

Perfect Tempering

M. Daghofer*, M. Konegger*, H. G. Evertz* and W. von der Linden*

**Institute for Theoretical and Computational Physics, TU Graz, Austria*

Abstract. Multimodal structures in the sampling density (e.g. two competing phases) can be a serious problem for traditional Markov Chain Monte Carlo (MCMC), because correct sampling of the different structures can only be guaranteed for infinite sampling time. Samples may not decouple from the initial configuration for a long time and autocorrelation times may be hard to determine.

We analyze a suitable modification [1] of the simulated tempering idea [2], which has orders of magnitude smaller autocorrelation times for multimodal sampling densities and which samples all peaks of multimodal structures according to their weight. The method generates **exact**, i. e. uncorrelated, samples and thus gives access to reliable error estimates. **Exact tempering** is applicable to arbitrary (continuous or discrete) sampling densities and moreover presents a possibility to calculate integrals over the density (e.g. the partition function for the Boltzmann distribution), which are not accessible by usual MCMC.

INTRODUCTION

Simulated Tempering was introduced in Ref. [2], parallel tempering, also known as Replica Exchange Monte Carlo, in Ref. [3, 4] and both have been widely used (see e. g. Refs. [5, 6]) to make Markov chain Monte Carlo faster. For an introduction to both methods see Ref. [7].

Propp and Wilson introduced the coupling from the past (CFTP) method to draw exact samples, i.e. samples which are guaranteed to be uncorrelated and to obey the desired distribution in Ref. [8]. Applications of this method to continuous degrees of freedom and cluster algorithms exist, see Refs. [9] and [10].

A small modification of the Simulated Tempering algorithm likewise allows to obtain uncorrelated samples, see Ref. [1]. The first two sections give an introduction to this procedure and the resulting error estimates. We then examine the tempering Markov matrix and the autocorrelation time in and give indications about the needed parameters. We investigate the tempering algorithm for multidimensional continuous distributions and find a polynomial dependence on the dimension. Finally, we compare exact sampling with simulated tempering to the CFTP method for the two dimensional Ising model and find a quadratic dependence on the system size for simulated tempering, while CFTP needs exponential time for cold but finite temperatures.

EXACT SAMPLING WITH SIMULATED TEMPERING

Besides speeding simulations up, Simulated Tempering provides an alternative way to obtain exact samples from arbitrary probability density functions [1]. Fig. 1 shows the

principle for a multi-modal distribution $p_1(X)$ consisting of two Gaussians, but it does not depend on the specified example and can thus be applied to a variety of probability distributions. We want to draw Exact samples from the distribution $p_1(X)$, which we can not sample directly, where X can be a discrete or continuous quantity of arbitrary dimension. In order to do so, we introduce an additional parameter β and the joint probability $p(X, \beta) = p(X|\beta)p(\beta)$. We have large freedom in choosing $p(X, \beta)$, for the simulation depicted in fig. 1, we chose:

$$p(X, \beta_m) = \frac{1}{Z} \frac{1}{Z_m} p_1(X) \beta_m p_0(X)^{1-\beta_m}, \quad (1)$$

where Z is the overall normalization, Z_m is a constant depending on β_m , which determines $p(\beta_m)$. The additional variable β was allowed to take $M + 1$ discrete values β_m with $\beta_0 = 0$ and $\beta_M = 1$. $p_0(X)$ should be chosen in a way to allow generating Exact samples easily; in our example, it was a single broad Gaussian peak. Furthermore, its range in X -space should be broad enough to cover all structures of $p_1(X)$. For application to physical systems, p_1 would of course be chosen as $e^{-E(X)}$ (with $E(X)$ denoting the energy) and $p_0 = \text{const.}$, which yields $p(X|\beta) = e^{-\beta E(X)}$. In this case, β_M is not chosen to be one, but can take any other value β_{max} .

We then do Markov chain Monte Carlo in the $\{X, \beta\}$ -space, where we alternate a couple of sweeps in X -space with moves in β -direction. In β -direction, $\beta_{m'}$ with $m' = m \pm 1$ is proposed with equal probability and accepted according to the Metropolis scheme. In X -space and for $\beta_m \neq 0$, usual Metropolis updates are employed. A special case arises for X -moves at $\beta_m = 0$. In this case $p(X|\beta = 0) \propto p_0(X)$ and we are able to draw a new exact sample X' distributed according to $p_0(X)$, which gives us a sample X' *uncorrelated* from X .

An example of the resulting random walk is depicted on the 'floor' of Fig. 1. Whenever this random walk reaches $\beta = 0$, a new exact sample from p_0 is drawn independent from the current state of the Markov chain so that the walk forgets its past. The MC time needed for one exact sample is thus given by the time needed by the Markov chain to travel from $\beta_0 = 0$ to $\beta_M = 1$ and back again.

A plain MCMC run would instead be trapped in one of the two peaks and rarely tunnel to the other. Repeating several plain MCMC runs and taking their average would give the wrong expectation value $\bar{x} = 0$, because the different weight of the peaks would not be accounted for.

EXPECTATION VALUES AND ERROR ESTIMATES

As the $\{X, \beta\}$ -samples obtained by the simulation obey $p(X, \beta)$, the X drawn at a given temperature β_m obeys $p(X|\beta_m)$. Expectation values for $\beta_M = 1$ are therefore calculated




FIGURE 1. Example for a Simulated Tempering run. On the ‘floor’, the Markov chain travels through the $\{x, \beta\}$ -space, the larger dots are the obtained samples, the dotted lines show the way the Markov process has taken. Via $\beta_0 = 0$, the walk reaches both peaks at $\beta_M = 1$, although no direct tunneling between them occurs. The peaks (solid lines) are the probabilities $p(x|\beta)$ for the various discrete β -values. The samples drawn at a certain temperature obey this distribution. On the right hand ‘wall’, the vertical axis is the time axis of the simulation; one sees the wandering of the random walk through the temperatures. The thick lines are inserted where the walk reaches $\beta_0 = 0$, i. e. where an independent exact sample is drawn from $p_0 = p(x|\beta = 0)$ (chosen as a single broad Gaussian peak). At these points, the walk forgets its past and a new uncorrelated bin starts.

from all (correlated and uncorrelated) samples obtained at a given temperature:

$$\begin{aligned} \bar{X}(\beta_M = 1) &= \frac{1}{N_M} \sum_{j=1}^{N_M} X_j = \frac{1}{N_M} \sum_{i=1}^{N_{M,ind}} \sum_{v=1}^{N_i} X_{i,v} \\ \langle \bar{X}(\beta_M = 1) \rangle &= \frac{1}{N_M} \sum_j \langle X_j \rangle = \langle X \rangle \end{aligned} \quad (2)$$

The X_j are the measurements obtained at the desired temperature $\beta_M = 1$, their index j was broken into i and v with i denoting the independent and uncorrelated bins and v labeling the correlated measurements within one bin, see Fig. 1. $N_{M,ind}$ is the number of independent bins which contain at least one sample drawn at β_M , and N_i the number of measurements within the i -th bin. $N_M = \sum_{i=1}^{N_{M,ind}} N_i$ is the total number of times the simulation has visited the desired temperature $\beta_M = 1$. A bar denotes the sample mean obtained in the Monte Carlo run, $\langle \dots \rangle$ denotes an expectation value over all samples. The sample mean is obviously unbiased.

It is worth noting that measuring the bin averages does not give the same result, because the probability for a move in β -direction, and thus the number of measurements

(N_i) taken in a bin before the walk returns to $\beta = 0$, is a random variable and depends on the current sample X :

$$\left\langle \frac{1}{N_{ind}} \sum_{i=1}^{N_{ind}} \frac{\sum_{v=1}^{N_i} X_{i,v}}{N_i} \right\rangle \neq \left\langle \frac{1}{N_{ind}} \sum_{i=1}^{N_{ind}} \frac{\sum_{v=1}^{N_i} X_{i,v}}{\bar{N}_{bin}} \right\rangle = \underbrace{\left\langle \frac{\sum_{i,v} X_{i,v}}{N_{ind} \bar{N}_{bin}} \right\rangle}_{=N} = \langle X \rangle \quad (3)$$

Here, $\bar{N}_{bin} = \frac{1}{N_{ind}} \sum_{i=1}^{N_{ind}} N_i$ is the average number of measurements per bin. For the same reason, taking only the first sample of each bin does not give correct results. For a multimodal $p_1(X)$ with a different height (and/or width) of the peaks as in Fig. 1, the Markov Chain may visit the smaller peak very often, but it will stay at the larger one longer.

The independent samples provide a way to analyze correlations and to calculate reliable error estimates [1]. When calculating the variance of the estimate \bar{X} , the new labels i and v become useful as it is now important to distinguish between correlated and uncorrelated samples:

$$\begin{aligned} \langle \bar{X}^2 \rangle &= \frac{1}{N^2} \sum_{i,j} \left\langle \sum_{v=1}^{N_i} \sum_{\mu=1}^{N_j} X_{i,v} X_{j,\mu} \right\rangle = \\ &= \frac{1}{N^2} \sum_{i,j} \left\langle \sum_{v=1}^{N_i} \sum_{\mu=1}^{N_j} \Delta X_{i,v} \Delta X_{j,\mu} \right\rangle + \frac{2\langle X \rangle}{N^2} \underbrace{\left\langle \sum_j N_j \sum_{i,v} \Delta X_{i,v} \right\rangle}_{=N} + \frac{\langle X \rangle^2}{N^2} \underbrace{\left\langle \sum_{i,j} N_i N_j \right\rangle}_{=N^2} = \\ &= \frac{1}{N^2} \sum_i \left\langle \sum_{v,\mu=1}^{N_i} \Delta X_{i,v} \Delta X_{i,\mu} \right\rangle + \frac{1}{N^2} \sum_{i \neq j} \underbrace{\left\langle \sum_{v=1}^{N_i} \Delta X_{i,v} \right\rangle}_{=0} \underbrace{\left\langle \sum_{\mu=1}^{N_j} \Delta X_{j,\mu} \right\rangle}_{=0} + \frac{2\langle X \rangle}{N} \underbrace{\left\langle \sum_{i,v} \Delta X_{i,v} \right\rangle}_{=0} + \langle X \rangle^2 \end{aligned} \quad (4)$$

where $\langle \sum_{v,\mu} \Delta X_{i,v} \Delta X_{j,\mu} \rangle = \langle \sum_v \Delta X_{i,v} \rangle \langle \sum_\mu \Delta X_{j,\mu} \rangle$ for $i \neq j$, because the measurements are from different bins, $\langle \sum_{v,\mu=1}^{N_i} \Delta X_{i,v} \Delta X_{i,\mu} \rangle$ is independent of i , because all bins are equivalent. From Eq. 4, it follows for the variance

$$\text{var}(\bar{X}) = \langle \bar{X}^2 \rangle - \langle \bar{X} \rangle^2 = \langle \bar{X}^2 \rangle - \langle X \rangle^2 = \frac{N_{ind}}{N^2} \left\langle \sum_{v,\mu=1}^{N_i} \Delta X_{i,v} \Delta X_{i,\mu} \right\rangle. \quad (5)$$

The unknown expectation value $\langle \sum_{v,\mu=1}^{N_i} \Delta X_{i,v} \Delta X_{i,\mu} \rangle$ is estimated from the Monte Carlo run, thus $\langle \sum_{v,\mu=1}^{N_i} \Delta X_{i,v} \Delta X_{i,\mu} \rangle_{\text{est}} \approx \frac{1}{N_{ind}} \sum_{i=1}^{N_{ind}} \sum_{v,\mu=1}^{N_i} \Delta X_{i,v} \Delta X_{i,\mu}$. However, the variance depends on the determination of the above expectation value, so it can only be correct, if all modes of p_1 have been sampled sufficiently. Similar formulae can be derived for the expectation values and error estimates of more complex observables (e.g. of the covariance), where correlations between the measured parameters can thus be taken into account.




FIGURE 2. Autocorrelation time for the Simulated Tempering algorithm in 1D depending on the distance between the two peaks for two to five β -slices. The distance is measured in multiples of the width σ of the Gaussian.

BEHAVIOR IN ONE DIMENSION

Although nobody would think of using Monte Carlo simulation for one dimensional problems, as much more efficient approaches are available, it is interesting to examine the Markov matrix for a Simulated Tempering simulation in the two-dimensional X - β -space with discretized X . The probability density $p_1(x)$ for $\beta = 1$ was chosen to consist of two Gaussians well separated from each other and $p_0(x)$ was chosen to be constant. For Simulated Tempering, the number of β -slices was varied from two (just $\beta = 0$ with $p(X|\beta = 0) = p_0$ and $\beta = 1$ with $p(X|\beta = 1) = p_1$) to five. The intermediate β -values were chosen so as to give approximately the same transition rate between all pairs of adjacent β -values. Autocorrelation and thermalization are largely determined by the second largest eigenvalue (e_2) of the Markov matrix, i. e. the one with magnitude closest to one. The autocorrelation time was approximately calculated as $\tau_{AC} \approx 1/(1 - |e_2|)$.

Fig. 2 shows this autocorrelation time as a function of the distance of the two peaks. One sees that more β -slices become necessary as the distance increases. For plain Markov chain Monte Carlo in the one-dimensional discrete X -space, the autocorrelation time far exceeded the range plotted in Fig. 2 even for a distance of $d = 12$ ($\tau_{AC} \approx 2.6499e + 03$) and its calculation is numerically unstable for larger distances.

The columns of the Tempering Markov matrix which correspond to $\beta = 0$ are identical, which means just that whenever the current state of the chain is at $\beta = 0$, the outcome of the next move will not depend on the current position in X -space.

PARALLEL TEMPERING

Another method similar to Simulated Tempering Parallel Tempering, also called Exchange Monte Carlo, see Refs. [3, 7]. In this method, we have M copies of X at the M values for β . Instead of the space $\{X, \beta_m\}$ as in Simulated Tempering, we now consider the product space $\{X_0, X_1, \dots, X_m, \dots, X_M\}$ where the configuration X_m is at the temperature β_m . At every β_m , there is *exactly one* configuration X , denoted by X_m and obeying the distribution $p(X|\beta_m)$. The probability for the total product space is given by the product of the individual probabilities. We now do Markov chain Monte Carlo again with this product probability. In X -space, Metropolis Monte Carlo updates are performed for all β s independently. New configurations X'_m are obtained with the usual Metropolis random walk for $\beta > 0$, while a new sample is drawn directly from $p_0(X)$ for $\beta = 0$. Alternated with the updates in X -space, Metropolis moves to swap configurations X_m and X_{m+1} at adjacent β -values are performed.

During the Monte Carlo run, all $X - m$ will eventually get swapped to $\beta = 0$, where a new sample is drawn. This time, however, the random walk does not completely forget its past, which can be inferred from the Markov matrix for a similar toy situation as for Simulated Tempering above. Suppose, we have three β -values $\beta_0 = 0$, β_1 , $\beta_2 = 1$, and the following temperature swaps occur in the Markov chain Monte Carlo:

$$\begin{array}{ccccccc} 2 & 2 & \tilde{0} & \tilde{0} & \tilde{0} & & \\ 1 & \rightarrow & \tilde{0} & \rightarrow & 2 & \rightarrow & \tilde{1} & \rightarrow & \tilde{1} & , \\ 0 & & 1 & & \tilde{1} & & 2 & & \tilde{2} & \end{array}$$

where a tilde means that an exact sample is drawn from p_0 . All Configurations have now been at $\beta = 0$, but the columns of the matrix corresponding to the above sequence of swaps are still not equal, which means that the current state of the Markov chain still depends on its initial state. However, these correlations are small after an initial thermalization and autocorrelation times are short.

NEEDED PARAMETERS

In order to do Simulated or Parallel Tempering, we have to adjust the values for the β_m and the Z_m , see Eq.(1). The β_m -values have to be dense enough to give a considerable overlap of $p(X|\beta_m)$ and $p(X|\beta_{m\pm 1})$. On the other hand, we want to have as few β -values as possible between $\beta = 1$ and $\beta = 0$. The β -values can be adjusted in a Parallel Tempering prerun, where a new value is inserted whenever the swapping rate between adjacent β s is too low.

The ideal Z_m needed for Simulated Tempering would make all β -values equally likely. This prevents the Markov chain from spending too much time at on single temperature and thus speeds travel from $\beta = 0$ to $\beta = 1$ and back again. This leads to:

$$Z_m \propto \int_X dX p_1(X)^{\beta_m} p_0(X)^{1-\beta_m}$$

For physical systems, the weight $Z(\beta)$ gives the partition function, which can only be determined in terms of $Z(\beta = 0)$. For problems in data analysis, it is the model evidence, i.e. the probability for the chosen model integrated over all possible parameter values. The weights can be obtained from the visiting frequency for the β -values in Simulated Tempering preruns, but this is rather difficult, because they may differ by orders of magnitude. They are not needed for Parallel Tempering, where they cancel out, but the integral can still be calculated with a procedure similar to thermodynamical integration, see Ref. [6]. With the random samples produced at β_m , we can estimate Z_{m+1} for β_{m+1} :

$$\frac{Z_{m+1}}{Z_m} = \left\langle \frac{p_1(X)^{\beta_{m+1}} p_0(X)^{1-\beta_{m+1}}}{p_1(X)^{\beta_m} p_0(X)^{1-\beta_m}} \right\rangle_{\beta_m}, \quad (6)$$

where $\langle \dots \rangle_{\beta_m}$ denotes an expectation value calculated at β_m . The integral $Z(\beta = 1)$ is the product of all the measured ratios:

$$Z(\beta = 1) = Z_M = Z_0 \cdot \prod_{m=0}^{M-1} \frac{Z_{m+1}}{Z_m}. \quad (7)$$

Care must be taken in evaluating this quantity, because the configurations are interchanged between β -values and the measurements obtained for the different β -values are therefore heavily correlated and the same applies to using parallel tempering data for multihistogramming (Ref. [11]), as was proposed in Ref. [12].

BEHAVIOR IN HIGHER DIMENSIONS

In this section, we examine the behavior of the Tempering algorithm in higher dimensions. We chose p_0 as one single broad Gaussian with width $\sigma_0 = 1$ centered at $X = 0$ and the wanted probability p_1 consisted of two Gaussians of width $\sigma = 0.04$ centered at $X = (0.3, 0.3, \dots)$ and $X = (0.8, 0.8, \dots)$, which were multiplied by 5000, so as to yield a norm $n = 10000$. 100 sweeps were performed between β -moves, the β_m and Z_m were adjusted in a parallel tempering prerun. The geometric mean was used to insert new slices, except for finding the second lowest $\beta_1 > 0$, where the old value was just divided by a constant, if the swapping rate was too low. As the number of needed β values N_β depends on the logarithm of the ratio of the volumes of p_1 and p_0 $N_\beta \propto -\log((\sigma_1/\sigma_0)^D)$, the dependence on the dimension D is expected to be linear, which is indeed approximately the case, as can be seen in Fig. 3.

Figure 4 shows the number of MC updates needed for one independent sample. One sees that the increase in needed samples with the dimension of the problem approximately obeys a power law in contrast to the behavior found for the CFTP algorithm (Ref. [9]), which has an exponential dependence on the dimension and generally similar performance as the rejection method, see Ref. [13]. For all presented dimensions, the results for the norm were consistent with the errorbars (see Fig. 5) and likewise the average for X , i. e. the simulation found both peaks.




FIGURE 3. Number N_β of needed β -values needed for various dimensions D .




FIGURE 4. Number of needed MC updates per independent sample for various dimensions D .

APPLICATION TO THE ISING MODEL

Although other choices could be more promising [14], we chose constant transition (Simulated Tempering) and swapping rates (Parallel Tempering, see Ref. [3]) between




FIGURE 5. Logarithm of the obtained norm $\log(Z(\beta = 1))$ for various dimensions D .

adjacent temperatures. This leads to:

$$\frac{1}{(\Delta\beta)^2} \propto \frac{d}{d\beta} \langle E \rangle_\beta \propto \langle E^2 \rangle_\beta - \langle E \rangle_\beta^2 = \frac{C_v}{k_B \beta^2}, \quad (8)$$

where E is the energy, C_v the specific heat and k_B the Boltzmann constant. This relation shows that we need denser β -values where C_v is large, i. e. near a (second order) phase transition. As the specific heat is small for low temperatures again, further cooling below the phase transition is easy. The specific heat is an extensive quantity, the number of needed temperatures is therefore expected to grow as $N_\beta \propto \sqrt{N_{spins}}$ with N_{spins} the number of spins. In this case, we used the arithmetic mean to insert new β -values. As expected, the number of needed temperatures scales proportional to the linear system size $N_\beta \propto L = \sqrt{N_{spins}}$, see fig. 6.




FIGURE 6. Number N_β of needed β -values needed for the two dimensional Ising model depending on the linear system size $L = \sqrt{N_{spins}}$.




FIGURE 7. Time per independent sample for the two dimensional Ising model depending on the linear system size $L = \sqrt{N_{spins}}$. Solid lines: CFTP Method and single spin flip algorithm Dotted lines: Simulated Tempering and Swendsen Wang algorithm

Figure 7 shows the time per independent sample for Exact Tempering and for the Coupling From The Past (CFTP) method with single spin flips introduced by Propp and Wilson in Ref. [8]. For CFTP, the time needed for an independent sample grows on a logarithmic scale with the system sizes for temperatures above the critical T_C , obeys a

power law at T_C and grows exponentially for the ordered phase below T_C . There are CFTP schemes for the Swendsen Wang algorithm, but its runtime scales exponentially with the system size except at $\beta = 0, T = \infty$ and it is *much* slower than the single spin flip algorithm around T_C , see Ref. [13].

Exact Tempering with the Swendsen Wang algorithm, on the other hand, gives linear scaling for all temperatures. Although this is slower than CFTP for high temperatures, one does in fact get the high temperature results for free, because they are sampled anyway in a Tempering run for low temperatures. The critical exponent for Exact Tempering is two for both the Swendsen Wang (1.92 ± 0.06 at the lowest temperature) and the Wolff (1.98 ± 0.07) algorithm. The reason for this is, that the time for an independent sample is determined by the number of steps needed to go from $\beta = 0$ to $\beta = \beta_{max}$ and back again and not by the algorithm used for the spin updates. This random walk in the temperatures scales proportional to the square of their number N_β , which gives

$$\tau_{ind} \propto N_\beta^2 \propto L^2 = N_{spins} . \quad (9)$$

This dependence $\tau_{ind} \propto N_{spins}$ breaks down for the single spin flip algorithm. Exact Tempering then becomes much slower, because the spin configurations at the critical temperature cannot be sampled as efficiently. This effect becomes more severe for first order transitions, where the algorithm does not manage the transition from the disordered to the ordered phase at T_C . ‘Tempering’ of a model parameter which carries the transition from first to second order might then be a solution. This was introduced in Ref. [15] for the Swendsen-Wang algorithm applied to the Potts model, where the variable ‘tempering’ parameter was the number of states q . Exact Tempering is also applicable to this variant, because the percolation problem for $q = 1$ can be sampled exactly.

CONCLUSIONS

This paper provides a discussion of Exact Sampling with Simulated Tempering [1] and compares it to Exact Sampling via the Propp-Wilson method [8]. The former is found to be advantageous in most cases. Simulated Tempering provides a way to draw exact, i.e. completely uncorrelated samples from arbitrary distributions in high dimensions. The peaks of multimodal densities are sampled with their respective weights. The parameters β_m and Z_m needed for the Simulated Tempering run can be adjusted in a Parallel Tempering prerun. While the Parallel Tempering algorithm itself does not provide perfectly uncorrelated samples, its autocorrelation time is small. For practical purposes, it is a robust alternative, because it does not need the parameters Z_m . Both methods allow to calculate the integral over the probability density, i.e. partition functions or model evidences.

ACKNOWLEDGMENT

This work has been supported by the Austrian Science Fund (FWF), project no. P15834-PHY. This document has been typeset using the $\mathcal{A}\mathcal{M}\mathcal{S}$ - $\mathcal{L}\mathcal{A}\mathcal{T}\mathcal{E}\mathcal{X}$ packages.

REFERENCES

1. C. J. Geyer, and E. A. Thompson, *J. Amer. Statist. Assoc.*, **90**, 909 (1995).
2. E. Marinari, and G. Parisi, *Europhys. Lett.*, **19**, 451–458 (1992).
3. K. Hukushima, and K. Nemoto, *J. Phys. Soc. Japan*, **65**, 1604 (1996).
4. K. Hukushima, H. Takayama, and K. Nemoto, *Int. J. Mod. Phys. C*, **7**, 337 (1996).
5. W. Kerler, and P. Rehberg, *Phys. Rev. E*, **50**, 4220–4225 (1994).
6. K. Pinn, and C. Wierczkowski, *Int. J. Mod. Phys. C*, **9**, 541 (1998).
7. E. Marinari, “Optimized Monte Carlo Methods,” in *Advances in Computer Simulation*, edited by J. Kertesz and I. Kondor, Springer, Berlin, 1997, URL `cond-mat/9612010`.
8. J.G. Propp, and D.B. Wilson, *Random Structures and Algorithms*, **9**, 223–252 (1996).
9. Andrew M. Childs, Ryan B. Patterson, and David J. C. MacKay, *Phys. Rev. E*, **63** (2001).
10. Huber, M., Efficient exact sampling from the Ising model using Swendsen-Wang, Research Report, University of Cornell, to be published (1998).
11. A. M. Ferrenberg, and R. H. Swendsen, *Phys. Rev. Lett.*, **63**, 1195–1198 (1989).
12. A. Mitsutake, and Y. Okamoto, *Chem. Phys. Lett.*, **332**, 131–138 (2000).
13. Konegger, M., *diploma thesis, TU Graz* (2000).
14. Helmut G. Katzgraber, Simon Trebst, David A Huse, and Matthias Troyer, *J. Stat. Mech.*, p. P03018 (2006).
15. W. Kerler, and A. Weber, *Phys. Rev. B*, **47**, 11563–11566 (1993).