

Ensemble Optimization Techniques for the Simulation of Slowly Equilibrating Systems

S. Trebst¹, D.A. Huse², E. Gull³, H.G. Katzgraber³, U.H.E. Hansmann^{4,5},
and M. Troyer³

¹ Microsoft Research and Kavli Institute for Theoretical Physics,
University of California, Santa Barbara, CA 93106, USA

² Department of Physics, Princeton University, Princeton, NJ 08544, USA

³ Theoretische Physik, Eidgenössische Technische Hochschule Zürich,
CH-8093 Zürich, Switzerland

⁴ Department of Physics, Michigan Technological University,
Houghton, MI 49931, USA

⁵ John-von-Neumann Institute for Computing, Forschungszentrum Jülich,
D-52425 Jülich, Germany

Abstract. Competing phases or interactions in complex many-particle systems can result in free energy barriers that strongly suppress thermal equilibration. Here we discuss how extended ensemble Monte Carlo simulations can be used to study the equilibrium behavior of such systems. Special focus will be given to a recently developed adaptive Monte Carlo technique that is capable to explore and overcome the entropic barriers which cause the slow-down. We discuss this technique in the context of broad-histogram Monte Carlo algorithms as well as its application to replica-exchange methods such as parallel tempering. We briefly discuss a number of examples including low-temperature states of magnetic systems with competing interactions and dense liquids.

1 Introduction

The free energy landscapes of complex many-body systems with competing phases or interactions are often characterized by many local minima that are separated by entropic barriers. The simulation of such systems with conventional Monte Carlo [1] or molecular dynamics [2] methods is slowed down by long relaxation times due to the suppressed tunneling through these barriers. While at second order phase transitions this slow-down can be overcome by improved updating techniques, such as cluster updates [3, 4], this is not the case for systems which undergo a first-order phase transition or for systems that exhibit frustration or disorder. For these systems one instead aims at improving the way that relatively simple, local updates are accepted or rejected in the sampling process by introducing artificial statistical ensembles such that tunneling times through barriers are reduced and autocorrelation effects minimized. In the following we discuss recently developed techniques to find statistical ensembles that optimize the performance of Monte Carlo sampling,

first in the context of broad-histogram Monte Carlo algorithms and then outline how these methods can be applied in the context of replica-exchange or parallel-tempering algorithms.

2 Extended Ensemble Methods

Let us consider a first-order phase transition, such as in a two-dimensional q -state Potts model [5] with a Hamiltonian

$$H = -J \sum_{\langle i,j \rangle} \delta_{\sigma_i \sigma_j}, \quad (1)$$

where the spins σ_i can take the integer values $1, \dots, q$. For $q > 4$ this model exhibits a first-order phase transition, accompanied by exponential slowing down of conventional local-update algorithms. The exponential slow-down is caused by the free-energy barrier between the two coexisting meta-stable states at the first-order phase transition.

This barrier can be quantified by considering the energy histogram

$$H_{\text{canonical}}(E) \propto g(E) P_{\text{Boltzmann}}(E) = g(E) \exp(-\beta E), \quad (2)$$

which is the probability of encountering a configuration with energy E during the Monte Carlo simulation. The density of states is given by

$$g(E) = \sum_c \delta_{E, E(c)}, \quad (3)$$

where the sum runs over all configurations c . Away from first-order phase transitions, $H_{\text{canonical}}(E)$ has approximately Gaussian shape, centered around the mean energy. At first-order phase transitions, where the energy jumps discontinuously, the histogram $H_{\text{canonical}}(E)$ develops a double-peak structure. The minimum of $H_{\text{canonical}}(E)$ between these two peaks, which the simulation has to cross in order to go from one phase to the other, becomes exponentially small upon increasing the system size. This leads to exponentially large tunneling and autocorrelation times.

This tunneling problem at first-order phase transitions can be alleviated by extended ensemble techniques which aim at broadening the sampled energy space. Instead of weighting a configuration c with energy $E = E(c)$ using the Boltzmann weight $P_{\text{Boltzmann}}(E) = \exp(-\beta E)$ more general weights $P_{\text{extended}}(E)$ are introduced which define the extended ensemble. The configuration space is explored by generating a Markov chain of configurations

$$c_1 \rightarrow c_2 \rightarrow \dots \rightarrow c_i \rightarrow c_{i+1} \rightarrow \dots, \quad (4)$$

where a move from configuration c_1 to c_2 is accepted with probability

$$P_{\text{acc}}(c_1 \rightarrow c_2) = \min \left[1, \frac{P(c_2)}{P(c_1)} \right] = \min \left[1, \frac{W_{\text{extended}}(E_2)}{W_{\text{extended}}(E_1)} \right]. \quad (5)$$

In general, the extended weights are defined in a single coordinate, such as the energy, thereby projecting the random walk in configuration space to a random walk in energy space

$$E_1 = E(c_1) \rightarrow E_2 \rightarrow \dots \rightarrow E_i \rightarrow E_{i+1} \rightarrow \dots \quad (6)$$

For this random walk in energy space a histogram can be recorded which has the characteristic form

$$H_{\text{extended}}(E) \propto g(E)W_{\text{extended}}(E), \quad (7)$$

where the density of states $g(E)$ is fixed for the simulated system.

One choice of generalized weights is the multicanonical ensemble [6] where the weight of a configuration c is defined as $W_{\text{multicanonical}}(c) \propto 1/g(E(c))$. The multicanonical ensemble then leads to a flat histogram in energy space

$$H_{\text{multicanonical}}(E) \propto g(E)W_{\text{multicanonical}}(E) = g(E)\frac{1}{g(E)} = \text{const.} \quad (8)$$

removing the exponentially small minimum in the canonical distribution. After performing a simulation, measurements in the multicanonical ensemble are reweighted by a factor $W_{\text{Boltzmann}}(E)/W_{\text{multicanonical}}(E)$ to obtain averages in the canonical ensemble.

Since the density of states and thus the multicanonical weights are not known initially, a scalable algorithm to estimate these quantities is needed. The Wang-Landau algorithm [7] is a simple but efficient iterative method to obtain good approximates of the density of states $g(E)$ and the multicanonical weights $W_{\text{multicanonical}}(E) \propto 1/g(E)$. Besides overcoming the exponentially suppressed tunneling problem at first-order phase transitions, the Wang-Landau algorithm calculates the generalized density of states $g(E)$ in an iterative procedure. The knowledge of the density of states $g(E)$ then allows the direct calculation of the free energy from the partition function, $Z = \sum_E g(E)e^{-\beta E}$. The internal energy, entropy, specific heat and other thermal properties are easily obtained as well, by differentiating the free energy. By additionally measuring the averages $A(E)$ of other observables A as a function of the energy E , thermal expectation values can be obtained at arbitrary inverse temperatures β by performing just a single simulation:

$$\langle A(\beta) \rangle = \frac{\sum_E A(E)g(E)e^{-\beta E}}{\sum_E g(E)e^{-\beta E}}. \quad (9)$$

3 Markov Chains and Random Walks in Energy Space

The multicanonical ensemble and Wang-Landau algorithm both project a random walk in high-dimensional configuration space to a one-dimensional

random walk in energy space where all energy levels are sampled equally often. It is important to note that the random walk in configuration space, (4), is a biased Markovian random walk, while the projected random walk in energy space, (6), is non-Markovian, as memory is stored in the configuration. This becomes evident as the system approaches a phase transition in the random walk: While the energy no longer reflects from which side the phase transition is approached, the current configuration may still reflect the actual phase the system has visited most recently. In the case of the ferromagnetic Ising model, the order parameter for a given configuration at the critical energy $E_c \sim -1.41N$ (in two space dimensions) will reveal whether the system is approaching the transition from the magnetically ordered (lower energies) or disordered side (higher energies).

This loss of information in the projection of the random walk in configuration space has important consequences for the random walk in energy space. Most strikingly, the local diffusivity of a random walker in energy space, which for a diffusion time t_D is defined as

$$D(E, t_D) = \langle (E(t) - E(t + t_D))^2 \rangle / t_D \quad (10)$$

is *not* independent of the location in energy space. This is illustrated in Fig. 1 for the two-dimensional Ising ferromagnet. Below the phase transition around $E_c \sim -1.41N$ a clear minimum evolves in the local diffusivity. In this region large ordered domains are formed and by moving the domain boundaries through local spin flips only small energy changes are induced resulting in a suppressed local diffusivity in energy space.

Because of the strong energy dependence of the local diffusivity the simulation of a multicanonical ensemble sampling all energy levels equally often turns out to be suboptimal [8]. The performance of flat-histogram algorithms can be quantified for classical spin models such as the ferromagnet where the number of energy levels is given by $[-2N, +2N]$ and thereby scales with the number of spins N in the system. When measuring the typical round-trip time between the two extremal energies for multicanonical simulations, these round-trip times τ are found to scale like

$$\tau \sim N^2 L^z, \quad (11)$$

showing a power-law deviation from the N^2 -scaling behavior of a completely unbiased random walk. Here L is the linear system size and z a critical exponent describing the slow-down of a multicanonical simulation in the proximity of a phase transition [8, 9]. The value of z strongly depends on the simulated model and the dimensionality of the problem. In two dimensions the exponent increases from $z = 0.74$ for the ferromagnet as one introduces competing interactions leading to frustration and disorder. The exponent becomes $z = 1.73$ for the two-dimensional fully frustrated Ising model which is defined by a Hamiltonian

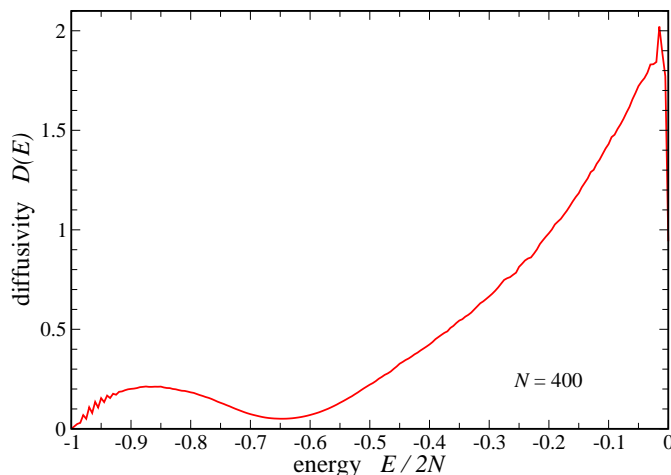


Fig. 1. Local diffusivity $D(E)$ of a random walk sampling a flat histogram in energy space for the two-dimensional ferromagnetic Ising model with $N = 20 \times 20$ spins. The local diffusivity strongly depends on the energy with a strong suppression around the critical energy $E_c \approx -1.41N$ and the ground-state energy $E_0 = -2N$

$$H = \sum_{\langle i,j \rangle} J_{ij} \sigma_i \sigma_j , \quad (12)$$

where the spins around any given plaquette of four spins are frustrated, e.g. by choosing the couplings along three bonds to be $J_{ij} = -1$ (ferromagnetic) and $J_{ij} = +1$ (antiferromagnetic) for the remaining bond. For the spin glass, where the couplings J_{ij} are randomly chosen to be $+1$ or -1 , exponential scaling ($z = \infty$) is found [8, 10]. Increasing the spatial dimension of the ferromagnet the exponent is found to decrease as $z \approx 1.81, 0.74$ and 0.44 for dimension $d = 1, 2$ and 3 and z vanishes for the mean-field model in the limit of infinite dimensions [9].

4 Optimized Ensembles

The observed polynomial slow-down of the multicanonical ensemble poses the question whether for a given model there is an optimal choice of the histogram $H_{\text{optimal}}(E)$ and corresponding weights $W_{\text{optimal}}(E)$, which eliminates the slow-down. To address this question an adaptive feedback algorithm has recently been introduced that iteratively improves the weights in an extended ensemble simulation leading to further improvements in the efficiency of the algorithm by several orders of magnitude [11]. The scaling for the optimized ensemble is found to scale like $O([N \ln N]^2)$ thereby reproducing the behavior of an unbiased Markovian random walk up to a logarithmic correction.

At the heart of the algorithm lies the idea to maximize a current j of walkers that move from the lowest energy level, E_- , to the highest energy level, E_+ , or vice versa, in an extended ensemble simulation by varying the weights $W_{\text{extended}}(E)$. To measure the current a label is added to the walker that indicates which of the two extremal energies the walker has visited most recently. The two extrema act as “reflecting” and “absorbing” boundaries for the labeled walker: e.g., if the label is plus, a visit to E_+ does not change the label, so this is a “reflecting” boundary. However, a visit to E_- does change the label, so the plus walker is absorbed at that boundary. The behavior of the labeled walker is *not* affected by its label except when it visits one of the extrema and the label changes.

For the random walk in energy space, two histograms are recorded, $H_+(E)$ and $H_-(E)$, which for sufficiently long simulations converge to steady-state distributions which satisfy $H_+(E) + H_-(E) = H(E) = W(E)g(E)$. For each energy level the fraction of random walkers which have label “plus” is then given by $f(E) = H_+(E)/H(E)$. The above-discussed boundary conditions dictate $f(E_-) = 0$ and $f(E_+) = 1$.

The steady-state current to first-order in the derivative is

$$j = D(E)H(E)\frac{df}{dE}, \quad (13)$$

where $D(E)$ is the walker’s diffusivity at energy E . There is no current if $f(E)$ is constant, since this corresponds to the equilibrium state. Therefore the current is to leading order proportional to df/dE . Rearranging the above equation and integrating on both sides, noting that j is a constant and f runs from 0 to 1, one obtains

$$\frac{1}{j} = \int_{E_-}^{E_+} \frac{dE}{D(E)H(E)}. \quad (14)$$

To maximize the current and thus the round-trip rate, this integral must be minimized. However, there is a constraint: $H(E)$ is a probability distribution and must remain normalized which can be enforced with a Lagrange multiplier:

$$\int_{E_-}^{E_+} dE \left(\frac{1}{D(E)H(E)} + \lambda H(E) \right). \quad (15)$$

To minimize this integrand, the ensemble, that is the weights $W(E)$ and thus the histogram $H(E)$ are varied. At this point it is assumed that the dependence of $D(E)$ on the weights can be neglected.

The optimal histogram, $H_{\text{optimal}}(E)$, which minimizes the above integrand and thereby maximizes the current j is then found to be

$$H_{\text{optimal}}(E) \propto \frac{1}{\sqrt{D(E)}}. \quad (16)$$

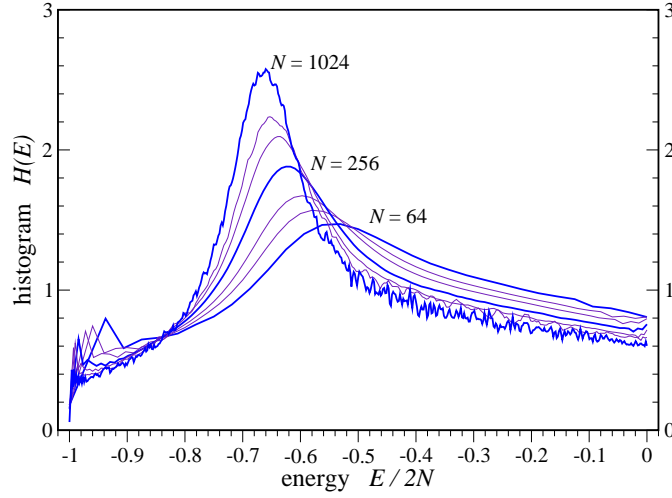


Fig. 2. Optimized histograms for the two-dimensional ferromagnetic Ising model for various system sizes. After the feedback of the local diffusivity a peak evolves near the critical energy of the transition $E_c \approx -1.41N$. The feedback thereby shifts additional resources towards the bottleneck of the simulation which are identified by a suppressed local diffusivity

Thus for the optimal ensemble, the probability distribution of sampled energy levels is simply inversely proportional to the square root of the local diffusivity.

The optimal histogram can be approximated in a feedback loop of the form

- Start with some trial weights $W(E)$, e.g. $W(E) = 1/g(E)$.
- Repeat
 - Reset the histograms $H(E) = H_+(E) = H_-(E) = 0$.
 - Simulate the system with the current weights for N sweeps:
 - Updates are accepted with probability $\min[1, W(E')/W(E)]$.
 - Record the histograms $H_+(E)$ and $H_-(E)$.
 - From the recorded histogram an estimate of the local diffusivity is obtained as

$$D(E) \propto \frac{1}{H(E) \frac{df}{dE}},$$

where the fraction $f(E)$ is given by $f(E) = H_+(E)/H(E)$ and $H(E)$ is the histogram $H(E) = H_+(E) + H_-(E)$.

- Define new weights as

$$W_{\text{optimized}}(E) = W(E) \sqrt{\frac{1}{H(E)} \cdot \frac{df}{dE}}.$$

- Increase the number of sweeps for the next iteration

$$N_{\text{sweeps}} \leftarrow 2N_{\text{sweeps}}.$$

- Stop once the histogram $H(E)$ has converged.

The implementation of this feedback algorithm requires to change only a few lines of code in the original local-update algorithm for the Ising model. Some additional remarks are useful:

1. In contrast to the Wang-Landau algorithm, the weights $W(E)$ are modified only after a batch of N_{sweeps} sweeps, thereby ensuring detailed balance between successive moves at all times.
2. The initial value of sweeps N_{sweeps} should be chosen large enough that a couple of round trips are recorded, thereby ensuring that steady state data for $H_+(E)$ and $H_-(E)$ are measured.
3. The derivative df/dE can be determined by a linear regression, where the number of regression points is flexible. Initial batches with the limited statistics of only a few round trips may require a larger number of regression points than subsequent batches with smaller round-trip times and better statistics.
4. Similar to the multicanonical ensemble, the weights $W(E)$ can become very large and storing the logarithms may be advantageous. The reweighting then becomes $\ln W_{\text{optimized}}(E) = \ln W(E) + [\ln \frac{df}{dE} - \ln H(E)]/2$.

At the end of the simulation, the density of states can be estimated from the recorded histogram as $g(E) = H_{\text{optimized}}(E)/W_{\text{optimized}}(E)$ and normalized as described above.

Figure 2 shows the optimized histogram for the two-dimensional ferromagnetic Ising model. The optimized histogram is no longer flat, but a peak evolves at the critical region around $E_c \approx -1.41N$ of the transition. The feedback of the local diffusivity reallocates resources towards the bottlenecks of the simulation which have been identified by a suppressed local diffusivity.

When analyzing the scaling of round-trip times for the optimized ensemble one finds a considerable speedup: The power-law slow-down of round-trip times for the flat-histogram ensemble $O(N^2 L^z)$ is reduced to $O([N \ln N]^2)$ for the optimized ensemble, e.g. there is only a logarithmic correction to the scaling of a completely unbiased random walk with $O(N^2)$ -scaling. For the two-dimensional fully frustrated Ising model the scaling of round-trip times is shown in Fig. 3. This scaling improvement results in a speedup by a nearly two orders of magnitude already for a system with some 128×128 spins.

5 Simulation of Dense Fluids

Extended ensembles cannot only be defined as a function of energy, but in arbitrary reaction coordinates \mathbf{R} onto which a random walk in configuration

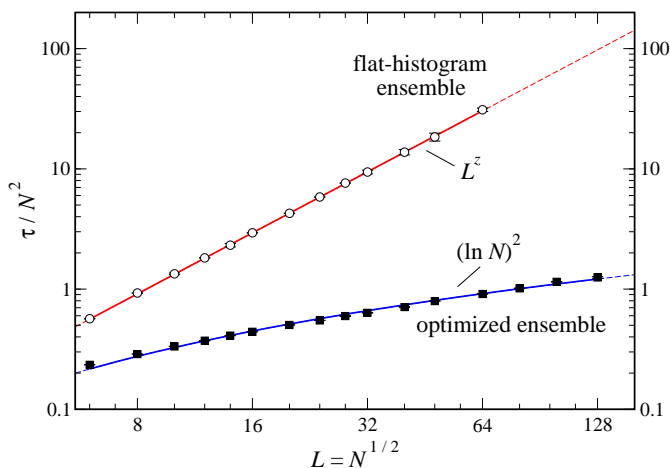


Fig. 3. Scaling of round-trip times for a random walk in energy space sampling a flat histogram (*open squares*) and the optimized histogram (*solid circles*) for the two-dimensional fully frustrated Ising model. While for the multicanonical simulation a power-law slow-down of the round-trip times $O(N^2 L^z)$ is observed, the round-trip times for the optimized ensemble scale like $O([N \ln N]^2)$ thereby approaching the ideal $O(N^2)$ -scaling of an unbiased Markovian random walk up to a logarithmic correction

space can be projected. The generalized weights in these reaction coordinates $W_{\text{extended}}(\mathbf{R})$ are then used to bias the random walk along the reaction coordinate by accepting moves from a configuration c_1 with reaction coordinate \mathbf{R}_1 to a configuration c_2 with reaction coordinate \mathbf{R}_2 with probability

$$p_{\text{acc}}(c_1 \rightarrow c_2) = p_{\text{acc}}(\mathbf{R}_1 \rightarrow \mathbf{R}_2) = \min \left(1, \frac{W_{\text{extended}}(\mathbf{R}_2)}{W_{\text{extended}}(\mathbf{R}_1)} \right). \quad (17)$$

The generalized weights $W_{\text{extended}}(\mathbf{R})$ can again be chosen in such a way that similar to a multicanonical simulation a flat histogram is sampled along the reaction coordinate by setting the weights to be inversely proportional to the density of states defined in the reaction coordinates, that is $W_{\text{extended}}(\mathbf{R}) \propto 1/g(\mathbf{R})$.

An optimal choice of weights can be found by measuring the local diffusivity of a random walk along the reaction coordinates and by applying the feedback method to shift weight towards the bottlenecks in the simulation. This generalized ensemble optimization approach has recently been illustrated for the simulation of dense Lennard-Jones fluids close to the vapor-liquid equilibrium [12]. The interaction between particles in the fluid is described by a pairwise Lennard-Jones potential of the form

$$\Phi_{\text{LJ}}(R) = 4\epsilon \left[\left(\frac{\sigma}{R} \right)^{12} - \left(\frac{\sigma}{R} \right)^6 \right], \quad (18)$$

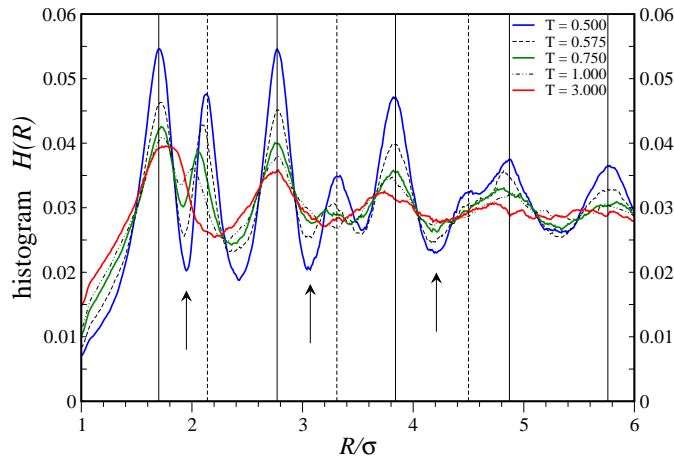


Fig. 4. Optimized histograms for a dense two-dimensional Lennard–Jones fluid after feedback of the local diffusivity in a radial coordinate for varying temperatures. For the optimized ensemble peaks evolve at the free energy barriers between the shells of the liquid which proliferate at lower temperatures (*solid lines*). Additional peaks emerge at low temperatures (*dashed lines*) revealing interstitial states (*arrows*) between the shells of the liquid

where ϵ is the interaction strength, σ a length parameter, and R the distance between two particles. It is this distance R between two arbitrarily chosen particles in the fluid that one can use as a new reaction coordinate for a projected random walk. For a given temperature defining an extended ensemble with weights $W_{\text{extended}}(R)$ and recording a histogram $H(R)$ during a simulation will then allow to calculate the pair distribution function $g(R) = H(R)/W_{\text{extended}}(R)$. The pair distribution function $g(R)$ is closely related to the potential of mean force (PMF)

$$\Phi_{\text{PMF}}(R) = -\frac{1}{\beta} \ln g(R), \quad (19)$$

which describes the average interaction between two particles in the fluid in the presence of many surrounding particles.

For high particle densities and low enough temperatures shell structures form in the fluid which are reminiscent of the hexagonal lattice of the solid structure at very low temperatures. These shell structures are revealed by a sinusoidal modulation in the PMF. Thermal equilibration between the shells is suppressed by entropic barriers which form between the shells. Again, one can ask what probability distribution, or histogram, should be sampled along the reaction coordinate, in this case the radial distance R , so that equilibration between the shells is improved. Measuring the local diffusivity for a random walk along the radial distance R in an interval $[R_{\text{min}}, R_{\text{max}}]$ and subsequently applying the feedback algorithm described above optimized histograms $H(R)$

are found which are plotted in Fig. 4 for varying temperatures [12]. The feedback algorithm again shifts additional weight in the histogram towards the bottleneck of the simulation, in this case towards the barriers between the shells. Interestingly, additional peaks emerge in the optimized histogram as the temperature is lowered towards the vapor-liquid equilibrium. The minima between these peaks point to additional meta-stable configurations which occur at these low temperatures, namely interstitial states which occur as the shells around two particles merge as detailed in [12].

This example illustrates that for some simulations the local diffusivity and optimized histogram *themselves* are very sensitive measures that can reveal interesting underlying physical phenomena which are otherwise hard to detect in a numerical simulation. In general, a strong modulation of the local diffusivity for the random walk along a given reaction coordinate is a good indicator that the reaction coordinate itself is a good choice that captures some interesting physics of the problem.

6 Parallel Tempering / Replica-Exchange Methods

The simulation of frustrated and/or disordered systems suffers from a similar tunneling problem than the simulation of first-order phase transitions: local minima in energy space are separated by barriers that grow with system size. While the multicanonical or optimized ensembles do not help with the NP-hard problems faced by spin glasses, they are efficient in speeding up simulations of frustrated magnets without disorder [11].

An alternative to these extended ensembles for the simulation of frustrated magnets is the “parallel tempering” or “replica-exchange” Monte Carlo method [13–16]. Instead of performing a single simulation at a fixed temperature, simulations are performed for M replicas at a set of temperatures T_1, T_2, \dots, T_M . In addition to standard Monte Carlo updates at a fixed temperature, exchange moves are proposed to swap two replicas between adjacent temperatures. These swaps are accepted with a probability

$$\min[1, \exp(\Delta\beta\Delta E)], \quad (20)$$

where $\Delta\beta = \beta_j - \beta_i$ is the difference in inverse temperatures and $\Delta E = E_j - E_i$ the difference in energy between the two replicas i and j .

The effect of these exchange moves is that a replica can drift from a local free energy minimum at low temperatures to higher temperatures, where it is easier to cross energy barriers and equilibration is fast. Upon cooling (by another sequence of exchanges) it can end up in a different local minimum on time scales that are much shorter compared to a single simulation at a fixed low temperature. This random walk of a single replica in temperature space is the conjugate analog of the random walk in energy space discussed for the extended ensemble techniques. The complement of the statistical ensemble, defined by the weights $W_{\text{extended}}(E)$, is the particular choice

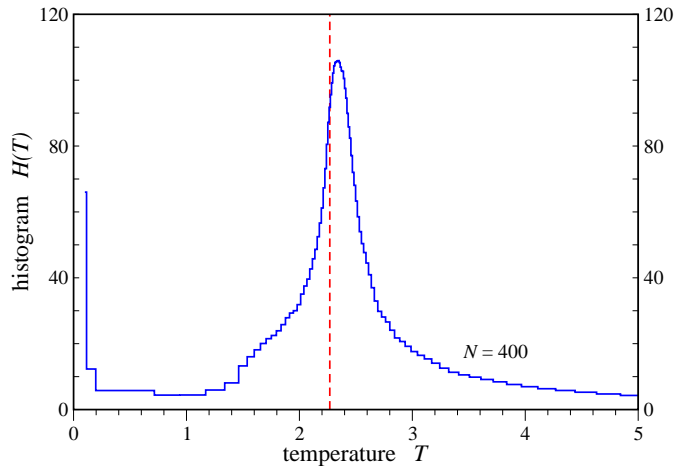


Fig. 5. Optimized temperature distribution $H(T)$ for a two-dimensional Ising ferromagnet with $N = 20 \times 20$ spins and 100 replicas/temperature points. After the feedback of the local diffusivity a peak evolves near the critical temperature of the transition $T_c \approx 2.269$. The feedback thereby shifts additional resources towards the bottleneck of the simulation which are identified by a suppressed local diffusivity. Note the similarity to Fig. 2

of temperature points in the temperature set $\{T_1, T_2, \dots, T_M\}$ for the parallel tempering simulation. The probability of sampling any given temperature T in an interval $T_i < T < T_{i+1}$ can then be approximated by $H(T) \propto 1/\Delta T$, where $\Delta T = T_{i+1} - T_i$ is the length of the temperature interval around the temperature T . This probability distribution $H(T)$ is the equivalent to the histogram $H(E)$ in the extended ensemble simulations. The ensemble optimization technique discussed above can thus be reformulated to optimize the temperature set in a parallel-tempering simulation in such a way that the rate of round trips between the two extremal temperatures, T_1 and T_M respectively, is maized [17, 18].

Starting with an initial temperature set $\{T_1, T_2, \dots, T_M\}$ a parallel tempering simulation is performed where each replica is labeled either “plus” or “minus” indicating which of the two extremal temperatures the respective replica has visited most recently. This allows to measure a current of replicas diffusing from the highest to the lowest temperature by recording two histograms, $h_+(T)$ and $h_-(T)$ for each temperature point. The current j is then given by

$$j = D(T)H(T)\frac{df}{dT}, \quad (21)$$

where $D(T)$ is the local diffusivity for the random walk in temperature space, and $f(T) = h_+(T)/[h_+(T) + h_-(T)]$ is the fraction of random walkers which have visited the highest temperature T_M most recently. The probability dis-

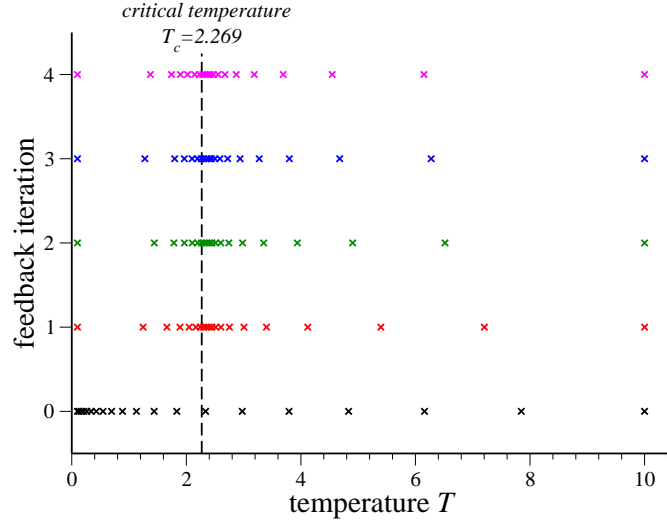


Fig. 6. Optimized temperature sets for a two-dimensional Ising ferromagnet with $N = 20 \times 20$ spins. The initial temperature set with 20 temperature points is determined by a geometric progression [17,19] for the temperature interval $[0.1, 10]$. After feedback of the local diffusivity the temperature points accumulate near the critical temperature $T_c = 2.269$ of the phase transition (*dashed line*). Similar to the ensemble optimization in energy space the feedback of the local diffusivity relocates resources towards the bottleneck of the simulation

tribution $H(T)$ is normalized, that is

$$\int_{T_1}^{T_M} H(T) dT = C \int_{T_1}^{T_M} \frac{dT}{\Delta T} = 1, \quad (22)$$

where C is a normalization constant. Rearranging (21), the local diffusivity $D(T)$ of the random walk in temperature space can be estimated as

$$D(T) \propto \frac{\Delta T}{df/dT}. \quad (23)$$

In analogy to the argument for the extended ensemble in energy space the current j is maximized by choosing a probability distribution

$$H_{\text{optimal}}(T) \propto \frac{1}{\sqrt{D(T)}} \propto \sqrt{\frac{1}{\Delta T} \frac{df}{dT}}, \quad (24)$$

which is inversely proportional to the square root of the local diffusivity. The optimized temperature set $\{T'_1, T'_2, \dots, T'_M\}$ is then found by choosing the n -th temperature point T'_n such that

$$\int_{T'_1}^{T'_n} H_{\text{optimal}}(T) dT = \frac{n}{M}, \quad (25)$$

where M is the number of temperature points in the original temperature set, and the two extremal temperatures $T'_1 = T_1$ and $T'_M = T_M$ remain unchanged. Similarly to the algorithm for the ensemble optimization this feedback of the local diffusivity should be iterated until the temperature set is converged.

Figures 5 and 6 illustrate the so-optimized temperature sets for the Ising ferromagnet obtained by several iterations of the above feedback loop. After the feedback of the local diffusivity, temperature points accumulate near the critical temperature $T_c = 2.269$ of the transition. This is in full analogy to the optimized histograms for the extended ensemble simulations where resources are shifted towards the critical energy of the transition, for comparison see Figs. 2 and 5.

It is interesting to note that for the so-optimized temperature set the acceptance rates for swap moves are not independent of the temperature [17]. Around the critical temperature, where temperature points are accumulated by the feedback algorithm, the acceptance rates are higher than at higher/lower temperatures, where the density of temperature points becomes considerably smaller after feedback [17,18]. The almost Markovian scaling behavior for the optimized random walks in either energy or temperature space is thus generated by a problem-specific statistical ensemble which is characterized neither by a flat histogram nor flat acceptance rates for exchange moves, but by a characteristic probability distribution which concentrates resources at the minima of the measured local diffusivity.

7 Outlook

The ensemble optimization technique reviewed in this chapter should be broadly applicable to a wide range of applications – possibly speeding up existing uniform sampling techniques by orders of magnitude. It has recently been used to improve broad-histogram Monte Carlo techniques [11] as well as parallel-tempering Monte Carlo simulations [17], with applications to frustrated and disordered spin systems [11,17], dense fluids [12], as well as folded proteins [18]. It also holds promise to improve the simulation of quantum systems close to a phase transition when optimizing the extended ensemble introduced for the quantum Wang–Landau algorithm outlined in [20].

References

1. D.P. Landau, K. Binder, *A guide to Monte Carlo Simulations in Statistical Physics*. (Cambridge University Press, 2000)
2. D. Frenkel, B. Smit, *Understanding Molecular Simulation*. (Academic Press, 1996)
3. R.H. Swendsen, J.-S. Wang, Phys. Rev. Lett. **58**, 86 (1987)

4. U. Wolff, Phys. Rev. Lett. **62**, 361 (1989)
5. F.Y. Wu, Rev. Mod. Phys. **54**, 235 (1982)
6. B.A. Berg, T. Neuhaus, Phys. Lett. B **267**, 249 (1991); Phys. Rev. Lett. **68**, 9 (1992)
7. F. Wang, D.P. Landau, Phys. Rev. Lett. **86**, 2050 (2001); Phys. Rev. E **64**, 056101 (2001)
8. P. Dayal, S. Trebst, S. Wessel, D. Würtz, M. Troyer, S. Sabhapandit, S.N. Coppersmith, Phys. Rev. Lett. **92**, 097201 (2004)
9. Y. Wu, M. Körner, L. Colonna-Romano, S. Trebst, H. Gould, J. Machta, M. Troyer, Phys. Rev. E **72**, 046704 (2005)
10. S. Alder, S. Trebst, A.K. Hartmann, M. Troyer, J. Stat. Mech. P07008 (2004)
11. S. Trebst, D.A. Huse, M. Troyer, Phys. Rev. E **70**, 046701 (2004)
12. S. Trebst, E. Gull, M. Troyer, J. Chem. Phys. **123**, 204501 (2005)
13. R.H. Swendsen, J. Wang, Phys. Rev. Lett. **57**, 2607 (1986)
14. E. Marinari, G. Parisi, Europhys. Lett. **19**, 451 (1992)
15. A.P. Lyubartsev, A.A. Martsinovski, S.V. Shevkunov, P.N. Vorontsov-Velyaminov, J. Chem. Phys. **96**, 1776 (1992)
16. K. Hukushima, Y. Nemoto, J. Phys. Soc. Jpn. **65**, 1604 (1996)
17. H.G. Katzgraber, S. Trebst, D.A. Huse, M. Troyer, J. Stat. Mech. P03018 (2006)
18. S. Trebst, M. Troyer, U.H.E. Hansmann, J. Chem. Phys. **124**, 174903 (2006)
19. C. Predescu, M. Predescu, C.V. Ciobanu, J. Chem. Phys. **120**, 4119 (2004)
20. M. Troyer, S. Wessel, F. Alet, Phys. Rev. Lett. **90**, 120201 (2003)