

Performance Limitations of Flat Histogram Methods and Optimality of Wang-Landau Sampling

P. Dayal¹, S. Trebst^{1,2}, S. Wessel¹, D. Würtz¹, M. Troyer^{1,2}, S. Sabhapandit³, S.N. Coppersmith³

⁽¹⁾ *Theoretische Physik, Eidgenössische Technische Hochschule Zürich, CH-8093 Zürich, Switzerland*

⁽²⁾ *Computational Laboratory, Eidgenössische Technische Hochschule Zürich, CH-8092 Zürich, Switzerland and*

⁽³⁾ *Department of Physics, University of Wisconsin, Madison, WI 53706 USA*

(Dated: June 5, 2003)

We determine the optimal scaling of local-update flat-histogram methods with system size by using a perfect flat-histogram scheme based on the exact density of states of 2D Ising models. The typical tunneling time needed to sample the entire bandwidth does not scale with the number of spins N as the minimal N^2 of an unbiased random walk in energy space. While the scaling is power law for the ferromagnetic and fully frustrated Ising model, for the $\pm J$ nearest-neighbor spin glass the distribution of tunneling times is governed by a fat-tailed Fréchet extremal value distribution that obeys exponential scaling. We find that the Wang-Landau algorithm shows the same scaling as the perfect scheme and is thus optimal.

PACS numbers: 02.70.Rr, 75.10.Hk, 64.60.Cn

Monte Carlo methods are well-suited for the simulation of large many body problems, since the complexity for a single Monte Carlo update step scales only polynomially and often linearly in the system size, while the configuration space grows exponentially with the system size. The performance of a Monte Carlo method is then determined by how many update steps are needed to efficiently sample the configuration space. For second order phase transitions in unfrustrated systems the problem of “critical slowing down” – a rapid divergence of the number of Monte Carlo steps needed to obtain a subsequent uncorrelated configuration – was solved more than a decade ago by cluster update algorithms [1]. At first order phase transitions and in systems with many local minima of the free energy such as frustrated magnets or spin glasses, there is the similar problem of long tunneling times between local minima. With energy barriers ΔE scaling linearly with the linear system size L , the tunneling times τ at an inverse temperature $\beta = 1/k_B T$ scale exponentially with the system size, $\tau \sim \exp(\beta \Delta E) \propto \exp(\text{const} \times L)$. Several methods were developed to overcome this tunneling problem, such as the multicanonical method [2], broad histograms [4], simulated and parallel tempering [3], and Wang-Landau sampling [5]. The common aim of all these methods is to broaden the range of energies sampled within Monte Carlo simulations from the sharply peaked distribution of canonical sampling at fixed temperature in order to ease the tunneling through barriers.

Ideally, all relevant energy levels are sampled equally often during a simulation, thus producing a “flat histogram” in energy space. Some methods approach this goal by variations and generalizations of canonical distributions [2, 3], while others [4, 5] discard the notion of temperature completely and instead are formulated in terms of the density of states. With a probability $p(E)$ for a single configuration with energy E , the probability of sampling an arbitrary configuration with energy E is

given as $P_E = \rho(E)p(E)$, where the density of states $\rho(E)$ counts the number of states with energy E . Upon choosing $p(E) \propto 1/\rho(E)$ instead of $p(E) \propto \exp(-\beta E)$ one obtains a constant probability P_E for visiting each energy level E , and hence a flat histogram. Wang and Landau [5] proposed a simple and elegant flat histogram algorithm that iteratively improves approximations to the initially unknown density of states $\rho(E)$. Once $\rho(E)$ is determined with sufficient accuracy, the Monte Carlo algorithm just performs a random walk in energy space. Within two years of publication this algorithm has been applied to a large number of problems [6, 7, 8] and extended to quantum systems [9].

In this Letter we investigate the performance of flat histogram algorithms in general, and the Wang-Landau algorithm in particular, for three systems for which the density of states $\rho(E)$ is known exactly on finite two-dimensional (2D) lattices: the Ising ferromagnet as the simplest example, the fully frustrated Ising model as a prototype for frustrated systems, and the $\pm J$ Ising spin glass. For each of these models we construct a *perfect flat histogram method* by simulating a random walk in configuration space where we employ the *known density of states* for these models to set $p(E) \propto 1/\rho(E)$.

As a measure of performance we use the average tunneling time τ to get from a ground state (lowest energy configuration) to an anti-ground state (configuration of highest energy), which is the relevant time scale for sampling the whole phase space [10]. Since the number of energy levels in a d -dimensional system with linear size L scales with the number of spins $N = L^d$, the tunneling time for a pure random walk in energy space is

$$\tau \propto N^2 = L^{2d}. \quad (1)$$

This scaling was in fact found for first order phase transitions [2, 5]. However, none of the systems we study exhibit this scaling. While the ferromagnetic and fully-




FIG. 1: Scaling of tunneling times τ from the ground state to the anti-ground state as a function of system size L for a perfect flat histogram method that samples using the exact density of states. Shown are results for the ferromagnet (FM) and fully frustrated (FF) 2D Ising models with both local and N -fold way updates; the inset illustrates the frustrated couplings. In all cases polynomial scaling $\tau \propto L^{2d+z}$ is found, with $z_{\text{local}}^{\text{FM}} = 0.743 \pm 0.007$, $z_{N\text{-fold}}^{\text{FM}} = 0.729 \pm 0.011$, $z_{\text{local}}^{\text{FF}} = 1.727 \pm 0.004$, and $z_{N\text{-fold}}^{\text{FF}} = 1.692 \pm 0.004$.

frustrated models exhibit power law scaling, for the spin glass the distribution of characteristic tunneling times is extremely broad and appears to diverge exponentially with system size. For all three models the scaling of the performance of the Wang-Landau algorithm is the same as that of the perfect flat histogram method.

We first look at homogeneous systems, with both ferromagnetic and fully frustrated couplings (see inset of Fig. 1). Exact densities of states $\rho(E)$ were calculated using the program of Beale [11] for the ferromagnet and the algorithm of Saul and Kardar [12] for the fully frustrated Ising model. In Fig. 1 the measured tunneling times are plotted versus system size L for the perfect flat histogram method using both local and N -fold way updates. Instead of the expected scaling Eq. (1), we find a more rapid increase following $\tau \propto L^{2d+z}$ with $z_{\text{local}}^{\text{FM}} = 0.743 \pm 0.007$ for the ferromagnetic (FM) model and $z_{\text{local}}^{\text{FF}} = 1.727 \pm 0.004$ for the fully frustrated (FF) model. Frustration significantly increases the scaling exponent but still conserves power law scaling.

A previous study [7] suggests that N -fold way updates [13] speed up Wang-Landau sampling. We find identical scaling exponents for N -fold way and local updates within our error bars, see Fig. 1, implying that any performance improvement remains constant with system size. As can be seen from Fig. 2, N -fold way updates reduce the tunneling time by roughly a factor of 2, independent of system size. In practice, the reduction of tunneling times is offset by the added expense of N -fold way updates. The slight decrease of tunneling ratios in the fully frustrated model compared to the ferromagnetic




FIG. 2: Scaling of the ratios of tunneling times measured using local and N -fold way updates. All simulation are done with single-spin flip updates weighted by the exact density of states. For each system size of the ferromagnet and the fully frustrated 2D Ising models 50,000 tunneling events and at least 25 tunneling events for the 2D spin glass 1000 randomly generated realizations for $L = 6, 8 \dots 18$ and 350 realizations for $L = 20$ are used for averaging.




FIG. 3: Normalized distributions of tunneling times (left panels) and the ratio of the number of first excited states to the number of ground states $\rho(E_1)/\rho(E_0)$ (right panels). For both system sizes, $L = 6$ and $L = 16$, 1000 randomly generated 2D $\pm J$ spin glass realizations are sampled. The measured histograms of tunneling times (using local updates and the exact density of states) and $\rho(E_1)/\rho(E_0)$ both follow fat tailed Fréchet distributions (solid lines).

model can be explained by the increasing degeneracy of ground states in the fully frustrated model, which decreases the ground state lifetime and makes N -fold way updates less effective.

To determine the performance in more complex energy landscapes, we study the 2D $\pm J$ Ising spin glass, where we find exponential scaling. We measured tunneling times of the perfect flat histogram method for 1000




FIG. 4: Scaling of the parameters of Fréchet distribution of the tunneling times (see Eq. 2) of the 2D $\pm J$ Ising spin glass using perfect flat histogram sampling as a function of system size L . Solid and dashed lines show least square fits assuming exponential and algebraic (power law) scaling respectively. The exponential fits of the data with $L \geq 8$ for μ and β give $\chi^2 = 12$ and 7.6, respectively, being significantly better than algebraic fits with $\chi^2 = 130$ and 61, respectively. The inset shows the scaling of the shape parameter.

realizations for the system sizes $L = 6, 8, 10, 12, 14, 16, 18$, and 350 realizations for $L = 20$ using the exact density of states obtained by the algorithm of Saul and Kardar [12]. For fixed system size the tunneling times are scattered over several orders of magnitude for the various realizations, as shown in Figs. 3 a) and 3 b). To analyze the underlying distributions we use extremal value theory [14]. The central limit theorem for extremal values [15] states that the extrema of large samples are distributed according to one of only three distributions, depending on whether the tails of the original distribution are fat-tailed (algebraic), exponential or thin-tailed (decaying faster than exponential). This theorem is successfully applied in the analysis of tails in diverse fields such as hydrology, insurance and finance [16]. Surprisingly, here we find that *not only the extrema*, but *all of the measured tunneling times* [see Figs. 3 a) and 3 b)] are distributed according to the Fréchet extremal value distribution for fat-tailed distributions:

$$H_{\xi;\mu;\beta}(\tau) = \exp \left[- \left(1 - \xi \frac{\tau - \mu}{\beta} \right)^{1/\xi} \right], \quad (2)$$

with $\xi < 0$. The parameters of the distribution are determined by a maximum likelihood estimator. Fig. 4 shows that the location parameter μ specifying the maximum of the distribution and the scale parameter β determining the width of the distribution scale exponentially with linear system size L :

$$\mu \propto \exp(L/(4.21 \pm 0.04)), \quad (3)$$

$$\beta \propto \exp(L/(3.37 \pm 0.05)). \quad (4)$$




FIG. 5: Correlation between tunneling times and ratios of the number of first excited states to the number of ground states $\rho(E_1)/\rho(E_0)$. Shown are data from 1000 randomly generated 2D $\pm J$ spin glass realizations of fixed system size $L = 16$.

The shape parameter ξ , shown in the inset of Fig. 4, determines the power law decay of the fat tails of the distribution

$$\frac{dH_{\xi;\mu;\beta}}{d\tau} \xrightarrow{\tau \rightarrow \infty} \tau^{-(1-1/\xi)}. \quad (5)$$

From this asymptotic behavior one can see that the m -th moment of a fat tailed Fréchet distribution (with $\xi < 0$) is well defined only if $|\xi| < 1/m$. We find that $|\xi| > 1/2$ for $L \geq 10$, which implies that the variance ($m = 2$) does not exist and the central limit theorem for mean values does not apply. Any direct estimate of the mean tunneling time – as opposed to the most likely time given by the Fréchet location parameter – then has an infinite error. This breadth may explain differences in our conclusion from those in Refs. 17, 18. It also looks plausible, although we cannot go to large enough systems, that $|\xi|$ increases monotonically with system size and could become larger than 1, in which case even the mean tunneling time ($m = 1$) becomes ill-defined. The most likely tunneling time, given by the location parameter μ , would still remain well defined and finite.

Since the tunneling time is to a large extent dominated by the energy landscape at low energies, we study, as a qualitative measure for the complexity of the energy landscape, the distribution of $\rho(E_1)/\rho(E_0)$ where $\rho(E_1)$ is the number of first excited states and $\rho(E_0)$ the number of ground states. Again we find fat tailed Fréchet distributions, as shown in Figs. 3 c) and 3 d), suggesting that intrinsic properties of the 2D $\pm J$ spin glass account for the observed distribution of the measured tunneling times. In Fig. 5 we show the measured tunneling times versus the ratio $\rho(E_1)/\rho(E_0)$. A strong correlation over five orders of magnitude is found, supporting the argument.




FIG. 6: Convergence of the tunneling times τ during the calculation of the density of states using Wang Landau sampling versus Wang-Landau parameter f . Main panel: 2D Ising ferromagnet model; inset: typical samples for the 2D $\pm J$ spin glass. For each system size the tunneling times are averaged over 50 (FM) and 500 (SG) independent runs, and are given in units of the average tunneling time τ_{exact} measured during random walks in the exact density of states using Wang-Landau (local update) sampling.

The question arises why the scaling behaviors of the fully frustrated model and the spin glass are different. Both models have an exponentially large number of ground states and an extensive ground state entropy, but the tunneling time scaling is algebraic for the fully frustrated model and exponential for the spin glass. We believe that the reason is the difference in complexity of the energy landscapes in the two models. While the energy landscape above the large number of ground states of a frustrated model can be simple, the energy landscape of the spin glass is more complex with an extremely large number of local minima: The number of first excited states $\rho(E_1)$ that can be reached from the $\rho(E_0)$ ground states by one single spin flip is at most $L^d \rho(E_0)$. The number of local minima that are not connected to the ground state by a single spin flip can thus be estimated as $\rho(E_1) - L^d \rho(E_0)$. For the 2D ferromagnetic model the ratio $\rho(E_1)/\rho(E_0)$ is exactly L^2 , and $\rho(E_1)/\rho(E_0)$ is of order L^2 for the fully frustrated model. In contrast, in the 2D spin glass the ratio $\rho(E_1)/\rho(E_0)$ exceeds L^2 by several orders of magnitude for some realizations, as can be seen from Figs. 3 c), 3 d) and 5.

Finally we compare the tunneling times measured in the Wang-Landau algorithm to the perfect flat histogram method. The Wang-Landau algorithm approaches the exact density of states $\rho(E)$ by multiplying the current estimate at each visited level by a factor f that is reduced towards 1 over time as the algorithm converges. In Fig. 6 we show tunneling times for Wang-Landau sampling as a function of this correction factor f for the Ising ferromagnet (main panel) and for the spin glass (inset). Results for the fully frustrated model (not shown) are

qualitatively similar. In the initial stages of the simulation ($\ln f \gtrsim 10^{-6}$) the tunneling times are shorter than for exact sampling, since the random walk is biased — it is always driven away from the last region visited (due to the increased $\rho(E)$ there). Eventually the tunneling times converge to exactly the same times as for the perfect flat histogram method, indicating convergence of the Wang-Landau algorithm. The Wang-Landau algorithm is thus optimal in the sense that it performs identically to a perfect flat histogram method. Unlike in recent applications to continuum systems [19] no convergence problems are observed for these lattice models.

Our benchmarks of a perfect flat histogram method provide a lower bound for the tunneling times of other “flat histogram” methods such as multicanonical [2], tempering [3], broad histograms [4] or Wang-Landau sampling [5]. From our analysis of tunneling times we find that the Wang Landau algorithm scales identically as the perfect flat histogram method and is thus optimal. We expect that the other methods will perform similarly when well-tuned. However, whether one uses local or N -fold way updates, the scaling is not the N^2 scaling of a random walk in a system with N energy levels, but slower, namely $N^2 L^z$ for both the ferromagnetic ($z = 0.743 \pm 0.007$) and the fully frustrated Ising model ($z = 1.727 \pm 0.004$). The power law scaling for the frustrated model is very encouraging and demonstrates, for the first time, that the Wang-Landau algorithm is well suited for frustrated models. A combination with alternative sampling schemes, such as in Ref. [8], can further improve the performance. The observation of a power law with an exponent larger than 2 is not trivially explained in the context of a random walk in energy space and the subject of further investigations.

The exponential scaling for the $\pm J$ spin glass even for the perfect method shows a limitation for any flat histogram method. Here the distribution of tunneling times follows a fat tailed Fréchet extremal value distribution. The origin of this extremal character of the 2D $\pm J$ Ising spin glass remains an interesting open question. Further studies are in progress to investigate this issue as well as three-dimensional classical spin and quantum spin glasses.

We thank F. Alet, N. Kawashima, W. Krauth, M. Pelikan, A. Sandvik, and B. J. Schulz for helpful discussions and acknowledge the support of the Aspen Center for Physics where this project was started. PD, ST, and SW acknowledge support by the Swiss National Science Foundation, and SC and SS acknowledge support from the Petroleum Research Fund of the American Chemical Society and by NSF-DMR.

-
- [1] R. H. Swendsen and J. S. Wang, Phys. Rev. Lett. **58**, 86 (1987); U. Wolff, Phys. Rev. Lett. **62**, 361 (1989).
- [2] B. A. Berg and T. Neuhaus, Phys. Rev. Lett. **68**, 9 (1992); Phys. Lett. B. **267**, 249 (1991).
- [3] K. N. K. Hukushima, J. Phys. Soc. Japan **65**, 1604 (1996); G. P. E. Marinari, Europhys. Lett. **19**, 451 (1992); A.P.Lyubartsev et al., J. Chem. Phys. **96**, 1776 (1992).
- [4] P. M. C. de Oliveira et al., Braz. J. Phys. **26**, 677 (1996).
- [5] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001); Phys. Rev. E **64**, 056101 (2001).
- [6] C. Yamaguchi and Y. Okabe, J. Phys. A **34**, 8781 (2001); F. Calvo, Molec. Phys **100**, 3421 (2002); Q. Yan et al., J. Chem. Phys. **116**, 8745 (2002); M. S. Shell et al., Phys. Rev. E **66**, 056703 (2002); T. S. Jain and J. J. de Pablo, J. Chem. Phys. **116**, 7238 (2002); N. Rathore and J. J. de Pablo, J. Chem. Phys. **116**, 7225 (2002); T.J.H.Vlugt, Molecular Physics **100**, 2763 (2002); E. B. Kim et al., J. Chem. Phys. **117**, 7781 (2002).
- [7] B. J. Schulz et al., Int. J. Mod. Phys. **C13**, 477 (2002).
- [8] C. Yamaguchi and N. Kawashima, Phys. Rev. E **65**, 056710 (2002).
- [9] M. Troyer et al., Phys. Rev. Lett. **90**, 120201 (2003).
- [10] We have checked that the behavior of τ is not sensitive to details of its definition.
- [11] P. D. Beale, Phys. Rev. Lett. **76**, 78 (1995).
- [12] L. K. Saul and M. Kardar, Nuclear Physics B **432**, 641 (1994).
- [13] A. B. Bortz et al., J. Comput. Phys. **17**, 10 (1975).
- [14] P. Embrechts, C. Klüpelberg, and T. Mikosch, *Modelling Extremal Events* (Springer, 1997).
- [15] R. A. Fisher and L. H. C. Tippett, Proc. Cambridge Philos. Soc. **24**, 180 (1928).
- [16] R.-D. Reiss and M. Thomas, *Statistical Analysis of Extreme Values* (Birkhäuser, 1997).
- [17] B. A. Berg and T. Celik, Int. J. Mod. Phys. **C3**, 1251 (1992).
- [18] Z. F. Zhan et al., Physica A **285**, 239 (2000).
- [19] Q. Yan and J. J. de Pablo, Phys. Rev. Lett. **90**, 035701 (2003); M.S. Shell et al., cond-mat/0305210.