

A brief history of the introduction of generalized ensembles to Markov chain Monte Carlo simulations

Bernd A. Berg^a

Department of Physics, Florida State University, Tallahassee, FL 32306, USA

Received 26 July 2016 / Received in final form 30 September 2016
Published online 5 April 2017

Abstract. The most efficient weights for Markov chain Monte Carlo calculations of physical observables are not necessarily those of the canonical ensemble. Generalized ensembles, which do not exist in nature but can be simulated on computers, lead often to a much faster convergence. In particular, they have been used for simulations of first order phase transitions and for simulations of complex systems in which conflicting constraints lead to a rugged free energy landscape. Starting off with the Metropolis algorithm and Hastings' extension, I present a minireview which focuses on the explosive use of generalized ensembles in the early 1990s. Illustrations are given, which range from spin models to peptides.

1 Introduction

In this minireview I give a personal account of the birth of generalized ensemble simulations, which happened in the early 1990s. Though essential ideas go well back into the 1970s, they were not built on, but rediscovered around 1990. Later papers were often unaware of their predecessors. Nevertheless, I make here an attempt to present developments in a chronological order. Starting with the Metropolis algorithm, subsequent early work on the methodology, and Hastings' extension to a general Markov chain Monte Carlo (MCMC) scheme, I come then to the introduction of generalized ensembles: Umbrella sampling and more than a decade later the almost simultaneous emergence of multicanonical (MUCA) ensembles, expanded ensembles and replica exchange simulations (parallel tempering).

My account of the last two methods remains rather sketchy as I was not involved into their development. My emphasis is on publications in which I participated and those which influenced my own work or crossed its path by one or another reason. No attempt of an "objective" overview or covering the field beyond the early developments is made. So, my focus is on MUCA simulations for which calculations of interface tensions of first order phase transitions via Binder's method were the point of departure. Subsequently, I discuss applications to spin glasses and peptides

^a e-mail: bberg@fsu.edu

(small proteins) as examples of complex systems with conflicting constraints. Concerning the performance, it is outlined that even for first order phase transitions a residual exponential slowing down remains. Wang-Landau sampling is recommended for calculating the MUCA weights. Summary and conclusion are given in the final section.

2 Metropolis algorithm and Hastings' extension

MCMC calculations started in earnest with the 1953 paper by Metropolis, Rosenbluth² and Teller² [1]. In their own words: "A general method, suitable for fast computing machines, for investigating such properties as the equations of state for substances consisting of interacting molecules is described. The method consists of a modified Monte Carlo integration over configuration space." It relies on calculating the change in energy ΔE , which is caused by a stochastic move of a particle position: "If $\Delta E < 0$, i.e., if the move would bring the system to a state of lower energy, we allow the move and put the particle in its new position. If $\Delta E > 0$, we allow the move with probability $\exp(-\Delta E/kT)$; i.e., we take a random number ξ_3 between 0 and 1, and if $\xi_3 < \exp(-\Delta E/kT)$, we move the particle to its new position. If $\xi_3 > \exp(-\Delta E/kT)$, we return it to its old position. Then, whether the move has been allowed or not, i.e., whether we are in a different configuration or in the original configuration, we consider that we are in a new configuration for the purpose of taking our averages."

Whether to count a configuration after a rejected move again or not caused some discussion at its time and may still be a stumbling block for newcomers. I remember that I found the false choice in a lattice gauge theory (LGT) Metropolis programs of a postdoc in the 1990s, who had been running them for quite a while on supercomputers.

2.1 Notable early developments

In 1959 Salsburg et al. [2] made an attempt to calculate the partition function by MCMC using a method called in the modern language "reweighting":

$$\exp(-\beta E) \rightarrow \exp(-\beta E - \Delta\beta E) = \exp(-\beta' E). \quad (1)$$

As noticed by the authors their method is restricted to very small lattices when one wants to cover the entire β range from 0 to ∞ as needed for the partition function, because for a fixed, possibly large, statistics and away from critical points $\Delta\beta \sim 1/\sqrt{V}$ (V volume) holds for the admissible $\Delta\beta$ range.

In the next almost thirty years reweighting was several times forgotten and re-discovered (for instance, McDonald and Singer [3] used reweighting to evaluate physical quantities over a small range of temperatures) until finally the time was right and the MCMC community had grown large enough so that reweighting never became extinct again after the 1988 paper by Ferrenberg and Swendsen [4, 5].

Continuing with the MCMC time line, in his 1963 article Glauber [6] transcends the Metropolis dynamics by introducing the heatbath and other dynamical approaches for the Ising model.

2.2 Hastings

In Hastings own words [7]: "A generalization of the sampling method introduced by Metropolis et al. (1953) is presented along with an exposition of the relevant theory,

techniques of application and methods and difficulties of assessing the error in Monte Carlo estimates.” Unfortunately, the mathematician Hastings renamed everything the chemists and physicists had introduced before: “Equilibrium” to “stationary”, “detailed balance” to “reversibility” and “ergodicity” to “irreducibility”.

Instead of focusing on the Boltzmann distribution, Hastings constructs a Markov chain matrix P with a general stationary distribution π . Following his paper, the matrix P is made to satisfy the detailed balance (reversibility) condition

$$\pi_i p_{ij} = \pi_j p_{ji} \quad (2)$$

for all pairs of states i and j . This property ensures that $\sum_i \pi_i p_{ij} = \pi_j$ for all j , and hence that π is an equilibrium (stationary) distribution of P . Ergodicity (irreducibility) of P must be checked in each specific application. For this it is only necessary that there is a positive probability of going from state i to state j in some finite number of transitions.

It is then assumed that p_{ij} has the form

$$\pi_i p_{ij} = q_{ij} \alpha_{ij} \quad (i \neq j) \quad \text{with} \quad p_{ij} = 1 - \sum_{j \neq i} p_{ij}. \quad (3)$$

Here $Q = \{q_{ij}\}$ is the transition matrix of an arbitrary Markov chain on the states $0, 1, \dots, S$ and α_{ij} is given by

$$\alpha_{ij} = \frac{s_{ij}}{1 + (\pi_i q_{ij})/(\pi_j q_{ji})}, \quad (4)$$

where it is noticeable that the normalization of the distribution π drops out. Metropolis updating (s_{ij}^M) is a special case:

$$s_{ij}^M = \begin{cases} 1 + (\pi_i q_{ij})/(\pi_j q_{ji}) & \text{for } (\pi_j q_{ji})/(\pi_i q_{ij}) \geq 1 \Rightarrow \alpha_{ij}^M = 1, \\ 1 + (\pi_j q_{ji})/(\pi_i q_{ij}) & \text{for } (\pi_j q_{ji})/(\pi_i q_{ij}) < 1 \Rightarrow \alpha_{ij}^M = (\pi_j q_{ji})/(\pi_i q_{ij}). \end{cases} \quad (5)$$

It goes for $q_{ij} \neq q_{ji}$ sometimes under the name biased Metropolis updating. Requiring symmetry:

$$q_{ij} = q_{ji} \Rightarrow \alpha_{ij}^M = \begin{cases} 1 & \text{for } \pi_j/\pi_i \geq 1, \\ \pi_j/\pi_i & \text{for } \pi_j/\pi_i < 1, \end{cases} \quad (6)$$

and assuming the Boltzmann distribution $\pi_i = \exp(-\beta E_i)$, we obtain the original Metropolis algorithm back.

However, Hastings did not make progress in identifying suitable distributions π for important applications and error estimates of MCMC simulations are nowadays much better under control [8]. Besides, Hastings appeared to be unaware of Glauber’s earlier work [6] generalizing the Metropolis approach.

3 Umbrella sampling

Until about 1990 the paradigm of MCMC calculations in statistical physics appeared to be that the use of Boltzmann weights will optimize the simulation. However, there had been notable exceptions.

Already in 1972 Valleau and Card [9] used overlapping bridging distributions to estimate entropy and free energy by what they called “multistage sampling”. Nowadays we call this multi-histogram reweighting. Torrie and Valleau [10] continued




Fig. 1. Probability density $P(E)$, E energy, for a umbrella sample (solid line) together with canonical probabilities at two temperatures, T_1 and T_2 , as obtained by reweighting of the umbrella sample (sketched according to [12]).

with estimation of free energies from distributions designed for this purpose. Patey and Valleau [11] sampled ionic potentials with weights $w(r_{12}) \exp(-U/kT)$, where the weighting functions $w(r_{12})$ was chosen to spread the sampling over a desired range. In 1977 this line of work culminated in the proposal of “Nonphysical Sampling Distributions in Monte Carlo Free-Energy Estimation: Umbrella Sampling” by Torrie and Valleau [12].

Markov chains, which sample these distributions, are special cases of Hastings framework, but the real problem is the engineering of sampling weights that work for a particular purpose. In any case, no reference to Hastings is given, though his [7] and the Valleau et al. papers [9–12] have in common to be products of Toronto University.

The basic idea of umbrella sampling is sketched in Figure 1: The umbrella distribution is much broader than canonical distributions in its range, which are then not obtained by direct simulations, but by reweighting of a umbrella distribution. This leads us to the question “How to get the umbrella distribution in the first place?” and we find in the paper of Torrie and Valleau [12]: “Superficially, the most serious limitation of the sampling techniques described here may appear to be the lack of a direct and straightforward way of determining the weighting function to use for a given problem. Instead $w(\mathbf{q}^N)$ must be determined by a trial-and-error procedure for each case, often beginning with the information available from the distribution in a very short Boltzmann-weighted experiment which is then broadened in stages through subsequent short test runs with successively greater bias of the sampling. What this rather inelegant procedure lacks aesthetically is more than compensated by the efficiency of the ultimate umbrella-sampling experiment.”

Umbrella sampling survived mainly in a niche of computational chemistry, e.g., [13]. But ten years later the conclusion by Lie and Scheraga [14] was still: “The difficulty of finding such weighting factors has prevented wide applications of umbrella sampling”.

What is really gained by using umbrella distributions instead of a series of canonical samples and possibly patching them? The documentation at its time remained rudimentary and gains appear to be no larger than, maybe, factors of 10 to 20. An exception is the improvement one can find in chapter 6.3 of the 1987 statistical mechanics book by Chandler [15]. He considered a 20×20 Ising model below the

(critical) Curie temperature, where visitation of states with magnetization $m = 0$ is a rare event. By patching 10 umbrella windows he can reach $m = 0$ and sample configurations which are in canonical Boltzmann simulations suppressed by 6 to 7 orders of magnitude. Unfortunately, this impressive demonstration of the numerical method missed a physical focus. Naturally this would have been a calculation of the interface tension, known exactly since Onsager's [16] 1944 solution of the 2D Ising model. However, this requires Binder's [17] relation (7), which is discussed in the next section (presumably Chandler was not aware of it).

Independently of Chandler, and unaware of umbrella sampling, Bhanot and collaborators [18] developed a similar method of patching microcanonical simulations, then defined as simulations restricted to some small energy range $E_{\min} \leq E \leq E_{\max}$, and applied it to a number of physical topics.

4 Multicanonical and other generalized ensembles

The spread of the multicanonical approach is an example of an interdisciplinary success story with unexpected turns of events. In 1988 Parisi et al. [19] claimed that the deconfining phase transition of SU(3) LGT is second order, what appeared rather weird, because a large number of previous publications had agreed on a weak first order transition. However, as this was Georgio Parisi, everyone felt obliged to take this very seriously and immediately, as well as over the subsequent years, a large number of papers appeared, including one by myself and collaborators [20], which all re-confirmed the first order nature of this transition.

This also stimulated work on simpler models than SU(3) LGT with the aim to develop better methods to identify first order transitions. Examples of simple systems with first order phase transitions are 2D and 3D Ising models below the Curie temperature for which the transition is at $h = 0$ in an applied magnetic field and q -state Potts models for which there is a first order transition in the temperature ($q \geq 4$ in 2D, $q \geq 3$ in 3D required). One focus was on calculations of their interface tensions. In particular, Potvin and Rebbi [21] as well as Kajanti et al. [22], used the 7-state 2D Potts model as laboratory for a newly developed approach, which estimated the interface tension by driving half a lattice "adiabatically" from one phase into the other. Thomas Neuhaus and I stumbled in this context over Binder's method for estimating interface tensions.

4.1 Binder's method

Figure 2 sketches the probability density of the magnetization m for a 2D Ising model with periodic boundary conditions at a temperature below its critical and at external magnetic field $h = 0$. For the interface tension f^s Binder [17] derived the relation

$$f^s = \lim_{L \rightarrow \infty} f_L^s \quad \text{with} \quad f_L^s = -\frac{1}{L^{(D-1)}} \ln \left(\frac{P_L^{\min}}{P_L^{\max}} \right), \quad (7)$$

where P_L^{\max} is the maximum and P_L^{\min} the minimum of the probability density of the magnetization on a finite lattice of size L^D with periodic boundary conditions. In case of the order-order transition of Figure 2 at $h = 0$ it is obvious that the probability density is symmetric in the magnetization m . For a temperature driven first order transition there is no such symmetry. The energy, instead of the magnetization, is recorded on the abscissa and it becomes part of the calculations to determine suitable pseudotransition temperatures $\beta_t(L)$.




Fig. 2. Ising model with periodic boundary conditions below the Curie temperature: Order parameter (magnetization) distribution (qualitatively drawn after reference [17]).

On finite lattices P_L^{\max} and P_L^{\min} can be calculated by MCMC, but initial attempts [17] remained pitiful, because they were carried out in the canonical ensemble and estimates deteriorate then quickly with increasing lattice size as configurations for P_L^{\min} are exponentially suppressed according to $P_L^{\min} \sim \exp(-f^s L^{D-1})$. So, Binder's method remained dormant for almost ten years.

4.2 Multicanonical ensemble

For a temperature driven first order transition the MUCA approach [23, 24] calculates P_L^{\min} by sampling no longer with Boltzmann weights but with the inverse number of states,

$$W(E) = \text{const}/n(E), \quad (8)$$

as weights. Subsequently, one reweights the thus obtained MUCA ensemble to the canonical ensemble at an effective transition temperature, $\beta_t(L)$, for which equal heights of the maxima of the spectral probability density are realized (equal weights have also been suggested [25]).

In contrast to umbrella sampling the MUCA weights are precisely defined, but in principle MUCA appears to suffer from the same shortcoming: The weights are a-priori unknown (in contrast to the Boltzmann weights). However, well defined target weights led to a clear focus on the problem of getting them. Especially, for the envisioned calculations of interface tensions weights can be obtained by finite size extrapolations from small to increasingly larger systems: For very small systems canonical simulations work and provide the starting point. Then one extrapolates weights for the next larger system, corrects them by simulation results from this system, proceeds to the next larger system, and so on.

For the 2D 10-state Potts model lattices up to size 100^2 were covered in this way. Figure 3 shows for a 70^2 lattice the energy histogram obtained by using approximate MUCA weights and at the pseudotransition temperature the canonically reweighted histogram. Improvements by a factor 10^5 over canonical simulations are easily achieved. Figure 4 gives an example of a configuration with interfaces as found for P_L^{\min} in the canonical ensemble. Equation (7) yields for each lattice size




Fig. 3. Reweighting of a MUCA simulation at the first order phase transition of the 2D 10-state Potts model to the canonical ensemble [24].

Table 1. Interface tensions f^s for some 2D Potts models.

q	MUCA [24,26]	Potvin and Rebbi [21]	Exact (Borgs and Janke [27])
7	0.0241 (10)	0.1886 (12)	0.020792...
10	0.09781 (75)	—	0.094701...

the effective interface tension and in Figure 5 these estimates are extrapolated to the infinite volume result ($1/L \rightarrow 0$). The thus obtained f^s is the second entry of the third row of Table 1.

However, there was a problem now: The MUCA estimate for the interface tension of the 2D 10-state Potts model is lower than this estimate for the 2D 7-state Potts model by Potvin and Rebbi [21], the third entry of the second row of Table 1 (Kajanti et al. [22] agree and give $2f^s \approx 0.2$ without error bar). As the strength of the 2D Potts model transitions increases with the number of states, this should be the other way round. The discrepancy was confirmed by a direct MUCA calculation of the 2D 7-state Potts model [26]¹ for which f^s is listed as the second entry of the second row of Table 1. What now? As nobody will easily admit that his or her numerical method does not work, an endless debate was looming. But, see the next section.

4.3 Miracles happen (thanks to Borgs and Janke)

After the numerical results were published Borgs and Janke [27] realized, based on previous work by Buffenoir and Wallon [28] and others, that there is actually an exact solution for the interface tension of 2D q -state Potts models. While numerical calculations tend to reproduce known exact results accurately, this was not so much the case for these previously unknown exact results, which are for the 2D 7-state and

¹ This paper was actually my first collaboration with W. Janke, who was then on a joint postdoc appointment with the Embry Riddle Aeronautical University at Daytona Beach and the Florida State University.




Fig. 4. 2D 10-state Potts model: A configuration with disorder-order interfaces from the P_L^{\min} configurations of the canonical ensemble.

10-state models collected in the 4th column of Table 1. The MUCA estimate for the 10-state Potts model agrees within less than 4%, but only within five sigma. For the 7-state model the MUCA estimate needs three error bars to include the exact value. However, this is very good when compared to the other numerical estimate for the 7-state model, which is an entire order of magnitude too large. So, due to the paper by Borgs and Janke the looming endless discussion was cut short (the interfaces are perceived as too stiff in the method of [21,22]). Looking back, the controversy helped a lot to draw attention on MUCA simulations as a reliable method. Almost needless to say, once the exact results are known, the discrepancies with MUCA estimates can be resolved by applying more sophisticated fits [29] than the one of Figure 5.

4.4 Other early multicanonical simulations and new horizons

Soon MUCA calculations were also carried out for the interface tensions of the order-order transitions of 2D and 3D Ising models below their critical temperatures and MUCA is in this application called multimagnetical (MUMA). The 2D case allows for




Fig. 5. Extrapolation of the interface tension for the 2D 10-state Potts model [24].

comparison with Onsager's [16] exact result and MUMA calculations were well consistent [30]. This was certainly of interest when this paper was submitted in September 1991 and the controversy about the estimates for Potts models was still unsettled. However, due to the malice of some referees, which were finally overruled by the editor, the paper appeared only in January 1993. By then subsequent work [31], which included estimates for 3D Ising model interface tensions, appeared almost simultaneously in print. Figure 6 shows the reweighted canonical probability densities of one of our MUMA simulations of the 3D Ising model and it is clear that an astronomically large improvement over canonical simulation is reached: Configurations for the minimum are sampled, which are suppressed by a factor of approximately $e^{-150} \approx 10^{-60}$ in the canonical ensemble.

Somewhat later, extensions of the MUCA method to cluster algorithms, multi-bondic (MUBO) simulations, were introduced and studied by Janke and collaborators [32,33].

Realizing the importance of the MUCA approach beyond first order phase transitions, Celik and I [34] published in 1992 an application to complex systems with a rugged free energy landscape due to frustrated interactions: The 2D Edwards-Anderson Ising spin glass. By sampling with MUCA weights (8) one can move in and out of valleys and find their relative heights by connecting them through the disordered phase. This is illustrated in Figure 7: The mountain in the background is sampled from configurations of the disordered high temperature phase and the MUCA updating process travels in and out of low temperature valleys (this is where the histogram entries are), which are separated by high barriers (means no histogram entries) in the probability density of the overlap parameter. Our spin glass investigations with MUCA went on for quite a while [35] and the final paper with Billoire and Janke [36] was about overlap barriers in the 3D and 4D Edwards-Anderson Ising spin glass model. In contrast to first order phase transitions finite size extrapolations of the MUCA weights are not possible for spin glasses, because each replica is unique. So, reference [34] introduced a recursive approach to estimate them. Progress has since then been achieved with respect to this and is shortly reviewed in Section 4.7.




Fig. 6. Probability densities of the magnetization for the 3D Ising model at a temperature below its critical value [31].

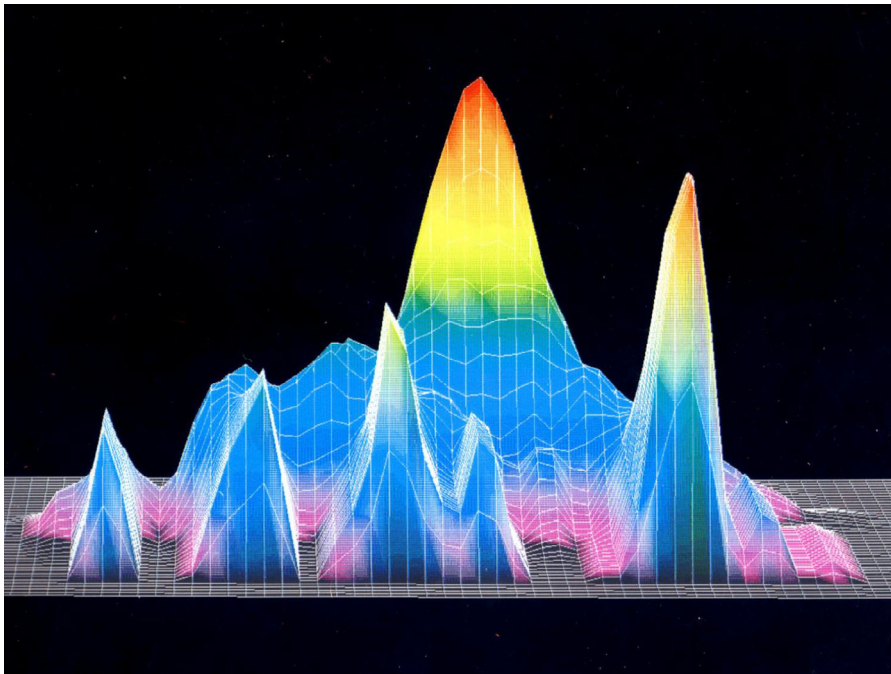


Fig. 7. Snapshot from a MUCA simulation of the 2D Edwards-Anderson Ising spin glass, which exhibits barriers in the overlap parameter [34].




Fig. 8. Configurations from folding of poly-alanine into its helix groundstate [38,39].

The concluding remarks in [34] are: “The similarities of spin glasses to other problems with conflicting constraints suggest that MUCA simulations may be of value for a wide range of investigations: optimization problems like the traveling salesman, neural networks, protein folding, and others.” Hansmann and Okamoto made the application to small proteins reality [37]. Figure 8 pictures the folding of poly-alanine configurations into their helix groundstates in a MUCA simulation that connects the native groundstate with a disordered initial configuration [38,39]. Their review article [40] coined the generic name “Generalized Ensembles” for MUCA ensembles, umbrella sampling, expanded ensembles and the replica exchange method.

4.5 Expanded ensembles and the replica exchange method

Expanded ensembles [41,42] enlarge the configuration space by one or more additional variables. If this is the temperature, the method is also called “simulated tempering”. In the replica exchange method [43,44] several identical replica of the systems are simulated with parameter values chosen so that transitions between the replica, which are accepted according to the Metropolis method, become possible, i.e., have reasonable acceptance rates. This is a special case of the approach of reference [45], which was not focused enough to provide guidance for practical applications. The replica exchange method allows for easy parallelization, because the simulation of each replica is independent from that of the other replica and the exchanges requires only transfer of a few variables. See [47] for work by Janke and collaborators on this subject.

If the replica differ in the temperature variable, the replica exchange method is often called “parallel tempering” (PT), which should not be confused with simulated tempering. Presumably due to the easy use of many processors on supercomputers, PT has become quite popular, in particular for protein simulations to which it was introduced in reference [46] within the MCMC framework. However, in biophysics molecular dynamics (MD) simulations are most widely used. To them PT was adapted a few years later in reference [48].




Fig. 9. Energy optimized droplet versus percolated configuration as discussed in [49].

That all these methods flourished since the early 1990s, and were not forgotten again, shows that by now the time was right to create a sufficiently large and sophisticated user community.

4.6 Performance of generalized ensemble simulations

The number of steps a random walk (diffusive process) needs to bridge a distance N is proportional to N^2 . Therefore, the optimal performance of a MUCA simulation to bridge an extensive energy gap is proportional to V^2 , where V is the volume of the system. In the early work [24] it was assumed that for first order transitions the performance would be close to this optimum, something like $V^{2.3}$ was consistent with the data. This turned out to be too optimistic. The slowing down is still exponential in system size L . Only the worst exponential, $\sim \exp(-f^s L^{D-1})$, has been eliminated. The argument due to Neuhaus and Hager [49] is illustrated in Figure 9 for an order-order transition. The interface energy minimized shape for the green within the red phase is a droplet as long the green phase region is sufficiently small. However, once the green phase is sufficiently large, the optimal shape becomes the percolated green rectangle and transitions from the droplet to the rectangle have to go through exponentially suppressed configurations. This is an example of hidden barriers for which a good reweighting variable is difficult to identify. In practice the coefficient of this subleading exponential appears to be rather small, so that the diffusive process dominates the slowing down for some while and effectively a power law in L is observed as long as the systems are not too large.

This argument does not apply directly to PT, because each replica is then simulated in the canonical ensemble. As a consequence rare configurations of free energy barriers like the one of the right side of Figure 9 are in practice not sampled but at best jumped over and it is wrong to use this method when one is interested in sampling barrier configurations. On the other hand, when one is interested in connecting distinct groundstate branches, like those of our spin glass Figure 7, through the disordered phase, PT appears to be well suited.

4.7 Wang-Landau recursion and parallelization

Ever since [34] recursions for MUCA weights received appropriate attention. The proposed schemes, like [50], relied on manipulations of additive histograms of the already covered energy ranges. This changed when Wang and Landau [51] introduced a new sampling method that relies on a multiplicative instead of a additive recursion of weights:

$$\ln W(E) \rightarrow \ln W(E) - a, \quad a > 0, \quad (9)$$

whenever the (discrete) energy E is sampled. Then $a \rightarrow a/2$ after a sufficiently flat histogram is obtained. Though proposed as a sampling method, in the MUCA context

it is best used as a recursion where the weights get frozen once they are working estimates of the MUCA weights. Here “working estimate” means that the energy range of interest is cycled with the frozen weights (instead of getting stuck in part of that range). In tests, which I performed iterating the MUCA weights this way, it has turned out to be a robust way of getting working estimates within a small fraction of the CPU time needed for the subsequent calculation with fixed weights. So, there is little incentive to hunt for further improvement, although it is not finally settled whether additive recursions can do as well or even better (a poster [52] comparing Wang-Landau and MUCA recursions was presented at this conference).

With Bazavov [53] I have worked out a straightforward implementation of the Wang-Landau recursion for the case of a continuous energy variable. The basic steps are:

1. Bin into histograms just large enough so that a single updating step cannot jump a bin.
2. Track the mean values within each bin.
3. Do not use constant weights over a bin, but interpolate logarithmic weights (9) linearly between mean values of neighbor bins.

The performance was as good as the one achieved in models with discrete energies.

For simulations on supercomputers parallelization is crucial. Here I like to point the reader towards two relatively recent papers. Massively parallel Wang-Landau sampling is discussed by Landau and collaborators in [54] and parallelization of MUCA simulations by Janke and collaborators in [55].

5 Summary and conclusions

A message is that mastering simulational methods is not the trivial part of a physical, chemical or biophysical MCMC study. Astronomically large efficiency factors can float around between getting a simulation optimized or not. There have been hundreds of papers refining and improving the methods sketched here. Within this limited review I cannot follow up on them. An efficient way to track such investigations is to look for citations of the papers quoted here and work from there up to the latest developments.

Obviously, for a number of applications generalized ensembles allowed great leaps in the efficiency of MCMC simulations, but ultimately it is unavoidable that one reaches some limits where one is stuck again. The two main problems appear to me:

1. There are often hidden barriers and we have been unable to find reaction coordinates in which these barriers become explicit. Recommendation: Keep trying.
2. The optimal performance of the discussed methods is limited by the V^2 power law due to being a diffusive process. For large systems that can still be far too slow. See [56] for some recent ideas to improve on the diffusive process.

At second order phase transitions collective updating with cluster algorithms provides great improvements, but they have a rather limited applications range. Non-equilibrium simulations could be a way out if one could relate the generated configurations to those of canonical equilibrium ensembles. Putting some dynamics into simulations by combinations of MCMC with MD may help. Most important, we are awaiting the not yet existing ideas of the next generation of computational physicists.

This article is dedicated to Wolfhard Janke on the occasion of his 60th birthday.

I would like to thank Martin Weigel and the other organizers for inviting me to Wolfhard's belated birthday party.

References

1. N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller, *J. Chem. Phys.* **21**, 1087 (1953)
2. Z.W. Salsburg, J.D. Jacobson, W.S. Fickett, W.W. Wood, *J. Chem. Phys.* **30**, 65 (1959)
3. I.R. McDonald, K. Singer, *Disc. Faraday Soc.* **43**, 40 (1967)
4. A.M. Ferrenberg, R.H. Swendsen, *Phys. Rev. Lett.* **61**, 2635 (1988)
5. A.M. Ferrenberg, R.H. Swendsen, *Phys. Rev. Lett.* **63**, 1658 (1988)
6. R.J. Glauber, *J. Math. Phys.* **4**, 294 (1963)
7. W.K. Hastings, *Biometrika* **57**, 97 (1970)
8. B.A. Berg, *Markov Chain Monte Carlo Simulations and their Statistical Analysis* (World Scientific, 2004)
9. J.P. Valleau, D.N. Card, *J. Chem. Phys.* **57**, 5457 (1972)
10. G.M. Torrie, J.P. Valleau, *Chem. Phys. Lett.* **28**, 578 (1974)
11. G.N. Patey, J.P. Valleau, *J. Chem. Phys.* **63**, 2334 (1975)
12. G.M. Torrie, J.P. Valleau, *J. Comp. Phys.* **23**, 187 (1977)
13. J. Chandrasekar, S.F. Smith, W.K. Jorgensen, *J. Am. Chem. Soc.* **107**, 154 (1985)
14. Z. Lie, H.A. Scheraga, *J. Mol. Struct.* **179**, 333 (1978)
15. D. Chandler, *Introduction to Modern Statistical Mechanics* (Oxford University Press, 1987)
16. L. Onsager, *Phys. Rev.* **65**, 117 (1944)
17. K. Binder, *Phys. Rev. A* **25**, 1699 (1982)
18. G. Bhanot, R. Salvador, S. Black, P. Carter, R. Toral, *Phys. Rev. Lett.* **59**, 803 (1987)
19. G. Parisi et al., *Phys. Rev. Lett.* **61**, 1545 (1988)
20. N.A. Alves, B.A. Berg, S. Sanielevici, *Phys. Rev. Lett.* **64**, 3107 (1990)
21. J. Potvin, C. Rebbi, *Phys. Rev. Lett.* **62**, 3062 (1989)
22. K. Kajanti, L. Kärkkäinen, K. Rummukainen, *Phys. Lett. B* **223**, 213 (1989)
23. B.A. Berg, T. Neuhaus, *Phys. Lett. B* **267**, 249 (1991)
24. B.A. Berg, T. Neuhaus, *Phys. Rev. Lett.* **68**, 9 (1992)
25. C. Borgs, S. Kappler, *Phys. Lett. A* **171**, 37 (1992)
26. W. Janke, B.A. Berg, M. Katoot, *Nucl. Phys. B* **382**, 649 (1992)
27. C. Borgs, W. Janke, *J. Phys. I France* **2**, 2011 (1992)
28. E. Buffenoir, S. Wallon, Saclay preprint SPHT/92-077
29. A. Billoire, T. Neuhaus, B.A. Berg, *Nucl. Phys. B* **413**, 795 (1994)
30. B.A. Berg, U. Hansmann, T. Neuhaus, *Phys. Rev. B* **47**, 497 (1993)
31. B.A. Berg, U. Hansmann, T. Neuhaus, *Z. Phys. B* **90**, 229 (1993)
32. W. Janke, S. Kappler, *Phys. Rev. Lett.* **74**, 212 (1995)
33. B.A. Berg, W. Janke, *Phys. Rev. Lett.* **98**, 040602 (2007)
34. B.A. Berg, T. Celik, *Phys. Rev. Lett.* **69**, 2292 (1992)
35. B.A. Berg, T. Celik, U. Hansmann, *Phys. Rev. B* **50**, 16444 (1994)
36. B.A. Berg, A. Billoire, W. Janke, *Phys. Rev. B* **61**, 12143 (2000)
37. U. Hansmann, Y. Okamoto, *J. Comp. Chem.* **14**, 1333 (1993)
38. U. Hansmann, Y. Okamoto, *J. Chem. Phys.* **110**, 1267 (1999)
39. U. Hansmann, Y. Okamoto, *J. Chem. Phys.* **111**, 1339 (1999)
40. U. Hansmann, Y. Okamoto, *Ann. Rev. Comp. Phys.*, edited by D. Stauffer (World Scientific, 1999) vol VI, pp. 129–157
41. A.P. Lyubartsev, A.A. Martsinovski, S.V. Shevkanov, P.N. Vorontsov-Velyaminov, *J. Chem. Phys.* **96**, 1776 (1992)
42. E. Marinari, G. Parisi, *Europhys. Lett.* **19**, 451 (1992)

43. G.J. Geyer, in *Computing Science and Statistics, Proceedings of the 23rd Symposium on the Interface*, edited by E.M. Keramidas, S.M. Kaufman (Interface Foundation, Fairfax, VA, 1991) pp. 156–163
44. K. Hukusima, K. Nemoto, J. Phys. Soc. Japan **65**, 1604 (1996)
45. R.H. Swendsen, J.-S. Wang, Phys. Rev. Lett. **57**, 2607 (1986)
46. U.H. Hansmann, Chem. Phys. Lett. **281**, 140 (1997)
47. E. Bittner, A. Nußbaumer, W. Janke, Phys. Rev. Lett. **101**, 130603 (2008)
48. Y. Sugita, Y. Okamoto, Chem. Phys. Lett. **314**, 141 (1999)
49. T. Neuhaus, J.S. Hager, J. Stat. Phys. **113**, 47 (2003)
50. B.A. Berg, J. Stat. Phys. **82**, 323 (1996)
51. F. Wang, D.P. Landau, Phys. Rev. Lett. **86**, 2050 (2001)
52. E. Bittner, W. Janke, in preparation.
53. A. Bazavov, B.A. Berg, Comp. Phys. Comm. **180**, 2339 (2009)
54. T. Vogel, Y.W. Li, T. Wüst, D.P. Landau, Phys. Rev. Lett. **110**, 210603 (2013)
55. J. Zierenberg, M. Marenz, W. Janke, Comp. Phys. Comm. **184**, 1155 (2013)
56. T. Vogel, D. Perez, Phys. Rev. Lett. **115**, 190602 (2015)