Supplemental Material: Phase-separation kinetics in the two-dimensional long-range Ising model

Fabio Müller,* Henrik Christiansen,† and Wolfhard Janke‡
Institut für Theoretische Physik, Universität Leipzig, IPF 231101, 04081 Leipzig, Germany
(Dated: October 17, 2022)

TECHICAL DETAILS OF THE SIMULATION

The simulations are performed on finite $L \times L$ square lattices with periodic boundary conditions. Unavoidable finite-size effects can be mitigated by implementing the periodic boundary condition via Ewald summation [1]. We replace the couplings which describe the infinite system,

$$J_{i,j} = J(\mathbf{r}_{i,j}) = |\mathbf{r}_j - \mathbf{r}_i|^{-(d+\sigma)} = r_{i,j}^{-(d+\sigma)}, \tag{1}$$

from the Hamiltonian (1) of the main article by couplings which directly incorporate the Ewald summation [2],

$$J^{\mathrm{PBC}}(\mathbf{r}_{i,j}) = \sum_{\mu,\nu=-\infty}^{\infty} J\left(\mathbf{r}_{i,j} + \mu L\hat{\mathbf{e}}_x + \nu L\hat{\mathbf{e}}_y\right), \tag{2}$$

where $\hat{\mathbf{e}}_x$ and $\hat{\mathbf{e}}_y$ are the unit vectors in x and y direction. Since Ewald summation preserves translational symmetry the couplings do not need to be stored for each pair s_i, s_j but it is sufficient to calculate the couplings once at the beginning of a simulation for each possible distance vector $\mathbf{r}_{i,j}$. Thus, these couplings can simply be used in conjunction with the minimum image convention. In the following we drop the superscript "PBC" to ease the notation.

The evolution of the system takes place only through exchanges of randomly drawn directly neighboring spins, which constitute the most local configurational changes conserving the order parameter. The simulation process is repeated for at least 1000 realizations for $L = 256$, 400 realizations for $L = 512$, and 200 realizations for $L = 1024$ with different initial configurations and thermal noise. The simulations proceed by proposing nearest-neighbor exchanges and accepting them according to the Metropolis criterion

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

where only the change in energy ΔE and the inverse temperature β enter. The energy difference is measured between proposed new and the old configuration,

$$\Delta E = E^{\mathrm{new}} - E^{\mathrm{old}}. \tag{4}$$

For the energy difference for each single spin pair we have

$$\Delta E(s_i, s_j) = J_{i,j}s_i^{\mathrm{new}}s_j^{\mathrm{new}} - J_{i,j}s_i^{\mathrm{old}}s_j^{\mathrm{old}}, \tag{5}$$

which is nonzero only in the case where *one* of the two spins is flipped. Exchanging two spins which point in the same direction does of course not lead to a change in the configuration and hence neither in the energy of the system. We discuss the role of parallel spin pairs in the next section, but now first focus on the exchange of spins pointing in the opposite direction.

For an exchange of two such spins s_i and s_j pointing in opposite direction we can write the change in energy as

$$\Delta E = 2\sum_{k\notin\{i,j\}} J(\mathbf{r}_{i,k})s_i^{\mathrm{old}}s_k + 2\sum_{k\notin\{i,j\}} J(\mathbf{r}_{j,k})s_j^{\mathrm{old}}s_k. \tag{6}$$

In the summation the two exchanged spins are neglected because the interaction between the two spins which are flipped does not change during the exchange, as shown in Eq. (5). Accounting for $s_i^{\mathrm{old}} = -s_j^{\mathrm{old}}$ we obtain

$$\Delta E = 2s_i^{\mathrm{old}} \sum_{k\notin\{i,j\}} \left[J(\mathbf{r}_{i,k}) - J(\mathbf{r}_{j,k})\right] s_k. \tag{7}$$

Using $\mathbf{r}_{j,k} = \mathbf{r}_{i,k} - \mathbf{r}_{i,j}$ yields

$$\begin{aligned}\Delta E &= 2s_i^{\mathrm{old}} \sum_{k\notin\{i,j\}} \left[J(\mathbf{r}_{i,k}) - J(\mathbf{r}_{i,k} - \mathbf{r}_{i,j})\right] s_k \\ &= 2s_i^{\mathrm{old}} \sum_{k\notin\{i,j\}} J^{\mathrm{eff}}_{\mathbf{r}_{i,j}}(\mathbf{r}_{i,k})s_k \end{aligned} \tag{8}$$

where

$$J^{\mathrm{eff}}_{\mathbf{r}_{i,j}}(\mathbf{r}_{i,k}) = J(\mathbf{r}_{i,k}) - J(\mathbf{r}_{i,k} - \mathbf{r}_{i,j}). \tag{9}$$

These effective couplings now have the form of a generalized dipole interaction which in contrast to the original couplings loose the strict positivity, decay faster, and carry a strong anisotropy. As mentioned before, we only propose spin exchanges along the principal axis of the lattice between two spins at distance 1. Thus, the vector $\mathbf{r}_{ij}$ can be generated from $\hat{\mathbf{e}}_x$ by a rotation with an angle $\phi \in \{0, \pi/2, \pi, 3\pi/2\}$,

$$\mathbf{r}_{i,j} = O_{(\phi(j))}\hat{\mathbf{e}}_x \equiv O_{(j)}\hat{\mathbf{e}}_x \tag{10}$$

where the subscript (j) on the rotation matrix O indicates the dependence of the rotation on the specific $\mathbf{r}_{i,j}$ (the index i is anyway fixed throughout). Then

$$\begin{aligned} J^{\mathrm{eff}}_{\mathbf{r}_{i,j}}(\mathbf{r}_{i,k}) &= J(\mathbf{r}_{i,k}) - J(\mathbf{r}_{i,k} - \mathbf{r}_{i,j}) \\ &= J(\mathbf{r}_{i,k}) - J(\mathbf{r}_{i,k} - O_{(j)}\hat{\mathbf{e}}_x) \\ &= J(O_{(j)}O_{(j)}^{-1}\mathbf{r}_{i,k}) - J(O_{(j)}(O_{(j)}^{-1}\mathbf{r}_{i,k} - \hat{\mathbf{e}}_x)). \end{aligned} \tag{11}$$

Even after the Ewald summation the couplings J remain invariant under the above considered rotations

$$J(O_{(j)}\mathbf{r}) = J(\mathbf{r}). \tag{12}$$

Hence

$$\begin{aligned} J^{\text{eff}}_{\mathbf{r}_{i,j}}(\mathbf{r}_{i,k}) &= J(\mathbf{r}_{i,k}) - J(\mathbf{r}_{i,k} - \mathbf{r}_{i,j}) \\ &= J(O^{-1}_{(j)}\mathbf{r}_{i,k}) - J(O^{-1}_{(j)}\mathbf{r}_{i,k} - \hat{\mathbf{e}}_x) \\ &= J^{\text{eff}}_{\hat{\mathbf{e}}_x}(O^{-1}_{(j)}\mathbf{r}_{i,k}) \end{aligned} \tag{13}$$

which expresses the effective couplings for any of the 4 possible nearest neighbors $\mathbf{r}_{i,j}$ in terms of the effective couplings for $\mathbf{r}_{i,j} = \hat{\mathbf{e}}_x$ subject to an (inverse) rotation of $\mathbf{r}_{i,k}$. With theses couplings, we can rewrite (7) as

$$\Delta E = 2s_i^{\text{old}} \sum_{k\neq i,j} J^{\text{eff}}_{\mathbf{r}_{i,j}}(\mathbf{r}_{i,k}) s_k \tag{14}$$

$$= 2s_i^{\text{old}} \sum_{k\neq i,j} J^{\text{eff}}_{\hat{\mathbf{e}}_x}(O^{-1}_{(j)}\mathbf{r}_{i,k}) s_k, \tag{15}$$

which allows us to compute the energy difference ΔE with a single set of effective couplings

$$J^{\text{eff}}_{\hat{\mathbf{e}}_x}(\mathbf{r}_{i,k}) = J(\mathbf{r}_{i,k}) - J(\mathbf{r}_{i,k} - \hat{\mathbf{e}}_x). \tag{16}$$

This helps to minimize the memory requirement of the simulation which becomes important for the simulation of very large systems where the couplings and the quantities which are derived from them become very large in memory.

With this formulation of ΔE in (15), we are now able to efficiently use the new algorithm [3] for the here considered conserved long-range Ising model, which was originally optimized for the nonconserved long-range Ising model. It uses the observation, that the energy difference ΔE involved in a proposed spin exchange does not need to be known exactly and its direct calculation can be avoided in favor of a decision based on exact lower and upper bounds [4]. This is achieved by means of a dynamical tree-like hierarchical data structure, which enables the collective treatment of many interactions at a time. While the original implementation was optimized and tested for the nonconserved long-range Ising model, we are able to carry over all optimizations which were conceived there. However, for the sake of an efficient implementation, instead of excluding s_i and s_j from the summation in (15), it is more practical to set the two effective couplings $J^{\text{eff}}_{\hat{\mathbf{e}}_x}(0) = J^{\text{eff}}_{\hat{\mathbf{e}}_x}(\hat{\mathbf{e}}_x) = 0$ in (16) which represent the self-interaction of the two exchanged spins and the interaction between them. This has the additional advantage of matching the conventions of the algorithm in Ref. [3].

THE ROLE OF PARALLEL SPIN-PAIRS

During the process of phase separation the domains of aligned spins grow steadily as is evident from the growth

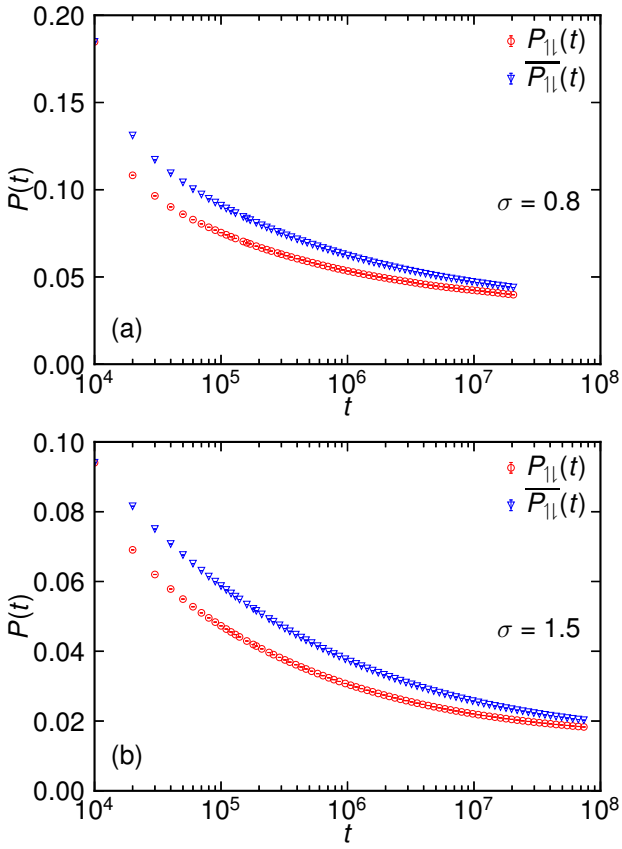

FIG. 1. Evolution of the fraction of anti-parallel spin pairs during the phase separation process for (a) $\sigma = 0.8$ and (b) $\sigma = 1.5$ and $L = 1024$.

of the characteristic length $\ell(t)$. As a direct consequence, the probability of finding two aligned spins when picking a random pair of two neighboring spins grows with proceeding simulation time. The decision about the exchange of such a parallel spin pair is, of course, trivial since it changes neither the configuration nor the energy of the system. Instead of randomly picking pairs of spin to exchange, we have opted for an alternative approach where we maintain a list of anti-parallel neighbor pairs and propose only spin exchanges from this list [5]. In this formalism the Monte Carlo time is incremented at each update by $1/m$ where m is the length of the list. Thus, one Monte Carlo sweep corresponds to $2N = 2L^2$ spin-exchange attempts, which sets the time scale of the simulations. With this simple algorithmic trick the overhead of the trivial spin exchanges is almost completely eliminated. Especially for the usual nearest-neighbor Ising model, the approach speeds up the simulation considerably, particularly in the late stage of the phase-ordering process.

For a parallel implementation, for example on GPUs as proposed in Ref. [6], the role of aligned spin pairs is very different. There, the parallel efficiency relies on the fact that all updates during one sweep can be treated in

parallel. By minimizing the penalty of the parallelization with a clever protocol a 440-fold speedup for the simulation in a very dilute regime was reported. The algorithm was originally presented for an off-lattice particle system, but could also be applied to our simulation setting where a spin exchange in the Ising language would be mapped to a particle hopping in the particle picture. In the particle language a spin exchange of two aligned spins representing particles, corresponds to a particle hopping onto an occupied lattice site, which is of course forbidden. Also in the particle picture a pre-rejection scheme would be trivial in a sequential implementation: Do not calculate the energy and reject the update (or do not even propose it as described above), and move to the next update. Conversely, in a parallel algorithm all the updates need to be performed simultaneously to exploit all available resources. Even if the decision about some proposed updates is trivial, the algorithm can only continue to run once the decisions about all the updates are taken, rendering the trivial updates as expensive as the non-trivial ones. This issue is also mentioned in Ref. [6], where it is stated clearly that a "naive pre-rejection strategy on the GPU implementation cannot improve the performance". To investigate the influence of the pre-rejection probability on the performance of a parallel implementation without pre-rejection strategy we need to look at the fraction of neighbors which point into opposite direction defined as the fraction of the number of anti-parallel and total number of neighboring spin pairs

$$P_{\uparrow\downarrow}(t) = \frac{N_{\uparrow\downarrow}(t)}{N_{\uparrow\downarrow}(t) + N_{\uparrow\uparrow}(t)}, \tag{17}$$

where $N_{\uparrow\downarrow}$ is the number of anti-parallel neighbor pairs and $N_{\uparrow\uparrow}$ the number of parallel neighbor pairs. Additionally, we want to investigate the running average of this quantity

$$\overline{P_{\uparrow\downarrow}}(t) = \frac{1}{t}\int_0^t P_{\uparrow\downarrow}(t')dt'. \tag{18}$$

In Fig. 1 we show $P_{\uparrow\downarrow}(t)$ (red circles) and $\overline{P_{\uparrow\downarrow}}(t)$ (blue triangles) for $\sigma = 0.8$ and $\sigma = 1.5$ for $L = 1024$. Most relevant to the average performance of the algorithm is the average fraction of anti-parallel spins over the course of the entire simulation $\overline{P_{\uparrow\downarrow}}(t_{\max})$ which shows that during the complete time evolution picking a random pair of neighboring spins yields two anti-parallel spins in only 5% of the cases for $\sigma = 0.8$ and 2% of the cases for $\sigma = 1.5$. While for a sequential algorithm, where the cost of the energy evaluation dominates, this translates into a speedup of $\approx 20 - 50$, a parallel implementation would need the same amount of time as for the rare case where all the proposed pairs of neighbors are anti-parallel, thus loosing a large factor in its effective parallel efficiency.

We attribute the dependence of $\overline{P_{\uparrow\downarrow}}(t)$ on σ (growing number of anti-parallel spin pairs with decreasing σ) to the increasing *absolute* quench temperatures with decreasing σ. On average for all σ we find $\overline{P_{\uparrow\downarrow}}(t_{\max}) \approx 0.04$ which translates to a loss in parallel efficiency by a factor of ≈ 25 for the GPU implementation and thus results in a remaining speedup of < 20 (instead of ≈ 440) in comparison to a sequential implementation with a pre-rejections scheme. Our speedup factor of $150 - 450$ in the main text is reported with respect to a sequential implementation with the same pre-rejection strategy as used for our new algorithm, which altogether yields a more than 7-fold speedup of our simulation running on a *single* core of a CPU compared to a simulation utilizing a *full* GPU with the method introduced in Ref. [6].

* fabio.mueller@itp.uni-leipzig.de
† henrik.christiansen@itp.uni-leipzig.de; Current address: NEC Laboratories Europe GmbH, Kurfürsten-Anlage 36, 69115 Heidelberg, Germany
‡ wolfhard.janke@itp.uni-leipzig.de

[1] P. Ewald, Die Berechnung optischer und elektrostatischer Gitterpotentiale, Ann. Phys. (Berl.) **369**, 253 (1921).
[2] 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).
[3] F. Müller, H. Christiansen, S. Schnabel, and W. Janke, Fast, hierarchical, and adaptive algorithm for Metropolis Monte Carlo simulations of long-range interacting systems, arXiv:2207.14670 [physics.comp-ph], under review at Phys. Rev. X (2022).
[4] 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).
[5] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, New York, 1999).
[6] Y. Liang, X. Xing, and Y. Li, A GPU-based large-scale Monte Carlo simulation method for systems with long-range interactions, J. Comput. Phys. **338**, 252 (2017).