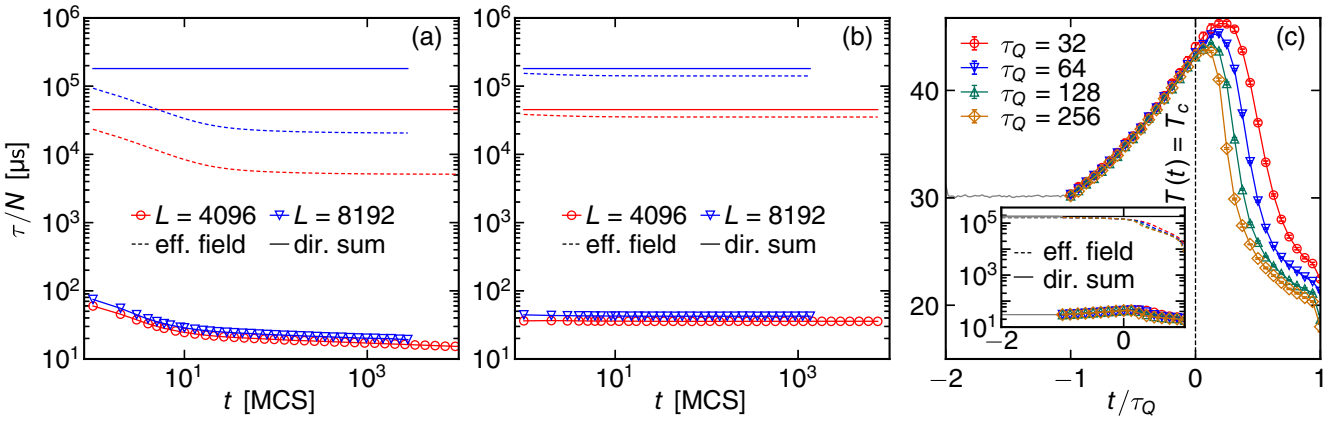

FIG. 8. Run times per spin update τ/N for the LRXYM versus simulation time t for $\sigma = 0.8$ in quenches from a disordered start to (a) $T_q = 0.1T_c$ and (b) $T_q = T_c$ on large lattices with $L = 4096$ and $L = 8192$ and (c) run times per spin update for simulation with $L = 8192$ using a Kibble-Zurek protocol. The dashed lines in the same color show the run times calculated for the effective field (eff. field) algorithm and the solid lines (without markers) are the run times of the direct summation (dir. sum).

of ≈1000 as compared to the effective field (and ≈8500 as compared to the direct summation). The corresponding speedup factors for the critical quench are approximately 3300 (4300) and 3500 (5500) for the KZM.

IV. LENNARD-JONES SYSTEM WITH FULL INTERACTION RANGE

To highlight the power of our algorithm also for continuous spatial degrees of freedom, we finally demonstrate its applicability to a Lennard-Jones particle system [58] with potential

$$V_{\mathrm{LJ}} = 4\epsilon\left[\left(\frac{\sigma}{r_{ij}}\right)^{12} - \left(\frac{\sigma}{r_{ij}}\right)^{6}\right]. \tag{30}$$

Here, we keep the full interaction range, i.e., do not truncate (and shift) the potential at the often employed cutoff $r_c = 2.5\sigma$. It is well known that, e.g., the critical temperature and critical density of a LJ system do depend on r_c [59–61].

We consider N interacting particles in a volume L^d whose linear extent L can be adjusted to yield the desired density $\rho = N/L^d$. Periodic boundary conditions are applied which, due the fast decay of $V_{\mathrm{LJ}} \propto -r_{ij}^{-6}$ can easily be realized by the minimum image convention; i.e., in this application Ewald summation is not necessary.

For the here presented simulations, we use the most general bound introduced in Sec. II; i.e., we virtually collect all particles of a box at the points of minimal or maximal interaction, respectively. This is analogous to bounds 2 introduced in Sec. III B for the LRIM. In equilibrium simulations of particles the proposed MC moves can be freely chosen. Here we perform 90% local displacements within radius $r = \sigma$ and 10% nonlocal moves where the particle's potential new position is chosen randomly in the whole simulation box. In Fig. 9 we show for $d = 2$ the run times per update τ/N for different N but fixed density $\rho = 0.35$ and varying temperature T, covering both the (oversaturated) vapor and vapor-liquid phase. We find clear evidence for $O(N \log N)$ complexity. The run times appear largely independent of the temperature, which is in contrast to the results for the LRIM in equilibrium. For smaller densities, we generally find faster run times for the same number of particles. An implementation for $d = 3$ is straightforward, too, and will be presented elsewhere [62].

V. APPLICABILITY AND LIMITATIONS

Although in this paper we only present in detail the application of the algorithm to three different models, it can be used for other lattice spin and off-lattice particle systems as well. Of course, we are not restricted to two spatial

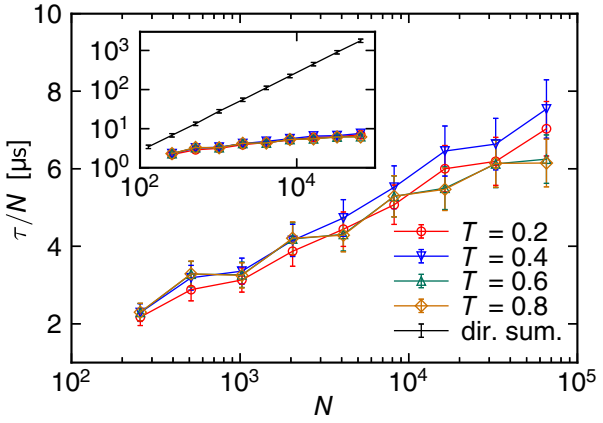

FIG. 9. Run times per update τ/N versus system size N plotted on a log-normal scale for a Lennard-Jones system with untruncated interactions at constant density $\rho = 0.35$. The inset shows the same data on a log-log scale together with the run times obtained from a standard Metropolis simulation using a direct summation of all interactions.

dimensions nor to hypercubic lattices. Systems with quenched disorder [63] are not excluded either: Random field models [64] are trivially accommodated within the framework introduced above since the extra field simply enters as an offset to ΔE.

As mentioned in Sec. III D 1, the scaling of the run times of the algorithm seems to be unaffected by the tightness of the employed bounds. A reduction of the involved prefactor, however, can be achieved. Such constant speedups can be decisive for the feasibility of large-scale simulations which require massive compute resources. Besides the above considered O(n) models with $n \geq 2$, site-diluted spin systems which model crystal defects through unoccupied lattice sites [65,66] constitute another example where bounds 3 [Eq. (23)] cannot be applied directly. For this system modified bounds 3 can be used instead,

$$\begin{aligned}\Delta E_B^{\min} &= 2\max[J_\Gamma^{\mathrm{int}}(\mathbf{R}) - (2N_B^- + N_B^0)J_\Gamma^{\max}(\mathbf{R}),\\ &\qquad (2N_B^+ + N_B^0)J_\Gamma^{\min}(\mathbf{R}) - J_\Gamma^{\mathrm{int}}(\mathbf{R})],\\ \Delta E_B^{\max} &= 2\ \min[J_\Gamma^{\mathrm{int}}(\mathbf{R}) - (2N_B^- + N_B^0)J_\Gamma^{\min}(\mathbf{R}),\\ &\qquad (2N_B^+ + N_B^0)J_\Gamma^{\max}(\mathbf{R}) - J_\Gamma^{\mathrm{int}}(\mathbf{R})],\end{aligned} \tag{31}$$

where N_B^0 is the number of vacancies in box B. This opens a way to treat q-state Potts models as well [67], where the components which are inert to both the old and the newly proposed state would be treated as vacancies. Now, we need to store the population of each of the q spin states for all boxes.

Random field and site dilution are forms of disorder that can easily be managed since they affect individual spins and their interaction with the environment as a whole. More challenging are models where disorder manifests as variation of the interaction of pairs of spins such as systems with bond dilution [68] or the Edwards-Anderson spin-glass model [69,70]. Here, the order parameter is not as closely related to the spin-box interaction energy as the magnetization in the case of the pure Ising model. This implies that it is difficult to formulate an estimator similar to bounds 3 or bounds 2. Also in this more difficult case we have checked that a reduction in complexity is achieved in simulations using the basic bounds 1, although the speedup is less pronounced as compared to the pure Ising case. Another large class of problems consists of spin systems with competing interactions, e.g., antiferromagnetic short-range and ferromagnetic long-range interactions or vice versa [71,72]. Here, one could, for example, evaluate the short-range interactions directly and treat the long-range interactions using our algorithm.

The LJ system considered above can easily be generalized to two-body potentials of other functional forms and extended to multispecies systems. Depending on the specific details of the different interactions between the particles, many scenarios can be imagined. A general approach would be to use a separate decomposition of the system for each particle type.

Yet, the question about possible limitations of the algorithm arises. Cases which might pose problems are systems with interactions that do not decay sufficiently fast (and thereby cannot be grouped together with decaying interaction strength). For mean-field models the algorithm may thus not be used (efficiently).

VI. CONCLUSION AND OUTLOOK

We have presented a general, hierarchical, and adaptive algorithm for Metropolis Monte Carlo simulations of long-range interacting systems. The range of possible applications of the algorithm is very broad. The formulation does not depend on the lattice structure and is thus valid for both general lattice spin models and systems with long-range interacting particles in continuous space as long as the interaction decays sufficiently fast with distance. In the two applications considered here, viz. the nonconserved long-range Ising model and a Lennard-Jones system in two dimensions, we observe run times that support an average asymptotic complexity of $O(N \log N)$ (where N is the number of spins or particles), but the existing data for the LRIM may also be described by a power law with a small exponent. However, the scenario of a logarithmic scaling seems more likely due to the hierarchical, treelike nature of the algorithm and is also strongly supported by the scaling of the run times for the LRXYM and the Lennard-Jones particle system.

Importantly, our method has small prefactors for the asymptotic scaling of the run times, resulting in speedup factors which exceed 10 000 in relevant physical scenarios. In a single day, we can perform simulations which before would have taken ≈30 years with any of the established methods, enabling the exploration of parameter ranges that were hitherto not accessible.

Until recently it was only possible to investigate the nonequilibrium properties of the long-range Ising model during quenches to low temperatures where the low acceptance rates allow an efficient simulation via the effective field approach [35]. The application of the new algorithm is not limited to this scenario, proving very efficient also in the case of large acceptance rates, as encountered, e.g., during quenches to the critical temperature. While here we exemplify our algorithm for the long-range Ising model and XY with nonconserved order parameter and a Lennard-Jones system, our method can easily be applied to other spin and off-lattice systems. One application where we have already successfully employed the algorithm is the phase separation in a conserved order parameter simulation of the long-range Ising model where the system evolves at the quench temperature through spin exchanges [73]. Other nonequilibrium simulation settings where long-range interactions are of interest are field-driven hysteresis [74] and Kibble-Zurek processes, driving

systems through phase transitions with different cooling rates [26–28], for which we have explicitly shown that they are now accessible with our new algorithm. Especially interesting is the extension to quantum systems where the Suzuki-Trotter mapping [75,76] of a d-dimensional quantum system to the corresponding $(d+1)$-dimensional classical system allows the application of our algorithm, which is designed for general d. Very recently, motivated by D-Wave experiments [77], Bando and Nishimori [78] investigated the generalized quantum Kibble-Zurek mechanism in the transverse-field Ising model coupled to an external bath [79] where long-range interactions arise naturally in Trotter direction. Also of great interest are models where the quantum spins themselves interact via a tunable long-range potential [80,81], which describes many experimental situations [82,83]. The algorithm we present here constitutes an important step toward an efficient simulation of the corresponding classical system.

ACKNOWLEDGMENTS

This project was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Project No. 189 853 844—SFB/TRR 102 (Project No. B04), and the Deutsch-Französische Hochschule (DFH-UFA) through the Doctoral College "$\mathbb{L}^4$" under Grant No. CDFA-02-07. We further acknowledge support by the Leipzig Graduate School of Natural Sciences "BuildMoNa."

[1] G. L. Eyink and K. R. Sreenivasan, *Onsager and the Theory of Hydrodynamic Turbulence*, Rev. Mod. Phys. **78**, 87 (2006).

[2] A. Campa, T. Dauxois, and S. Ruffo, *Statistical Mechanics and Dynamics of Solvable Models with Long-Range Interactions*, Phys. Rep. **480**, 57 (2009).

[3] R. H. French, V. A. Parsegian, R. Podgornik, R. F. Rajter, A. Jagota, J. Luo, D. Asthagiri, M. K. Chaudhury, Y.-M. Chiang, S. Granick *et al.*, *Long Range Interactions in Nanoscale Science*, Rev. Mod. Phys. **82**, 1887 (2010).

[4] Y. Levin, R. Pakter, F. B. Rizzato, T. N. Teles, and F. P. C. Benetti, *Nonequilibrium Statistical Mechanics of Systems with Long-Range Interactions*, Phys. Rep. **535**, 1 (2014).

[5] A. Campa, T. Dauxois, D. Fanelli, and S. Ruffo, *Physics of Long-Range Interacting Systems* (Oxford University Press, Oxford, 2014).

[6] J. S. Douglas, H. Habibian, C.-L. Hung, A. V. Gorshkov, H. J. Kimble, and D. E. Chang, *Quantum Many-Body Models with Cold Atoms Coupled to Photonic Crystals*, Nat. Photonics **9**, 326 (2015).

[7] B. Neyenhuis, J. Zhang, P. W. Hess, J. Smith, A. C. Lee, P. Richerme, Z. Gong, A. V. Gorshkov, and C. Monroe, *Observation of Prethermalization in Long-Range Interacting Spin Chains*, Sci. Adv. **3**, e1700672 (2017).

[8] Y.-C. Zhang, V. Walther, and T. Pohl, *Long-Range Interactions and Symmetry Breaking in Quantum Gases through Optical Feedback*, Phys. Rev. Lett. **121**, 073604 (2018).

[9] J. W. Perram, H. G. Petersen, and S. W. De Leeuw, *An Algorithm for the Simulation of Condensed Matter which Grows as the 3/2 Power of the Number of Particles*, Mol. Phys. **65**, 875 (1988).

[10] J. Barnes and P. Hut, *A Hierarchical $O(N \log N)$ Force-Calculation Algorithm*, Nature (London) **324**, 446 (1986).

[11] *Fast Methods for Long-Range Interactions in Complex Systems*, Institute for Advanced Simulation Series, Vol. 6, edited by G. Sutmann, P. Gibbon, and T. Lippert (Forschungszentrum Jülich, Jülich, 2011).

[12] A. Arnold, F. Fahrenberger, C. Holm, O. Lenz, M. Bolten, H. Dachsel, R. Halver, I. Kabadshow, F. Gähler, F. Heber *et al.*, *Comparison of Scalable Fast Methods for Long-Range Interactions*, Phys. Rev. E **88**, 063308 (2013).

[13] E. Luijten and H. Blöte, *Monte Carlo Method for Spin Models with Long-Range Interactions*, Int. J. Mod. Phys. C **06**, 359 (1995).

[14] K. Fukui and S. Todo, *Order-N Cluster Monte Carlo Method for Spin Systems with Long-Range Interactions*, J. Comput. Phys. **228**, 2629 (2009).

[15] E. Flores-Sola, M. Weigel, R. Kenna, and B. Berche, *Cluster Monte Carlo and Dynamical Scaling for Long-Range Interactions*, Eur. Phys. J. Special Topics **226**, 581 (2017).

[16] E. P. Bernard, W. Krauth, and D. B. Wilson, *Event-Chain Monte Carlo Algorithms for Hard-Sphere Systems*, Phys. Rev. E **80**, 056704 (2009).

[17] M. Michel, S. C. Kapfer, and W. Krauth, *Generalized Event-Chain Monte Carlo: Constructing Rejection-Free Global-Balance Algorithms from Infinitesimal Steps*, J. Chem. Phys. **140**, 054116 (2014).

[18] W. Krauth, *Event-Chain Monte Carlo: Foundations, Applications, and Prospects*, Front. Phys. **9**, 229 (2021).

[19] A. Hucht, *Nonequilibrium Phase Transition in an Exactly Solvable Driven Ising Model with Friction*, Phys. Rev. E **80**, 061138 (2009).

[20] M. Michel, X. Tan, and Y. Deng, *Clock Monte Carlo Methods*, Phys. Rev. E **99**, 010105(R) (2019).

[21] G. Giachetti, N. Defenu, S. Ruffo, and A. Trombettoni, *Berezinskii-Kosterlitz-Thouless Phase Transitions with Long-Range Couplings*, Phys. Rev. Lett. **127**, 156801 (2021).

[22] A. J. Bray, *Theory of Phase-Ordering Kinetics*, Adv. Phys. **51**, 481 (2002).

[23] *Kinetics of Phase Transitions*, edited by S. Puri and V. Wadhawan (CRC Press, Boca Raton, FL, 2009).

[24] M. Henkel and M. Pleimling, *Non-Equilibrium Phase Transitions: Ageing and Dynamical Scaling far from Equilibrium*, Vol. 2 (Springer, Heidelberg, 2010).

[25] P. Calabrese and A. Gambassi, *Ageing Properties of Critical Systems*, J. Phys. A **38**, R133 (2005).

[26] T. W. B. Kibble, *Some Implications of a Cosmological Phase Transition*, Phys. Rep. **67**, 183 (1980).

[27] W. H. Zurek, *Cosmological Experiments in Superfluid Helium?*, Nature (London) **317**, 505 (1985).

[28] W. H. Zurek, *Cosmological Experiments in Condensed Matter Systems*, Phys. Rep. **276**, 177 (1996).

[29] N. Clisby, *Accurate Estimate of the Critical Exponent ν for Self-Avoiding Walks via a Fast Implementation of the Pivot Algorithm*, Phys. Rev. Lett. **104**, 055702 (2010).
[30] N. Clisby, *Efficient Implementation of the Pivot Algorithm for Self-Avoiding Walks*, J. Stat. Phys. **140**, 349 (2010).
[31] S. Schnabel and W. Janke, *Accelerating Polymer Simulation by Means of Tree Data-Structures and a Parsimonious Metropolis Algorithm*, Comput. Phys. Commun. **256**, 107414 (2020).
[32] P. Ewald, *Die Berechnung Optischer und Elektrostatischer Gitterpotentiale*, Ann. Phys. (Berlin) **369**, 253 (1921).
[33] T. Horita, H. Suwa, and S. Todo, *Upper and Lower Critical Decay Exponents of Ising Ferromagnets with Long-Range Interaction*, Phys. Rev. E **95**, 012143 (2017).
[34] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevX.13.031006 for an alternative derivation of the bounds.
[35] H. Christiansen, S. Majumder, and W. Janke, *Phase Ordering Kinetics of the Long-Range Ising Model*, Phys. Rev. E **99**, 011301(R) (2019).
[36] For the effective field simulation times, we performed simulations having different acceptance rates α (corresponding to the different temperatures) for a range of smaller system sizes and phenomenologically fitted the resulting run times with the function $R(\alpha) = AN\alpha + B$ giving the runtime per update, where N is the number of spins and A and B are free fitting parameters. It is thus assumed explicitly that the complexity is $O(N^2)$. Further, it is recognized that there is a constant, system size independent term B involved. All the run times for the effective field are not measured directly but calculated from the system size and the acceptance rates measured during the simulation with the new algorithm.
[37] E. Burovski, W. Janke, M. Guskova, and L. Shchur, *Acceptance Rate Is a Thermodynamic Function in Local Monte Carlo Algorithms*, Phys. Rev. E **100**, 063303 (2019).
[38] H. Christiansen, S. Majumder, M. Henkel, and W. Janke, *Aging in the Long-Range Ising Model*, Phys. Rev. Lett. **125**, 180601 (2020).
[39] R. Agrawal, F. Corberi, E. Lippiello, P. Politi, and S. Puri, *Kinetics of the Two-Dimensional Long-Range Ising Model at Low Temperatures*, Phys. Rev. E **103**, 012108 (2021).
[40] H. Christiansen, S. Majumder, and W. Janke, *Zero-Temperature Coarsening in the Two-Dimensional Long-Range Ising Model*, Phys. Rev. E **103**, 052122 (2021).
[41] H. Janssen, B. Schaub, and B. Schmittmann, *New Universal Short-Time Scaling Behaviour of Critical Relaxation Processes*, Z. Phys. B Condens. Matter **73**, 539 (1989).
[42] R. Islam, C. Senko, W. C. Campbell, S. Korenblit, J. Smith, A. Lee, E. E. Edwards, C.-C. J. Wang, J. K. Freericks, and C. Monroe, *Emergence and Frustration of Magnetism with Variable-Range Interactions in a Quantum Simulator*, Science **340**, 583 (2013).
[43] P. Richerme, Z.-X. Gong, A. Lee, C. Senko, J. Smith, M. Foss-Feig, S. Michalakis, A. V. Gorshkov, and C. Monroe, *Non-Local Propagation of Correlations in Quantum Systems with Long-Range Interactions*, Nature (London) **511**, 198 (2014).
[44] P. Jurcevic, B. P. Lanyon, P. Hauke, C. Hempel, P. Zoller, R. Blatt, and C. F. Roos, *Quasiparticle Engineering and Entanglement Propagation in a Quantum Many-Body System*, Nature (London) **511**, 202 (2014).
[45] B. Gardas, J. Dziarmaga, W. H. Zurek, and M. Zwolak, *Defects in Quantum Computers*, Sci. Rep. **8**, 4539 (2018).
[46] A. Keesling, A. Omran, H. Levine, H. Bernien, H. Pichler, S. Choi, R. Samajdar, S. Schwartz, P. Silvi, S. Sachdev, P. Zoller, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, *Quantum Kibble-Zurek Mechanism and Critical Dynamics on a Programmable Rydberg Simulator*, Nature (London) **568**, 207 (2019).
[47] R. Puebla, S. Deffner, and S. Campbell, *Kibble-Zurek Scaling in Quantum Speed Limits for Shortcuts to Adiabaticity*, Phys. Rev. Res. **2**, 032020(R) (2020).
[48] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk *et al.*, *Quantum Annealing with Manufactured Spins*, Nature (London) **473**, 194 (2011).
[49] P. I. Bunyk, E. M. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, A. J. Przybysz, and J. Whittaker, *Architectural Considerations in the Design of a Superconducting Quantum Annealing Processor*, IEEE Trans. Appl. Supercond. **24**, 1 (2014).
[50] A. Jelić and L. F. Cugliandolo, *Quench Dynamics of the 2d XY Model*, J. Stat. Mech. (2011) P02032.
[51] S. Schnabel, F. Müller, H. Christiansen, and W. Janke (to be published).
[52] M. I. Berganza and L. Leuzzi, *Critical Behavior of the XY Model in Complex Topologies*, Phys. Rev. B **88**, 144104 (2013).
[53] F. Cescatti, M. Ibáñez-Berganza, A. Vezzani, and R. Burioni, *Analysis of the Low-Temperature Phase in the Two-Dimensional Long-Range Diluted XY Model*, Phys. Rev. B **100**, 054203 (2019).
[54] A. J. Bray and A. D. Rutenberg, *Growth Laws for Phase Ordering*, Phys. Rev. E **49**, R27 (1994).
[55] B. Zheng, M. Schulz, and S. Trimper, *Dynamic Simulations of the Kosterlitz-Thouless Phase Transition*, Phys. Rev. E **59**, R1351 (1999).
[56] Y. Ozeki and N. Ito, *Nonequilibrium Relaxation Method*, J. Phys. A **40**, R149 (2007).
[57] This may be due to cache effects which are shifted to smaller system sizes for the LRXYM, since storing the continuous spin variables occupies a larger amount of memory. Once the needed amount of storage exceeds the cache size, these variables need to be moved to the RAM, which has considerably longer access times.
[58] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Elsevier, Amsterdam, 2001).
[59] B. Smit and D. Frenkel, *Vapor-Liquid Equilibria of the Two-Dimensional Lennard-Jones Fluid(s)*, J. Chem. Phys. **94**, 5663 (1991).
[60] B. Smit, *Phase Diagrams of Lennard-Jones Fluids*, J. Chem. Phys. **96**, 8639 (1992).
[61] A. Trokhymchuk and J. Alejandre, *Computer Simulations of Liquid/Vapor Interface in Lennard-Jones Fluids: Some Questions and Answers*, J. Chem. Phys. **111**, 8510 (1999).
[62] S. Schnabel and W. Janke (to be published).

[63] A. P. Young, *Spin Glasses and Random Fields* (World Scientific, Singapore, 1997).
[64] A. J. Bray and M. A. Moore, *Scaling Theory of the Random-Field Ising Model*, J. Phys. C **18**, L927 (1985).
[65] A. B. Harris, *Effect of Random Defects on the Critical Behaviour of Ising Models*, J. Phys. C **7**, 1671 (1974).
[66] S. Kazmin and W. Janke, *Critical Exponent ν of the Ising Model in Three Dimensions with Long-Range Correlated Site Disorder Analyzed with Monte Carlo Techniques*, Phys. Rev. B **102**, 174206 (2020).
[67] S. Majumder, S. K. Das, and W. Janke, *Universal Finite-Size Scaling Function for Kinetics of Phase Separation in Mixtures with Varying Number of Components*, Phys. Rev. E **98**, 042142 (2018).
[68] A. H. Kole, G. T. Barkema, and L. Fritz, *Comparison of Cluster Algorithms for the Bond-Diluted Ising Model*, Phys. Rev. E **105**, 015313 (2022).
[69] S. F. Edwards and P. W. Anderson, *Theory of Spin Glasses*, J. Phys. F **5**, 965 (1975).
[70] H. G. Katzgraber and A. P. Young, *Monte Carlo Studies of the One-Dimensional Ising Spin Glass with Power-Law Interactions*, Phys. Rev. B **67**, 134410 (2003).
[71] A. Giuliani, J. L. Lebowitz, and E. H. Lieb, *Ising Models with Long-Range Antiferromagnetic and Short-Range Ferromagnetic Interactions*, Phys. Rev. B **74**, 064420 (2006).
[72] C. M. Horowitz, M. A. Bab, M. Mazzini, M. L. Rubio Puzzo, and G. P. Saracco, *Phase Transitions and Critical Phenomena in the Two-Dimensional Ising Model with Dipole Interactions: A Short-Time Dynamics Study*, Phys. Rev. E **92**, 042127 (2015).
[73] F. Müller, H. Christiansen, and W. Janke, *Phase-Separation Kinetics in the Two-Dimensional Long-Range Ising Model*, Phys. Rev. Lett. **129**, 240601 (2022).
[74] B. K. Chakrabarti and M. Acharyya, *Dynamic Transitions and Hysteresis*, Rev. Mod. Phys. **71**, 847 (1999).
[75] M. Suzuki, *Relationship between d-Dimensional Quantal Spin Systems and $(d+1)$-Dimensional Ising Systems: Equivalence, Critical Exponents and Systematic Approximants of the Partition Function and Spin Correlations*, Prog. Theor. Phys. **56**, 1454 (1976).
[76] M. Suzuki, *Quantum Monte Carlo Methods—Recent Developments*, Physica (Amsterdam) **194A**, 432 (1993).
[77] Y. Bando, Y. Susa, H. Oshiyama, N. Shibata, M. Ohzeki, F. J. Gómez-Ruiz, D. A. Lidar, S. Suzuki, A. del Campo, and H. Nishimori, *Probing the Universality of Topological Defect Formation in a Quantum Annealer: Kibble-Zurek Mechanism and Beyond*, Phys. Rev. Res. **2**, 033369 (2020).
[78] Y. Bando and H. Nishimori, *Simulated Quantum Annealing as a Simulator of Nonequilibrium Quantum Dynamics*, Phys. Rev. A **104**, 022607 (2021).
[79] P. Werner, K. Völker, M. Troyer, and S. Chakravarty, *Phase Diagram and Critical Exponents of a Dissipative Ising Spin Chain in a Transverse Magnetic Field*, Phys. Rev. Lett. **94**, 047201 (2005).
[80] D. Jaschke, K. Maeda, J. D. Whalen, M. L. Wall, and L. D. Carr, *Critical Phenomena and Kibble-Zurek Scaling in the Long-Range Quantum Ising Chain*, New J. Phys. **19**, 033032 (2017).
[81] R. Puebla, O. Marty, and M. B. Plenio, *Quantum Kibble-Zurek Physics in Long-Range Transverse-Field Ising Models*, Phys. Rev. A **100**, 032115 (2019).
[82] K. Kim, M.-S. Chang, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G.-D. Lin, L.-M. Duan, and C. Monroe, *Quantum Simulation of Frustrated Ising Spins with Trapped Ions*, Nature (London) **465**, 590 (2010).
[83] H. Labuhn, D. Barredo, S. Ravets, S. De Léséleuc, T. Macrì, T. Lahaye, and A. Browaeys, *Tunable Two-Dimensional Arrays of Single Rydberg Atoms for Realizing Quantum Ising Models*, Nature (London) **534**, 667 (2016).