

Multicanonical methods, molecular dynamics, and Monte Carlo methods: Comparison for Lennard-Jones glasses

Kamal K. Bhattacharya* and James P. Sethna†

Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, New York 14853-2501

(Received 30 June 1997)

We applied a multicanonical algorithm to a two-dimensional and a three-dimensional Lennard-Jones system with quasicrystalline and glassy ground states. Focusing on the ability of the algorithm to locate low-lying energy states, we compared the results of the multicanonical simulations with standard Monte Carlo simulated annealing and molecular-dynamics methods. We find slight benefits to using multicanonical sampling in small systems (less than 80 particles), which disappear with larger systems. This is disappointing as the multicanonical methods are designed to surmount energy barriers to relaxation. We analyze this failure theoretically and show that (i) the multicanonical method is reduced in the thermodynamic limit (large systems) to an effective Monte Carlo simulated annealing with a random temperature vs time and (ii) the multicanonical method gets trapped by unphysical entropy barriers in the same metastable states whose energy barriers trap the traditional quenches. The performance of Monte Carlo and molecular-dynamics quenches were remarkably similar. [S1063-651X(98)09102-8]

PACS number(s): 02.50.Ng, 61.43.Bn, 02.70.Ns, 02.70.Lq

I. INTRODUCTION

In the past decade there has been an outpouring of interest in accelerating statistical mechanics simulations. This started with the work of Swendsen and collaborators: Swendsen and Wang introduced a cluster-flip method for accelerating non-disordered spin systems [1,2] and Widom, Strandburg, and Swendsen introduced a cluster-flip for finding quasicrystalline ground states in a two-dimensional atomic simulation [3]. These methods all gain a major speedup by introducing mostly nonlocal update rules and often prove capable of bypassing critical slowing down problems [4].

During the same period, different accelerating approaches were introduced, which are based on efficient schemes to analyze data from traditional Monte Carlo simulations [4–6] and are frequently called histogram methods. These methods have expanded the applicability of various kinds of critical phenomenon simulations, although they are not necessarily designed to bypass critical slowing down problems as efficiently as, e.g., cluster algorithms. Nevertheless, substantial progress can be achieved combining histogram and cluster-flip algorithms (see, e.g., Ref. [7]).

More recently, so-called reweighting techniques have been introduced, which are based on an early approach by Torrie and Valleau [8]. They proposed a method to enlarge the sampling range of a Monte Carlo algorithm by using nonphysical weighting functions. The general idea in the more recent approaches is to change the relative weights of different configurations to sample equally in all ranges of energy rather than focusing on a narrow temperature range. The most frequently used reweighting method is multicanonical sampling [9–13], which represents the most general method as other reweighting methods, e.g., entropic sam-

pling [14], can be directly mapped onto this approach [15]. In systems with a strongly double-peaked probability distribution of magnetization or energy states (a situation often found in systems exhibiting a first-order phase transition), the multicanonical approach has been proven to be a powerful tool. Simple reweighting schemes allow to overcome the “supercritical slowing down” [12] known from canonical Monte Carlo simulations at a fixed temperature. For example, in non-disordered spin systems with a field-driven first-order phase transition (e.g., the Ising model) or a temperature-driven first-order phase transition (e.g., the q -state Potts model) the supercritical slowing down of canonical Monte Carlo is due to the low Boltzmann weight of the domain-wall states. Multicanonical sampling approaches the problem by introducing a weight function, which weights all magnetization states (Ising model) or energy states (Potts model) equally and therefore ensures that domain-wall states are sampled with the same likelihood as all other accessible states. The canonical distribution function at a fixed temperature, which contains all the thermodynamic information, can be reconstructed. Usually multicanonical sampling uses local update schemes along the lines of the Metropolis algorithm [16]; variations using cluster-flip or other methods are feasible and have proved useful (for a review consult [11,12] and references therein).

Of course, acceleration methods are most crucial for glassy systems, which otherwise can be inaccessible to numerical simulations. (Even the experiments fall out of equilibrium. Think of an experiment as 10^{23} parallel atomistic processors with picosecond clock times.) Whether one believes that glasses are sluggish because of large energy barriers to relaxation (rates $\sim e^{B/T}$) or that the free-energy barriers are due to tortuous entropically difficult routes between the metastable configurations, a clever algorithm could in principle jump directly between the glassy states. Instead of relative rates that grow as a power of $T - T_c$, acceleration could gain us exponential speedups.

*Electronic address: kamal@theorie.Physik.uni-goettingen.de

†Electronic address: sethna@msc.cornell.edu

Acceleration methods have been extensively applied to disordered spin models. These studies have focused less on understanding the performance of the algorithms because the physics of the systems is less thoroughly understood. The multicanonical methods have been applied to spin glasses in two and three dimensions to calculate the zero-temperature entropy, ground-state energies, distribution of overlaps, etc. (see Refs. [17–25]). The authors succeed in evaluating these properties with remarkable accuracy; nevertheless, whether the replica theory [26] or the droplet scaling ansatz [27] is the more appropriate picture in describing the ground-state properties of glasses could not be resolved. The performance of multicanonical sampling for glassy systems is clearly worse than for system with a less rugged landscape. We believe this failure is systematic and cannot be avoided within the framework of multicanonical sampling.

Berg and Celik [17] argued that multicanonical methods should be superior to simulated annealing (gradual cooling) [28], although a direct comparison was made only to canonical sampling (quenches to a fixed low temperature). Lee and Choi [29] applied multicanonical sampling to the “traveling salesman” problem and claimed to achieve a dramatic improvement over the traditional Monte Carlo simulated annealing approach. Newman [30] has used both cluster methods [7] and multicanonical (entropic) sampling [14] to study the random-field Ising model. He finds dramatic speedups from both methods, often reaching equilibrium in a few passes through the lattice. Newman has focused on small systems (mostly 24^3 , with a few runs for systems up to 64^3) and simultaneously used histogram methods to measure critical exponents and phase boundaries for a range of disorders and temperatures. He confirmed results of the related “simulated tempering” approach, developed by Marinari and Parisi [31,32]. Simulated tempering proved very useful in spin-glass simulations [33] and is similar in spirit to the multicanonical approach [34].

Acceleration methods have been used little in continuum atomic simulations, perhaps because of the widespread reliance on molecular-dynamics methods. Straightforward, direct molecular-dynamics simulations of the equations of motion do not converge to an equilibrium state much faster than the Monte Carlo simulated annealing methods, but they are also not noticeably worse [35] and they have a direct physical interpretation. Shumway and Sethna [36] studied a one-dimensional atomic system in an incommensurate sinusoidal potential and developed an evolutionary algorithm that generated optimal cluster moves as the system was quenched to lower temperatures; later attempts to generalize these ideas to higher dimensions have not been successful so far [37]. Hansmann and Okamoto [38] used multicanonical sampling and Monte Carlo simulated annealing to study the folding of the peptide met-enkephalin; the multicanonical method found the ground state more consistently using the same amount of computer time. This result underlines the general belief [17,39] that simulations in the multicanonical ensemble are in many ways superior to traditional simulated annealing.

In this paper we apply multicanonical sampling to two-component Lennard–Jones systems and compare the performance with traditional simulated annealing and straightforward molecular dynamics in finding low-energy

configurations. We search for low-energy states of a three-dimensional Lennard–Jones glass, one of the prototype glassy systems [40–43], and use the set of parameters recently introduced by Kob and Andersen [44–46]. In addition to that, we apply multicanonical sampling and simulated annealing to a two-dimensional Lennard–Jones system with quasicrystalline ground states, using the parameters of Ref. [3]. We find that multicanonical sampling offers little benefit for the study of either. We argue that this is likely a general effect, applicable to all simulation methods applied to glassy systems in the thermodynamic limit.

II. INTRODUCTION TO THE METHODS: MULTICANONICAL METHODS AND THE ENTROPIC SAMPLING VARIANT

The standard way of implementing a Monte Carlo algorithm is using importance sampling. The idea behind this approach is simple. Rather than weighting each sample in phase space equally, each state is weighted with a sample probability distribution $\Gamma(x)$, where x denotes the sampled configuration of the system. To estimate the thermal average of an observable A , one calculates

$$\langle A \rangle = \frac{\sum_x A(x) \exp[-\beta H(x)] \Gamma^{-1}(x)}{\sum_x \exp[-\beta H(x)] \Gamma^{-1}(x)}, \quad (1)$$

where H is the Hamiltonian of the system [so $H(x)$ is the energy E for the state x] and $\beta = 1/k_B T$. Choosing $\Gamma(x)$ non-uniformly ensures that states with important contributions to the partition sum are preferentially sampled and therefore the number of states needed to be sampled to provide a reasonable estimate of A is significantly reduced.

In standard Monte Carlo methods, i.e., canonical Monte Carlo or simulated annealing, the weighting distribution is the Boltzmann distribution $\Gamma = \exp[-\beta H(x)]$. This has the advantage of a direct physical interpretation: The computer is doing the same thermal average as an equilibrium system at temperature $1/k_B \beta$. It has an important disadvantage that configurations and events that are rare in the physical system are also rare in the simulation. In particular, if the system has a “rugged energy landscape” with large free-energy barriers B separating physically important metastable states, the system will cross between these states with the same slow Arrhenius rate $\nu \exp(-B/T)$ that is found experimentally.

The idea of multicanonical sampling is to circumvent this problem by choosing $\Gamma(x)$ so that the distribution of states $P(H(x)) \sim \Omega(H(x)) \Gamma(x)$ is approximately flat in energy (or some other variable, such as magnetization [12]). In principle, we want to choose $\Gamma(E) = 1/\Omega(E) = \exp[-S(E)]$, where $\Omega(E)$ is the density of states at energy E and $S(E)$ is the entropy. Of course, we do not begin the simulation knowing the entropy as a function of energy.

In our work we use the entropic sampling algorithm [14], which is a numerical and mathematical equivalent variant of the multicanonical approach [15]. The only difference between entropic and multicanonical sampling is the way by which one generates estimates $J(E)$ of $S(E)$. The entropic sampling algorithm uses a quite straightforward recursive

updating method: (i) Initialize to zero an array $H(E)$, which will keep track of the energy of the visited states, (ii) sample states according to the current $J_i(E)$ and add the energy of each sampled state to the histogram H for a reasonably long time, and (iii) set the new J_{i+1} according to the rule

$$J_{i+1}(E) = \begin{cases} J_i(E) + \ln[H(E)] & \text{if } H(E) \neq 0 \\ J_i(E) & \text{if } H(E) = 0. \end{cases} \quad (2)$$

The multicanonical sampling update scheme differs from entropic sampling only in the treatment of the histogram bins with few entries (for an analysis of various schemes see Refs. [13,39]). The original approach introduces a constant slope for $J(E)$ below a cutoff energy, corresponding to a small constant temperature. These extra parameters are annoying [39] in the implementation; however, they do tend to keep the system from being trapped in energy regions that have not hitherto been sampled frequently. As we will argue, in glassy systems both algorithms will tend to get trapped in low-energy metastable states even when the statistics are fine. (In any case our simulations spend around half the time at high energies, so any algorithmic improvements can bring at best a factor of 2 in computer time.) In this paper we use the simpler entropic sampling method of Eq. (2); nevertheless, we will refer to this method as multicanonical sampling.

III. THEORETICAL EXPECTATIONS FOR RELATIVE PERFORMANCE

What makes people think that multicanonical sampling should be an improvement over simulated annealing or molecular dynamics? We consider three possible reasons.

(i) Perhaps the multicanonical method is better because it allows the system to cross energy barriers (as is mentioned frequently [9–14]). This is indeed an improvement over canonical sampling at a fixed temperature; however, a simulated annealing method also runs at a variety of temperatures.

Indeed, the two methods are *identical* in the thermodynamic limit. The acceptance ratio for a given single-atom Monte Carlo move for multicanonical sampling is $P(E) = \exp[S(E) - S(E')]$. In a large system with N atoms, the entropy density $S(E)/N$ is a smooth function of the energy density E/N ; since the energy density change for a single-atom move $(E' - E)/N$ is small, we may expand $S(E)$ to first order in $E' - E$. Using the relation $\partial S(E)/\partial E = 1/T$, the acceptance ratio becomes $P(E) = \exp[-(E' - E)/T]$. Thus multicanonical sampling at the energy E has exactly the same acceptance ratio as simulated annealing at a temperature $T(E) = [\partial S(E)/\partial E]^{-1}$.

Thus the local behavior, the acceptance ratio for Monte Carlo moves from a given state, is virtually the same for multicanonical and canonical sampling. Near first-order transitions, canonical quenches produce large changes in the state for small changes in temperature and thus behave quite differently from the multicanonical approaches (which by varying the energy explore the interface states directly). This is one of the major applications of multicanonical sampling methods. We expect that the multicanonical methods will perform for these systems rather similarly to microcanonical quenches that conserve the energy; see Ref. [47]. In our

glassy simulations, this distinction is presumably not important. The differences between the two methods near a given state should be similar in magnitude and type to the differences between the microcanonical (fixed-energy) simulations and the canonical (fixed-temperature) simulations: Differences can be seen for small systems, but disappear as the system gets larger. To be explicit, for a large system the final state of an multicanonical sampling run for which the time-dependent energy is $E(t)$ should be statistically equivalent to a simulated annealing run with randomly fluctuating temperature $T(E(t))$. One notes also that the quench rate is not tunable for the multicanonical method: The “diffusion constant” in energy space depends on the atomic step size and on the number of particles. This is potentially a serious handicap, as changing the quench rate is the primary tool used in glassy systems to find lower energy states. Thus the power of multicanonical methods to vary the energy to facilitate barrier crossing is, for large systems at least, no different from repeatedly heating and cooling the entire system.

(ii) Perhaps the multicanonical method might be picking the heating and cooling schedule intelligently in order to escape from local minima. Indeed, since the effective temperature becomes lower as the energy decreases, a multicanonical sampling system stuck in a high-energy metastable state will have larger thermal excitations (bigger acceptance of upward moves in energy) than one in a low-energy state and will depart faster. This is the explanation, we believe, for the substantial success of the multicanonical sampling method seen in the past.

This preferential escape from high-energy metastable states will unfortunately also become unimportant for large systems. One can see this most easily by considering a local region trapped in a high-energy configuration with local energy e' and with a lower energy configuration e nearby, separated by a barrier b . For a small system, where the local energy difference $e' - e$ is important, the effective temperatures in states e' and e will differ, but for a large system of size N this temperature difference (from the differing acceptance ratios from the two states) will vanish as $1/N$. There are glassy systems in which the energy barriers and energy differences are not all local: mean-field spin glasses, for example, have energy barriers that grow as powers of the number of spins N [26]. However, the maximum energy barrier (and presumably the maximum energy asymmetry $e' - e$) scales with a power N^α with α strictly less than one (at least in finite dimensions), so the change in effective temperature N^α/N still vanishes as $N \rightarrow \infty$ [48].

(iii) Perhaps the multicanonical method is exploring different energies more effectively than an externally chosen cooling schedule for simulated annealing. For example, the multicanonical method is guaranteed to converge to an equilibrium density of states at each energy. The same is true for a simulated annealing run at infinitely slow cooling, but is not true for repeated coolings at a fixed rate, which would be expected to generate metastable states repeatedly. On the one hand, the theorems suggest that the ground state should be occupied as often as any other energy; on the other hand, it is hard to see how a multicanonical quench to low energies can bypass the metastable states that trap simulated annealing runs of comparable computer time.

To address this issue, let us consider the characteristics of

the random walk $E(t)$ that the system performs in energy space as a function of time within the multicanonical approach. For some systems, such as the Ising model, multicanonical sampling does indeed produce a roughly unbiased random walk (if one starts from a good estimate of the density of states). As the system becomes larger, the energy range scales with N and the step size of the energy stays fixed, so the time scale for diffusing from high energies to near the ground state scales as N^2 (distance scales with the square root of time). This behavior is confirmed in simulations [9,11] studying, e.g., the first-order phase transition below T_c changing the Ising model from pointing up to down, where traditional canonical methods would suffer from the surface tension barrier $\sigma L^{d-1} = \sigma N^{(d-1)/d}$ and so the time scales as $N^2 \exp \sigma N^{(d-1)/d}$. Bypassing this “supercritical slowing down” [12] is an important application for multicanonical methods. [It should be mentioned that the optimal random walk time scale for diffusion through energy space has not been observed in simulations of, e.g., the two dimensional Potts model using *local* updates; usually the system diffuses on a time scale $\sim N^{2.5}$ (see, e.g., Refs. [9,11]). Although the supercritical slowing down is bypassed, the optimal random walk behavior has only been observed with the aid of *nonlocal* update moves.]

It is known numerically that this simple argument breaks down in simulations of spin glasses [17]. The typical time to cover the energy range (called the ergodicity or tunneling time in the literature) for spin glasses scales as N^4 [19] or perhaps e^N [11] instead of N^2 . Why should the random walk argument not work for glasses?

We assert that the answer to why the random walk argument breaks down can be found in the trapping of multicanonical sampling due to the same metastable states explored in thermal coolings. Indeed, it is the metastable states that the system is *not able to explore* that trap the multicanonical sampling algorithm. The states of glassy systems are often described in a caricature treelike structure (sideways in Fig. 1). The horizontal axis of the tree can be thought of as either energy or temperature: the branches represent mutually inaccessible ergodic components. For the Ising model, there are two major ergodic components (corresponding to the two directions for the magnetization) and a few domain-wall states. For glasses, the ergodic components are sometimes thought of as regions of configuration space separated by infinite free-energy barriers (as in the mean-field spin-glass models [26]) and sometimes as regions separated by energy barriers that are too large to cross in the time scale of the experiment or simulation.

The key point is that the accessible density of states for a glass can be very different from the total density of states. In Fig. 1 we note that the ergodic component containing the ground state has a density of states that differs from the density of states for the system as a whole, starting at the energy of the first accessible metastable states. The number Σ of these inaccessible metastable states is related to the density of tunneling states in configurational glasses [49–51] and is thought to increase exponentially with the size of the system (M independent two-state systems with uncrossable barriers generate $\Sigma = 2^M$ states). In spin glasses, the number of components separated by infinite free-energy barriers




FIG. 1. Glassy metastable states are often thought to form a treelike structure. We anticipate that the inaccessible metastable states will form a barrier to leaving the ground state within the multicanonical approach because the total density of states grows much faster than the accessible density of states.

(ones that diverge as $N \rightarrow \infty$) diverges with a power of N [48].

The multicanonical sampling method is guaranteed to sample the ground-state energy just as much as any other energy. It is easy to see, however, that once the system is in the ground state, it will stay there for a long time. If the density of states within the ground-state ergodic component is $\tilde{\Omega}_g(E)$ and the density of states in the entire system is $\Omega(E)$, then the acceptance ratio for a multicanonical sampling move from E to E' is $\Omega(E)/\Omega(E')$, while the probability of a random move raising the energy is $\tilde{\Omega}_g(E')/\tilde{\Omega}_g(E)$. Hence the likelihood of sampling high-energy states E within the ground-state component will fall as $\tilde{\Omega}_g(E)/\Omega(E)$. Consider the very crude model where all ergodic components are similar and stay completely separate until the glass transition energy E_g (the energy at the glass transition temperature), at which point they merge (Fig. 2). For this model, escaping from the ground-state component will take a time that scales as the total number of ergodic components and hence diverges as $N \rightarrow \infty$. Since multicanonical sampling spends the same amount of time in each energy range, the time between independent visits to the true ground state will scale as the time needed to escape from the ground-state ergodic component. So we begin our exploration with the expectation that multicanonical sampling should be useful for small systems, but will not provide significant advantages for large system sizes.

IV. IMPLEMENTATION

We argued in the preceding section that multicanonical sampling will not be a fundamental improvement over repeated coolings using simulated annealing, at least in large systems. On the other hand, entropic sampling and other multicanonical methods have been reported to lead to sub-




FIG. 2. Grossly oversimplified picture of glassy states, assuming that all components merge at the glass transition at energy E_g . The parameter ξ denotes the fraction of ergodic components at each energy or temperature. At the glass transition, the fraction of accessible ergodic components is of order $1/\Sigma$, thus the time to find the ground-state ergodic component (or leave it once found) scales with Σ as $\Sigma \rightarrow \infty$.

stantial gains in equilibration times for small glassy spin systems [17–25]. We are interested in simulations of structural glasses: collections of atoms that typically form metastable, glassy configurations when slowly cooled. In this section we give a detailed description of our implementation of multicanonical sampling, simulated annealing, and molecular dynamics. In order to ensure a fair comparison, we have, as often as possible, taken cooling schedules and time and spatial step sizes from standard references in the literature.

A. Three-dimensional system

In our three-dimensional simulations, we applied the three algorithms to a binary mixture of large (L) and small (S) particles with the same mass, interacting via the Lennard-Jones potential of the form $V_{\alpha\beta}(r) = 4\epsilon_{\alpha\beta}[(\sigma_{\alpha\beta}/r)^{12} - (\sigma_{\alpha\beta}/r)^6]$. The values of ϵ and σ were chosen as follows: $\epsilon_{LL} = 1.0$, $\sigma_{LL} = 1.0$, $\epsilon_{LS} = 1.5$, $\sigma_{LS} = 0.8$, $\epsilon_{SS} = 0.5$, and $\sigma_{SS} = 0.88$. All results are given in reduced units, where σ_{LL} was used as the length unit and ϵ_{LL} as the energy unit. The systems were kept at a fixed density ($\rho \approx 1.2$), periodic boundary conditions have been applied, and the potential has been truncated appropriately according to the minimum image rule [52] and shifted to zero at the respective cutoff. The minimum image rule prevents a particle from using the periodic boundary conditions to “see” more than one copy of its neighbors: Using the conventional cutoff at $r = 2.5\sigma$ would demand at least 160 particles. The choice of parameters follows recent simulations of Lennard-Jones glasses [44–46,53]; this choice suppresses recrystallization of the system on molecular-dynamics time scales. This potential, together with this set of parameters, mimics the potential for $\text{Ni}_{80}\text{P}_{20}$. We looked at five different system sizes

($N = 20, 40, 60, 80, 100$). For each N we generated 30 low-energy configurations. The initial configurations were random in the case of simulated annealing and multicanonical sampling and high-temperature equilibrium configurations in the molecular-dynamics case. To compare the three methods, we defined one run length to be 10^6 sweeps through the system for the two Monte Carlo methods. The molecular-dynamics runs are quenched at a rate that consumes the same CPU time as used by the Monte Carlo sampling.

B. Two-dimensional system

In our two-dimensional simulations, we did not apply molecular dynamics (our hard-wall boundary conditions made it inconvenient). We again used a binary Lennard-Jones system, introduced by Widom, Strandburg, and Swendsen [3] with a slightly unconventional form for the potential: $V_{\alpha\beta}(r) = \epsilon_{\alpha\beta}[(\sigma_{\alpha\beta}/r)^{12} - 2(\sigma_{\alpha\beta}/r)^6]$. The Lennard-Jones parameters are chosen to favor configurations of decagonal order ($\epsilon_{LL} = 0.5$, $\sigma_{LL} = 1.176$, $\epsilon_{LS} = 1.0$, $\sigma_{LS} = 1.0$, $\epsilon_{SS} = 0.5$, and $\sigma_{SS} = 0.618$) and the system is known to have a quasicrystalline ground state. The particles are initially randomly distributed in a large cylindrical box with infinitely high walls. The potential was truncated at $r_{\text{cutoff}} = 2.5\sigma_{\alpha\beta}$ and shifted to zero at this point. All results are in reduced units with ϵ_{LS} and σ_{LS} as fundamental units. Here four different system sizes ($N = 31, 66, 101, 160$) were used, where the numbers of different particles were chosen to keep the ratio fixed close to the value of 1.06 large atoms per small atom. Widom, Strandburg, and Swendsen [3] found that this ratio led to defect-free ground states. This system provides an excellent testing ground for multicanonical sampling for various reasons. The ground state is known to be quasicrystalline, a state with a strong bond orientational order without a long-range-order periodicity. Defective configurations are easily recognized, as the typical defects consist of triangles of like particles. There are plenty of metastable states with high-energy barriers, as it takes rearrangement of a large number of particles to disentangle the triangular defects. Widom, Strandburg, and Swendsen [3] have shown that simulated annealing fails to locate ground states and always gets trapped in a long-lived metastable states, a problem they circumvented using three-particle cluster flips.

The final minimum-energy configurations for runs of all three methods were optimized by starting from the lowest-energy configurations found and quenching down to $T = 0$ using a conjugate gradient method. The resulting $T = 0$ energies are compared in Sec. V.

C. Multicanonical sampling

The multicanonical sampling method was implemented generally as described in Sec. II. The Metropolis algorithm [16] was used for local updates. We developed an initial estimate $J(E)$ for the entropy $S(E)$ with a long run (approximately 10^7 sweeps) starting from a flat distribution: This initial estimate was used as the starting distribution for the subsequent runs. In three dimensions we redid this initialization for each system size; in two dimensions we initialized in this way for the 101 particle system and used finite-size scaling [13] from this distribution for initialization at other system sizes. Finite-size scaling does not appear to work for

glassy systems with quenched disorder [17,19]. We did not include this initial computer time in the comparisons; thus we err on the side of multicanonical sampling. Notice that in continuous systems it is not obvious how to set the optimal bin size for the histogram (unlike in spin systems, where the smallest energy step determines the bin size). We therefore tested several bin sizes. For the two-dimensional systems a fixed bin size of 0.001 was used and in the three-dimensional systems we found it useful to set the bin size to $0.01/N$ energy units. Our investigations suggest that larger bin sizes can introduce artificial barriers in the low-energy range and smaller bins lead to more noise. To set a context, the typical successful energy step in a 20-particle simulation in three dimensions varied from around one at high temperatures to around 0.01 near the ground state. We explored energy-dependent step sizes for the single-atom moves, but they did not improve performance. Multicanonical sampling demands an upper cutoff for the energy: We use zero for the upper limit for both two and three dimensions.

D. Simulated annealing

Simulated annealing uses locally the Metropolis update scheme with the Boltzmann factor as the sample probability distribution. The cooling schedule implemented here is similar to the one used in Refs. [3,38], where the temperature is repeatedly lowered by a small factor and then annealed. In our runs, we choose 50 annealing steps of 20 000 sweeps each, with an initial temperature of one and a final temperature of 0.05; each temperature is thus cooled down by a factor of 0.942. The initial configurations were set at random.

E. Molecular dynamics

The molecular-dynamics (MD) routine used the velocity form of the Verlet algorithm [52]. The unit of time is given by $(m\sigma_{LL}^2/48\epsilon_{LL})^{1/2}$, where m is the mass of the particle; the Verlet time step in these units is $\delta t = 0.01$. The system was coupled to a heat bath and the temperature was reduced linearly in time according to $T_{\text{bath}} = T_{\text{start}} - \gamma_{\text{MD}}t$. Note that the cooling here is linear in time, as is traditional in the molecular dynamics of Lennard-Jones glasses [54]. The cooling rate $\gamma_{\text{MD}} = 1.0 \times 10^{-4}$ was chosen so that the MD runs consume an amount of computer time similar to that of the Monte Carlo algorithms. This cooling rate is in the middle of the range explored in recent simulations, although our system sizes are much smaller [53,54]. The initial configurations were equilibrated at a temperature $T_{\text{start}} = 1.0$ (at a small cost of computer time, which we did not factor into the comparisons) and cooled to the final temperature $T_{\text{final}} = 0.05$, yielding approximately 2.0×10^6 molecular-dynamics steps.

V. COMPARISON

In this section we will first compare the performance of the three methods in locating low-energy states of the two- and three-dimensional Lennard-Jones systems. The performance of the three methods is remarkably similar. Second, we will compare low-energy configurations of the two-dimensional system to show that the algorithms get trapped

TABLE I. $T=0$ energy per particle for the lowest-energy configuration found with multicanonical sampling, simulated annealing, and molecular dynamics.

N	Multicanonical sampling	Simulated annealing	Molecular dynamics
20	-0.89	-0.89	-0.89
40	-4.18	-4.15	-4.16
60	-5.38	-5.36	-5.37
80	-6.52	-6.56	-6.57
100	-6.85	-6.86	-6.87

in similar metastable states. Third, we will quantitatively analyze the trapping of the multicanonical sampling algorithm in a metastable state.

We present the $T=0$ energies of the lowest-energy configurations for the three-dimensional Lennard-Jones systems in Table I. For $N=20$ particles each algorithm is able to locate the same lowest-energy state, presumably the ground state. For $N=40$ and $N=60$ particles the lowest-energy state is found by multicanonical sampling. The gain in energy ΔE_{40} over simulated annealing is around 0.03 and the gain is 0.02 over molecular dynamics. For $N=80$ and $N=100$ particles the lowest energies are found by molecular dynamics, and the gain over multicanonical sampling is $\Delta E_{80} \sim 0.05$ and $\Delta E_{100} \sim 0.02$.

There are three things to notice about this table. First, the dramatic energy difference with increasing system size is due to the change in the cutoff in the potential given by the minimum image rule. Using the 20-particle cutoff in the larger system sizes, we found energies that hardly varied with system size. Second, the fact that these energies differ in the third decimal place does not mean that the differences are negligible. In Refs. [53,54] the dependence of the final energy on the cooling rate for exactly this system was studied using molecular dynamics: To gain an energy of 0.03 starting from the cooling rate we are using, they had to decrease the cooling rate by a factor of 10. Third, as we argued in Sec. III, any gains given by multicanonical sampling disappear as the system size grows.

In Table II we list the mean and the standard deviation of the energies from the 30 runs at each system size with each algorithm. It has been found in the literature that the fluctuations for simulated annealing are much larger than for multicanonical sampling [38]. We find this to be true both for simulated annealing and for molecular dynamics. Indeed, the

TABLE II. Mean energy per particle and the standard deviation evaluated using all low-energy configurations found by multicanonical sampling, simulated annealing and molecular dynamics.

N	Multicanonical sampling	Simulated annealing	Molecular dynamics
20	-0.84 ± 0.03	-0.79 ± 0.06	-0.81 ± 0.04
40	-4.10 ± 0.03	-4.07 ± 0.05	-4.08 ± 0.04
60	-5.34 ± 0.01	-5.33 ± 0.03	-5.33 ± 0.03
80	-6.50 ± 0.01	-6.51 ± 0.02	-6.51 ± 0.03
100	-6.83 ± 0.01	-6.84 ± 0.02	-6.82 ± 0.02

TABLE III. Energy per particle for the lowest-energy configuration found with multicanonical sampling and simulated annealing.

N	Multicanonical sampling	Simulated annealing
31	-2.1	-2.1
66	-2.27	-2.22
101	-2.34	-2.33
160	-2.37	-2.38

average performance of multicanonical sampling remains comparable to that of the other two methods, even at the larger system sizes (where the extremal performance was worse).

The “ultimate goal” of this field is to accelerate three-dimensional glass simulations, bypassing barriers to relaxation. Perhaps this is too high a standard: Nobody has such an algorithm. We now will apply multicanonical sampling to a two-dimensional Lennard-Jones system, where an effective cluster-flip acceleration method has been developed [3]. In Tables III and IV we show the extremal and the mean $T=0$ energies for a variety of system sizes. Again, multicanonical sampling is slightly better for the smaller systems, but the advantage disappears for the largest system.

It is remarkable how similarly the three different methods perform. Although it seems to be known that molecular dynamics and simulated annealing are comparable [35], we are not aware of any reference providing a direct comparison. Of course, comparisons of efficiency are highly implementation dependent. The two Monte Carlo methods could benefit from a temperature-dependent step size (although we did experiment with it without finding any substantial improvement). One could refine the cooling schedule for the two traditional methods. One could introduce a temperature cutoff (like in the original multicanonical approach) or use variable bin sizes to improve the multicanonical sampling method. Again, our experiments with bin size and cutoff were not encouraging. Our main conclusion is that the choice of methods is a matter of taste. In particular we are encouraged by the fact that Monte Carlo methods are competitive, especially as they adapt easily to cluster acceleration methods.

All three methods suffer from the large number of metastable states prevalent in the configuration space of the two- and three-dimensional systems and thus are not capable of locating ground states. To show that they find similar metastable states, we plot in Figs. 3 and 4 the lowest-energy configuration found by multicanonical sampling and by simulated annealing. These configurations are typical representatives of metastable states for the two-dimensional

TABLE IV. Mean energy per particle and the standard deviation evaluated using all low-energy configurations found by multicanonical sampling and simulated annealing.

N	Multicanonical sampling	Simulated annealing
31	-1.97 ± 0.1	-1.87 ± 0.17
66	-2.24 ± 0.04	-2.12 ± 0.07
101	-2.29 ± 0.07	-2.23 ± 0.06
160	-2.34 ± 0.05	-2.34 ± 0.02




FIG. 3. Lowest-energy configuration generated with multicanonical sampling. The gray particles forming a triangle represent a defect.

system. The defects are clusters of three large particles, which are shown in gray in Figs. 3 and 4.

VI. CHARACTERIZATION

Why is multicanonical sampling not bypassing the free-energy barriers to relaxation? In this section we present a vivid illustration of how the multicanonical sampling algorithm gets trapped in a metastable state. In the bottom half of Fig. 5 we plot the energy as a function of time: At very short times it performs a random walk in energy space as advertised, but it rapidly gets trapped in a low-energy metastable state.

The simulation shown in Fig. 5 is the same as the runs for $N=100$ particles tabulated in Tables I and II except for two important differences. (i) The runs of Tables I and II ran for 10^6 sweeps; here we ran for 10^7 sweeps. (ii) The entropy estimate $J(E)$ for the runs in Tables I and II was dynamically updated every 10^5 sweeps using the recursive updating scheme equation (2). Here we calculated a best estimate $\overline{J(E)}$ from the 30 runs in the tables and used this function as




FIG. 4. Lowest-energy configuration generated with simulated annealing. The gray particles forming a triangle represent a defect.




FIG. 5. Multicanonical sampling run, showing trapping into a metastable state. In this run, we used a fixed $\overline{J(E)}$ (i.e., no dynamical updating) gleaned from several previous runs of the same system. Energy vs time is shown in the bottom panel. Only the first quarter of the time series is shown. At short times, we observe a random walk in energy. At times below 2×10^5 sweeps the system randomly walks through energies above the glass transition. The system then gets trapped in two successively lower metastable free-energy wells (see Fig. 6). The histogram $\ln[H(E)] = \Delta J(E)$ of visited states measured during the simulation is shown in the upper panel. Note the large peak associated with the trapping. The second and third peaks are above the glass transition and correspond to more transient states.

a fixed entropy estimate. The best estimate $\overline{J(E)}$ (comprising information from 40×10^6 sweeps) is a sufficiently smooth function that we do not expect (or observe) the system to be trapped in some artificial well resulting from statistical fluctuations in $J(E)$.

The top half of Fig. 5 shows the logarithm of the histogram $H(E)$ tabulating the visited states as a function of energy. This function is important as it is used in the recursive updating scheme $\ln[H(E)] = \Delta J(E) = J(E)_{\text{update}} - J(E)_{\text{estimate}}$ [see Eq. (2)].

Figure 6 shows that the system is trapped in a single-harmonic metastable state. The upper panel shows an expanded view of the first peak in $\Delta J(E)$. For times after 1.3×10^6 the system exclusively samples in a single well: Repeated quenches yield the same minimum energy $E_1 = -6.8241$. This is a metastable state: As seen in Table I, the true ground state has an energy ≤ -6.87 .

In the harmonic approximation we can analytically calculate the density of states and compare the contribution from the metastable state directly to the measured data. The harmonic density of states has the form

$$\Omega_{\text{harmonic}}(E) \propto \left(\frac{2(E - E_1)}{K} \right)^{3N/2 - 1}, \quad (3)$$

where K involves the geometric mean of the phonon frequencies and can be thought of as a typical spring constant. In the multicanonical sampling algorithm the probability of sampling a state in this harmonic well is given by the ratio of the density of states in the single well divided by the estimated density of states $\exp[\overline{J(E)}]$. This probability is com-




FIG. 6. Expanded view of Fig. 5 at low energies. Energy vs time is shown in the bottom panel. From 2×10^5 to 1.3×10^6 sweeps, the system seems trapped in a number of states: Repeated quenches yield metastable energies clustering around $E_2 = -6.82 \pm 0.002$. The system then falls into a slightly lower state, with energy $E_1 = -6.8241$; repeated quenches show that it stays in this single well for the remainder of the simulation. The equilibration within this lowest well is excellent: The acceptance ratio is near 50% and the system exhibits an efficient random walk in energy within the single well. The histogram $\ln[H(E)] = \Delta J(E)$ of visited states is shown in the upper panel as a dark line. The light line is the theoretical prediction assuming a single, harmonic well at energy $E_1 = -6.8241$: $\Delta J_{\text{harmonic}}(E) = (3N/2 - 1) \ln(E - E_1) - \overline{J(E)} + C$. The bumps in the theoretical curve are due to the irregularities in our initial estimate $\overline{J(E)}$; the identical-looking bumps in the measured data reflect the effects of $\overline{J(E)}$ in weighting the histogram. Our theoretical prediction describes the measured data well for the energy range explored during the last portions of the simulation. The range above $E/N \sim -6.78$ is underestimated as in this region the data are composed largely of states corresponding to the minima clustered around E_2 .

pared directly to the histogram of sampled states in the upper half of Fig. 6. The system is trapped in a single harmonic well.

By running for shorter times and by dynamically updating the entropy estimate during each run, we have substantially mitigated the trapping problem of multicanonical sampling shown in Figs. 5 and 6. Imagine the entropy estimate $J(E)$ after updating it by adding $\ln[H(E)]$. The acceptance ratio to leave the region given by $1/J(E)$ will dramatically increase and thus the trapping will be bypassed, as indeed we observed in practice. Lingering near a state increases the estimate entropy in that region and eventually push the system out. But dynamical updating should not be an essential ingredient of the algorithm, which is formulated presuming an *a priori* knowledge of the entropy as a function of energy. Formally dynamical updating violates the Markovian character of the algorithm and convergence to the equilibrium state is no longer guaranteed. In practical terms it is very distressing that the algorithm needs to produce a noticeable bump in the density of states to escape from a metastable state. The comparison against molecular dynamics and simulated annealing would be substantially more unfavorable for long runs without dynamical updating.

Figures 5 and 6 are a tangible illustration of the trapping mechanism depicted in Figs. 1 and 2. The inaccessible metastable states contributing to $J(E)$ form a strange type of entropic barrier around the metastable state E_1 . Leaving E_1 via a saddle point at $E_2 > E_{\text{peak}} \sim -6.8$ is suppressed by roughly the exponential of $\Delta J(E_2) - \Delta J(E_{\text{peak}})$ shown in Fig. 6.

Slow cooling in molecular dynamics or simulated annealing can lead to trapping in metastable states due to large energy barriers. Entropic sampling and the other multicanonical methods get trapped in metastable states because of large entropic barriers imposed by the algorithm. In both case the algorithms are sabotaged by the large number of low-lying metastable states. Multicanonical sampling provides insight into this problem, but does not provide a solution.

VII. CONCLUSIONS

In this study we applied the multicanonical method using the specific entropic sampling scheme to Lennard-Jones systems. We focused on the ability of the algorithm to find ground states of these glassy systems and compared the performance to the two traditional glassy simulation methods simulated annealing and molecular dynamics. The use of multicanonical sampling did not reveal any new insights into the ground-state properties of Lennard-Jones glasses. We explain these results on the basis of the following observations.

First, in the thermodynamic limit multicanonical methods are locally equivalent to simulated annealing. Furthermore, the global dynamics of multicanonical sampling resembles a random heating and cooling of the sample. Thus, for large systems simulated annealing and multicanonical sampling must have the same properties. In principle, multicanonical sampling has the advantage of providing the density of states, which allows one to evaluate the canonical distribution function. In glasses this feature is not necessarily helpful, as the multicanonical methods sample phase space as slowly as the annealing methods; thus, in practice, multicanonical sampling will not be able to extract any equilibrium

expectation values better than simulated annealing.

Second, the large number of inaccessible metastable states imposes a bizarre entropy barrier to the multicanonical method. The algorithm simply gets stuck in a metastable state, as it might using molecular dynamics and simulated annealing. We underlined this point by comparing the probability distribution estimated by the algorithm inside the metastable state with a theoretical expression derived in the harmonic approximation.

Furthermore, our results emphasize the known fact that simulated annealing and molecular dynamics have similar performances in glassy systems. As a consequence, one should acknowledge the importance of averaging over many molecular-dynamics trajectories, especially for glassy systems. Averages over an ensemble of trajectories are a basic concept in Monte Carlo simulations; the striking similarity in performance to molecular-dynamics simulations is a hint at the importance of similar averages in glassy molecular-dynamics simulations.

Finally, the goal of finding a method that gains an exponential speedup of glassy simulations still remains. Our study clearly indicates that standard reweighting techniques will presumably be no substantial help in tackling this problem. The complicated structure of the glassy configuration space needs more “intelligent” algorithms, which are able not only to bypass energy barriers but also to find an efficient path through the rugged energy landscape.

ACKNOWLEDGMENTS

We would like to thank M. E. J. Newman, J. Jacobsen, K. W. Jacobsen, G. Chester, and R. Kree for many useful and illuminating discussions and B. Berg for very helpful and elucidating comments on various points. The work of J.P.S. was supported by NSF Grant No. DMR-9419506. K.K.B. is grateful for support by the German Academic Exchange Service (Doktorandenstipendium HSP II/AUFE). Computational support was provided by the Cornell Theory Center.

-
- [1] R. H. Swendsen and J. S. Wang, Phys. Rev. Lett. **57**, 2607 (1986).
 - [2] R. H. Swendsen and J. S. Wang, Phys. Rev. Lett. **58**, 86 (1987).
 - [3] M. Widom, K. J. Strandburg, and R. H. Swendsen, Phys. Rev. Lett. **58**, 706 (1987).
 - [4] R. H. Swendsen, J. S. Wang, and A. M. Ferrenberg, in *The Monte Carlo Method in Condensed Matter Physics*, edited by K. Binder (Springer, Berlin, 1992), p 75.
 - [5] A. M. Ferrenberg and R. H. Swendsen, Phys. Rev. Lett. **61**, 2635 (1988); **63**, 1658 (1989).
 - [6] J. Lee and J. M. Kosterlitz, Phys. Rev. Lett. **65**, 137 (1990).
 - [7] M. E. J. Newman and G. T. Barkema, Phys. Rev. E **53**, 393 (1996).
 - [8] G. M. Torrie and J. P. Valleau, J. Comput. Phys. **23**, 187 (1977).
 - [9] B. A. Berg and T. Neuhaus, Phys. Lett. B **267**, 249 (1991).
 - [10] B. A. Berg and T. Neuhaus, Phys. Rev. Lett. **68**, 451 (1992).
 - [11] B. A. Berg, e-print, cond-mat/9707011, and references therein.
 - [12] W. Janke in *Computer Simulation Studies in Condensed Matter Physics VII*, edited by D. P. Landau, K. K. Mon, and H. B. Schüttler, Lecture Notes in Physics Vol. 78 (Springer, Berlin, 1994), p. 29.
 - [13] G. R. Smith and A. D. Bruce, J. Phys. A **28**, 6623 (1995), and references therein; Europhys. Lett. **34**, 91 (1996).
 - [14] J. Lee, Phys. Rev. Lett. **71**, 211 (1993); **71**, 2353 (1993).
 - [15] B. A. Berg, U. H. E. Hansmann, and Y. Okamoto, J. Phys. Chem. **99**, 2236 (1995).
 - [16] A. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).
 - [17] B. A. Berg and T. Celik, Phys. Rev. Lett. **69**, 2292 (1992).
 - [18] B. A. Berg and T. Celik, Int. J. Mod. Phys. C **3**, 1251 (1992).
 - [19] B. A. Berg, T. Celik, and U. Hansmann, Europhys. Lett. **22**, 63 (1993).
 - [20] T. Celik, Nucl. Phys. B **30**, 908 (1993).

- [21] T. Celik, U. H. E. Hansmann, and M. Katoot, *J. Stat. Phys.* **73**, 775 (1993).
- [22] B. A. Berg and U. E. Hansmann, *Nucl. Phys. B* **34**, 664 (1994).
- [23] U. E. Hansmann and B. A. Berg, *Int. J. Mod. Phys. C* **5**, 85 (1994).
- [24] B. A. Berg, U. E. Hansmann, and T. Celik, *Phys. Rev. B* **50**, 16 444 (1994).
- [25] B. A. Berg, U. H. E. Hansmann, and T. Celik, *Nucl. Phys. B* **42**, 905 (1995).
- [26] M. Mezard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
- [27] D. S. Fisher and D. A. Huse, *Phys. Rev. B* **38**, 386 (1988).
- [28] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, *Science* **220**, 671 (1983).
- [29] J. Lee and M. Y. Choi, *Phys. Rev. E* **50**, R651 (1994).
- [30] M. E. J. Newman (private communication).
- [31] E. Marinari and G. Parisi, *Europhys. Lett.* **19**, 451 (1992).
- [32] E. Marinari, e-print, cond-mat/9612010.
- [33] W. Kerler and P. Rehberg, *Phys. Rev. E* **50**, 4220 (1994).
- [34] U. H. E. Hansmann and Y. Okamoto, *Phys. Rev. E* **54**, 5863 (1996).
- [35] G. V. Chester (private communication).
- [36] S. L. Shumway and J. P. Sethna, *Phys. Rev. Lett.* **67**, 995 (1991).
- [37] S. L. Shumway (private communication).
- [38] U. H. E. Hansmann and Y. Okamoto, *Physica A* **212**, 415 (1994).
- [39] B. A. Berg, *J. Stat. Phys.* **82**, 323 (1996).
- [40] F. H. Stillinger and T. A. Weber, *J. Chem. Phys.* **80**, 4434 (1984).
- [41] F. H. Stillinger and T. A. Weber, *J. Chem. Phys.* **81**, 5089 (1989).
- [42] F. H. Stillinger and T. A. Weber, *Phys. Rev. B* **31**, 5262 (1985).
- [43] T. A. Weber and F. H. Stillinger, *Phys. Rev. B* **32**, 5402 (1985).
- [44] W. Kob and H. C. Andersen, *Phys. Rev. Lett.* **73**, 13 476 (1994).
- [45] W. Kob and H. C. Andersen, *Phys. Rev. E* **51**, 4626 (1995).
- [46] W. Kob and H. C. Andersen, *Phys. Rev. E* **52**, 4134 (1995).
- [47] Ole H. Nielsen, James P. Sethna, Per Stoltze, Karsten W. Jacobsen, and Jens K. Nørskov, *Europhys. Lett.* **26**, 51 (1994).
- [48] H. Sompolinski, *Phys. Rev. Lett.* **47**, 935 (1981).
- [49] R. C. Zeller and R. O. Pohl, *Phys. Rev. B* **2**, 2029 (1971).
- [50] P. W. Anderson, B. I. Halperin, and C. M. Varma, *Philos. Mag.* **25**, 1 (1972); W. A. Phillips, *J. Low Temp. Phys.* **7**, 351 (1972).
- [51] W. B. Phillips, *Amorphous Solids: Low Temperature Properties* (Springer-Verlag, Berlin, 1981).
- [52] M. P. Allen und D. J. Tildesley, *Computer Simulation of Liquids* (Oxford, Oxford, 1986).
- [53] K. Vollmayr, W. Kob, and K. Binder, *J. Chem. Phys.* **105**, 4714 (1996).
- [54] K. Vollmayr, Ph.D. thesis, Johannes-Gutenberg University, 1996 (unpublished).