Simulating thermal boundary conditions of spin-lattice models with weighted averages

Wenlong Wang*

*Department of Physics, University of Massachusetts, Amherst, Massachusetts 01003 USA and
Department of Physics and Astronomy, Texas A&M University, College Station, Texas 77843-4242, USA*

Thermal boundary conditions has played an increasingly important role in revealing the nature of short-range spin glasses and is likely to be relevant also for other disordered systems. Diffusion method initializing each replica with a random boundary condition at the infinite temperature using population annealing has been used in recent large-scale simulations. However, the efficiency of this method can be greatly suppressed because of temperature chaos. For example, most samples have some boundary conditions that are completely eliminated from the population in the process of annealing at low temperatures. In this work, I study a weighted average method to solve this problem by simulating each boundary conditions separately and collect data using weighted averages. The efficiency of the two methods are studied using both population annealing and parallel tempering, showing that the weighted average method is more efficient and accurate.

I. INTRODUCTION

Thermal boundary conditions (TBC) includes the set of all combinations of periodic/antiperiodic boundary conditions in each spatial dimension [1–3]. For example in three dimensions ($d = 3$), there are $2^d = 8$ possible choices. Each boundary condition i occurs in the thermal ensemble with weight p_i depends on its free energy F_i as $p_i = \exp[-\beta(F_i - F)]$, where F is the total free energy of the system in TBC and β is the inverse temperature. TBC was initially motivated to reduce domain-wall effects in spin glasses [1] and was later shown to be useful in studying temperature chaos and bond chaos via boundary condition crossings, i.e, the weights $\{p_i\}$ change chaotically as a function of temperature or couplings [2, 3]. Previously, thermal boundary conditions was used with exact algorithms for finding ground states of two-dimensional spin glasses [4, 5] (referred to as “extended” boundary conditions). A restricted version of thermal boundary conditions using periodic and antiperiodic boundary conditions in a single direction (a subset of TBC) was also used in Refs. [6–9].

Simulating the full TBC ensemble in Refs. [1–3] was done using population annealing (PA) [10–13]. See Ref. [14] for a recent discussion of the algorithm. Population annealing works with a population of replicas and aims to maintain thermal equilibrium while lowering the temperature. When temperature is decreased, the population is resampled according to the Boltzmann weights of the replicas, followed by regular Monte Carlo sweeps to all replicas in the population. In this work, the Metropolis algorithm is used. Simulating TBC using the diffusion method with PA is very simple. One can initialize each replica with a random boundary condition at $\beta = 0$ and then boundary conditions are resampled along with the replicas. The word “diffusion” would become apparent when one implements the method in parallel tempering,

as discussed in Sec. II C.

It was noticed in Refs. [1, 2] that the efficiency of the diffusion method can be greatly suppressed by temperature chaos. Some boundary conditions can be totally removed from the population. This is not satisfactory as there is no mechanism to recover these lost boundary conditions once they are eliminated from the population. Furthermore, temperature chaos predicts that these boundary conditions could become important again at lower temperatures. Therefore, it is worth to study a new method that does not have this problem. The most natural way is to simulate each boundary condition separately and combine the observables using weighted averages with free energy. How to weight different kinds of observables is discussed in Sec. II D. Since it is also interesting to study the performance of parallel tempering (PT) in simulating the full set of thermal boundary conditions, the two methods are therefore studied in both PA and PT.

The paper is organized as follows: Section II introduces the Edwards-Anderson model, measured quantities, the diffusion method and the weighted average method. Numerical results are shown in Sec. III, followed by concluding remarks and future challenges in Sec. IV.

II. MODELS, OBSERVABLES AND METHODS

A. The Edwards-Anderson model

The Edwards-Anderson (EA) Hamiltonian is defined as

$$H = - \sum_{\langle ij \rangle} J_{ij} S_i S_j, \quad (1)$$

where $\{S_i = \pm 1\}$ are the spin degrees of freedom defined on a three-dimensional cubic lattice. The sum over $\langle ij \rangle$ means sum over all nearest neighbours. J_{ij} is the coupling between spins S_i and S_j and all couplings are independently chosen from the standard Gaussian distribution with mean 0 and variance 1. Thermal boundary

*Electronic address: wenlong@physics.umass.edu

conditions is applied to all instances.

B. Observables

There are three classes of observables that need to be treated differently using weighted averages. The first class are observables that are functions of a single replica like the energy E or the magnetization m . The second class are the thermodynamic observables of the free energy F and the entropy S . There is a third class of observables in spin glasses due to the nature of the symmetry breaking which are functions of two replicas like the spin overlap q defined as

$$q_{ab} = \frac{1}{N} \sum_i S_i^a S_i^b, \quad (2)$$

where micro-states a, b are chosen independently from the Boltzmann distribution. Note that a, b can be chosen from the same boundary conditions or different boundary conditions. Another quantity in this class is the link overlap which needs some care in the definition due to the change of boundary conditions, the link overlap q_ℓ is defined as

$$q_{\ell ab} = \frac{1}{dN} \sum_{\langle ij \rangle} S_i^a \text{sign}(J_{ij}^a) S_j^a S_i^b \text{sign}(J_{ij}^b) S_j^b, \quad (3)$$

where the sign function is ± 1 depends on whether the argument is positive or negative respectively. The definition has nothing to do with the weighted average method, but this is a more useful definition in the TBC setting for studying the nature of spin glasses. In this way, the difference of the link overlap of two different boundary conditions arises only from the domain walls, not from a mixture of domain walls and boundary condition differences. In the following, I will focus on the spin overlap function and will refer to it as the overlap function where no confusion arises. Note also that there is a 8×8 overlap function matrix in the TBC ensemble. An important statistic I , which quantifies the weight of an overlap distribution near $q = 0$, is defined as

$$I = \int_{-0.2}^{0.2} P(q) dq, \quad (4)$$

where $P(q)$ is an overlap distribution function.

C. The diffusion method

We have already discussed the diffusion method in population annealing, which initializes each replica with a random boundary condition at $\beta = 0$. For completeness, I also introduce and study the diffusion method in parallel tempering because parallel tempering is widely used in spin-glass simulations. Following the same strategy, one can simulate the TBC using parallel tempering by

generating random states with random boundary conditions at $\beta = 0$. Then the boundary conditions diffuse along with the replicas under the swap moves of parallel tempering, hence the name the diffusion method. The convenience of working with $\beta = 0$ is because proposing a boundary condition change at finite temperatures can be costly as many bonds are affected when boundary condition changes. Furthermore, this is also essential to measure the absolute free energy.

The implementation of this method is simple and detailed balance is satisfied. However, the efficiency of the method still needs to be studied. Since in this method, PA and PT work in a similar way, one may expect that PT should suffer from temperature chaos too. This turns out to be the case as discussed in Secs. III A and III B. The effect of boundary condition eliminations in PA is replaced by one or more diffusion bottlenecks in PT. In the next, I discuss the weighted average method, which does not have this problem.

D. The weighted average method

It was shown that the absolute free energy can be measured very accurately using the free energy perturbation method in both population annealing and parallel tempering [11, 15]. It can be shown from statistical mechanics that the average energy, entropy, free energy and the spin overlap distribution should be averaged as

$$E = \sum_i E_i p_i, \quad (5)$$

$$S = \sum_i S_i p_i - \sum_i p_i \log p_i, \quad (6)$$

$$F = \sum_i F_i p_i + T \sum_i p_i \log p_i, \quad (7)$$

$$P(q) = \sum_{ij} P_{ij}(q) p_i p_j, \quad (8)$$

where $p_i = \frac{e^{-\beta F_i}}{\sum_j e^{-\beta F_j}}$, E_i, S_i and F_i are the energy, entropy and free energy of boundary condition i and $P_{ij}(q)$ is the overlap distribution between the boundary conditions i and j .

It is worth noting that the weighted average method generates a lot more information than the diffusion method. For example, the overlap matrix is briefly discussed in Sec. IV. For now, we discuss briefly of the implementation of the weighted average method in PA and PT. In the weighted average method, each boundary condition is simulated independently. It is only when computing the spin overlaps that one needs communications between replicas with different boundary conditions. Therefore, in this way, PT can be simulated using parallel computing too. Since it is usually a practice to simulate two independent sets of replicas of each bound-

any condition, one can use 8 or 16 threads in the simulation. This is not doable in the diffusion method of PT. On the other hand, PA is intrinsically parallel and is also much more flexible with the number of threads. In my simulations, I have used OpenMP parallel computing for both PA and PT. For the equilibration measure of PA, one can either use the weighted average or the minimum of the entropy of families [14] of all boundary conditions.

III. RESULTS

In this section, I investigate the performance of the diffusion method and the weighted average method. I first compare PA and PT for the diffusion method in Sec. III A. The conclusion is that the diffusion method, while works but can be inefficient and the limitations of the method are discussed in Sec. III B. Finally, I study the performance of the weighted average method, compare PA and PT and also with the diffusion method, showing that the weighted average method works better than the diffusion method in Sec. III C.

A. PA and PT: the diffusion method

In this section, I compare the efficiency of the diffusion method of PA and PT in simulating thermal boundary conditions. The large-scale population annealing data is taken from a recent work on temperature chaos [2]. In the implementation of parallel tempering, I used N_β temperatures evenly distributed in β for the high temperature part and N_T temperatures evenly distributed in T for the low temperature part. No Monte Carlo sweeps were done at $\beta = 0$, it is sufficient to generate random states and boundary conditions at the infinite temperature. The amount of work in this work is counted in Monte Carlo sweeps and one Monte Carlo sweep is a sequential update of all the spins for a replica at a temperature once. The work is distributed evenly between thermalization and data collection. The simulation parameters of PA and PT are summarized in Tables I and II respectively. The main conclusion of this section is that PA and PT has similar efficiency in the diffusion method. The limitations of the method are discussed in Sec. III B.

1. Detailed comparison of a single hard sample

In this section, I make a detailed comparison for the hardest instance of $L = 8$ chosen from a total of 2001 instances in Ref. [2]. I will focus on the systematic error of $-\beta F$, the q distribution at a low temperature $\beta = 5$ as well as the evolution of $\{p_i\}$ as a function of β . One can see from the evolution of $\{p_i\}$ that this sample is rather chaotic with multiply boundary condition crossings.

The convergence of the dimensionless quantity $-\beta F$ at a low temperature $\beta = 5$ is studied as a function

TABLE I: The simulation parameters of PA [2] using the diffusion method (DF) for different system sizes L with thermal boundary conditions. R represents the total number of replicas, $1/\beta_{\max}$ is the lowest temperature simulated, N_T the number of temperatures in the annealing schedule evenly distributed in β , N_S the number of sweeps per replica per temperature and M the number of samples. The simulation parameters using the weighted average method (WA) are also shown in the table, where R represents the population size of each boundary condition.

Method	L	R	$1/\beta_{\max}$	N_T	N_S	M
DF	4	$5 \cdot 10^4$	0.2	101	10	101
DF	6	$2 \cdot 10^5$	0.2	101	10	101
DF	8	$5 \cdot 10^5$	0.2	201	10	101
DF	10	$2 \cdot 10^6$	0.2	301	10	101
WA	4	$5 \cdot 10^3$	0.2	101	10	101
WA	6	$2 \cdot 10^4$	0.2	101	10	101
WA	8	$5 \cdot 10^4$	0.2	201	10	101
WA	10	$2 \cdot 10^5$	0.2	301	10	101
WA	12	$2 \cdot 10^5$	1/3	301	10	101

TABLE II: The simulation parameters of PT using the diffusion method (DF) for different system sizes L with thermal boundary conditions. N_β represents the number of temperatures in the high temperature part uniformly distributed in β , N_T the number of temperatures in the low temperature part uniformly distributed in T , $T_{\min}$ the lowest temperature simulated, $T_{\max}$ the highest temperature simulated in the low temperature part, N_S the number of sweeps per temperature and M the number of samples. The work is evenly distributed in two independent sets of replicas. The simulation parameters using the weighted average method (WA) are also shown in the table, where N_S represents the number of sweeps of each boundary condition.

Method	L	N_β	N_T	$T_{\min}$	$T_{\max}$	N_S	M
DF	4	5	10	0.2	2.0	$2 \cdot 10^6$	101
DF	6	5	20	0.2	2.0	$4 \cdot 10^6$	101
DF	8	10	30	0.2	2.0	$1.2 \cdot 10^7$	101
DF	10	10	40	0.2	2.0	$6 \cdot 10^7$	101
WA	4	5	10	0.2	2.0	$2 \cdot 10^5$	101
WA	6	5	20	0.2	2.0	$4 \cdot 10^5$	101
WA	8	10	30	0.2	2.0	$1.2 \cdot 10^6$	101
WA	10	10	40	0.2	2.0	$6 \cdot 10^6$	101
WA	12	20	40	1/3	2.0	$5 \cdot 10^6$	101

of the amount of work W . Even though two independent sets of replicas are simulated in parallel tempering to compute the overlap distribution and data were collected from both sets, when studying the convergence of the free energy, this factor of 2 is not counted since this only improves statistical errors, not systematic errors. Indeed, measuring free energy does not require two sets of independent runs. One can say that it is an advantage of population annealing that can effectively compute the

overlap distribution from one run due to the ensemble nature of the algorithm. The same parameters in Tables I and II are used for this instance and the amount of work is varied by changing the population size in PA and the number of sweeps in PT. The convergence of $-\beta F$ is shown in the top panel of Fig. 1. The standard is taken from a weighted average [11] of the largest data set of PA, and the errorbars in the plot are estimated from multiple independent runs. The overlap distribution and the evolution of $\{p_i\}$ as a function of β are shown in Fig. 2. The results are randomly chosen from the largest runs of PA and PT in Fig. 1. We see that PA and PT have similar performance in the diffusion method.

FIG. 1: (Color online) Linear-log plot of the systematic error $\Delta(-\beta F)$ of PA and PT as a function of the amount of work W for a chaotic instance of $L = 8$ at $T = 0.2$. Top panel: the diffusion method (DF). Bottom panel: the weighted average method (WA). The errorbars are the standard deviation of the $-\beta F$ distribution computed from multiple independent runs, not the errorbar of the sample mean of $-\beta F$. Note that even though the weighted average simulates all the 8 boundary conditions independently, it still converges faster than the diffusion method and is therefore more efficient.

2. A large-scale comparison

Now we look at a large-scale comparison of PA and PT using the diffusion method. Two important quantities for spin glasses are compared: the free energy per spin f and the statistic I of the overlap distribution. The results for f and I are shown in Figs. 3 and 4 respectively, again showing the similar performance of PA and PT.

B. Limitations of the diffusion method

We have shown in the last section that one can use the diffusion method to simulate thermal boundary conditions. This method is simple to implement and yields fair result. However, the method has its own limitations due to the diffusion nature of the method, in particular, the problem of the boundary condition dying out.

FIG. 2: (Color online) The comparison of the evolution of $\{p_i\}$ as a function of β of PA and PT using the diffusion method. The solid lines are for population annealing and the points are for parallel tempering. Each color represents a boundary condition. The inset is the comparison of the overlap distribution at $\beta = 5$. The results are randomly chosen from the largest runs of Fig. 1.

FIG. 3: (Color online) Comparison of the free energy per spin f at $\beta = 5$ of population annealing and parallel tempering using the diffusion method.

The problem of boundary condition dying out of the diffusion method is actually a very severe one. In the state-of-the-art simulation of Ref. [2], where most samples if not all are well equilibrated with decent simulation parameters, it is striking that most of the samples suffer from this problem as shown in Table III and Fig. 5. One can see that about 77% of instances have some boundary conditions eliminated at $T = 1/3$ and the number quickly grows to about 98% at $T = 0.2$. The distribution of the number of boundary conditions that are eliminated n_D also quickly shifts from a bias towards $n_D = 0$ to $n_D = 7$.

This problem is a resolution problem as the resolution of the weight of a boundary condition is $1/R$. A bound-

FIG. 4: (Color online) Comparison of the statistic I at $\beta = 5$ of population annealing and parallel tempering using the diffusion method.

TABLE III: The fraction of instances with boundary condition dying out $f_{\mathcal{J}}$ in the large-scale simulation of Ref. [2] with $M = 2001$ samples of each system size L . In contrast, no boundary condition dies out in the weighted average method.

L	4	6	8	10	10	12
T	0.2	0.2	0.2	0.2	1/3	1/3
$f_{\mathcal{J}}$	0.9835	0.9840	0.9820	0.9720	0.7701	0.7731

ary condition i is likely to be eliminated when p_i becomes close to $1/R$: $p_i \sim 1/R$. When $p_i < 1/R$, it is not even expected to be part of the ensemble in most independent runs. Unluckily, there is no mechanism in the diffusion method to reintroduce a boundary condition once it is eliminated. To “confirm” that the simulation is fine, one can only try a much larger R , like a factor of 10 as used in Ref.[2], and see if the results remain the same. But this is not satisfactory as the runs can be consistently wrong and this is also not efficient. This is especially the case for large system sizes as temperature chaos asserts that the dominate boundary conditions switches more frequently as system size increases. The problem of the total absence of some boundary conditions also occurs in the parallel tempering implementation. Bottlenecks in the diffusion channel prevents some boundary conditions to reach the lowest temperature. This is not too surprising considering the similar performance of PA and PT. Therefore, this is a limitation of the diffusion method itself, independent of which algorithm one uses to implement it. In the next section, I study the weighted average method, its validity as well as its efficiency.

FIG. 5: (Color online) Two typical distribution of the number of boundary conditions that die out n_D of $L = 10$ at $T = 1/3$ and 0.2 respectively in the diffusion method. Note that as temperature is lowered a little in the low-temperature phase, a dramatic change occurs in the distribution as many boundary conditions are eliminated in the annealing process. In contrast, no boundary condition dies out in the weighted average method.

C. PA and PT: the weighted average method

In this section, I present a detailed study of the same hard instance as well as a large-scale simulation using the weighted average method with PA and PT. At first sight, this might be less efficient as one has to simulate all the 8 boundary conditions independently. However, as we shall see in this section that this is not the case. For the hard instance studied, it is even more efficient than the diffusion method. In addition, the problem of boundary conditions elimination is completely removed in the weighted average method.

1. Detailed comparison of the single hard sample

The convergence of $-\beta F$ is shown in the bottom panel of Fig. 1. Compare with the diffusion method in the top panel, one can see that it is more efficient. This is interesting and therefore the weighted average method is both more efficient and more accurate. PA and PT are still comparable in performance. One may have noticed that PT converges slightly faster, but there are more factors to consider here. One is that PT often requires two independent sets of replicas, which are not counted here, to compute the spin overlaps. Take this into account, PA and PT are similar in performance. Another factor is that PT also has slightly more overhead in parallel computing due to the frequent breaks to measure the overlap functions. However, one can choose to run the code in sequential like in the diffusion method. Therefore, it is more fair to conclude that PA and PT are comparable in efficiency. The overlap distribution and the evolution of $\{p_i\}$ as a function of β are shown in Fig. 6. The results

are randomly chosen from the largest runs of PA and PT in Fig. 1. The two algorithms give very similar result.

FIG. 6: (Color online) The comparison of the evolution of $\{p_i\}$ as a function of β of PA and PT using the weighted average method. The solid lines are for population annealing and the points are for parallel tempering. The results are randomly chosen from the largest runs of Fig. 1. The inset is the comparison of the overlap distribution at $\beta = 5$. Note that even with slightly less computational work, the agreement is better than that of Fig. 2 using the diffusion method, showing that the weighted average method is more accurate.

2. A large-scale simulation

Finally, we do a large-scale comparison of PA and PT. We have studied the EA model up to $L = 10$ at $T = 0.2$ and $L = 12$ at $T = 1/3$ (current state-of-the-art) using the weighted average method, demonstrating that the method can be used for an accurate large-scale study of spin glasses with thermal boundary conditions. The comparison of the free energy per spin f and the statistic I of the overlap distribution are shown in Figs. 7 and 8 respectively, again showing the similar performance of PA and PT. It also worth noting that the agreement in the measurements of I is better than that the diffusion method even though with slightly less amount of work, in line with the results of the single hard instance.

IV. CONCLUSIONS AND FUTURE CHALLENGES

In this paper, I studied two methods to simulate thermal boundary conditions, the diffusion method and the weighted average method using both population annealing and parallel tempering. I have illustrated the weaknesses of the diffusion method, in particular, the elimination of boundary conditions. The weighted average method is studied in detail, showing that it is more efficient than the diffusion method and solves the boundary condition elimination problem. Population annealing

FIG. 7: (Color online) Comparison of the free energy per spin f of population annealing and parallel tempering using the weighted average method. $\beta = 5$ for $L = 6, 8$ and 10 and $\beta = 3$ for $L = 12$.

FIG. 8: (Color online) Comparison of the statistic I of population annealing and parallel tempering using the weighted average method. $\beta = 5$ for $L = 6, 8$ and 10 and $\beta = 3$ for $L = 12$.

and parallel tempering on the other hand are similar in their performance.

Furthermore, the weighted method method generates a lot more information than the diffusion method and has opened new ways in doing spin-glass research in thermal boundary conditions. Detailed information of each boundary condition is accessible, and the spin and link overlap distributions between any two boundary conditions are also readily available. For example, it is interesting to study the overlap distributions in the whole in-

stance as well as in a window between the same boundary conditions and two different boundary conditions. This can be very useful in studying the nature of domain walls. i.e. whether the low-lying excitations are space filling or have fractal surfaces. It is also interesting to study the structure of the spin overlap matrix. Except studying the size dependence of the traditional I in periodic boundary conditions and thermal boundary conditions, one can also study the weighted or minimum I of the diagonal part of the overlap matrix. In this way, one can remove effects of the overlaps between different boundary conditions, which are generally larger than the diagonal ones near the origin, and therefore may further reduce finite-size effects. Efforts along these directions are currently

in progress and will be presented in future publications.

Acknowledgments

I gratefully acknowledge support from NSF (Grant No. DMR-1208046 and DMR-1151387). I thank Jon Machta for helpful discussions and suggestions, and also for the careful reading of the manuscript. I thank the anonymous referee for insightful suggestions for this manuscript. Finally, I thank the Texas A&M University for access to their Ada and Curie clusters.

-
- [1] W. Wang, J. Machta, and H. G. Katzgraber, Phys. Rev. B **90**, 184412 (2014).
 - [2] W. Wang, J. Machta, and H. G. Katzgraber, Phys. Rev. B **92**, 094410 (2015).
 - [3] W. Wang, J. Machta, and H. G. Katzgraber, ArXiv:1603.00543 (2016).
 - [4] J. W. Landry and S. N. Coppersmith, Phys. Rev. B **65**, 134404 (2002).
 - [5] C. K. Thomas and A. A. Middleton, Phys. Rev. B **76**, 220406(R) (2007).
 - [6] K. Hukushima, Phys. Rev. E **60**, 3606 (1999).
 - [7] M. Sasaki, K. Hukushima, H. Yoshino, and H. Takayama, Phys. Rev. Lett. **95**, 267203 (2005).
 - [8] M. Sasaki, K. Hukushima, H. Yoshino, and H. Takayama, Phys. Rev. Lett. **99**, 137202 (2007).
 - [9] M. Hasenbusch, Physica A **197**, 423 (1993).
 - [10] K. Hukushima and Y. Iba, in *The Monte Carlo method in the physical sciences: celebrating the 50th anniversary of the Metropolis algorithm*, edited by J. E. Gubernatis (AIP, 2003), vol. 690, p. 200.
 - [11] J. Machta, Phys. Rev. E **82**, 026704 (2010).
 - [12] J. Machta and R. Ellis, J. Stat. Phys. **144**, 541 (2011).
 - [13] E. Zhou and X. Chen, in *Proceedings of the 2010 Winter Simulation Conference (WSC)* (2010), p. 1211.
 - [14] W. Wang, J. Machta, and H. G. Katzgraber, Phys. Rev. E **92**, 063307 (2015).
 - [15] W. Wang, Phys. Rev. E **91**, 053303 (2015).