

## Efficient Cluster Algorithm for Spin Glasses in Any Space Dimension

Zheng Zhu,<sup>1</sup> Andrew J. Ochoa,<sup>1</sup> and Helmut G. Katzgraber<sup>1,2,3</sup>

<sup>1</sup>*Department of Physics and Astronomy, Texas A&M University, College Station, Texas 77843-4242, USA*

<sup>2</sup>*Materials Science and Engineering Program, Texas A&M University, College Station, Texas 77843, USA*

<sup>3</sup>*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501, USA*

(Received 22 January 2015; published 14 August 2015)

Spin systems with frustration and disorder are notoriously difficult to study, both analytically and numerically. While the simulation of ferromagnetic statistical mechanical models benefits greatly from cluster algorithms, these accelerated dynamics methods remain elusive for generic spin-glass-like systems. Here, we present a cluster algorithm for Ising spin glasses that works in any space dimension and speeds up thermalization by at least one order of magnitude at temperatures where thermalization is typically difficult. Our isoenergetic cluster moves are based on the Houdayer cluster algorithm for two-dimensional spin glasses and lead to a speedup over conventional state-of-the-art methods that increases with the system size. We illustrate the benefits of the isoenergetic cluster moves in two and three space dimensions, as well as the nonplanar chimera topology found in the D-Wave Inc. quantum annealing machine.

DOI: [10.1103/PhysRevLett.115.077201](https://doi.org/10.1103/PhysRevLett.115.077201)

PACS numbers: 75.10.Nr, 05.50.+q, 64.60.-i, 75.40.Mg

A plethora of problems across disciplines map onto spin-glass-like Hamiltonians [1]. Despite decades of intense analytical and numerical scrutiny, a deep understanding of these paradigmatic models of disordered systems remains elusive. Given the inherent difficulties of studying these Hamiltonians analytically beyond mean-field theory as well as the continuous increase of computer power, progress in this field has benefited noticeably from numerical studies. The development of efficient Monte Carlo methods such as parallel tempering [2] and population annealing [3] has helped in understanding these systems at a much deeper level; however, most numerical studies are still plagued by corrections to finite-size scaling due to the small system sizes currently available [4].

In contrast, simulations of spin Hamiltonians without disorder and frustration are comparably simple: Ferromagnetic systems have greatly benefited from the development of cluster algorithms [5,6] that help in overcoming critical slowing down close to phase transitions. Therefore, the holy grail of spin-glass simulations is to introduce accelerated cluster dynamics that improve upon the benefits of efficient simulation methods such as population annealing or parallel tempering Monte Carlo. In 2001 Houdayer introduced a seminal rejection-free cluster algorithm tailored to work for two-dimensional Ising spin glasses [7]. The method updates large patches of spins at once, therefore effectively randomizing the configurations and efficiently overcoming large barriers in the free-energy landscape. Furthermore, the energy of the system remains unchanged when performing a cluster move. This means that the numerical overhead is very small because the rejection rate is zero and there is no need to, for example, compute any random numbers for a cluster update. The use of these cluster moves made it possible to

obtain a speedup of several orders of magnitude in two-dimensional systems, therefore allowing us to simulate considerably larger system sizes.

While cluster algorithms such as the Swendsen-Wang and Wolff ones [5,6] work well for ferromagnetic systems in any space dimension because the clusters reflect the spin correlations in the system, this is not the case for algorithms that build clusters like the Houdayer cluster algorithm. In this case, the clusters do not reflect overlap correlations [8,9] and cluster updates only have an accelerating effect on the dynamics if the clusters do not span the entire system or if they comprise single spins. This is the case either when temperatures are close to zero (small clusters) or when the underlying geometry of the problem has a percolation threshold below 50%—as is the case in three space dimensions. Updating such a system-spanning cluster amounts to swapping out both replicas, thereby not randomizing the configurations. This means that while the method works in principle, it does not really provide any simulational benefit. As such, Houdayer cluster moves work, in principle, only for models where the percolation threshold is above 50%, as is the case in two-dimensional Ising spin-glass Hamiltonians. One way to remedy this situation is to increase the percolation threshold artificially, e.g., by diluting the lattice [10]. However, this is often not desirable and is highly dependent on the problem to be studied.

Here, we show that Houdayer-like cluster moves can be applied to spin systems on topologies where the percolation threshold is below 50%, provided that the interplay of temperature and frustration prevents clusters from spanning the whole system. We therefore introduce *isoenergetic cluster moves* for spin-glass-like Hamiltonians in any space dimension. These rejection-free cluster moves accelerate

thermalization by several orders of magnitude even for systems with space dimensions larger than 2. We show that the inherent frustration present in spin-glass Hamiltonians prevents clusters from spanning the whole system for temperatures below the characteristic energy scale of the problem. As such, spin-glass simulations can be sped up considerably in the hard-to-reach low-temperature regime of interest in many numerical studies.

The fact that the isoenergetic cluster moves are rejection free and leave the energy of the system unchanged is also of great importance to any heuristic based on Monte Carlo updates to compute ground-state configurations of spin-glass-like Hamiltonians. For example, the convergence of simulated annealing [11] can be considerably improved by adding isoenergetic cluster moves at each temperature step. Because the moves change the spin configurations but leave the energy of the system intact, the approach has the potential to “tunnel” through energy barriers, thus improving overall convergence.

We first introduce the benchmark model, followed by a short description of the Houdayer cluster algorithm and an outline of our isoenergetic cluster algorithm. Results in two and three space dimensions, as well as on the nonplanar chimera topology [12] are presented.

*Benchmark model and observables.*—The Hamiltonian of a generic Ising spin glass is defined by  $\mathcal{H} = \sum_{i \neq j}^N J_{ij} s_i s_j$ , where  $s_i \in \{\pm 1\}$  represent Ising spins and  $N$  is the total number of spins. In this study the interactions  $J_{ij}$  are selected from a Gaussian distribution with mean zero and variance  $J^2 = 1$ . Because we are only interested in highlighting the improved thermalization by adding isoenergetic cluster moves, we measure the average energy per spin defined via  $[\langle \mathcal{H} \rangle]/N$ , as well as the link overlap  $q_\ell = (1/N_b) \sum_{ij}^N s_i^{(1)} s_j^{(1)} s_i^{(2)} s_j^{(2)}$ . Here,  $\langle \dots \rangle$  represents a Monte Carlo average, the superscripts represent two replicas of the system,  $[\dots]$  indicates an average over the disorder, and  $N_b$  is the number of bonds in the system. Using Gaussian disorder, one can equate the internal energy per spin to the internal energy computed from the link overlap [13],  $E(q_\ell)$ , i.e.,

$$E(q_\ell) = -\frac{J^2 N_b}{T N} (1 - q_\ell). \quad (1)$$

To test that the system is thermalized, we thus study the time-dependent behavior of

$$\Delta = [\langle E(q_\ell) \rangle - \langle \mathcal{H}/N \rangle]. \quad (2)$$

When  $\Delta \rightarrow 0$ , the bulk of the disorder instances are thermalized [14]. Simulation parameters are listed in Table I.

*Reminder: Houdayer cluster algorithm.*—The Houdayer cluster algorithm (HCA) [7] is an efficient algorithm to study *two-dimensional* Ising spin glasses at low temperatures where thermalization is slow. It is similar to replica Monte Carlo [19], but with the difference that both replicas are at the *same* temperature. By allowing large cluster rearrangements of configurations, the HCA improves thermalization by efficiently tunneling through configuration space.

The algorithm works as follows: In the HCA, two independent spin configurations (replicas) are simulated at the same temperature. The site overlap between replicas (1) and (2),  $q_i = s_i^{(1)} s_i^{(2)}$ , is calculated. This creates two domains in  $q$  space: sites with  $q_i = 1$  and  $q_i = -1$ . Clusters are defined as the connected parts of these domains in  $q$  space. One then randomly chooses one site with  $q_i = -1$  and builds the cluster by adding all of the connected spins in the domain with probability 1. When no more spins can be added to the cluster in  $q$  space, the spins in *both* replicas that correspond to cluster sites are flipped with probability 1, irrespective of their orientation. The method can be implemented in a very efficient way because sites are added to the cluster with probability 1 and the cluster updates are rejection free. To ensure ergodicity, the cluster move is combined with standard single-spin Monte Carlo updates. Summarizing, one simulation step using the HCA consists of the following steps: (1) Perform one Monte Carlo sweep ( $N$  Metropolis updates) in each replica. (2) Perform one Houdayer cluster move. (3) Perform one parallel tempering update for a pair of neighboring temperatures.

Note that the last step is not necessary; however, the combination of the HCA moves and parallel tempering

TABLE I. Parameters of the simulation in two space dimensions (2D), three space dimensions (3D), and on the chimera (Ch) topology. For each topology simulated and system sizes  $N$ , we compute  $N_{\text{sa}}$  disorder instances and measure over  $2^b$  Monte Carlo sweeps (and isoenergetic cluster moves) for each of the  $2N_T$  replicas.  $T_{\text{min}}$  ( $T_{\text{max}}$ ) is the lowest (highest) temperature simulated, and  $N_T$  is the total number of temperatures used in the parallel tempering Monte Carlo method. Isoenergetic cluster moves only occur for the lowest  $N_c$  temperatures simulated (determined from Fig. 1).

	$N$	$N_{\text{sa}}$	$b$	$T_{\text{min}}$	$T_{\text{max}}$	$N_T$	$N_c$
2D	256, 576, 1024	$10^4$	22	0.2120	1.6325	30	30
Ch	128, 288, 512, 800, 1152	$10^4$	22	0.2120	1.6325	30	19
3D	64, 216, 512, 1000, 1728	$1.510^4$	23	0.4200	1.8000	26	13

(PT) updates improves thermalization considerably and represents the standard *modus operandi*.

In theory, the efficiency of the HCA depends strongly on the percolation threshold of the desired topology to be simulated. Because spins are added to the cluster with probability 1, if the percolation threshold of the studied lattice is below 50%, then the cluster might span the entire system and an update will not yield a new configuration. This is the reason why the HCA is claimed to only work in two space dimensions [7] where the percolation threshold is above 50% (see also Fig. 1, top panel).

*Isoenergetic cluster algorithm.*—Our proposed isoenergetic cluster moves are closely related to the HCA. We begin by simulating two replicas with the same disorder at multiple temperatures. The cluster moves alone are not ergodic, so, again, these must be combined with simple Monte Carlo updates. One simulation step using isoenergetic cluster moves consists of the following steps: (1) Perform one Monte Carlo sweep ( $N$  Metropolis updates) in each replica. (2a) If the number of cluster sites with  $q_i = -1$  is greater than  $N/2$ , then all of the spins in one of the configurations can be flipped (because of spin-reversal symmetry), thus reducing the cluster size while leaving the energy unchanged. (2b) Perform one Houdayer cluster move for all temperatures  $T \lesssim J$ . (3) Perform one parallel tempering update for a pair of neighboring temperatures.

The main difference thus lies in applying cluster moves to a carefully selected set of temperatures where the isoenergetic cluster moves (ICMs) are efficient (steps 2a and 2b) because clusters do not percolate, as well as reducing cluster sizes and thus the numerical overhead by exploiting spin-reversal symmetry (step 2a) [20,21]. For example, in the case of the chimera lattice, the overhead of the ICM over PT is approximately 25% and is roughly independent of the system size for the studied  $N$ . However, the overhead for the HCA over PT is at least 50% and grows with increasing system size.

Figure 1 shows the fraction of spins with negative overlap (i.e., the fraction of potential cluster sites) as a function of temperature  $T$  for different system sizes  $N$  and on three different topologies. The top panel of Fig. 1 shows data in two space dimensions where the percolation threshold is  $p_c \approx 0.592$  [15] (the solid horizontal line). As such, for all temperatures simulated, the fraction of cluster sites is below the percolation threshold and saturates at 50% for  $T \rightarrow \infty$ . This means that isoenergetic cluster updates are efficient for all temperatures studied because the clusters never percolate. Naively, one would expect that in higher space dimensions clusters percolate for all  $T$ 's. This is, however, not the case due to the frustration present in spin glasses, as can be seen for the chimera topology (the center panel of Fig. 1) or in three space dimensions (the bottom panel of Fig. 1). For increasing system size the fraction of cluster sites converges to a limiting curve that crosses the

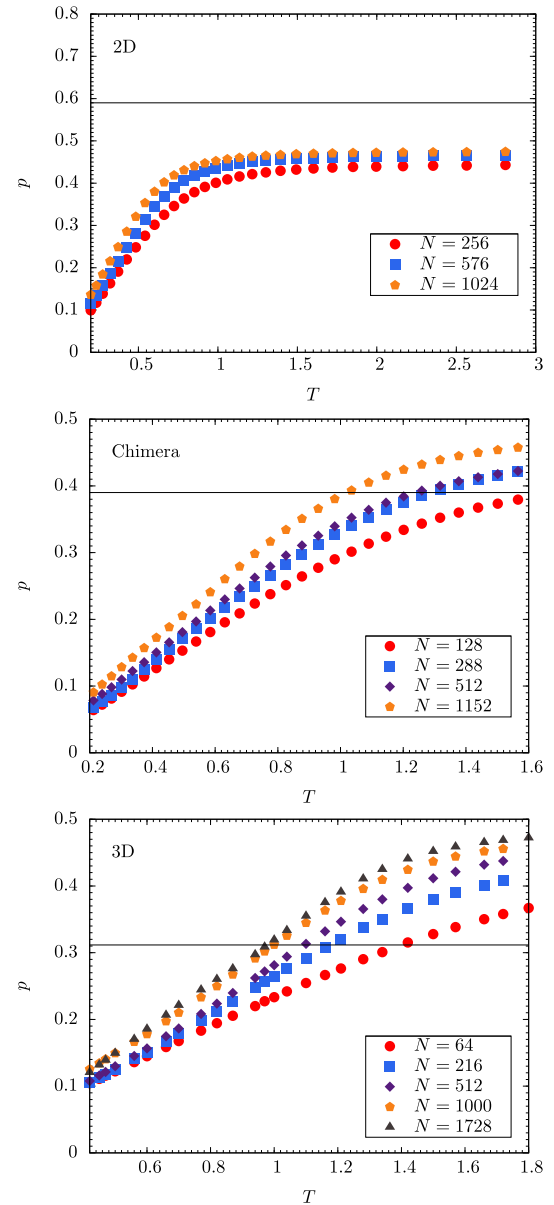


FIG. 1 (color online). (Top panel) Fraction of spins  $p$  of potential cluster sites as a function of temperature  $T$  for different system sizes  $N$  in two space dimensions (2D). The horizontal line represents the percolation threshold of a two-dimensional square lattice, i.e.,  $p_c \approx 0.592$  [15]. Because  $p \rightarrow 0.5$  for  $T \rightarrow \infty$ , for all  $T$ , clusters do not percolate, which is why the HCA is efficient in two-dimensional planar geometries. (Center panel)  $p$  as a function of temperature  $T$  for different system sizes  $N$  on the chimera topology. The horizontal line represents the percolation threshold of the nonplanar chimera topology, namely,  $p_c \approx 0.387$ , computed here using the approach developed in Ref. [16] (see the Supplemental Material [17]). For  $T \gtrsim J = 1$  clusters percolate and cluster updates provide no gain. (Bottom panel)  $p$  as a function of temperature  $T$  for different system sizes  $N$  in three space dimensions (3D). The horizontal line represents the percolation threshold of the three-dimensional cubic lattice ( $p_c \approx 0.311$  [18]). For  $T \gtrsim J = 1$  clusters percolate. In all panels, error bars are computed via a jackknife analysis over configurations and are smaller than the symbols.

percolation threshold (the horizontal solid lines) at approximately  $T \approx J = 1$ . This means that, for all  $T \gtrsim J$ , clusters percolate and the cluster updates are just numerical overhead without any advantage to the simulation. However, for  $T \lesssim J$  the fraction of cluster sites lies below the percolation threshold. This means that performing cluster moves in this temperature regime should improve thermalization. Note that it is a coincidental property that for three-dimensional Ising spin glasses  $T_c \sim 1$  [22], i.e., that cluster moves can be applied to any  $T \lesssim T_c$  [23].

When the interactions  $J_{ij}$  are drawn from a Gaussian distribution, the ground state is unique. As can be seen in Fig. 1, the fraction  $p$  of spins potentially in a cluster also approaches zero for  $T \rightarrow 0$ ; i.e., both replicas are in the ground state for low enough  $T$ . Therefore, the cluster is composed of no sites or the entire lattice. In the case of disorder distributions that yield a highly degenerate ground state, such as is the case for bimodal disorder, it is possible to continue to have clusters at zero temperature. It is thus possible to efficiently hop around the ground-state manifold by applying cluster moves to low-lying or even zero-temperature states, although this might not be ergodic. We do emphasize, however, that if clusters are too small, then the isoenergetic cluster moves also become ineffective. Therefore, plotting  $p$  as was done in Fig. 1 is essential in determining the efficiency and applicability of the method.

**Benchmarking results.**—Figure 2 shows  $\Delta$  [Eq. (2)] as a function of Monte Carlo time (measured in lattice sweeps)  $t = 2^b$ . The top panel of Fig. 2 shows data in two space dimensions for simulations using isoenergetic cluster moves (PT + ICM) and vanilla PT Monte Carlo for  $N = 1024$  spins at  $T = 0.212$ . Once  $\Delta \sim 0$ , we deem the system thermalized. Clearly, the inclusion of cluster moves—as can also be expected from the results of Houdayer—shows an improved thermalization. The center panel of Fig. 2 shows data on the chimera topology with  $N = 1152$  spins and  $T = 0.212$ , where the HCA is not expected to show any improvement over PT due to  $p_c < 0.5$ . As can be seen, our ICM clearly improve thermalization in comparison to PT by at least 2 orders of magnitude, an amount that grows with increasing system size. Finally, the bottom panel of Fig. 2 shows  $\Delta$  as a function of simulation time in three space dimensions with  $N = 1728$  spins and  $T = 0.42 \ll T_c$ . Although not as impressive as with the chimera topology, we see a speedup of approximately one order of magnitude—an amount that again grows with increasing system size.

Finally, Fig. 3 shows the ratio of the thermalization time using PT and using PT + ICM for different topologies at the lowest simulation temperature (see Table I) as a function of the system size  $N$ . In all cases, the speedup increases with increasing system size, therefore illustrating that the addition of isoenergetic cluster moves greatly improves thermalization.




FIG. 2 (color online). (Top panel)  $\Delta$  [Eq. (2)] as a function of simulation time  $t = 2^b$  measured in Monte Carlo sweeps in two space dimensions (2D) for  $N = 1024$  and  $T = 0.212$ . Simulations using vanilla PT thermalize at least  $2^{25}$  Monte Carlo sweeps, whereas with the addition of ICMs thermalization is reduced to approximately  $2^{16}$  Monte Carlo sweeps. This means approximately 2 orders of magnitude improvement. (Center panel)  $\Delta$  as a function of simulation time  $t = 2^b$  measured in Monte Carlo sweeps for an Ising spin glass on chimera with  $N = 1152$  spins at  $T = 0.212$ . Simulations using PT thermalize at approximately  $2^{25}$  Monte Carlo sweeps, whereas the addition of ICMs reduces thermalization to  $2^{18}$  Monte Carlo sweeps. Again, approximately 2 orders of magnitude speedup. (Bottom panel)  $\Delta$  as a function of simulation time  $t = 2^b$  measured in Monte Carlo sweeps in three space dimensions (3D) for  $N = 1728$  and  $T = 0.42 \sim 0.43T_c$ . Using standard PT, the system thermalizes approximately after  $2^{23}$  Monte Carlo sweeps. This time is reduced to  $\sim 2^{20}$  Monte Carlo sweeps when ICMs are added. In all panels, error bars are computed via a jackknife analysis over configurations.






FIG. 3 (color online). Ratio between the *approximate* average thermalization time of PT and PT + ICM for different topologies at the lowest simulation temperature (see Table I) as a function of system size  $N$ . In all cases the speedup increases with increasing system size. Note that thermalization times have been determined by eye.

*Summary.*—We have presented a rejection-free cluster algorithm for spin glasses in any space dimension that greatly improves thermalization. By restricting Houdayer cluster moves to temperatures where cluster percolation is hampered by the interplay of frustration and temperature, we are able to extend the Houdayer cluster algorithm for two-dimensional spin glasses to any topology or space dimension. Our standard implementation of the cluster updates represents only a minor overhead [21] compared to the thermalization time speedup obtained from the iso-energetic cluster moves—a speedup that increases with the system size [26].

We would like to thank F. Hamze, J. Machta, and M. Weigel for the fruitful discussions. H. G. K. acknowledges support from the NSF (Grant No. DMR-1151387) and would like to thank Suntory’s Hibiki and Yamazaki for inspiration. We thank Texas A&M University for the extensive CPU time on the Ada cluster. This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via MIT Lincoln Laboratory Air Force Contract No. FA8721-05-C-0002. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purpose.

[1] D. L. Stein and C. M. Newman, *Spin Glasses and Complexity*, Primers in Complex Systems (Princeton University Press, Princeton, NJ, 2013).

- [2] K. Hukushima and K. Nemoto, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).
- [3] J. Machta, *Phys. Rev. E* **82**, 026704 (2010).
- [4] P. Hohenberg and B. Halperin, *Rev. Mod. Phys.* **49**, 435 (1977).
- [5] R. H. Swendsen and J.-S. Wang, *Phys. Rev. Lett.* **58**, 86 (1987).
- [6] U. Wolff, *Phys. Rev. Lett.* **62**, 361 (1989).
- [7] J. Houdayer, *Eur. Phys. J. B* **22**, 479 (2001).
- [8] L. Chayes, J. Machta, and O. Redner, *J. Stat. Phys.* **93**, 17 (1998).
- [9] J. Machta, C. M. Newman, and D. L. Stein, *J. Stat. Phys.* **130**, 113 (2008).
- [10] T. Jörg, *Prog. Theor. Phys. Suppl.* **157**, 349 (2005).
- [11] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, *Science* **220**, 671 (1983).
- [12] P. Bunyk, *IEEE Trans. Appl. Supercond.* **24**, 1 (2014).
- [13] H. G. Katzgraber, M. Palassini, and A. P. Young, *Phys. Rev. B* **63**, 184422 (2001).
- [14] B. Yuceosoy, J. Machta, and H. G. Katzgraber, *Phys. Rev. E* **87**, 012104 (2013).
- [15] X. Feng, Y. Deng, and H. W. J. Blöte, *Phys. Rev. E* **78**, 031136 (2008).
- [16] O. Melchert, *Phys. Rev. E* **87**, 042106 (2013).
- [17] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.115.077201> for a calculation of the percolation threshold of the chimera lattice.
- [18] J. Wang, Z. Zhou, W. Zheng, T. M. Garoni, and Y. Deng, *Phys. Rev. E* **87**, 052107 (2013).
- [19] R. H. Swendsen and J.-S. Wang, *Phys. Rev. Lett.* **57**, 2607 (1986).
- [20] Jon Machta tells us that the method might also work for ferromagnetic systems at low temperatures. Here, the step of flipping the whole spin configuration of one replica if the two replicas are in opposite pure states is crucial. In this case both replicas will be in the same pure state (all spins up or all spins down), and the clusters will be small in a ferromagnet as well.
- [21] Note that in step (2a) of the ICM algorithm we flip the signs of the spins to reduce the overhead whenever the number of cluster sites exceeds  $N/2$ . To ensure that the measurement of observables is not affected by this operation, the signs of the flipped spins need to be kept track of.
- [22] H. G. Katzgraber, M. Körner, and A. P. Young, *Phys. Rev. B* **73**, 224432 (2006).
- [23] Note that for the three-dimensional case, one may be tempted to think that clusters do not percolate because the spins are “frozen” below the critical temperature  $T_c \sim 1$ . However, this is not the case. In four space dimensions,  $T_c \sim 1.8$  [24] and  $p_c = 0.197$  [25]. However, clusters still only stop percolating for  $T \lesssim 1$ .
- [24] T. Jörg and H. G. Katzgraber, *Phys. Rev. B* **77**, 214426 (2008).
- [25] G. Paul, R. M. Ziff, and H. E. Stanley, *Phys. Rev. E* **64**, 026115 (2001).
- [26] We note that our ICMs also slightly reduce autocorrelation times in addition to greatly improving thermalization times. However, note that the error due to configurational averages dominates in spin-glass simulations.