

Estimating the Density of States of Frustrated Spin Systems

Barash, L. Y., Marshall, J., Weigel, M. & Hen, I.

Published PDF deposited in Coventry University's Repository

Original citation:

Barash, LY, Marshall, J, Weigel, M & Hen, I 2019, 'Estimating the Density of States of Frustrated Spin Systems' New Journal of Physics, vol. 21, 073065.

<https://dx.doi.org/10.1088/1367-2630/ab2e39>

DOI 10.1088/1367-2630/ab2e39

ISSN 1367-2630

ESSN 1367-2630

Publisher: IOP Publishing:

CC-BY Original content from this work may be used under the terms of the Creative Commons Attribution 3.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

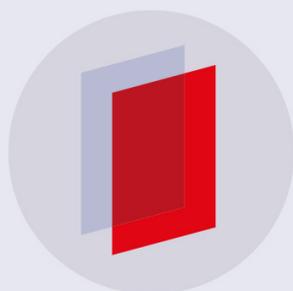
Copyright © and Moral Rights are retained by the author(s) and/ or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This item cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder(s). The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

PAPER • OPEN ACCESS

Estimating the density of states of frustrated spin systems

To cite this article: Lev Barash *et al* 2019 *New J. Phys.* **21** 073065

View the [article online](#) for updates and enhancements.



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.



PAPER

Estimating the density of states of frustrated spin systems

OPEN ACCESS

RECEIVED

19 September 2018

REVISED

24 May 2019

ACCEPTED FOR PUBLICATION

1 July 2019

PUBLISHED

30 July 2019

Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Lev Barash^{1,2,3} , Jeffrey Marshall⁴ , Martin Weigel⁵ and Itay Hen^{4,6} ¹ Landau Institute for Theoretical Physics, 142432 Chernogolovka, Russia² Science Center in Chernogolovka, 142432 Chernogolovka, Russia³ National Research University Higher School of Economics, 101000 Moscow, Russia⁴ Department of Physics and Astronomy, and Center for Quantum Information Science & Technology, University of Southern California, Los Angeles, CA 90089, United States of America⁵ Applied Mathematics Research Centre, Coventry University, Coventry, CV1 5FB, United Kingdom⁶ Information Sciences Institute, University of Southern California, Marina del Rey, CA 90292, United States of AmericaE-mail: barash@itp.ac.ru**Keywords:** entropic sampling methods, statistical mechanics, phase transitions, quantum computation, spin glasses, density of states

Abstract

Estimating the density of states (DOS) of systems with rugged free energy landscapes is a notoriously difficult task of the utmost importance in many areas of physics ranging from spin glasses to biopolymers. DOS estimation has also recently become an indispensable tool for the benchmarking of quantum annealers when these function as samplers. Some of the standard approaches suffer from a spurious convergence of the estimates to metastable minima, and these cases are particularly hard to detect. Here, we introduce a sampling technique based on population annealing enhanced with a multi-histogram analysis and report on its performance for spin glasses. We demonstrate its ability to overcome the pitfalls of other entropic samplers, resulting in some cases in large scaling advantages that can lead to the uncovering of new physics. The new technique avoids some inherent difficulties in established approaches and can be applied to a wide range of systems without relevant tailoring requirements. Benchmarking of the studied techniques is facilitated by the introduction of several schemes that allow us to achieve exact counts of the degeneracies of the tested instances.

1. Introduction

In statistical and condensed matter physics, the density of states (DOS) of a system describes the number of states at each energy level. The DOS, which is independent of temperature, represents a deep characterization of the system. In terms of thermodynamics, knowledge of the DOS allows one to calculate the partition function and hence all expectation values that can be derived from it, including the free and internal energies as well as the specific heat, as a function of temperature [1]. Alternatively, the DOS itself and closely related quantities are the center of interest in an analysis of the thermodynamics of phase transitions in the microcanonical ensemble [2]. DOS calculations can also determine the spacing between energy bands in semiconductors [3].

Whilst knowledge of the DOS, $\Omega(E)$, is extremely valuable, it cannot in general be easily acquired. Exact calculations are only possible in a few special cases such as Ising models on two-dimensional lattices [4, 5]. In general, the problem is exponentially hard as the system size increases (in computational complexity terms it is a #P-hard problem) [6, 7]. This difficulty notwithstanding, there exist a number of approximation techniques, mostly based on Monte Carlo methods, that allow one to estimate $\Omega(E)$. The most widely used approach of this type is the Wang–Landau (WL) algorithm [8, 9] and its variants [10–12], which is based on the multicanonical method [13].

As stochastic approximation techniques, these approaches are affected by statistical errors as well as systematic deviations (bias). Estimating statistical error is not easily possible from a single WL simulation alone and normally requires statistics over independent runs. The most relevant form of bias in WL simulations is that of a false convergence, where the DOS estimate settles down on a deceptively smooth shape that is however not a faithful representation of the actual DOS [14]. Naturally, such problems are notoriously hard to detect if no independent information about the actual DOS is available.

Such difficulties apply in particular to systems with complex free-energy landscapes that are typically accompanied by frustration in the interactions such as in the protein-folding problem [15] and in the spin-glass systems that result from a combination of frustration and quenched disorder [16]. The latter may be viewed as prototypical classically-hard optimization problems, and they are so challenging that specialized hardware has been built to simulate them [17]. Recently, DOS estimation, which provides the relative degeneracies of the energy levels of spin glasses, has become an indispensable tool in the context of the benchmarking of experimental quantum annealers [18, 19] and the attempts to demonstrate speedups over classical devices. The currently available commercial realization of this paradigm [20] effectively samples low-lying energy configurations of spin-glass samples that are coded into the couplers connecting an array of superconducting flux qubits. The properties of the resulting samples and the question of whether such devices indeed provide superior performance as compared to classical algorithms for some problem classes has been the subject of much recent debate [21–30]. The questions of reliable Boltzmann sampling aided by quantum annealers for machine learning applications [24–26] as well as the advantages that quantum annealing devices potentially hold when tasked with fairly sampling the ground-state manifolds of spin glasses with multiple minimizing configurations [27–29] are now a topic of considerable interest.

Our goal in this work is threefold. (i) We devise techniques for *verifiably* benchmarking algorithms for sampling the DOS, designed to overcome the pitfalls of misinterpreting false convergences of entropic samplers. (ii) Employing the above techniques, we demonstrate the difficulties in applying traditional algorithms for sampling the DOS to spin-glass instances. (iii) We introduce a population annealing (PA) algorithm for estimating the DOS that allows for the intrinsic control of statistical and systematic errors and demonstrate how it can outperform the standard approach without the associated problems of choosing energy windows and related parameters that occur for the latter.

2. Verifiable benchmarking of entropic samplers

As mentioned above, approximation algorithms for the DOS are not always reliable in converging to the correct answer, especially for the frustrated systems considered here. While there are some general results regarding the convergence of suitably modified WL type algorithms [11], convergence times can become astronomically large and convergence hard to assess from intrinsic indicators [31]. As we shall see below, for disordered systems convergence times can also fluctuate wildly between different realizations of the couplings. It is hence highly desirable for benchmarking purposes to have at hand sets of samples that are sufficiently challenging for the tested algorithms, but for which nevertheless the (exact) DOS is known from other considerations. In general, such samples are not readily available, but we present here two groups of examples that are extremely useful in this respect: locally planar lattices and samples with planted solutions.

For concreteness, we shall consider spin models of the Ising type, whose cost function (or Hamiltonian) is of the form

$$H = -\sum_{\langle i,j \rangle} J_{ij} s_i s_j - \sum_i h_i s_i, \quad (1)$$

where $\langle i, j \rangle$ denotes the set of edges of the underlying graph. Here, $s_i = \pm 1$ are the Ising spin variables and the quenched parameters J_{ij} and h_i denote the exchange couplings and random fields, respectively. In the following, we will focus on the zero-field case $h_i = 0$. While the problems of computing the DOS (or partition function) and of finding a ground state for this problem in at least three dimensions are NP hard [32], both tasks can be completed in polynomial time on any set of graphs of a genus bounded by a constant, which includes planar and toroidal two-dimensional lattices [5]. For such cases there exist efficient algorithms to solve the above problems. For ground states, these include minimum-weight perfect matching [33–35], while for the partition function or DOS the usual approaches are based on the counting of dimer coverings which can be achieved via an evaluation of Pfaffians [5, 36, 37].

Unfortunately, such techniques are restricted to locally planar graphs and so they do not apply, for example, to the Chimera graphs used in current implementations of quantum annealers, which have a genus that grows linearly in the number of sites. For such more general problems we propose here an approach that is based on generating problem instances for which the values and degeneracies of the ground- and first excited-states, $\Omega(E_0)$ and $\Omega(E_1)$, are *exactly* computable. Since the states of lowest energies are usually the most difficult to sample, the degeneracies of these two energy levels are the most difficult to ascertain, and a correct estimation of these serves as a good indicator of true convergence. We create such samples by considering problem instances with *planted solutions* [21, 38, 39]—an idea borrowed from constraint satisfaction (SAT) problems [40, 41] where the planted solution represents a ground-state configuration of equation (1) that minimizes the energy and is known in advance. Following [21], the Hamiltonian of a planted-solution spin glass is a sum of terms, each of which consists of a small number of connected spins, namely, $H = \sum_j H_j$. Each term H_j is chosen such that one of its ground states is the planted solution. It follows then that the planted solution is also a ground state of the total Hamiltonian. This class of instances has two attractive properties: (i) the ground-state energies of the generated problems are known in advance, and (ii) the *exact* degeneracies of the

ground and first excited states are computable [27, 30]. These in turn allow us to check how close entropic samplers come to these exact values. The computation of $\Omega(E_0)$ and $\Omega(E_1)$ is based on the fact that our generated instances consist of a sum of terms, each of which has all minimizing configurations of H as its ground state. To enumerate all ground states, we implement a form of the ‘bucket’ algorithm [42] designed to eliminate variables one at a time to perform an exhaustive search efficiently (for a detailed description of the algorithm, see the supplementary information of [27]). By noting that the lowest energy excited states are those configurations that violate precisely one clause, their degeneracy may also be calculated. We perform the same exhaustive search as above, but where now the configurations tested will correspond to first excited states of one of the H_j (and are still minimizing for $H_{i \neq j}$). This gives the number of configurations which are ground states of $H_{i \neq j}$ and first excited states of H_j ; by performing this calculation for each of the H_j , we get the total number of first excited states of H .

While this approach in principle works for arbitrary graphs, we focus here on Chimera lattices, i.e. two-dimensional arrays of unit cells of eight spins with a $K_{4,4}$ bipartite connectivity [43, 44], see for example [30]. Our choice is motivated by the attention these graphs have gained in recent years in the context of optimization as well as sampling via quantum annealing as the quantum annealers currently commercially available feature qubits connected with this topology [45–47]. While the Chimera graph is two-dimensional in nature [48], it is also non-planar and as such gives rise to difficult spin-glass problems [49]. We generated 625 planted-solution instances of 501 spins each⁷, following a technique described in detail in [21] wherein the clauses H_j are chosen to be ‘frustrated loops’ along the Chimera graph. For each sample we employed the bucket algorithm in order to obtain $\Omega(E_0)$ and $\Omega(E_1)$.

The combination of full exact DOS for samples on the square lattice and toroidal boundary conditions and of exact values for $\Omega(E_0)$ and $\Omega(E_1)$ for the Chimera samples allows us to carefully examine the reliability and performance of sampling schemes for estimating the DOS, avoiding the pitfalls provided by badly converged estimates of stochastic approximation schemes.

3. Sampling the DOS

The common approximation algorithms for the DOS are based on Markov chain Monte Carlo [8, 13]. In the following, we will use the most popular of these, the WL algorithm [8, 9], in a variant dubbed the WL-1/ t method [10] that in principle can be shown to converge to the correct answer if given infinite run time [11], as a reference and contender of the method introduced here, entropic population annealing (EPA).

3.1. WL sampling

In WL sampling as introduced in [8, 9] a running estimate $\hat{\Omega}(E)$ of the DOS (initialized as $\hat{\Omega}(E) = 1 \forall E$) is updated in a random walk through energy space by multiplying $\hat{\Omega}(E)$ at the current energy E by a modification factor f (initially chosen to equal Euler’s number e) at each step. A new configuration of energy E' is proposed according to the chosen move scheme (for a spin model typically through single spin flips) and accepted with probability

$$p_{\text{acc}} = \min \left[1, \frac{\hat{\Omega}(E)}{\hat{\Omega}(E')} \right]. \quad (2)$$

If, after some time, the histogram $\hat{H}(E)$ of all possible energies is found to be ‘sufficiently flat’ (typically interpreted as no histogram bin having less than 80% of the average number of entries [8], but see also [50] for a related discussion), the modification factor is reduced as $f \rightarrow \sqrt{f}$, and the histogram $\hat{H}(E)$ is reset to an empty state. The algorithm stops if f is ‘sufficiently small’, for example $f = f_{\text{final}} = \exp(10^{-8})$. While the approach was invented as a variant of Markov chain Monte Carlo, the fact that the transition probabilities according to equation (2) change constantly means that neither detailed nor global balance are satisfied, and it is more useful to think of the method as a ‘stochastic approximation algorithm’ [11].

It is well known that the original scheme of [8, 9] does not converge to the true DOS, but the error asymptotically saturates at a value determined by the protocol used for reducing f [51]. This shortcoming is remedied by choosing a different modification protocol for f , leading to a slower decay of f at late times. The so-called 1/ t algorithm proposed in [10] uses two phases. In the first phase the standard WL algorithm is used, with the only difference that the energy histograms are considered to be sufficiently flat already if $\hat{H}(E) \neq 0$ for all E . Once $\ln f$ falls below the moving threshold N_E/t , where t is the simulation time measured in spin-flip attempts and N_E is the number of energy levels, the simulation enters the second phase. There, the modification factor is adapted quasi continuously according to $\ln f(t) = N_E/t$ and histogram flatness is no longer tested. The simulation stops once $f(t) \leq f_{\text{final}}$.

While no saturation of error occurs in the 1/ t algorithm [10, 14], it is still necessary to know the permissible energy levels (including the ground state) beforehand to judge histogram flatness, which is a major drawback of the method for disordered systems. In practice, we therefore employ pre-runs of the WL type without any

⁷ The 501 spins correspond to the graph considered in [30].

reduction of the modification factor with the goal of discovering the available energy levels⁸. For large systems and problems with complex free-energy landscapes, it is usually necessary to divide the total energy range into several windows for which the algorithm is run separately to achieve convergence on realistic time scales for interesting system sizes [9]. The right choice of window sizes in such schemes is a difficult problem especially for disordered and frustrated systems [52], and we are not aware of any reliable systematic approach to solve it. As a consequence, we had to spend considerable time with trial and error to arrive at suitable setups for the problems studied below. A number of further generalizations of the method have been proposed, for instance a combination with parallel tempering [12] which uses progressively smaller windows at lower energies, but here again there is no general algorithm for determining the appropriate window sizes automatically.

3.2. Entropic population annealing

The new algorithm introduced here, which we call EPA, is not based on Markov chains but on the sequential Monte Carlo method. PA was first studied in [53, 54] and more recently developed further in [55–61]. It is based on the initialization of a population of replicas drawn from the equilibrium distribution at high temperatures, which is then subsequently cooled to lower and lower temperatures. During this process, a combination of population control and spin flips is used to ensure that the ensemble remains in equilibrium. The simulation entails the following steps [55, 58].

- (i) Set up an equilibrium ensemble of $R_0 = R$ independent copies (replicas) of the system at inverse temperature $\beta_0 = 1/k_B T_0$.
- (ii) Take a step to inverse temperature $\beta_i > \beta_{i-1}$ by resampling the configurations $j = 1, \dots, R_{i-1}$ with their relative Boltzmann weight $\hat{\tau}_i(E_j)$, leading to $R_i \neq R_{i-1}$ replicas, in general.
- (iii) Update each replica by θ rounds of an MCMC algorithm at inverse temperature β_i .
- (iv) Goto step 2 unless the inverse target temperature β_f has been reached.

During resampling, the expected number of copies is

$$\hat{\tau}_i(E_j) = \frac{R}{R_{i-1}} \frac{e^{-(\beta_i - \beta_{i-1})E_j}}{Q(\beta_{i-1}, \beta_i)}, \quad (3)$$

with a normalizing factor

$$Q(\beta_{i-1}, \beta_i) = \frac{1}{R_{i-1}} \sum_{j=1}^{R_{i-1}} e^{-(\beta_i - \beta_{i-1})E_j}. \quad (4)$$

The actual number of copies is taken to be the integer part $\lfloor \hat{\tau}_i(E_j) \rfloor$ plus an additional copy added with a probability corresponding to the fractional part, $\hat{\tau}_i(E_j) - \lfloor \hat{\tau}_i(E_j) \rfloor$. While initially constant (inverse) temperature steps were used on increasing $\beta_i > \beta_{i-1}$ [55], it turns out that a better, parameter-free method consists of choosing β_i to ensure a certain overlap of the energy distributions between the two temperatures [58]. This overlap can be computed from the resampling factors,

$$\alpha(\beta_{i-1}, \beta_i) = \frac{1}{R_{i-1}} \sum_{j=1}^{R_{i-1}} \min\left(1, \frac{R \exp[-(\beta_i - \beta_{i-1})E_j]}{R_{i-1} Q(\beta_{i-1}, \beta_i)}\right),$$

and β_i is adapted using a bisection search such as to ensure an overlap α^* of energy histograms. The method is not very sensitive to the precise value of α^* , and we choose $\alpha^* = 0.86$ in the runs below.

While the algorithm described above is just PA [55] improved by adaptive temperature steps [58, 60], the possibility of sampling the entropy arises from a combination of the method with multi-histogram techniques [62]. An estimator of the free energy follows directly from the resampling factors [55],

$$-\beta_i \hat{F}(\beta_i) = \ln Z_{\beta_0} + \sum_{k=1}^i \ln Q_k, \quad (5)$$

where Z_{β_0} is the partition function at the initial temperature β_0 . In the following, we always choose $\beta_0 = 0$, such that simply $Z_{\beta_0} = 2^N$, where N is the number of spins. We can then estimate the DOS by combining the histograms at all temperature steps. As we show in appendix A, a variance-optimized estimator is given by

⁸ We generally choose the length of pre-runs so as to make sure that all levels are discovered, but in a few cases the actual ground state is only found in the main run.

$$\hat{\Omega}(E) = \frac{\sum_{i=1}^{N_\beta} \hat{H}_{\beta_i}(E)}{\sum_{i=1}^{N_\beta} R_i \exp[\beta_i \hat{F}(\beta_i) - \beta_i E]}. \quad (6)$$

Here, N_β is the total number of temperatures, and the energy histogram $\hat{H}_{\beta_i}(E)$ at inverse temperature β_i is normalized such that $\sum_E \hat{H}_{\beta_i}(E) = R_i$. In equation (6), the free-energy estimate $\hat{F}(\beta_i)$ is taken from equation (5). More sophisticated estimators that can lead to improved results in some cases are discussed in appendix A.

This approach is naturally suited for (moderately or massively) parallel calculations as the R replicas are simulated independently of each other and the only interaction occurs during resampling. An efficient GPU implementation was discussed in [58]. Importantly for our application, EPA does not require any prior knowledge of the range of realized energies. Additionally, as we shall see below, EPA performs better at estimating $\Omega(E)$ for hard spin-glass samples than the WL-1/ t algorithm. A detailed analysis of systematic and statistical errors of PA can be found in [63]. Here it is worthwhile to note that statistical errors can be estimated from a single run by a jackknife blocking analysis over the population that is introduced in [63]. This is further discussed in appendix B. Also, note that it is possible to include histograms from independent runs in the overall estimate provided through equation (6) by extending the sums over i over the temperature steps of all runs. The relative weight of these contributions is automatically taken into account through the free-energy factors deduced from equation (5) and the population sizes R_i . This is a natural generalization of the weighted averages first proposed for more basic observables in [55]. This scheme makes it possible to determine the DOS with arbitrary accuracy in a fixed time given sufficient parallel resources.

It should be noted that, as it stands, EPA only visits energies in the physical region $E \lesssim 0$, which is in contrast to WL that naturally also explores energies $E > 0$. Should one be interested in this unphysical regime, however, it is possible to derive its DOS from EPA, too. For systems on bipartite graphs the DOS is always symmetric, $\Omega(-E) = \Omega(E)$, so it is easy to construct the full DOS while just actually sampling energies $E \leq 0$. This is the case for all examples discussed below. For more general situations, it is also possible to construct the full DOS by running EPA twice, once for the Hamiltonian H and once for $-H$ and combining the results.

4. Results

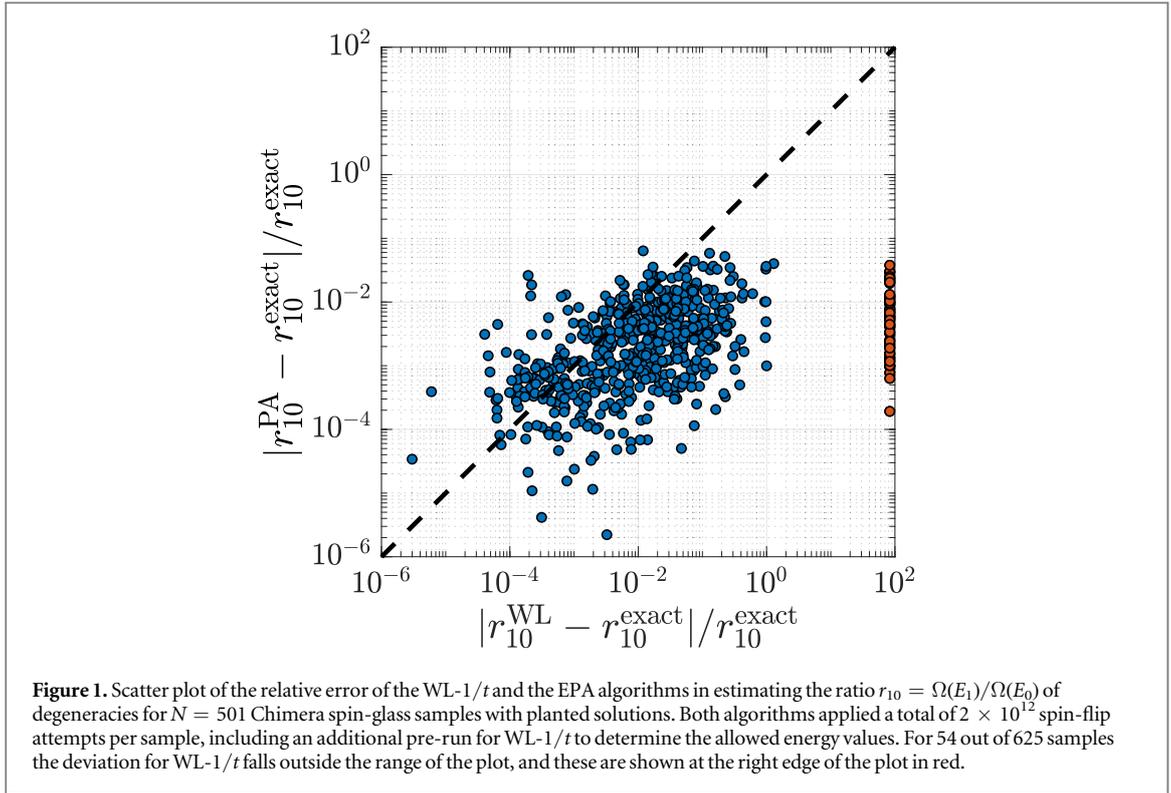
In order to test the efficiency of EPA against the WL-1/ t algorithm in an objective way that is unaffected by problems of false convergence, we applied the two algorithms to the planted solutions on Chimera graphs as well as to the stochastic $\pm J$ model on the square lattice with periodic boundaries, for both of which we have exact results. As a baseline for the comparison, we tested both methods for the case of the Ising ferromagnet on a square lattice for which extensive exact results are available. There, we find similar performance of the two techniques, see the discussion in appendix C.

4.1. Ising spin glasses on Chimera graphs

We first considered the Chimera spin-glass instances with planted solutions and $N = 501$ spins. In order to be able to compare the two algorithms on an equal footing, we could have directly considered them to be allowed the same runtime. This measure, however, is implementation and platform specific. Since both algorithms spend most of their time flipping spins, we compare them for simulations employing the same number 2×10^{12} of spin-flip attempts. For the used (serial) code for WL-1/ t this corresponded to a wallclock time of 37 h (on an Intel Xeon 2.4 GHz CPU), while for EPA we used a massively parallel GPU program [58] that took approximately 1.2 h (on an Nvidia Tesla K40 GPU) per realization.

As mentioned above, for WL it is required to know the allowed energy levels to decide about the flatness of histograms. This knowledge was here acquired by a pre-run of the WL type employing 2×10^{11} spin-flip attempts and with a fixed modification factor $\ln f = 1$ to explore the energy landscape (the corresponding run-time is included in the time estimate given above). This knowledge is not required for the EPA runs. With a single window covering the full energy range, WL-1/ t did not complete phase 1 for the vast majority of samples. We therefore used two windows with energy ranges $[E_0, E_0 + 1200]$ and $[E_0 + 1100, 50]$ in dimensionless units, respectively. The spin-flips were divided evenly between the two energy windows. Here, the disorder average of the ground state energy E_0 was found to be -3635 . The energy levels were determined during the pre-run, which found the ground state in the vast majority of cases. It is clear that for larger systems where it is much harder to find the ground state the determination of suitable windows for WL-1/ t becomes much harder. The simulation was started in a random configuration within the energy range of the window. With that restriction, 567 out of 625 samples completed phase 1 in the first window within the remaining 8×10^{11} flip attempts after the pre-run. No range restriction was required for EPA, and we used a population of size $R = 3992\,000$ with $\theta = 10$ rounds of spin flips per resampling step and a histogram overlap $\alpha^* = 0.86$, resulting in typically 100 temperature steps down to $\beta_f = 5 \times 10^4$.

⁹The unusual population size results from the attempt of matching the average number of spin-flip attempts between the two methods.



We note that both EPA and WL-1/ t intrinsically estimate entropy *differences*, i.e. ratios of degeneracies for neighboring energy values, and the absolute scale is only achieved through an additional normalization such as that given by Z_{β_0} in equation (5). It is therefore reasonable to study their performance in estimating

$$r_{10} = \frac{\Omega(E_1)}{\Omega(E_0)},$$

the ratio of degeneracies of first excited and ground states. In figure 1 we show the relative deviations of the ratios r_{10} from the exact values known through the planting, as estimated from WL-1/ t and EPA. WL-1/ t found the correct ground-state energy for 622 of the 625 samples. For some samples the relative deviations are so large that they exceed the scale of the plot of figure 1, some by many orders of magnitude¹⁰. These samples are shown at the boundary of the box and in a different color. It is clear that for most samples the deviations are substantially smaller for EPA than for WL-1/ t . In total, EPA outperforms WL-1/ t in 80% of the instances. The error of WL-1/ t is larger than 7% for 25% of samples and it is difficult to distinguish between the accurate and inaccurate WL results. In contrast, the EPA results are accurate to within 7% for all of the 625 samples.

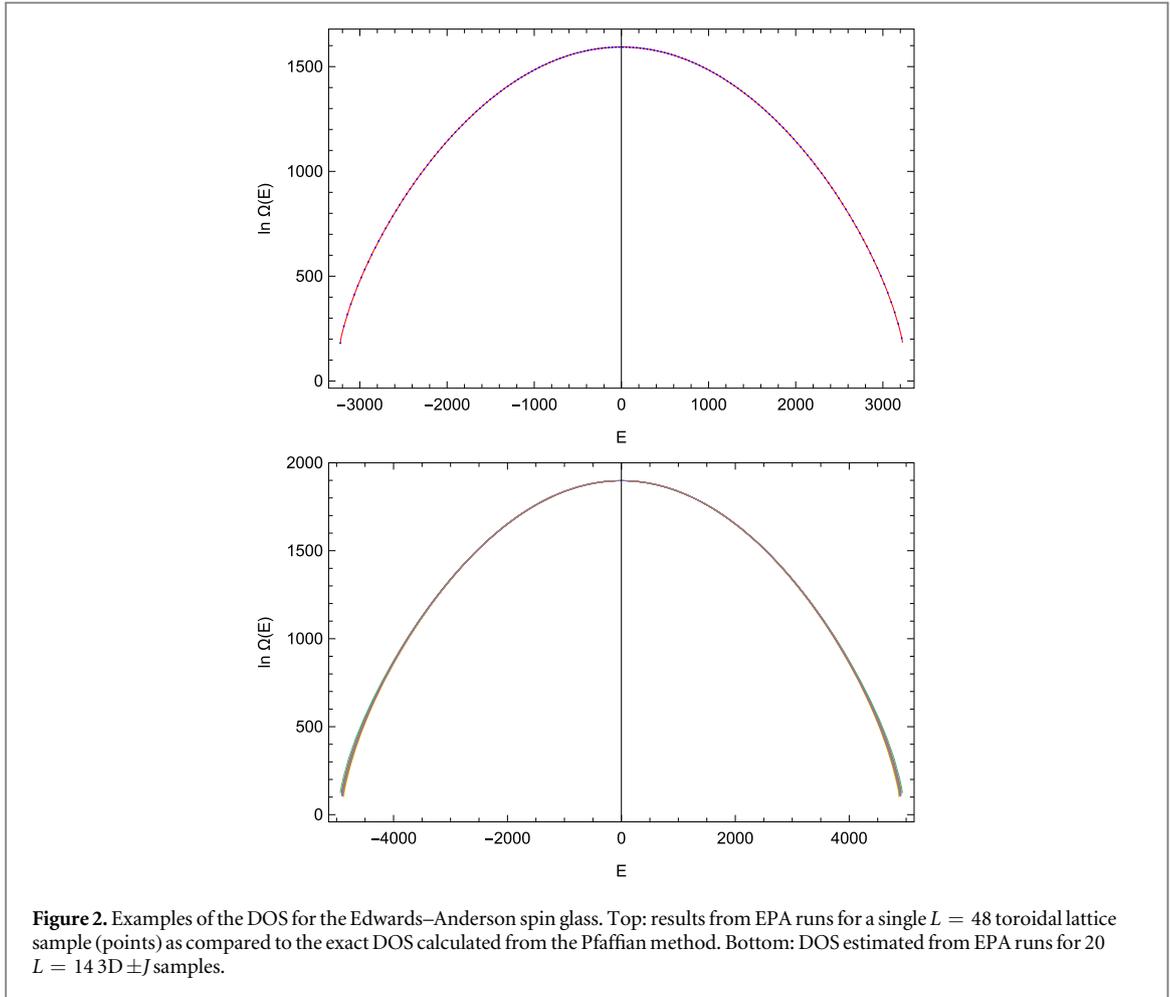
4.2. Ising spin glasses on toroidal graphs

For planar or otherwise two-dimensional lattices of a fixed genus, a counting of dimer coverings and the corresponding evaluation of Pfaffians can be used to determine the full DOS in polynomial time [5, 36, 37]. We studied toroidal graphs, i.e. $L \times L$ patches of the square lattice with periodic boundary conditions using the implementation proposed in [5] which has an asymptotic run-time scaling of $O(L^5)$. Using this approach, we evaluated 1000 samples with a standard $\pm J$ coupling distribution and 32×32 spins, and also 500 samples of size 48×48 . An example of $\ln \Omega(E)$ as estimated for a single sample of size $L = 48$ from EPA is shown in the top panel of figure 2. At this scale, the data are completely indistinguishable from the exact result also shown for comparison. As one reads off from the graph, the actual DOS $\Omega(E)$ spans about 700 orders of magnitude, and it is quite remarkable that it can be estimated so accurately from the simulations. To systematically assess the accuracy of the sampling for different disorder realizations, we considered the total deviation Δ of the simulation results from the exact DOS, where

$$\Delta = \frac{1}{N_E} \sum_{i=1}^{N_E} \left| \frac{\ln \Omega(E_i) - \ln \Omega_{\text{exact}}(E_i)}{\ln \Omega_{\text{exact}}(E_i)} \right|. \quad (7)$$

While for PA an absolute normalization of $\Omega(E)$ [64] follows from the free-energy estimator equation (5) in combination with equation (6), WL-1/ t as described above only yields the DOS up to an overall factor. To fix the

¹⁰ This includes the three samples for which WL-1/ t does not find the ground state, which effectively implies that r_{10} and the relative deviation are infinite.

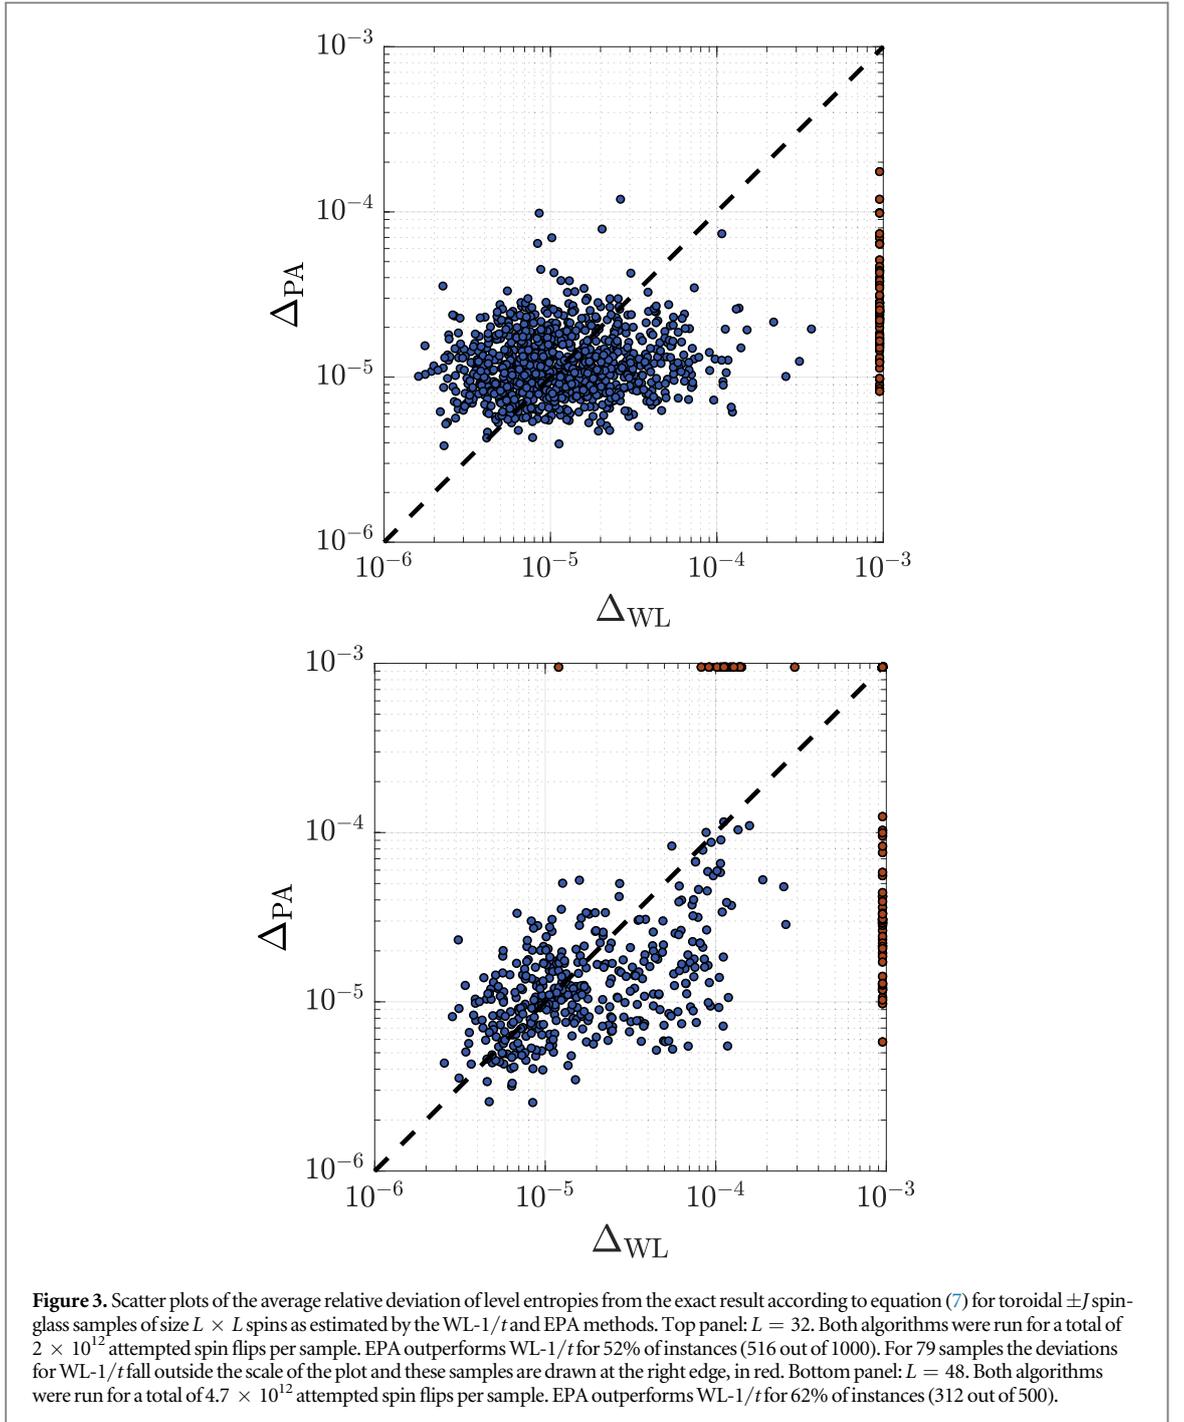


latter we used the fact that the sum $\sum \Omega(E)$ over all energy levels must equal the total number 2^N of states. Note that different ways of normalization of WL-1/ t lead to quite different fluctuations of the final DOS estimates, and the normalization via the total number of states used here leads to the best results, see the discussion in appendix D. Since EPA only samples states with energies $E \lesssim 0$, we restricted the energy range for WL-1/ t to $E \leq 50$ to ensure a fair comparison¹¹.

For 32×32 samples, we used 1.8×10^{12} spin-flip attempts in the main run of WL-1/ t employing a single energy window with $E \leq 50$. Just as for the Chimera samples, a pre-run was required to determine the range of possible energies, for which an additional 2×10^{11} updates were used. For the EPA algorithm, we used a population size $R = 2340\,000$, and performed $\theta = 19$ rounds of spin-flips between two resampling steps. The imposed histogram overlap of $\alpha^* = 0.55$ resulted in $N_\beta = 44$ temperature steps for most disorder realizations down to $\beta_f = 5 \times 10^4$. The total number of spin flips in these EPA runs is hence also approximately 2×10^{12} . For system size 48×48 , WL-1/ t required two energy windows to converge; these were chosen as $[E_0, E_0 + 64]$ and $[E_0 + 36, 50]$. After the pre-run of length 6×10^{11} across both windows, we used 3.1×10^{12} spin-flip attempts in the main run of the first window and 1.0×10^{12} updates in the second window. The two DOS segments obtained by WL-1/ t were sewn together by matching the estimates at a point in the intersection of the two windows. For the EPA algorithm, we used $R = 1\,019\,965$, $\theta = 10$ and $\alpha^* = 0.86$, which resulted in $N_\beta = 200$ ($\beta_f = 3$) and hence the total number of spin flips is 4.7×10^{12} as for the WL-1/ t runs.

The resulting values for the average relative deviation of the level entropies according to equation (7) are shown in the scatter plots provided in figure 3. The top panel corresponds to 1000 samples of size 32×32 . In some cases the ground states were not found or the first phase of WL-1/ t did not complete, leading to extremely large or infinite deviations; the corresponding samples are shown in red at the boundary of the plot. In this case we only find a moderate advantage for EPA, which outperforms WL-1/ t on 516 of the 1000 samples. Considering the larger system size $L = 48$, the advantage of EPA increases, leading to a smaller value of Δ for 291 samples out of the 451 samples where both methods found all energy levels. This observation is in line with a

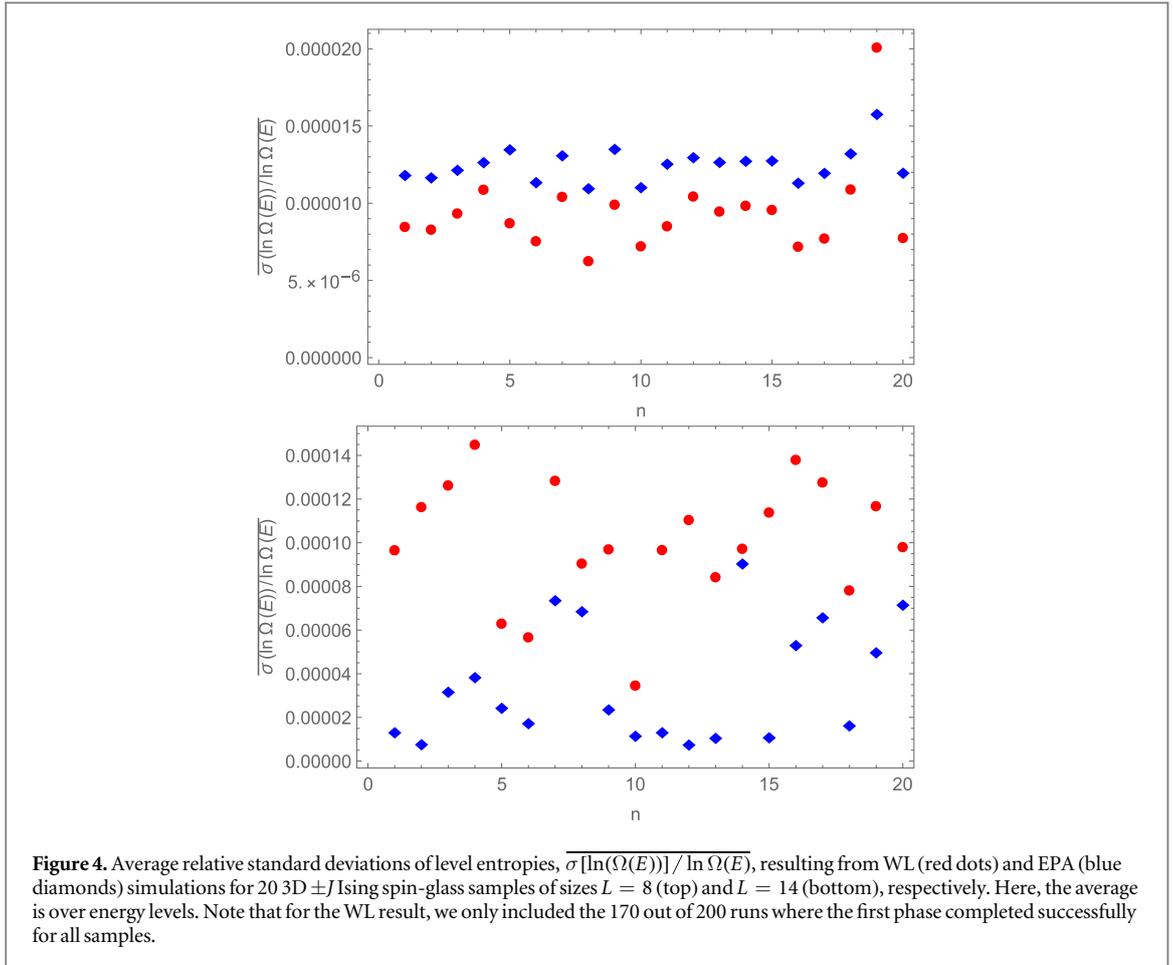
¹¹ For the normalization of the DOS we hence completed $\Omega(E) = \Omega(-E)$ for the positive energies. For consistency, we additionally applied the same normalization in EPA.



general trend of EPA faring relatively better for harder problems as compared to WL that we shall see confirmed for other examples below.

4.3. Three-dimensional Ising spin glasses

Concerning the trend of improving relative performance of EPA for harder problems, it is interesting to see how the two samplers perform on spin-glass instances in three dimensions, where the spin-glass problem is known to be NP-hard [32]. To this end, we studied samples of the $\pm J$ Edwards–Anderson model with $L = 8$ and $L = 14$. For $L = 8$ we were able to employ a single energy window for WL, and the parameters for EPA were $R = 1992\,984$, $\theta = 10$, $\alpha^* = 0.86$, and $\beta_f = 3$, using 10^{12} spin-flip attempts in both cases. For $L = 14$, 5.5×10^{12} spin flips were applied in each run. For the WL-1/ t method we used two windows, $[E_0, E_0 + 64]$ and $[E_0 + 36, 50]$, with 3.9×10^{12} and 1.0×10^{12} in the main run, respectively. The remaining 6×10^{11} spin-flip attempts were used in the pre-run. The parameters for EPA were $R = 867\,694$, $\theta = 10$, $\alpha^* = 0.86$, and $\beta_f = 10$. In the bottom panel of figure 2 we show $\ln \Omega(E)$ for the $L = 14$ samples. The sample-to-sample fluctuations in the DOS are in fact rather small and can only be seen in the very low-energy part of the spectrum (as well as its mirror image for large positive energies).



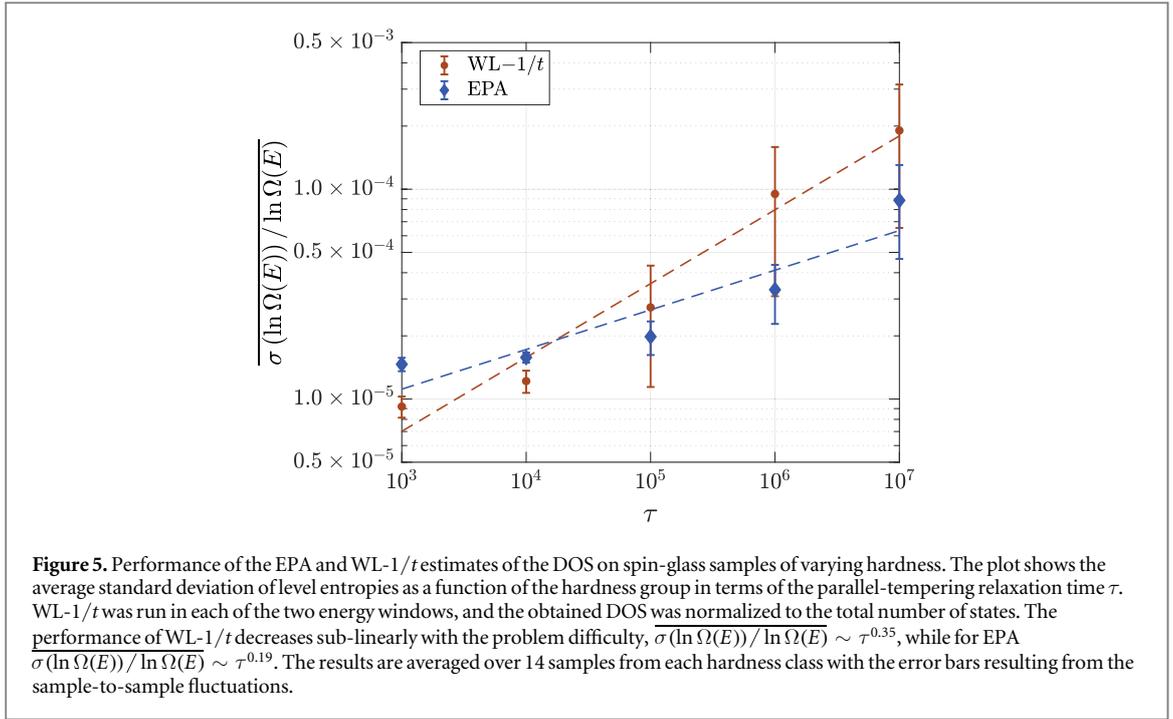
As the samples considered are neither planar nor planted, we do not have access to the exact DOS, and we hence quantify the success of the two algorithms in estimating the DOS by determining the level of fluctuation in the estimates of $\ln \Omega(E)$ between independent runs, both for the WL- $1/t$ and EPA methods. Specifically, we estimated the relative standard deviation $\sigma[\ln \Omega(E)]/\ln \Omega(E)$ from 200 independent runs for each disorder sample. In figure 4 we show this quantity, averaged over all energy levels, for 20 $\pm J$ 3D spin-glass samples of the two system sizes considered. While for $L = 8$ the WL runs yield slightly smaller error bars, for $L = 14$ the situation is reversed, with EPA resulting in 5 times smaller error bars, on average, corresponding to saving a factor of 25 in run-time.

4.4. Entropic sampling of problems of varying hardness

Having ascertained that the EPA method can yield significantly better approximations to the DOS of some hard problems in a given number of steps than the WL approach, it is interesting to analyze more closely the actual distribution of performances of these algorithms over the space of disorder realizations. The results above in figures 1 and 3 indicate the presence of larger fluctuations in the quality of approximation for WL- $1/t$ as compared to EPA across disorder realizations. To study this effect quantitatively, we used a set of disorder samples classified according to their *thermal hardness*. A well established measure of such hardness is the exponential autocorrelation or relaxation time in parallel tempering simulations [65, 66]. As very long simulations are required to determine these time scales accurately, a number of proxy quantities such as the so-called ‘tunneling time’ are frequently used in practical applications [67, 68]. Here, we rely on a method developed in the context of spin-glass simulations that analyses the dynamics of the random walk of replicas in temperature space [69] and extracts the corresponding relaxation times τ .

To benchmark the EPA algorithm against WL- $1/t$, we generated about 10^6 random instances on an $N = 512$ -spin Chimera graph (of which only 476 spins were used) and measured the relaxation times τ of each instance with parallel tempering¹². Next, we grouped together instances with similar classical hardness, i.e. similar relaxation times τ , $10^k \leq \tau \leq 3 \times 10^k$ for $k = 3, 4, 5, 6$ and 7. For each such ‘generation’ of τ , we randomly picked 100 representative

¹² Specifically, we chose a temperature grid of the PT simulations consisting of $N_T = 30$ temperatures. Temperatures with indices $i = 1, 2, \dots, 12$ were uniformly distributed in the range $T_{\min} = 0.045 \leq T_i \leq 0.2$, while the temperatures T_i with $i = 13, 14, \dots, N_T$ were spread evenly in the range $0.21 \leq T_i \leq T_{\max} = 1.632$ [70].



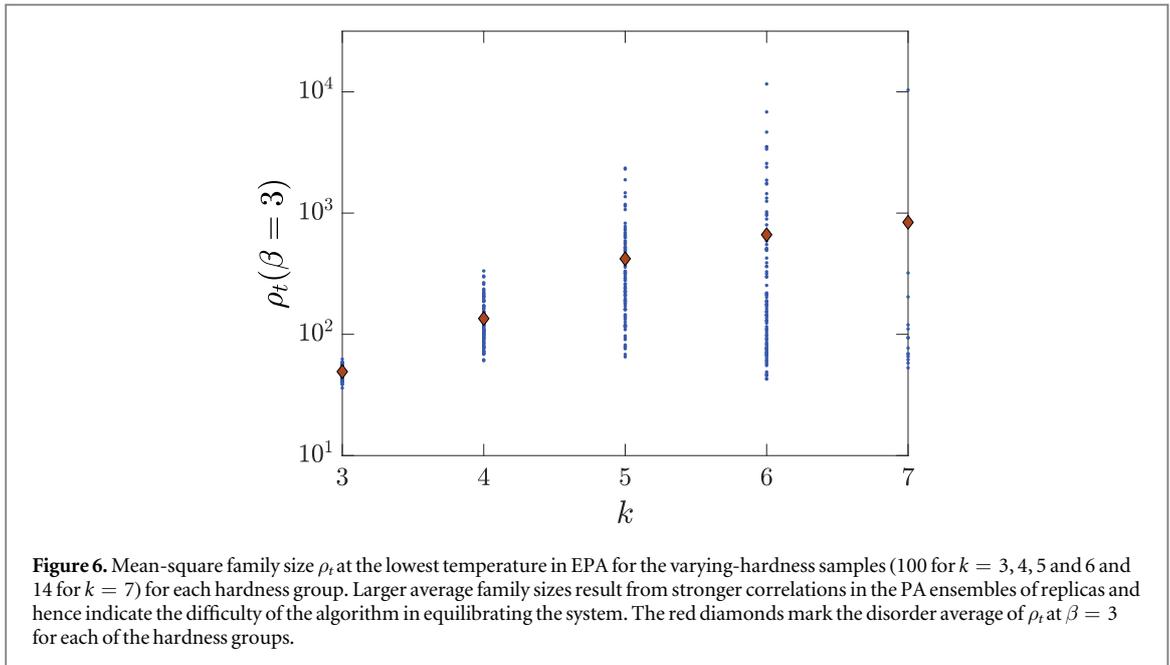
instances for the benchmarking of the algorithm (only 14 instances with $k = 7$ were found). We then performed WL-1/t simulations with a total of 10^{12} spin-flip attempts for all samples, evenly split between one simulation each restricted to the energy windows $[-900, -500]$ and $[-550, 50]$, respectively (the ground-state energy for these samples is roughly $E_0 \approx -800$). A pre-run of 2×10^{11} spin-flip attempts was again used to discover the range of possible energies for each sample. All runs completed the first phase of the simulation here, owing to the use of two energy windows. For the PA runs we used $R = 2.1 \times 10^6$, $\theta = 10$, $\alpha^* = 0.86$, corresponding to $N_\beta \approx 100$ temperature steps ($\beta_f = 5$) and 10^{12} spin flip attempts. The DOS estimates from both methods are only considered for $E \leq 0$ and $\Omega(E) = \Omega(-E)$ is used for $E > 0$. The resulting DOS estimate is normalized using the known total number of states 2^N .

As for the 3D samples, we compared the two algorithms by considering the relative standard deviations $\sigma[\ln \Omega(E)] / \ln \Omega(E)$, averaged over all energies. The resulting estimates are shown in figure 5 for the samples of the different hardness classes $k = 3, \dots, 7$. It is clear that EPA is less affected by sample hardness than WL-1/t, with the growth in fluctuation with sample hardness being much steeper for WL-1/t than for EPA. Note that this quantity only covers the effect of statistical errors, whereas the data in figures 1 and 3 considered the total deviation from the exact results that also includes bias effects. Note also that the sample-to-sample fluctuations, represented in the error bars of the data points in figure 5, are significantly larger for WL-1/t than for EPA. We find that WL-1/t and EPA have rather different behavior in sampling the DOS in different energy ranges, with WL-1/t being more focused on higher energies, for details see appendix E.

While in the present demonstration we used relaxation times from parallel tempering for classifying sample hardness, it is worthwhile to note that EPA can itself provide a hardness measure and thereby differentiate easy and hard samples. A few such quantities have been previously proposed for PA [56]. We consider here in particular the (temperature dependent) mean-square family size ρ_b , defined as [56]

$$\rho_t = R \sum_i n_i^2, \quad (8)$$

where n_i is the fraction of the current population that descends from the i th member of the initial population at $\beta_0 = 0$, while R corresponds to the initial population size. The quantity R/ρ_t can be understood as an effective population size, corresponding to the number of statistically independent replicas, such that $R/\rho_t = R$ corresponds to a perfectly uncorrelated population, while $R/\rho_t \rightarrow 0$ for the strongest correlations. These two limits hence represent the easiest and hardest samples, for which one would expect $\tau \rightarrow 0$ and $\tau \rightarrow \infty$, respectively, for parallel tempering. A related quantity that also takes the decorrelating effect of spin flips into account is the effective population size R_{eff} defined in [71]. In figure 6 we show a scatter plot of ρ_t for 100 samples of each of the hardness classes $k < 7$ and 14 samples for $k = 7$, respectively. The disorder average of ρ_t at $\beta = 3$ is found to be 49, 135, 420, 663 and 840 for $k = 3, 4, 5, 6$, and 7, respectively, indicating that while for the main part of the distribution the hardnesses in EPA and parallel tempering are strongly correlated, for the tails of the distribution the hardness in EPA increases more gently than that found in parallel tempering. As is demonstrated elsewhere, these intrinsic hardness measures can be used to make PA simulations adaptive to the sample



hardness [60, 63]. We note that the planted samples of section 4.1 have an average ρ_t of ≈ 2000 (see appendix F), indicating that planted samples of this type are much harder than random ones.

5. Summary and discussion

We have investigated the performance of sampling methods for estimating the DOS for systems with complex free-energy landscapes, focusing on spin glasses as the hardest problems among spin systems. We proposed a novel sampling technique based on sequential Monte Carlo on a large population of copies and demonstrated that it outperforms the most widely used entropic sampler, the WL algorithm, in the vast majority of cases. More importantly, the new approach shows better scaling as the hardness of problems is increased (for example through considering larger systems).

A notorious problem with benchmarking algorithms for estimating the DOS lies in safely assessing convergence. Here we address this issue by considering problems that are either of planar topology, in which case Pfaffian methods can be used to determine the DOS exactly for systems with more than 1000 spins, or which have planted solutions such that the exact degeneracies of the ground and first excited states can be calculated using a ‘bucket’ algorithm. Both classes provide hard optimization problems, implying a very non-trivial benchmark. In addition, we also considered more general problems such as stochastic spin-glass samples on Chimera graphs sorted by thermal hardness as well as the most challenging case of three-dimensional spin-glass instances of up to $14 \times 14 \times 14$ spins.

One essential advantage of the approach based on PA is that it does not require any prior knowledge about the energy spectrum, which in contrast needs to be acquired in an additional pre-run for the WL method. Furthermore, the well-known and delicate problem of dividing the energy range into windows that is the only way of making WL simulations converge for the more challenging cases, is completely absent for EPA. The difficulty of premature and false convergence that plagues WL and related methods is not so much of an issue for the newly introduced technique, where a re-distribution of weights can occur at all stages of the algorithm. In fact, for the EPA method it is easily possible to monitor equilibration from intrinsic properties and only output the DOS for energies where thermalization could be ensured. The main advantage of the approach, however, lies in the ideal suitability for massively parallel calculations, where given sufficient parallel resources the accuracy of the approximation can be arbitrarily improved at a constant wall-clock time by increasing the size of the population or combining the outcome of independent runs in a weighted average.

The specific spin system we study, namely spin glasses, is very relevant as current experimental quantum annealers attempt to solve precisely this type of problem. With questions still lingering about which distribution these devices sample from [30], it is important to have an accurate tool to estimate the DOS (for instance to understand thermalization [30]). For this problem we specifically consider instances with vastly different hardnesses, confirming that the accuracy of the technique proposed degrades significantly less for harder samples than previous approaches. We note that the EPA approach is in no way specific to the spin systems considered here, and it can be straightforwardly generalized to other problems such as lattice polymers and, with the help of binning or spectral methods [72], to cases with continuous degrees of freedom.

Our results are very promising as we could clearly show that (i) for a range of problems with complex energy landscapes existing sampling methods for the DOS are difficult to set up and do not converge very reliably, especially for hard samples; (ii) a dependable method for the benchmarking of entropic samplers on spin glasses is now available, which we hope will drive research forward in finding even better algorithms; and (iii) the EPA algorithm devised here allows for the reliable sampling of large-scale frustrated systems. We therefore trust that the algorithm will become a useful tool for DOS calculation in condensed matter physics, quantum computing and other areas of research.

Acknowledgments

Part of the computing resources were provided by the USC Center for High Performance Computing and Communications and the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the US Department of Energy under Contract No. DE-AC05-00OR22725. The research is based upon work (partially) supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the US Army Research Office contract W911NF-17-C-0050. 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 the ODNI, IARPA, or the US Government. This material is based on research sponsored by AFRL under agreement number FA8750-18-1-0109. The US Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. LB was supported by grant No. 14-21-00158 from the Russian Science Foundation. LB and MW acknowledge support by the European Commission through the IRSES network DIONICOS under Contract No. PIRSES-GA-2013-612707.

Appendix A. The DOS estimator

In this section we outline the derivation of the estimator (6) for the DOS. We initially follow the reasoning of [62]. In population annealing, the histogram of energies at temperature β_i , $\hat{H}_{\beta_i}(E)$ or in short \hat{H}_i , is an estimator of the equilibrium probability density of internal energies

$$\left\langle \frac{\hat{H}_i(E)}{R_i} \right\rangle = \frac{1}{Z_i} \Omega(E) e^{-\beta_i E}. \quad (\text{A.1})$$

Hence an estimate of the DOS from a single histogram is given by

$$\hat{\Omega}_i(E) = \frac{\hat{H}_i(E)}{R_i} \frac{Z_i}{e^{-\beta_i E}}. \quad (\text{A.2})$$

To make use of the histograms at different temperature steps, we take a weighted average

$$\hat{\Omega}(E) = \sum_i \alpha_i \hat{\Omega}_i(E). \quad (\text{A.3})$$

As a simple calculation shows [73], for independent individual estimates a minimum variance of the result is achieved when choosing the weights

$$\alpha_i = \frac{1/\sigma^2[\hat{\Omega}_i(E)]}{\sum_j 1/\sigma^2[\hat{\Omega}_j(E)]}.$$

From equation (A.2) we deduce that

$$\sigma^2[\hat{\Omega}_i(E)] = \frac{Z_i^2}{R_i^2} e^{2\beta_i E} \sigma^2[\hat{H}_i(E)], \quad (\text{A.4})$$

and hence the variance-optimized weighted average (A.3) becomes

$$\hat{\Omega}(E) = \frac{\sum_i R_i Z_i^{-1} e^{-\beta_i E} \hat{H}_i(E) \sigma^{-2}[\hat{H}_i(E)]}{\sum_j R_j^2 Z_j^{-2} e^{-2\beta_j E} \sigma^{-2}[\hat{H}_j(E)]}. \quad (\text{A.5})$$

Noting that $Z_i = \exp[-\beta_i F(\beta_i)]$ and that $\sum_E \hat{\Omega}(E) \exp[-\beta E]$ is an estimator of $Z = \exp[-\beta F]$, one can write the following equation,

$$e^{-\beta_k \hat{\mathcal{F}}_k} = \sum_E \frac{\sum_i R_i \exp[\beta_i \hat{\mathcal{F}}_i - \beta_i E] \hat{H}_i(E) \sigma^{-2}[\hat{H}_i(E)]}{\sum_j R_j^2 \exp[2\beta_j \hat{\mathcal{F}}_j - 2\beta_j E] \sigma^{-2}[\hat{H}_j(E)]} e^{-\beta_k E}, \quad (\text{A.6})$$

which can be read as a set of N_β self-consistency conditions for the parameters $\hat{\mathcal{F}}_k$ that represent the free energies at the inverse temperatures β_k . Equation (A.6) can be solved iteratively by evaluating the right hand side for each k to

receive updated values of $\hat{\mathcal{F}}_k$. Convergence is very slow (and might even fail) if starting with initial values for $\hat{\mathcal{F}}_k$ that are very far from the solution. If one starts with $\hat{\mathcal{F}}_k = \hat{F}(\beta_k)$ according to equation (5), however, which already provides very accurate estimates of the free energies, convergence is typically achieved in less than ten iterations.

It remains to discuss how to determine the variances $\sigma^2[\hat{H}_i(E)]$. Without further assumptions, these can be estimated via a jackknife analysis [74], i.e. by dividing the populations at all temperature steps into jackknife blocks. Calculating $\hat{H}_i(E)$ for each jackknife block then allows one to use the jackknife estimator of variance, energy by energy¹³. Simpler expressions can be derived with further assumptions, as we will discuss now. If all members of the population were independent of each other, the entries in a histogram would follow a Poisson distribution and hence the variance was

$$\sigma^2[\hat{H}_i(E)] = \langle \hat{H}_i(E) \rangle = R_i Z_i^{-1} \Omega(E) e^{-\beta_i E}, \quad (\text{A.7})$$

where the second equality follows from equation (A.1). Substituting this into equation (A.5) one finds

$$\hat{\Omega}(E) = \frac{\sum_i \hat{H}_i(E)}{\sum_j R_j \exp[\beta_j \hat{\mathcal{F}}_j - \beta_j E]}, \quad (\text{A.8})$$

which corresponds to the estimator (6) introduced in the main text when using the single-temperature estimate $\hat{\mathcal{F}}_j = \hat{F}(\beta_j)$. While this estimate can be further improved, in principle, via the iterations (A.6), we find that for the small temperature steps used in our runs the improvement is quite small, and the simpler approach of equation (A.8) with $\hat{\mathcal{F}}_j = \hat{F}(\beta_j)$ already yields excellent results.

In reality different replicas of the population are not independent of each other as the resampling introduces correlations. One might argue that as this reduces the effective population size by a factor $R_{\text{eff},i}/R_i$ [63], equation (A.7) should be replaced by

$$\sigma^2[\hat{H}_i(E)] = \frac{R_i}{R_{\text{eff},i}} \langle \hat{H}_i(E) \rangle = \frac{R_i^2}{R_{\text{eff},i}} Z_i^{-1} \Omega(E) e^{-\beta_i E}, \quad (\text{A.9})$$

such that equation (A.8) is replaced by

$$\hat{\Omega}(E) = \frac{\sum_i \frac{R_i}{R_{\text{eff},i}} \hat{H}_i(E)}{\sum_j \frac{R_j}{R_{\text{eff},j}} R_j \exp[\beta_j \hat{\mathcal{F}}_j - \beta_j E]}. \quad (\text{A.10})$$

In the absence of an estimate for $R_{\text{eff},i}$, it is also possible to revert to $R_i/\rho_{t,i}$ as defined in equation (8), which in general provides a lower bound for $R_{\text{eff},i}$ [71]. This approximation assumes, however, that the only effect of correlations is to reduce the effective number of events and, in particular, that this effect is independent of the energy level E . Preliminary tests of applying this variant of the DOS estimator yielded mixed results with decreases in Δ of equation (7) of at most 10%, so the effect appears to be rather weak.

Finally, we should also take into account the fact that the populations at different temperature steps are correlated¹⁴. In this case the average (A.3) should be performed with weights [73, 75, 76]

$$\alpha_i = \frac{\sum_j \text{cov}^{-1}(\hat{\Omega}_i, \hat{\Omega}_j)}{\sum_{ij} \text{cov}^{-1}(\hat{\Omega}_i, \hat{\Omega}_j)}.$$

Here,

$$\text{cov}(\hat{\Omega}_i, \hat{\Omega}_j) = \langle \hat{\Omega}_i(E) \hat{\Omega}_j(E) \rangle - \langle \hat{\Omega}_i(E) \rangle \langle \hat{\Omega}_j(E) \rangle$$

denotes the covariance matrix of the DOS estimates $\hat{\Omega}_i(E)$ from different temperature steps, and cov^{-1} is the corresponding inverse matrix. We note that, in generalization of equation (A.4), one has

$$\text{cov}[\hat{\Omega}_i(E), \hat{\Omega}_j(E)] = \frac{Z_i Z_j}{R_i R_j} e^{(\beta_i + \beta_j) E} \text{cov}[\hat{H}_i(E), \hat{H}_j(E)].$$

As for the variances, the covariances of the energy histograms can be estimated from a jackknife analysis over the populations [74, 76]. The corresponding expression for the optimized DOS estimate is then generalized to

$$\hat{\Omega}(E) = \frac{\sum_{i,j} R_j Z_j^{-1} e^{-\beta_j E} \hat{H}_j(E) \text{cov}^{-1}[\hat{H}_i(E), \hat{H}_j(E)]}{\sum_{i,j} R_i R_j Z_i^{-1} Z_j^{-1} e^{-(\beta_i + \beta_j) E} \text{cov}^{-1}[\hat{H}_i(E), \hat{H}_j(E)]}. \quad (\text{A.11})$$

¹³ Note that in order to apply the jackknife procedure one should consider an observable with a finite expectation value in the limit of an infinite number of measurements, such as $\hat{H}_i(E)/R_i$.

¹⁴ This is in contrast to the usual situation in multi-histogram reweighting, where the individual histograms belong to independent simulations.

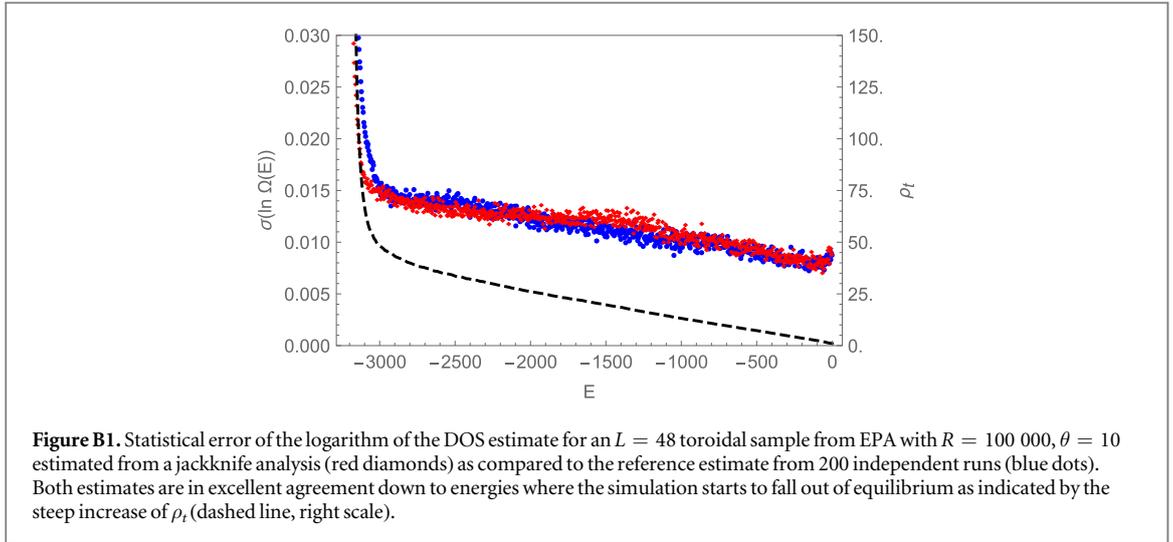


Figure B1. Statistical error of the logarithm of the DOS estimate for an $L = 48$ toroidal sample from EPA with $R = 100\,000$, $\theta = 10$ estimated from a jackknife analysis (red diamonds) as compared to the reference estimate from 200 independent runs (blue dots). Both estimates are in excellent agreement down to energies where the simulation starts to fall out of equilibrium as indicated by the steep increase of ρ_t (dashed line, right scale).

Appendix B. Estimating statistical errors and biases

Statistical errors in the DOS estimator $\hat{\Omega}(E)$ provided in equation (6) or the variants discussed in appendix A can be estimated by considering the statistics of different runs. As the correct way of combining different runs is through extending the summations on the right hand side of equation (6) over all the temperature points of all runs, combined with the corresponding free-energy estimates provided by equation (5) instead of taking a plain average over the estimates $\hat{\Omega}(E)$ of individual runs, it is important to estimate errors from a jackknife or bootstrap analysis over the runs instead of the standard estimator of the sample mean in order to minimize bias.

Alternatively, it is possible to estimate statistical errors from a single run of a sufficiently large population following the arguments of [71]. This amounts to a simultaneous jackknife analysis over the populations at all temperatures. As discussed in detail in [71], such an analysis can be justified if a linear order of the members of the population is assumed, and off-spring configurations in resampling are placed next to each other in the array of replicas. As one then generally observes an exponential decay of correlations with the index distance in replica space, sufficiently large blocks are statistically effectively independent of each other, and the jackknife estimator for the variance of the mean can be applied. For the case of the DOS estimator of equation (6) this implies that one divides the (linearly ordered) populations at all temperature steps into n blocks ($n = 100$ is often a good choice) and then applies equations (5) and (6) to all data apart from the replicas in block $s = 1, \dots, n$ to arrive at estimates $\ln \hat{\Omega}_{(s)}(E)$ ¹⁵. The variance of the mean (squared error bar) is then estimated by [74]

$$\hat{\sigma}^2[\ln \hat{\Omega}(E)] = \frac{n}{n-1} \sum_{i=1}^n [\ln \hat{\Omega}_{(s)}(E) - \ln \hat{\Omega}_{(i)}(E)]^2, \quad (\text{B.1})$$

where

$$\ln \hat{\Omega}_{(i)}(E) = \frac{1}{n} \sum_{s=1}^n \ln \hat{\Omega}_{(s)}(E).$$

As is illustrated in figure B1 this provides an accurate estimate of statistical errors as long as the total simulation remains in thermal equilibrium.

Regarding bias effects (i.e. lack of thermalization), we note that EPA as a sequential Monte Carlo method behaves differently to Markov chain samplers and the closely related WL approach. Whereas in the latter the results of a simulation that was not properly equilibrated will be affected by biases in its entirety, a PA simulation of a frustrated system will fall out of equilibrium at a certain threshold temperature, and consequently results derived from the populations at and below this temperature will be affected by biases. Conversely, however, results only incorporating data above this threshold are unaffected. For the DOS estimators studied here a strategy to avoid bias hence consists of only including histograms from the well-equilibrated regime in the estimate (6). Here, equilibration can be ensured by monitoring heuristic criteria such as $R_0 > 100\rho_t$ [56, 60] that can be evaluated on the fly to determine a stopping temperature.

¹⁵ As $\Omega(E)$ spans many orders of magnitude, it is normally much more reasonable to consider $\ln \Omega(E)$, and we hence here also estimate the error bars of this latter quantity.

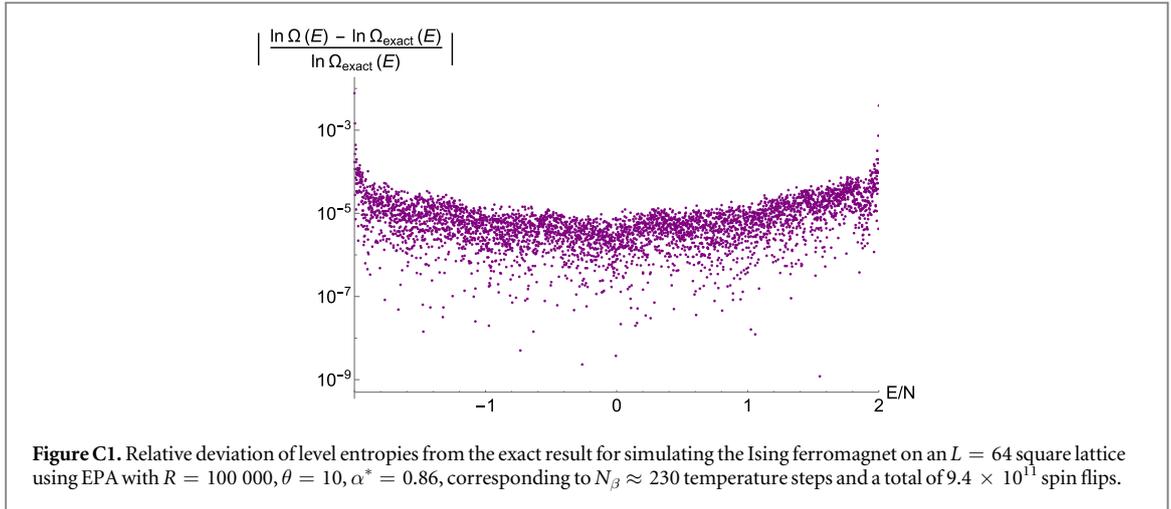


Figure C1. Relative deviation of level entropies from the exact result for simulating the Ising ferromagnet on an $L = 64$ square lattice using EPA with $R = 100\,000$, $\theta = 10$, $\alpha^* = 0.86$, corresponding to $N_\beta \approx 230$ temperature steps and a total of 9.4×10^{11} spin flips.

Table C1. Average deviation $\bar{\Delta}$ for the two-dimensional Ising model with $L = 32$. The deviation is averaged over 200 independent runs of the EPA and WL algorithms. Each of the runs performed $N = 9.82 \times 10^{11}$ spin flips. Also shown is the standard deviation of $\bar{\Delta}$.

Method	$\bar{\Delta}$	$\sigma(\bar{\Delta})$
WL-1/t	2.67×10^{-5}	1.55×10^{-5}
EPA, $\theta = 1$	8.55×10^{-5}	6.80×10^{-5}
EPA, $\theta = 5$	2.13×10^{-5}	1.29×10^{-5}
EPA, $\theta = 10$	2.66×10^{-5}	1.69×10^{-5}
EPA, $\theta = 20$	3.53×10^{-5}	1.82×10^{-5}
EPA, $\theta = 50$	5.43×10^{-5}	2.85×10^{-5}

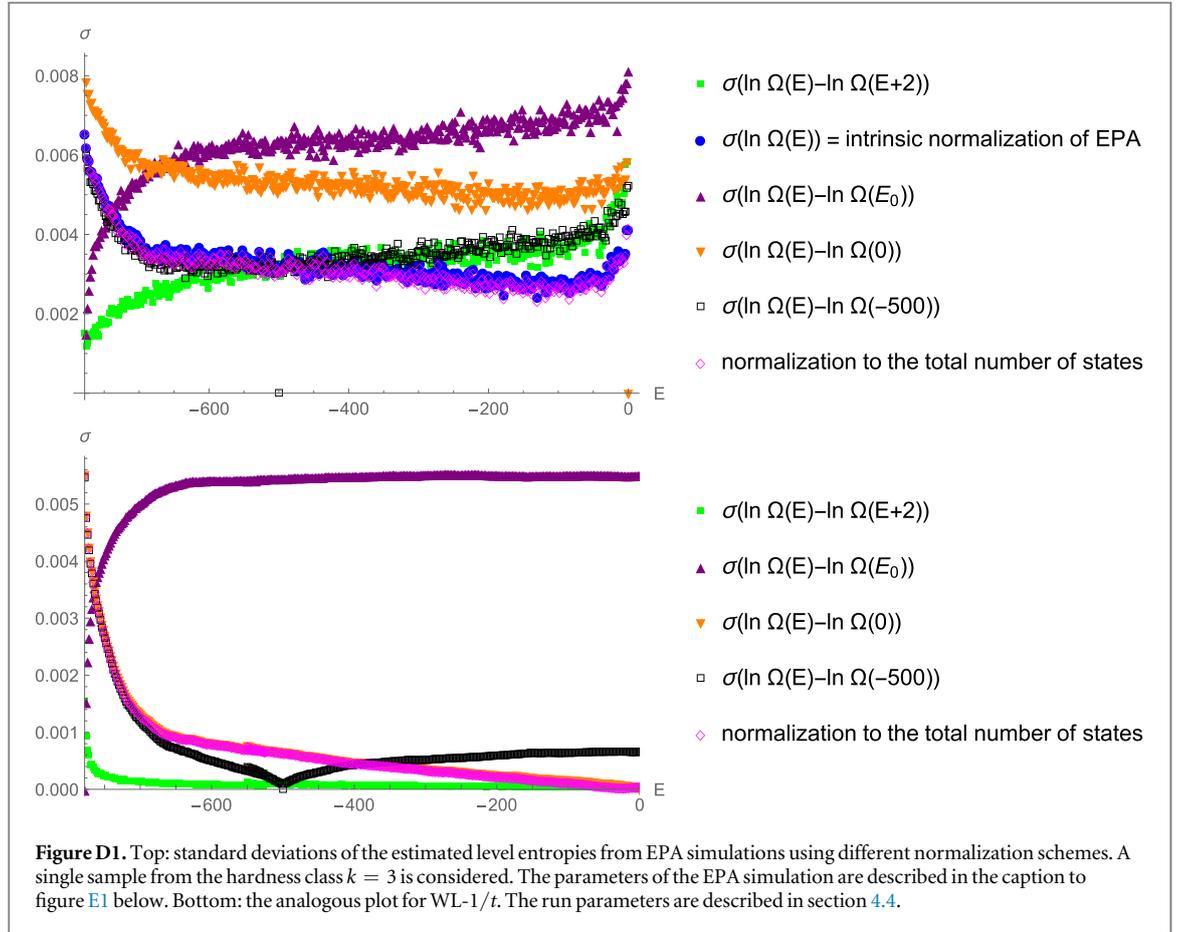
Appendix C. The square-lattice Ising ferromagnet

The case of the Ising ferromagnet on the square lattice, corresponding to the Hamiltonian (1) with $J_{ij} = 1$ for all nearest-neighbor pairs (i, j) , has served as a standard benchmark for entropic samplers since they were first considered [8]. In this case, the DOS can be exactly computed using methods that are somewhat simpler than those of [5], see [4]. Figure C1 shows the relative deviation of $\ln \Omega(E)$ obtained via EPA from the exact DOS for system size $L = 64$. Similar plots can be found for the WL method in [8, 9] and subsequently in many papers on improved methods. The parameters used for the EPA run are $R = 100\,000$, $\theta = 10$, and $\alpha^* = 0.86$, corresponding to $N_\beta \approx 230$ temperature steps and 9.4×10^{11} spin flips. It is apparent that a high accuracy is achieved across the whole energy range.

In table C1 we compare the average deviation Δ according to equation (7) for different runs of the WL-1/t and EPA algorithms. The WL-1/t runs obtained the full DOS, while EPA simulations obtained the degeneracies of energy levels with $E \lesssim 0$ and we exploited the symmetry $\Omega(-E) = \Omega(E)$. The accuracy of WL-1/t is approximately the same as the accuracy of EPA for $5 < \theta < 10$, for the same number of spin flips. Introducing an automatic adaptation of θ for EPA results in an additional approximately threefold reduction of $\bar{\Delta}$ [63].

Appendix D. Normalization of the DOS estimates

While EPA provides the DOS with its absolute normalization, this is not natively the case for WL. For a fair comparison hence different possible normalizations should be considered. Two main normalization schemes are known. The first one consists of fixing the value of $\Omega(E^*)$ at a specific energy, for example at the ground state $E^* = E_0$ or at $E^* = 0$. Alternatively, one can use the fact that the total number of states is 2^N , i.e.

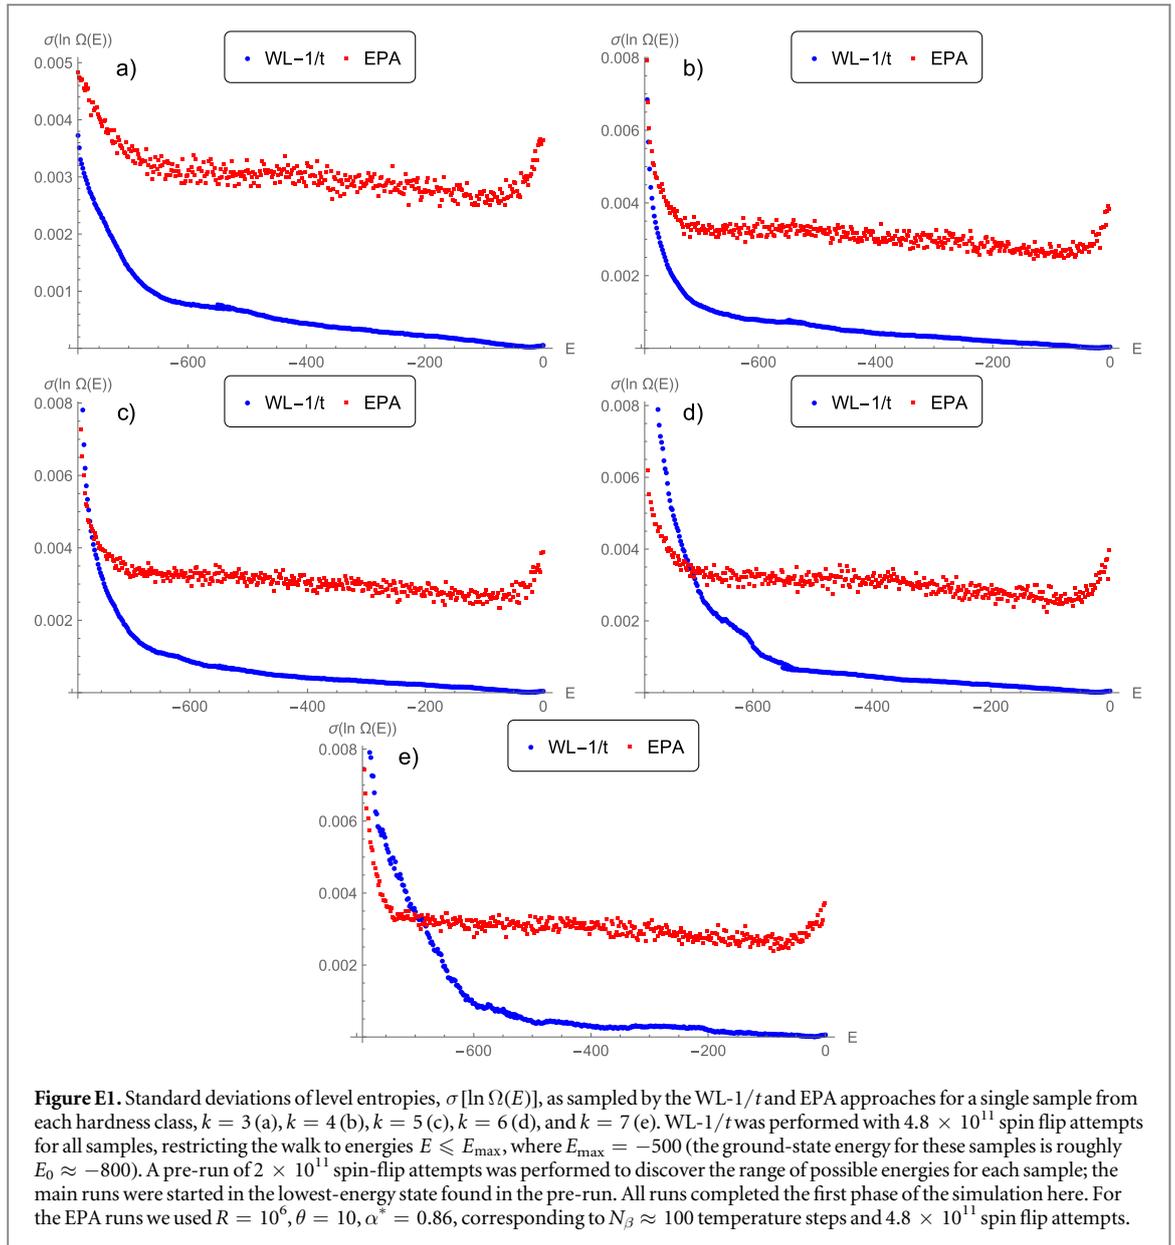


$$\sum_E \Omega(E) = 2^N. \quad (\text{D.1})$$

In figure D1 we show the effect of different normalizations applied to the DOS estimate resulting from an EPA simulation (top panel) and a WL-1/ t run (bottom panel) for a single realization (different realizations provide comparable relative performances). It is apparent that different normalizations lead to rather different statistical fluctuations for different energies, and that EPA and WL-1/ t behave quite differently in this respect. Clearly, fixing the DOS at a specific E^* leads to zero fluctuations at this point. Averaged over all energies, however, it is found that the normalization by the total number of states leads to the lowest fluctuations, and for EPA these are very similar to those found from the intrinsic normalization of EPA. It is also possible to consider entropy *differences* such as $\ln \Omega(E) - \ln \Omega(E + 2)$ that are intrinsically normalization independent and also feature relatively low levels of fluctuation. Two of the curves are almost identical in the bottom panel in figure D1, because the DOS has its largest value at $E = 0$, hence the degeneracy is most accurately estimated by WL-1/ t at $E = 0$.

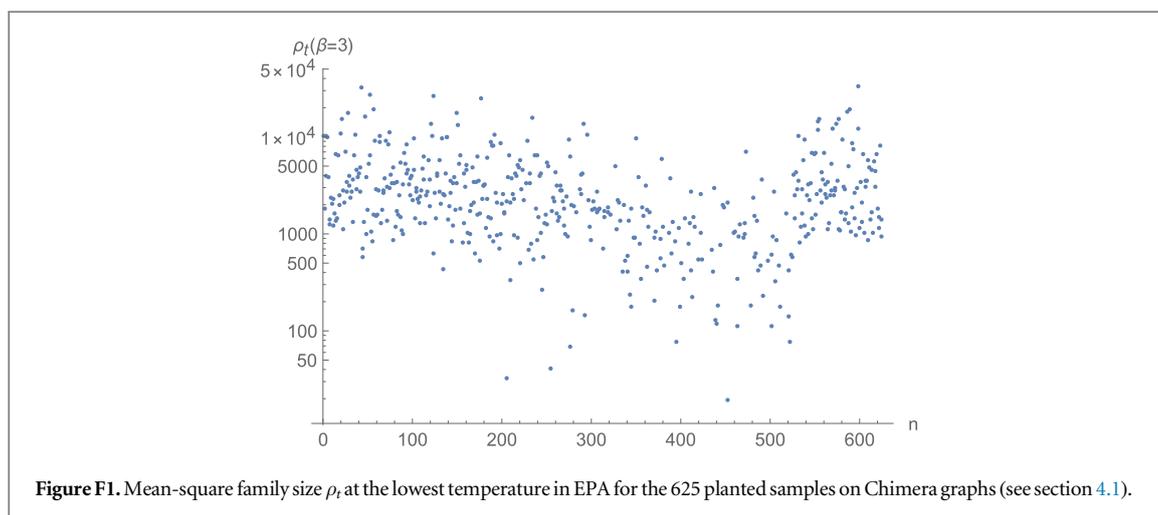
Appendix E. Precision in different energy ranges

A closer comparison of WL-1/ t and EPA is possible by considering the achieved precision with the same computational effort, but resolved by energies. To this end, we studied the behavior of $\sigma[\ln \Omega(E)]$, the standard deviation of the level entropies. Here, we used the normalization to the total number of states according to equation (D.1). Simulations were performed with both methods and the same set of samples, using parameters that ensure that the same number of spin flips is performed. Figure E1 shows the result of one sample of each hardness class $k = 3, 4, 5, 6$, and 7. While both methods have most problems for the immediate vicinity of the ground state, as expected, one can see that EPA results in a relatively small σ for low energies, especially for hard samples, and the results deteriorate only in the immediate vicinity of the ground state. In contrast, the results from WL-1/ t deteriorate already for $E \lesssim -600$, but it yields smaller fluctuations for higher energies, with the best performance at $E \approx 0$. This seems plausible as WL spends much more time at the higher energies with much larger entropies, whereas in PA the population is continuously moved from high to low energies.



Appendix F. Hardness of the planted ensemble

It is interesting to consider some of the hardness measures for the EPA method also for the other ensembles discussed here. Figure F1 shows the mean-square family size ρ_t at the lowest temperature in EPA for the 625 planted samples on Chimera graphs (see section 4.1). The average value is $\rho_t \approx 2000$, so the planted samples of this type are much harder than the random ones (see section 4.4).



ORCID iDs

Lev Barash  <https://orcid.org/0000-0002-2298-785X>

Jeffrey Marshall  <https://orcid.org/0000-0002-5011-1281>

Martin Weigel  <https://orcid.org/0000-0002-0914-1147>

Itay Hen  <https://orcid.org/0000-0002-7009-7739>

References

- [1] Sethna J P 2006 *Statistical Mechanics: Entropy, Order Parameters and Complexity* (Oxford: Oxford University Press)
- [2] Hüller A 1994 *Z. Phys. B* **93** 401–5
- [3] Sachs M 1963 *Solid State Theory* (New York: McGraw-Hill)
- [4] Beale P D 1996 *Phys. Rev. Lett.* **76** 78
- [5] Galluccio A, Loebl M and Vondrák J 2000 *Phys. Rev. Lett.* **84** 5924–7
- [6] Nourani Y and Andresen B 1999 *Theor. Comput. Sci.* **215** 51–68
- [7] Brown B, Flammia S T and Schuch N 2011 *Phys. Rev. Lett.* **107** 040501
- [8] Wang F and Landau D P 2001 *Phys. Rev. E* **64** 056101
- [9] Wang F and Landau D P 2001 *Phys. Rev. Lett.* **86** 2050–3
- [10] Belardinelli R E and Pereyra V D 2007 *Phys. Rev. E* **75** 046701
- [11] Liang F, Liu C and Carroll R J 2007 *J. Am. Stat. Assoc.* **102** 305–20
- [12] Vogel T, Li Y W, Wüst T and Landau D P 2013 *Phys. Rev. Lett.* **110** 210603
- [13] Berg B A and Neuhaus T 1992 *Phys. Rev. Lett.* **68** 9–12
- [14] Schneider S, Mueller M and Janke W 2017 *Comput. Phys. Commun.* **216** 1–7
- [15] Bryngelson J D, Onuchic J N, Succi N D and Wolynes P G 1995 *Proteins* **21** 167–95
- [16] Binder K and Young A P 1986 *Rev. Mod. Phys.* **58** 801
- [17] Belletti F et al 2009 *Comput. Sci. Eng.* **11** 48–58
- [18] Kadowaki T and Nishimori H 1998 *Phys. Rev. E* **58** 5355
- [19] Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A and Preda D 2001 *Science* **292** 472–5
- [20] Johnson M W et al 2011 *Nature* **473** 194–8
- [21] Hen I, Job J, Albash T, Rønnow T F, Troyer M and Lidar D A 2015 *Phys. Rev. A* **92** 042325
- [22] Katzgraber H G, Hamze F, Zhu Z, Ochoa A J and Muñoz-Bauza H 2015 *Phys. Rev. X* **5** 031026
- [23] Boixo S, Isakov S V, Smelyanskiy V N, Babbush R, Ding N, Jiang Z, Martinis J M and Neven H 2018 *Nat. Phys.* **14** 595–600
- [24] Benedetti M, Realpe-Gómez J, Biswas R and Perdomo-Ortiz A 2016 *Phys. Rev. A* **94** 022308
- [25] Adachi S H and Henderson M P 2015 arXiv:1510.06356
- [26] Amin M H, Andriyash E, Rolfe J, Kulchitskyy B and Melko R 2018 *Phys. Rev. X* **8** 021050
- [27] Zhang B H, Wagenbreth G, Martin-Mayor V and Hen I 2017 *Sci. Rep.* **7** 1044
- [28] Mandrà S, Zhu Z and Katzgraber H G 2017 *Phys. Rev. Lett.* **118** 070502
- [29] Könz M S, Mazzola G, Ochoa A J, Katzgraber H G and Troyer M 2018 arXiv:1806.06081
- [30] Marshall J, Rieffel E G and Hen I 2017 *Phys. Rev. Appl.* **8** 064025
- [31] Barash L Y, Fadeeva M A and Shchur L N 2017 *Phys. Rev. E* **96** 043307
- [32] Barahona F, Maynard R, Rammal R and Uhry J P 1982 *J. Phys. A: Math. Gen.* **15** 673
- [33] Bieche I, Maynard R, Rammal R and Uhry J P 1980 *J. Phys. A: Math. Gen.* **13** 2553
- [34] Pardella G and Liers F 2008 *Phys. Rev. E* **78** 056705
- [35] Khoshbakht H and Weigel M 2018 *Phys. Rev. B* **97** 064410
- [36] Blackman J A and Poulter J 1991 *Phys. Rev. B* **44** 4374–86
- [37] Saul L and Kardar M 1993 *Phys. Rev. E* **48** R3221
- [38] King A D, Lanting T and Harris R 2015 arXiv:1502.02098
- [39] Wang W, Mandrà S and Katzgraber H G 2017 *Phys. Rev. E* **96** 023312
- [40] Barthel W, Hartmann A K, Leone M, Ricci-Tersenghi F, Weigt M and Zecchina R 2002 *Phys. Rev. Lett.* **88** 188701
- [41] Krzakala F and Zdeborová L 2009 *Phys. Rev. Lett.* **102** 238701

- [42] Dechter R 1997 *Constraints* **2** 51–5
- [43] Choi V 2008 *Quantum Inf. Proc.* **7** 193–209
- [44] Choi V 2011 *Quantum Inf. Proc.* **10** 343–53
- [45] Johnson M W et al 2011 *Nature* **473** 194–8
- [46] Berkley A J et al 2013 *Phys. Rev. B* **87** 020502(R)
- [47] Bunyk P I et al 2014 *IEEE Trans. Appl. Supercond.* **24** 1–10
- [48] Weigel M, Katzgraber H G, Machta J, Hamze F and Andrist R S 2015 *Phys. Rev. X* **5** 019901(E)
- [49] Barahona F 1982 *J. Phys. A: Math. Gen.* **15** 3241
- [50] Gross J, Zierenberg J, Weigel M and Janke W 2018 *Comput. Phys. Commun.* **224** 387–95
- [51] Yan Q and de Pablo J J 2003 *Phys. Rev. Lett.* **90** 035701
- [52] Yin J and Landau D P 2012 *Comput. Phys. Commun.* **183** 1568–73
- [53] Iba Y 2001 *Trans. Japan. Soc. Artif. Intell.* **16** 279–86
- [54] Hukushima K and Iba Y 2003 *AIP Conf. Proc.* **690** 200–6
- [55] Machta J 2010 *Phys. Rev. E* **82** 026704
- [56] Wang W, Machta J and Katzgraber H G 2015 *Phys. Rev. E* **92** 063307
- [57] Borovský M, Weigel M, Barash L Y and Žukovič M 2016 *EPJ Web Conf.* **108** 02016
- [58] Barash L Y, Weigel M, Borovský M, Janke W and Shchur L N 2017 *Comput. Phys. Commun.* **220** 341–50
- [59] Barash L Y, Weigel M, Shchur L N and Janke W 2017 *Eur. Phys. J. Spec. Top.* **226** 595–604
- [60] Amey C and Machta J 2018 *Phys. Rev. E* **97** 033301
- [61] Barzegar A, Pattison C, Wang W and Katzgraber H G 2018 *Phys. Rev. E* **98** 053308
- [62] Ferrenberg A M and Swendsen R H 1989 *Phys. Rev. Lett.* **63** 1195
- [63] Weigel M, Barash L Y, Janke W and Shchur L N in preparation
- [64] Weigel M, Janke W and Hu C K 2002 *Phys. Rev. E* **65** 036109
- [65] Hukushima K and Nemoto K 1996 *J. Phys. Soc. Japan.* **65** 1604–8
- [66] Sokal A D 1997 Monte Carlo methods in statistical mechanics: foundations and new algorithms *Functional Integration: Basics and Applications (Proc. 1996 NATO Advanced Study Institute in Cargèse)* ed C DeWitt-Morette et al (New York: Plenum) pp 131–92
- [67] Katzgraber H G, Trebst S, Huse D A and Troyer M 2006 *J. Stat. Mech.* **P03018**
- [68] Bittner E, Nußbaumer A and Janke W 2008 *Phys. Rev. Lett.* **101** 130603
- [69] Alvarez Baños R et al 2010 *J. Stat. Mech.* **P06026**
- [70] Martin-Mayor V and Hen I 2015 *Sci. Rep.* **5** 15324
- [71] Weigel M, Barash L Y, Borovský M, Shchur L N and Janke W in preparation
- [72] Farris A C, Li Y W and Eisenbach M 2019 *Comput. Phys. Commun.* **235** 297–304
- [73] Brandt S 1998 *Data Analysis: Statistical and Computational Methods for Scientists and Engineers* 3rd edn (Berlin: Springer)
- [74] Efron B and Tibshirani R J 1994 *An Introduction to the Bootstrap* (London: Chapman and Hall)
- [75] Weigel M and Janke W 2009 *Phys. Rev. Lett.* **102** 100601
- [76] Weigel M and Janke W 2010 *Phys. Rev. E* **81** 066701