

Equilibrium microcanonical annealing for first-order phase transitions

Nathan Rose*

*Physics Department, University of Massachusetts, Amherst, Massachusetts 01003, USA
and IQB Information Technologies Inc., 458-550 Burrard Street, Vancouver, BC V6C 2B5, Canada*Jonathan Machta[†]*Physics Department, University of Massachusetts, Amherst, Massachusetts 01003, USA
and Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501, USA*

(Received 11 September 2019; published 12 December 2019)

A framework is presented for carrying out simulations of equilibrium systems in the microcanonical ensemble using annealing in an energy ceiling. The framework encompasses an equilibrium version of simulated annealing, population annealing, and hybrid algorithms that interpolate between these extremes. These equilibrium, microcanonical annealing algorithms are applied to the thermal first-order transition in the 20-state, two-dimensional Potts model. All of these algorithms are observed to perform well at the first-order transition though for the system sizes studied here, equilibrium simulated annealing is most efficient.

DOI: [10.1103/PhysRevE.100.063304](https://doi.org/10.1103/PhysRevE.100.063304)**I. INTRODUCTION**

First-order phase transitions occur in a variety of systems and are characterized by a discontinuous first derivative in the free energy and phase coexistence at the transition temperature. These characteristics pose difficulties for typical canonical ensemble Monte Carlo simulations because the transition rate between the coexisting phases decreases exponentially in system size at the transition temperature leading to exponential slowing.

Exponential slowing can be substantially reduced using multicanonical [1] or Wang-Landau [2,3] methods. A related means for reducing exponential slowing is to simulate the system in the microcanonical ensemble [4] rather than the canonical ensemble. For a temperature-driven, first-order phase transition, the coexisting phases will have distinct energies. Configurations with energies between the coexisting values of the energy possess domains of both phases separated by an interface and are very improbable in the canonical ensemble. In the microcanonical ensemble, however, a droplet of a minority phase is stable at energies between the energies of the two coexisting phases, and so simulations can be carried out relatively efficiently at any energy in the coexistence region and exponential slowing is much reduced. Nonetheless, in the microcanonical ensemble, there remains two or more barriers associated with condensation/evaporation [3,5,6] transitions at the ends of the co-existence regions and, depending on the geometry of the system, also transitions in the topology of the growing droplet [7] so that exponential slowing cannot be entirely eliminated.

In this paper we introduce a class of microcanonical Monte Carlo annealing algorithms and show that they are effective for simulating the strong first-order transition in large- q Potts models. These algorithms are related to the well-known simulated annealing [8] algorithm but with two important differences. First, annealing is carried out in energy rather than temperature, that is, we do microcanonical annealing. Second, we introduce an additional resampling step in the annealing process to properly sample the equilibrium ensemble at every energy. One of the algorithms in this class that is studied here is a microcanonical version of population annealing [9–11].

The purpose of annealing in the current setting differs from the usual objective of finding ground states or sampling thermal states for systems with rough free-energy landscapes. Here, microcanonical annealing is used to collect data at all energies to use reweighting to efficiently extract results at the transition temperature in the canonical ensemble.

The microcanonical ensemble is difficult to simulate using Markov chain Monte Carlo (MCMC) because of the constraint that proposed moves are accepted only if they do not change the energy. Here we circumvent that problem by simulating the ensemble of all equally weighted configurations with energies equal to *or below* a given energy ceiling, which we call the energy ceiling ensemble. It is straightforward to transform results from the energy ceiling ensemble to the microcanonical ensemble. An alternative approach to microcanonical Monte Carlo simulations is to introduce a kinetic energy term in the Hamiltonian and require that the total energy be fixed while the Monte Carlo moves change only the potential energy [4,6]. This “real” microcanonical ensemble achieves the same objective of allowing a continuous range of potential energies for a fixed value of the energy control parameter. Microcanonical methods with kinetic energy can be combined with replica exchange moves to improve the

*nathanrose68@gmail.com

†machta@physics.umass.edu

performance of replica exchange algorithms in molecular simulations [12].

Simulated annealing [8] consists of initializing a system at high temperature where equilibration using MCMC is easy, and gradually cooling the system following an annealing schedule that specifies a sequence of temperatures and number of MCMC sweeps at each temperature. The goal of simulated annealing is to find the ground state and, in its standard form, it is not a suitable algorithm for sampling equilibrium states. *Equilibrium* simulated annealing used here differs from standard simulated annealing in the choice of the starting configuration at the beginning of each annealing step. In standard simulated annealing, the initial configuration at a new temperature is the final configuration generated by the MCMC at the previous temperature. Note that this procedure is not well-defined in the energy ceiling ensemble since the last configuration generated by the MCMC may not be allowed at the new energy. In equilibrium simulated annealing, the new configuration is randomly sampled with appropriate weights for the new annealing step from the set of configurations generated by the MCMC at the previous annealing step. In the energy ceiling ensemble these weights are equal and the new state is randomly and uniformly chosen from the configurations generated at the previous energy that satisfy the new energy ceiling. Thus, to go from one annealing step to the next we “resample” (or, more correctly, “subsample”) from the configurations generated at the previous step.

If multiple replicas of the system are simultaneously annealed, then resampling can also be done among the set of replicas. The resulting scheme is population annealing [9,10,13–16] and is an example of a sequential Monte Carlo [17] algorithm. Population annealing in the canonical ensemble is an effective tool for simulating equilibrium systems with rough free-energy landscapes, such as spin glasses [14,15,18]. Canonical population annealing has been applied to the first-order transition in Potts models [19]. In microcanonical population annealing, resampling consists of discarding replicas that do not satisfy the new energy ceiling and making copies of those that do satisfy the new ceiling so that the population size remains constant. “Microcanonical” annealing in density rather than energy has been used to study hard sphere fluids [11]. Microcanonical population annealing is very similar to nested sampling [20]. Finally, we introduce a class of algorithms that interpolate between equilibrium simulated annealing and population annealing. In these “hybrid” annealing algorithms, multiple replicas are annealed and resampling is done from the union of time series of each replica.

Annealing in the energy ceiling ensemble provides a direct measurement of the entropy, which can be used to calculate the canonical free energy and partition function, thus giving access to canonical ensemble observables at any temperature. The entropy or free energy can also be used to perform weighted averages over several simulation runs, reducing both statistical and systematic errors of observables in the microcanonical or canonical ensembles [10].

The primary goal of the paper is to introduce the equilibrium microcanonical annealing framework and test the performance of these algorithms for strong first-order transitions. A secondary goal is to obtain high precision results for the finite-size, two-dimensional, 20-state Potts model. The paper

is organized as follows. In Sec. II we present the microcanonical annealing algorithms used in the simulations. In Sec. III we define the q -state Potts model and the observables studied in the simulations. Section IV gives details of the simulations and in Sec. V we present results from the simulations. The paper closes with a discussion in Sec. VI.

II. EQUILIBRIUM MICROCANONICAL ANNEALING

In this section we present a general framework for equilibrium microcanonical annealing. Two limiting cases of this framework are population annealing and equilibrium simulated annealing. In all cases we have an “annealing schedule” that consists of a sequence of energy ceilings $\{E^{(k)}\}$ and number of MCMC sweeps, $n_s(E^{(k)})$ that are carried out at each energy ceiling. We study these algorithms in the context of the Potts model, which has a bounded, discrete set of energy levels. For such systems the annealing schedule is naturally chosen as the set of energy levels of the system and annealing is carried out from the highest energy level to the ground state. In the following discussion the parenthetical superscript “ (k) ” is suppressed.

Microcanonical annealing directly simulates the energy ceiling ensemble in which all configurations with energies less than or equal to the energy ceiling, E are equally probable. In general, microcanonical annealing works with a population of replicas of the system that represent the energy ceiling ensemble. Each annealing step is carried out in two stages. In the first stage a population of R replicas of the system is updated using a MCMC updating procedure designed to equilibrate to the energy ceiling ensemble. During the course of this updating a pool of R^* replicas of the system is saved where, generally, $R^* \geq R \geq 1$.

The MCMC procedure used here is a single-spin-flip algorithm. The proposed update consists of flipping a randomly chosen spin. Suppose that the initial spin state is α and the final spin state is γ with energy E_γ . The acceptance probability, $P(\alpha \rightarrow \gamma)$, of the proposed move is given by

$$P(\alpha \rightarrow \gamma) = \begin{cases} 1 & E_\gamma \leq E \\ 0 & E_\gamma > E. \end{cases} \quad (1)$$

This is equivalent to the demon approach by Creutz [21] at energy E . The MCMC algorithm is carried out for $n_s(E)$ sweeps on each replica, where a sweep consists of N updates and N is the number of spins in the system. Clearly, this algorithm satisfies detailed balance with respect to the energy ceiling ensemble. If, in addition, the algorithm is ergodic, that is, all configurations of the system can be reached from one another via a series of updates, then the algorithm will converge to the desired energy ceiling ensemble. However, it is evident that this algorithm is *not* ergodic for all ceiling energies. For example, for Ising-Potts models with E sufficiently near the ground state, it is not possible to go from one ordered state to another and remain below the ceiling via single spin flip moves. Nonetheless, full coverage of configuration space and accurate equilibrium results can be obtained either by using a population of replicas of the system as is done in population annealing and/or by combining multiple independent runs as is done in equilibrium simulated annealing.

The second stage of microcanonical annealing consists of resampling with replacement from the pool of R^* replicas with energy ceiling E to generate a population of R replicas that is approximately in equilibrium at the next lower value of the energy ceiling, E' . Note that if R^* is not sufficiently large the pool may contain no configurations with energy less than or equal to E' , in which case the algorithm fails.

The entire microcanonical annealing procedure is as follows:

(1) A population of R independent replicas of the system is initialized in equilibrium with the energy ceiling E , set to the highest energy level of the system.

(2) $n_s(E)$ sweeps of the MCMC procedure are performed on each member of the population and R^* configurations generated from these MCMC sweeps are saved in the configuration pool.

(3) The energy ceiling is lowered to the next value E' in the annealing schedule ($E' < E$) and R configurations with energy less than or equal to E' are randomly drawn with replacement from the R^* configurations in the pool.

(4) Steps 2 and 3 are repeated until E is equal to the ground state energy of the system.

During the annealing step with ceiling energy E , the fraction $\epsilon(E)$ of replicas in the pool with energies greater than E is called the “culling fraction” and is used to estimate entropies.

The equilibrium microcanonical annealing framework described above can be easily generalized to an equilibrium canonical annealing framework with the following two modifications. First, the MCMC procedure must be chosen to equilibrate to a fixed temperature rather than a fixed energy ceiling. Second, the resampling step must select from the configuration pool with Boltzmann reweighting, $e^{-(\beta' - \beta)E_i}$, where E_i is the energy of configuration i and β' is the inverse temperature succeeding β in the annealing schedule.

A. Entropy estimators

There are two relevant entropies associated with microcanonical ensemble. The energy ceiling or “volume” entropy, S_E^c is defined as

$$S_E^c = \log \Sigma(E), \quad (2)$$

where $\Sigma(E)$ is the number of configurations of the system with energy less than or equal to E . If E' is the energy following E in the annealing schedule (here the energy spectrum, \mathcal{E}) and if $\epsilon(E)$ is the culling fraction at energy E , then a recursive estimator for the ceiling entropy is

$$\tilde{S}_{E'}^c = \tilde{S}_E^c + \log[1 - \epsilon(E)], \quad (3)$$

so that

$$\tilde{S}_E^c = \sum_{E_1 > E} \log[1 - \epsilon(E_1)] + S_\infty^c, \quad (4)$$

where the summation is over all energies in the annealing schedule greater than E and S_∞^c is the logarithm of the total number of system states. The standard definition of entropy in the microcanonical ensemble, S_E , sometimes referred to as

the surface entropy, is given by

$$S_E = \log[\Sigma(E) - \Sigma(E')], \quad (5)$$

$$\approx \log[\Sigma(E)\epsilon(E)],$$

so that the estimator for this entropy is related to the ceiling entropy estimator by

$$\tilde{S}_E = \tilde{S}_E^c + \log \epsilon(E). \quad (6)$$

In the thermodynamic limit the extensive parts of the two entropies are the same. In what follows we refer to \tilde{S}_E as the “microcanonical entropy.”

B. Estimators for microcanonical observables

Suppose A is an observable of the system such as the energy or magnetization. An estimator \tilde{A}_E^c for the equilibrium value of A in the energy ceiling ensemble is

$$\tilde{A}_E^c = \frac{1}{R^*} \sum_{j=1}^{R^*} A^j, \quad (7)$$

where A^j refers to the value of A in the j th replica in the configuration pool for energy ceiling E . We will generally use a superscript c to indicate a quantity in the energy ceiling ensemble.

An estimator \tilde{A}_E of the observable in the microcanonical ensemble at energy E is obtained by averaging over only those replicas in the configuration pool that have energies at the ceiling,

$$\tilde{A}_E = \frac{1}{\epsilon(E)R^*} \sum_{j=1}^{R^*} A^j \delta(E^j, E), \quad (8)$$

where δ denotes Kronecker function and E^j is the energy of replica j .

C. Estimators for canonical observables

The microcanonical entropy can be used to compute an observable in the canonical ensemble from the observable measured in microcanonical annealing. An estimator of the canonical partition function at inverse temperature β_c is given by

$$\tilde{Z}(\beta) = \sum_E e^{-\beta E + \tilde{S}_E}, \quad (9)$$

where the sum is taken over all the energy levels of the system. The canonical energy distribution, $\rho_\beta(E)$, is

$$\rho_\beta(E) = \frac{e^{-\beta E + \tilde{S}_E}}{\tilde{Z}(\beta)}. \quad (10)$$

An estimator $\tilde{A}(\beta)$ of the observable A in the canonical ensemble at inverse temperature β is given by

$$\tilde{A}(\beta) = \sum_E \rho_\beta(E) \tilde{A}_E, \quad (11)$$

and an estimator of the free energy is obtained from the definition

$$\beta \tilde{F} = -\log \tilde{Z}(\beta). \quad (12)$$

For systems that undergo thermal first-order phase transitions, phase coexistence is manifested as a two-peaked canonical energy distribution $\rho_{\beta_c}(E)$ at the transition temperature, β_c . In the thermodynamic limit, the two peaks are δ functions at the energies E_d and E_o of the coexisting disordered and ordered phases, respectively. For an observable A , estimators of the coexisting ordered and disordered values at the transition, \bar{A}_o and \bar{A}_d , are, respectively,

$$\bar{A}_o = \frac{\sum_{E < E_c} \rho_{\beta_c}(E) \tilde{A}_E}{\sum_{E < E_c} \rho_{\beta_c}(E)}, \quad (13a)$$

$$\bar{A}_d = \frac{\sum_{E \geq E_c} \rho_{\beta_c}(E) \tilde{A}_E}{\sum_{E \geq E_c} \rho_{\beta_c}(E)}, \quad (13b)$$

where the breakpoint energy, E_c must be chosen in the range $E_o < E_c < E_d$. In the thermodynamic limit, the breakpoint energy can be chosen arbitrarily in this range. For finite systems, an appropriate choice of E_c will help minimize finite-size corrections. Clearly, E_c should be chosen in the middle of the range to avoid significant overlap with the broadened peaks in the energy distribution.

D. Weighted averages

It has been shown [10,14] that for population annealing in the canonical ensemble, multiple independent runs can be combined using weighted averaging to reduce both statistical and systematic errors. In the case of the canonical ensemble the weight factor is proportional to the exponential of the free-energy estimator for each run. A similar result holds for equilibrium microcanonical annealing algorithms except that the weight factor is here proportional to the exponential of the entropy estimator [11].

Consider a collection of M independent microcanonical annealing runs, each with the same annealing schedule and the same population parameters R and R^* . Let $\tilde{A}_{E,m}$ denote the average value of an observable A in the microcanonical ensemble from the m th run. The best estimate \bar{A}_E from the collection of M runs is given by the weighted average,

$$\bar{A}_E = \sum_{m=1}^M \omega_{E,m} \tilde{A}_{E,m}, \quad (14)$$

where the weights are proportional to the exponential of the entropy estimator $\tilde{S}_{E,m}$ of run m ,

$$\omega_{E,m} = \frac{e^{\tilde{S}_{E,m}}}{\sum_{m'=1}^M e^{\tilde{S}_{E,m'}}}. \quad (15)$$

An analogous result holds for the energy ceiling ensemble,

$$\bar{A}_E^c = \sum_{m=1}^M \omega_{E,m}^c \tilde{A}_{E,m}^c, \quad (16)$$

where

$$\omega_{E,m}^c = \frac{e^{\tilde{S}_{E,m}^c}}{\sum_{m'=1}^M e^{\tilde{S}_{E,m'}^c}}. \quad (17)$$

For either ensemble, if a run fails at energy E , then for all energies $E' \geq E$ the corresponding weight vanishes.

It is most straightforward to first derive the weighted average formula for the energy ceiling result and then the microcanonical result. Consider the simple situation that at annealing step E the weights for all runs are the same, $\omega_{E,m}^c \equiv 1/M$. At the next annealing step with energy ceiling E' , the population to be averaged over in run m has a size $RR^*[1 - \epsilon_m(E)]$, where $\epsilon_m(E)$ is the culling fraction at energy ceiling E in run m . An unbiased average in the energy ceiling ensemble over many runs at energy ceiling E' should weight each run according to the size of its population, which requires that run m be weighted by a factor proportional to $[1 - \epsilon_m(E)]$. From Eq. (3) we have that $[1 - \epsilon_m(E)]$ is the exponential of the volume entropy change from E to E' .

In the more general case where runs already have differing weights at energy ceiling E we have the recursion relation,

$$\omega_{E',m}^c \propto [1 - \epsilon_m(E)] \omega_{E,m}^c, \quad (18)$$

where the constant of proportionality is independent of run and is set by the normalization of the weights. From Eq. (3) we see that $[1 - \epsilon_m(E)]$ is the exponential of the volume entropy change from E to E' . Collapsing the telescoping product of entropy changes yields the weight factor given in Eq. (17).

To obtain the analogous result for the microcanonical ensemble, note that the population that is averaged in the microcanonical ensemble is a factor $\epsilon_m(E)$ smaller than the population averaged in the ceiling ensemble. Thus, the microcanonical ensemble and ceiling ensemble weights are related by

$$\omega_{E,m} \propto \epsilon_m(E) \omega_{E,m}^c, \quad (19)$$

and from Eq. (6) we see that this transformation is precisely the transformation from \tilde{S}^c to \tilde{S} , thus verifying Eq. (15).

Formulas for the weighted average of the two entropies are slightly more complicated to derive because the entropy depends on a thermodynamic integration over all of the annealing steps, Eq. (4). Thus, the weighted average volume entropy can be written as

$$\bar{S}_E^c = \log \prod_{E_1 > E} \sum_{m=1}^M [1 - \epsilon_m(E_1)] \omega_{E_1,m}^c + \tilde{S}_\infty^c. \quad (20)$$

Expanding the definition Eq. (17) of the weight factor and using Eq. (3) to relate the culling factor to the change in entropy yields the telescoping product,

$$\bar{S}_E^c = \log \prod_{E_1 > E} \frac{\sum_{m=1}^M \exp(\tilde{S}_{E_1,m}^c)}{\sum_{m=1}^M \exp(\tilde{S}_{E_1,m}^c)} + \tilde{S}_\infty^c, \quad (21)$$

where E_1' is the successor to E_1 in the annealing schedule. After canceling terms in the numerator and denominator we have

$$\bar{S}_E^c = \log \frac{1}{M} \sum_{m=1}^M e^{\tilde{S}_{E,m}^c}, \quad (22)$$

and, analogously for the microcanonical entropy, we have

$$\bar{S}_E = \log \frac{1}{M} \sum_{m=1}^M e^{\tilde{S}_{E,m}}. \quad (23)$$

We emphasize that for fixed R and R^* , results from weighted averaging are exact in the limit $M \rightarrow \infty$ for any fixed number of MC sweeps per run.

Weighted averaging in the canonical ensemble is discussed in Refs. [10,14]. Given M independent simulations in the canonical ensemble, the best estimate $\bar{A}(\beta)$ of an observable A in the canonical ensemble at inverse temperature β is given by

$$\bar{A}(\beta) = \frac{\sum_{m=1}^M \tilde{A}_m(\beta) e^{-\beta \tilde{F}_m}}{\sum_{m=1}^M e^{-\beta \tilde{F}_m}}, \quad (24)$$

where $\beta \tilde{F}_m = -\log \tilde{Z}_m(\beta)$ is the free-energy estimator of the m th run. In analogy to Eq. (23), the best estimator of the free energy from multiple runs is given by

$$-\beta \bar{F} = \log \frac{1}{M} \sum_{m=1}^M e^{-\beta \tilde{F}_m}. \quad (25)$$

Since our simulations directly produce results in the microcanonical ensemble, it is important to show that transforming from the microcanonical to the canonical ensemble commutes with carrying out a weighted average in either ensemble. Specifically, we would like to show that

$$\bar{A}(\beta) = \frac{\sum_E \bar{A}_E e^{-\beta E + \bar{S}_E}}{\sum_E e^{-\beta E + \bar{S}_E}} \quad (26)$$

and

$$-\beta \bar{F} = \log \sum_E e^{-\beta E + \bar{S}_E}, \quad (27)$$

where the canonical weighted averages, $\bar{A}(\beta)$ and $\beta \bar{F}$, are defined in Eqs. (24) and (25), respectively, and where on the right-hand side (RHS) the barred quantities are microcanonical weighted averages defined in, Eqs. (14), (15), and (23). To verify Eq. (27), simply substitute the definition Eq. (23) of \bar{S}_E on the RHS, exchange the order of summation and then invoke the definitions, Eqs. (9) and (12). A similar but slightly more involved calculation verifies Eq. (26). Finally, we note that the weighted averages for observables, A_o and A_d , defined in Eqs. (13) are obtained from Eq. (26) with the summation over a restricted range of energies, $E < E_c$ and $E \geq E_c$, respectively. Alternatively, results analogous to Eq. (24) for A_o and A_d are obtained by replacing \tilde{F} by \tilde{F}_o and \tilde{F}_d , respectively, where these partial free energies are defined in the obvious way by limiting the range of summation in the partition function, Eq. (9).

E. Systematic errors

It is possible to quantify systematic errors in microcanonical annealing algorithms from the run-to-run variance of the entropy estimator. This result follows from the fact that weighted averages are exact in the limit $M \rightarrow \infty$. Systematic errors are then given by the difference between weighted and unweighted averages in this limit. The result, derived for the case of population annealing ($R \gg 1$ and $R^* = R$) in Refs. [11,14], is that the systematic error ΔA_E in the estimator

\bar{A}_E is given by the covariance of \bar{A}_E and \bar{S} ,

$$\Delta A_E = \text{cov}(\bar{A}_E, \bar{S}) = \text{var}(\bar{S}) \left[\frac{\text{cov}(\bar{A}_E, \bar{S})}{\text{var}(\bar{S})} \right]. \quad (28)$$

The second, trivial identity in Eq. (28) is useful because the ratio in the square brackets goes to a constant that depends on the observable A as the simulation size, R^* becomes large, while $\text{var}(\bar{S})$ is expected vanish as $1/R^*$ since the number of measurements going into estimating \bar{S} , also becomes large. Note that the scaling of the simulation size depends on the type of annealing algorithm: in equilibrium simulated annealing R^* becomes large with $R = 1$ and in population annealing, $R^* = R$, while various options are possible for hybrid annealing. In any case, systematic errors and their scaling with simulation size are characterized by the behavior of $\text{var}(\bar{S})$. Although this result was originally derived for population annealing, it holds for the whole class of equilibrium microcanonical annealing protocols discussed here since it depends on only the validity of weighted averaging. An analogous result holds for quantities in the energy ceiling ensemble.

The systematic error of an observable in the canonical ensemble, $\Delta A(\beta)$, is given by

$$\Delta A(\beta) = \text{cov}[\bar{A}(\beta), \beta \bar{F}] = \text{var}(\beta \bar{F}) \left[\frac{\text{cov}(\bar{A}(\beta), \beta \bar{F})}{\text{var}(\beta \bar{F})} \right]. \quad (29)$$

After transforming from the microcanonical to the canonical ensemble, this result can be applied to quantify errors in canonical observables obtained from microcanonical simulations. Results for systematic errors for A_o and A_d can be obtained from the variance of the partial free energies \tilde{F}_o and \tilde{F}_d , respectively, discussed at the end of Sec. IID.

The variance of entropy and free-energy estimators are therefore a useful tool for comparing the performance of different versions of equilibrium microcanonical annealing.

F. Relation to other annealing algorithms

The case $R = 1$ and $R^* \gg 1$ is a form of simulated annealing. Simulated annealing is conventionally carried out in the canonical ensemble with the last configuration generated by the MCMC procedure used as the starting point for the next annealing step. Thus, conventional simulated annealing is the case $R = R^* = 1$. If this conventional choice were applied in the microcanonical ensemble, then the algorithm would fail with probability ϵ at each annealing step and would be highly inefficient. The case $R \gg 1$ and $R^* = R$ is microcanonical population annealing, which has previously been applied to study hard sphere systems [11]. We refer to the intermediate cases $1 < R < R^*$ as hybrid annealing. In what follows we abbreviate population annealing, simulated annealing, and hybrid annealing as PA, SA, and HA, respectively. For convenience, brief descriptions of the three algorithms are included in Table I.

One of our goals is to understand the tradeoffs related to the population size R . One conclusion that seems evident is that a single large run of a hybrid algorithm with parameters $R = R_h$ and $R^* = R_h^*$ should outperform the weighted average of $M = R_h$ runs of simulated annealing with $R^* = R_h^*/M$ even

TABLE I. Descriptions of equilibrium simulated annealing (SA), hybrid annealing (HA), and population annealing (PA).

Algorithm	Description
SA	Single replica and many MCMC sweeps
HA	Intermediate number of replicas and MCMC sweeps
PA	Large number of replicas and few MCMC sweeps

though both simulations require nearly the same amount of computational work. The argument is that resampling from the larger pool in the hybrid algorithm will be more effective than simply reweighting independent runs. A related result was observed in comparing standard canonical simulated annealing with canonical population annealing for finding ground states of spin glasses [13].

III. MODEL AND OBSERVABLES

A. The q -state Potts model

The q -state Potts model [22] is a generalization of the Ising model consisting of N interacting spins $\{s_i | i = 1, \dots, N\}$ each taking values $s_i \in \{1, \dots, q\}$. The energy, E is given by

$$E = - \sum_{(i,j)} \delta(s_i, s_j), \quad (30)$$

where $\delta(s_i, s_j)$ is the Kronecker δ function and the summation is over all the interacting pairs of spins. For spins on a square (two-dimensional) lattice with nearest neighbor interactions and for $q > 4$, the Potts model undergoes a thermal first-order phase transition at inverse temperature $\beta_c = \log(1 + \sqrt{q})$. Higher q values result in stronger first-order phase transitions. In our simulations we consider the two-dimensional, 20-state Potts model on an $L \times L$ square lattice with periodic boundary conditions.

The first-order phase transition of the Potts model in the canonical ensemble is between a disordered phase with magnetization zero and energy per spin e_d and a magnetically ordered phase with magnetization per spin m_o and energy per spin e_o . Configurations with energy per spin between these values possess an interface between ordered and disordered domains, leading to a free-energy barrier proportional to the system size L .

In the microcanonical ensemble configurations with an interface are stable for energies per spin between e_d and e_o making it possible for an annealing algorithm to traverse the coexistence region nearly continuously. Nonetheless, there is a series of four discontinuous transitions in the micro-canonical ensemble, which introduce barriers between the disordered and ordered phases and increase equilibration times. First, at a length-scale dependent energy, $e_1(L) < e_d$ there is a transition corresponding to the condensation of an ordered droplet within a disordered background [3,5,6]. As the system size increases, $e_1(L)$ approaches e_d . As energy decreases, this droplet grows continuously until the system reaches energy e_2 where it can decrease its surface area by a discontinuous transition to a stripe configuration. Decreasing the energy further results in the widening of the ordered stripe until at energy e_3 there is another discontinuous transition




FIG. 1. Equilibrium configurations of the 20-state Potts model with $L = 40$ and periodic boundary conditions. From left to right and then bottom to top configurations are shown in the (1) homogeneous disordered phase, (2) disordered phase with an ordered droplet, (3) stripe phase, (4) ordered phase with disordered droplet, and (5) homogeneous ordered phase.

to a configuration with a disordered droplet in an ordered background [7]. Finally, for $e_4(L) > e_o$, there is an evaporation transition where the disordered droplet vanishes, leaving the homogeneous ordered phase. Mirroring $e_1(L)$, $e_4(L)$ depends on system size and approaches e_o from above. These transitions are illustrated in the sequence of snapshots shown in Fig. 1.

B. Observables

In this section we define the observables measured in our simulations. First, we measure the entropy using Eq. (6) and use it to calculate the free energy and the canonical energy distribution of Eq. (10) at the transition point β_c . From these quantities we calculate the following observables.

1. Energies of ordered and disordered phases at coexistence

Using the canonical energy distribution and Eqs. (13) we measure the ordered and disordered phase energies per spin, e_o and e_d , respectively.

2. Peak ratio and disordered phase excess

We measure two observables related to the relative weights of the ordered and disordered peaks in the energy distribution at β_c , the peak ratio and the disordered phase excess, both of which are defined below.

The peak ratio r_c at the transition temperature is defined as

$$r_c = \frac{\sum_{E < E_c} \rho_{\beta_c}(E)}{\sum_{E \geq E_c} \rho_{\beta_c}(E)}. \quad (31)$$

The peak ratio is related to the difference in the free energies of the ordered and disordered phases, which in the thermodynamic limit are equal at the transition temperature. However, since there are q different ordered phases, the exact value of the peak ratio in the thermodynamic limit is q at the transition temperature.

Since the peak ratio is not an observable of the form of Eq. (11), the weighted average cannot be computed simply from Eq. (24). We define another observable related to the peak ratio called the disordered phase excess X ,

$$X = \sum_{E > E_c} \rho_{\beta_c}(E) - \frac{1}{q+1}, \quad (32)$$

where the subtracted constant is chosen so the exact value of X in the thermodynamic limit is zero. The disordered phase excess is used as a measure of accuracy of a simulation at the transition point. Unlike the peak ratio, a weighted average of X using Eq. (24) gives the best estimate of X and is exact in the limit $M \rightarrow \infty$.

3. Wrapping fraction

To estimate the locations e_2 and e_3 of the droplet/stripe transitions we define the wrapping fraction, ω_w , as the average number of directions wrapped by connected paths of like spins. The number of wrapping directions takes values in $\{0, 1, 2\}$; in the disordered phase and the droplet state the wrapping number is zero, in the stripe state the wrapping number is one, while in the ordered phase the wrapping number is two.

4. Droplet fraction

To detect the evaporation transition between a configuration containing a disordered droplet and a homogeneous ordered phase, we consider the disordered cluster size histogram h_c , where here a disordered cluster is defined using the following procedure. First, identify the largest connected cluster of adjacent spins having the same value. Near the evaporation transition point this cluster will wrap both vertically and horizontally. Label all spins in this cluster “0” and label all other spins “1.” Group adjacent spins labeled “1” into connected clusters, and let h_c be the configuration averaged histogram of the sizes of these disordered clusters, where size is defined as the number of sites in the cluster.

The disordered cluster size distribution is shown in Fig. 6 for three energies. At energies lower than the evaporation transition point $e_4(L)$, all clusters consist of excitations within a homogeneous ordered phase, and h_c decays exponentially. At the transition point h_c is observed to be a power law distribution as the disordered droplet breaks down into smaller

disordered regions. At energies greater than the transition point, a finite fraction of spins exist within a disordered droplet, and thus, in addition to a power law, h_c contains a peak at some value greater than zero.

We define the droplet fraction, ω_c , as the probability that a configuration contains a large droplet of disordered phase. We use the following procedure to measure ω_c . First, a cutoff to the power-law at size of 35 was visually determined from h_c . Next, a least squares power law fit of h_c is made up to this cutoff, as shown in Fig. 6. After subtracting this power law from h_c , what remains is a single peak centered on the characteristic droplet size for the given energy and ω_c is obtained by numerically integrating this peak.

We expect ω_c to be one for energies sufficiently above the transition point. In other words, we expect every configuration to contain one droplet that is significantly larger than the droplet size distribution predicted by the power law. As the energy approaches e_4 from above, this peak will decrease in size and ω_c will drop to zero as the power law behavior of h_c crosses over to exponential decay. Below e_4 the power-law range of h_c disappears so ω_c is only well-defined near and above the evaporation transition.

5. Magnetization integrated autocorrelation time

The magnetization per spin, m for the Potts model is defined as

$$m = \frac{q(N_{\max}/N) - 1}{q - 1}, \quad (33)$$

where N_{\max} is the maximum number of spins taking the same value. In the disordered phase $N_{\max} \approx N/q$ and the magnetization vanishes, while in the ordered phase $N_{\max} > N/q$. As the temperature approaches zero, $N_{\max} \rightarrow N$ and the magnetization per spin approaches one.

The integrated autocorrelation time τ of the magnetization is defined by

$$\tau = \frac{1}{2} + \frac{\sum_{t=1}^{\infty} (\langle m(t_0)m(t_0+t) \rangle - \langle m \rangle^2)}{\langle m^2 \rangle - \langle m \rangle^2}, \quad (34)$$

where time is measured in Monte Carlo sweeps and the brackets refer either to an ensemble average or an average over starting times t_0 . The integrated autocorrelation time sets the time between independent measurements. Although the asymptotic decay of the correlation functions, i.e. the exponential autocorrelation time, is a better measure of the time to reach equilibrium, τ is usually also a good approximate measure of the equilibration time or, equivalently, the computational work required to equilibrate the system at a given energy.

The time series used to compute τ must be much larger than τ itself and the upper limit of the sum in Eq. (34) must be finite. Here we use an iterative method described in [23] that minimizes systematic errors and is used to set the upper limit of the sum defining τ .

It is important to note that the autocorrelation time measured in the coexistence region beyond the condensation transition is the timescale for generating independent samples that have the same ordered phase (droplet, stripe, or region). There

TABLE II. Simulation parameters for SA, HA, and PA. Values in the table are system size L , population size R , the parameter a_s , setting the number of MCMC sweeps [see Eq. (35)], number of runs M used for weighted averaging, and ratio of pool size to population size, R^*/R .

Algorithm	R	L	a_s	R^*/R	M
SA	1	30	1.2×10^6	1.2×10^6	30
		40	3.8×10^5	3.8×10^5	40
		50	1.6×10^5	1.6×10^5	50
		60	7.5×10^4	7.5×10^4	60
HA	100	30	1.2×10^4	1.2×10^3	30
		40	3.8×10^3	3.8×10^2	40
		50	1.6×10^3	1.6×10^2	50
		60	7.5×10^2	7.5×10	60
PA	8×10^5	30	2	1	30
	2×10^5	40	1.9	1	40
	10^5	50	1.6	1	50
	7.5×10^4	60	1	1	60

is a much longer timescale, not relevant for our measurements, to make transitions between different ordered phases.

IV. SIMULATION DETAILS

We perform SA, HA, and PA simulations for the 2D 20-state Potts model for systems sizes $L = 30, 40, 50,$ and 60 . For each size, the computational work, measured in sweeps, is kept the same for each algorithm. In each case, we obtain results by using weighted averaging over M simulation runs. The number of MC sweeps performed at each energy is equal to n_s given by

$$n_s(E) = \begin{cases} a_s, & \text{if } E > -\frac{N}{2} \\ 20a_s, & \text{if } -\frac{N}{2} \geq E > -\frac{3N}{2} \\ 5a_s, & \text{if } -\frac{3N}{2} > E, \end{cases} \quad (35)$$

where a_s is a parameter that depends on the algorithm and is chosen to make the computational work roughly equal between algorithms for the same system size.

This *ad hoc* sweep schedule concentrates the majority of MC sweeps in the range of energies where the integrated autocorrelation time is highest. We have found that this sweep schedule is an improvement compared to a uniform sweep schedule, but is not claimed to be optimal.

Simulation parameters for the three annealing algorithms are presented in Table II. Since both the number of proposed updates per MC sweep and the number of annealing steps in the simulation grow linearly with the system size N , the work per simulation grows quadratically for a given sweep parameter a_s . To keep the run times manageable, fewer MC sweeps (a_s decreased) are performed per annealing step for larger system sizes. To partially compensate for the decrease in the number of sweeps, reweighted averages are performed over more runs (M increased) for larger system sizes.

It is straightforward to take advantage of the massively parallel nature of the HA and PA algorithms by performing MC sweeps on independent replicas simultaneously. As the majority of computational work in these simulations is spent

performing MC sweeps and only a small fraction is spent during the resampling step, significant improvement in wall clock time can be gained for PA and HA through parallelization despite doing the resampling sequentially.

For PA the speedup due to parallelization is most dramatic, because the number of independent replicas is on the order of 10^5 . Population annealing simulations were implemented on a single NVIDIA Tesla C2075 GPU using the CUDA library using the methods described in Ref. [24]. One thread was assigned to each replica so that the total number of threads was equal to the population size. Despite the simplicity of this implementation, we still observed a notable speed-up with PA when compared to SA with the same number of sweeps running on a CPU. For instance, for our $L = 50$ simulations, the PA runs were about 85 times faster than SA runs.

For our HA simulations the population size of 100 was not large enough to benefit from using a GPU though a significant speedup over SA was obtained using OpenMP with 16 threads.

V. RESULTS

In this section we describe the results of our simulations. Our primary objective is comparing the three annealing algorithms but we also obtain precise finite-size results for the Potts model observables described above.

To obtain results in the canonical ensemble at the transition temperature, the microcanonical entropy is measured at every energy using each of the microcanonical annealing algorithms and the results are reweighted to obtain the free energy at the transition using Eqs. (9) and (12). The peak ratio observables, r_c and X are measured using Eqs. (13) and (32), respectively. Ordered and disordered phase energies per spin, e_o and e_d , are calculated from the canonical energy histogram, Eq. (10) using Eqs. (13). For the system sizes considered, the ordered and disordered peaks in the energy distribution at coexistence, shown in Fig. 2, are sufficiently separated that any ambiguity in defining breakpoint energy, E_c between ordered and disordered phases is negligible, and in practice E_c was chosen to be the point of minimum weight. The reported values of all observables are obtained from a weighted average over M




FIG. 2. Distribution $\rho_\beta(e)$ of the energy per spin, e at the transition temperature calculated with population annealing for system sizes $L = 30, 40, 50,$ and 60 (from top to bottom).

TABLE III. Energies of the disordered phase, e_d , ordered phase, e_o , the peak ratio, r_c , disordered phase excess X , and the variance of the free energy, $\text{var}(\beta F)$, at the transition temperature. Results are shown for the SA, HA, and PA simulations, and exact values in the thermodynamic limit [26].

L	Algorithm	e_d	e_o	r_c	X	$\text{var}(\beta F)$
30	SA	-0.626551(25)	-1.820723(12)	19.965(97)	$7.9 \times 10^{-5} \pm 2.2 \times 10^{-4}$	0.0013
	HA	-0.62678(73)	-1.81966(90)	18.1(1.6)	$4.9 \times 10^{-3} \pm 8.1 \times 10^{-3}$	0.38
	PA	-0.626603(48)	-1.820732(49)	20.17(29)	$-3.9 \times 10^{-4} \pm 6.6 \times 10^{-4}$	0.012
40	SA	-0.626514(34)	-1.820632(17)	19.84(15)	$3.6 \times 10^{-4} \pm 3.5 \times 10^{-4}$	0.0049
	HA	-0.62619(49)	-1.82017(32)	15.0(1.0)	$1.5 \times 10^{-2} \pm 7.1 \times 10^{-3}$	0.41
	PA	-0.62661(11)	1.82050(11)	19.39(69)	$1.4 \times 10^{-3} \pm 1.8 \times 10^{-3}$	0.064
50	SA	-0.626525(41)	-1.820649(20)	19.45(31)	$1.3 \times 10^{-3} \pm 7.6 \times 10^{-4}$	0.072
	HA	-0.62689(19)	-1.81991(30)	15.1(1.1)	$1.4 \times 10^{-2} \pm 7.8 \times 10^{-3}$	0.44
	PA	-0.62663(12)	-1.82072(15)	21.5(1.2)	$-3.5 \times 10^{-4} \pm 4.2 \times 10^{-3}$	0.25
60	SA	-0.626537(41)	-1.820753(31)	19.62(46)	$8.7 \times 10^{-4} \pm 1.1 \times 10^{-3}$	0.096
	HA	-0.62652(18)	-1.82050(36)	20.31(91)	$-7.0 \times 10^{-4} \pm 5.8 \times 10^{-3}$	0.70
	PA	-0.62660(18)	-1.82081(34)	22.93(83)	$-5.8 \times 10^{-3} \pm 1.7 \times 10^{-2}$	1.3
∞ (exact)	—	-0.626529...	-1.820684...	20	0	—

independent simulations as described in Sec. IID and error bars are obtained by bootstrapping over these runs.

Table III shows results from the three algorithms and four system sizes for e_d , e_o , r_c , X , and $\text{var}(\beta F)$. The results for SA are accurate to four or five significant digits and the results for the different algorithms are in agreement with one another within error bars. The SA results for size $L = 40$ are consistent with and more accurate than the results reported in Ref. [25].

Based on both the error bars, which quantify statistical errors, and on $\text{var}(\beta F)$, which quantifies systematic errors [see Eq. (29)], we see that SA yields significantly more accurate and well-equilibrated results than either PA or HA. Furthermore, except for size $L = 60$, PA is better than HA. The relative ranking of the three algorithms is also visible in Figs. 3 and 4, which are scatter plots of e_d and X , respectively, versus the free energy per spin, βf , for size $L = 40$. Each point represents one of the M independent runs. The larger the

variance of the value of the observable, the larger statistical errors for the observable while a large covariance with βf indicates a large systematic error according to Eq. (29).

While measuring observables at the transition can be done relatively efficiently using the microcanonical ensemble, the simulations become more difficult with increasing system size. It is believed that the major sources of hardness in these simulations are the four microcanonical phase transitions described in Sec. III A. To investigate this we used the integrated autocorrelation time of the magnetization τ as a measure of computational hardness.

We plot τ in Fig. 5 as function of the energy ceiling for the four system sizes considered. Indeed, τ displays sharp peaks that grow with system size. While only three peaks are visible for the system sizes considered, we believe that the largest, rightmost peak is associated with both the condensation transition and the first wrapping transition. For larger system sizes we expect that this peak would break into two distinct peaks.




FIG. 3. Scatter plot of the dimensionless free energy per spin, βf and the disordered phase energy per spin e_d for each of $M = 40$ runs for SA (dark blue triangles), HA (light green squares), and PA (red circles) at the transition temperature for system size $L = 40$.




FIG. 4. Scatter plot of the dimensionless free energy per spin, βf and disordered phase excess, X for each of $M = 40$ runs for SA (dark blue triangles), HA (light green squares), and PA (red circles) at the transition temperature for $L = 40$. The exact value in thermodynamic limit, $X = 0$, corresponds to peak ratio $r_c = q$.




FIG. 5. Integrated autocorrelation time τ in units of Monte Carlo sweeps, as a function of energy ceiling per spin e for system sizes $L = 30, 40, 50,$ and 60 , with peak heights increasing with L .

The evaporation transition is studied using methods described in Sec. III B. The disordered cluster size histogram h_c is plotted in Fig. 6 at three energies for size $L = 40$. As described in Sec. III B we observe a power law with an additional distinct peak for energies above the evaporation transition, power law behavior at the transition, and exponential decay below. The integral of the peak above the fitted power law in Fig. 6 gives a measurement of the ω_c . The power law behavior of h_c at the evaporation transition is unexpected and warrants further investigation.

In Fig. 7 we show the magnetization integrated autocorrelation time, τ , the wrapping fraction, ω_w , and droplet fraction, ω_c , for $L = 40$. The wrapping fraction is zero at high energies and displays two jumps as expected in the coexistence region




FIG. 6. Disordered cluster size histograms, h_c for ceiling energies per spin $e = -1.56$ (light green), -1.73 (red), and -1.84 (dark blue) showing the behavior of h_c above, near, and below the evaporation transition energy e_4 , respectively, for size system $L = 40$. Dashed lines represent fitted power laws to the histograms for cluster sizes below a cutoff of 35.




FIG. 7. The integrated autocorrelation time τ (red, left axis) is plotted against ceiling energy per spin e over the entire simulation energy range $-2 < e < 0$ for $L = 40$. The wrapping fraction ω_w (dark blue, right axis) is plotted from $-1.6 < e < .7$ and the droplet fraction ω_c (light green, right axis) is plotted from $-1.85 < e < -1.5$. The first two peaks in τ , from right to left, coincide with jumps in ω_w that signify the first and second wrapping transitions, respectively. The third peak in τ coincides with a rapid drop of ω_c signifying the evaporation transition.

$e_o < e < e_d$. These jumps coincide with the two rightmost peaks in τ , suggesting that they are in fact points where increased computational effort is needed.

We observe that $\omega_c \rightarrow 1$ for energies far above the condensation transition, and drops quickly to zero below the transition at energy e_4 . The point where ω_c begins to rapidly decrease marks the point where a droplet of disordered phase becomes unstable, which coincides in Fig. 7 with the final peak in τ .

The growth of the the autocorrelation time due to the wrapping and condensation/evaporation transitions poses the greatest challenge to any microcanonical algorithm including microcanonical annealing. These transitions limit the system sizes that can be feasibly studied and ultimately will lead to exponential growth in the resources required to perform simulations. However, in Fig. 5 we observe only a modest increase in τ with increasing system size. Compare this to the canonical energy histogram in Fig. 2, where the depth of the valley between ordered and disordered grows exponentially for the same sizes considered. The depth of this valley provides an estimate for the timescale needed by a canonical MCMC algorithm to tunnel between an ordered phase and the coexisting disordered phase. This suggests that for modest system sizes microcanonical annealing is not only more efficient than canonical MCMC but also achieves better scaling until exponential slowing eventually dominates.

VI. DISCUSSION

We have introduced a class of equilibrium annealing algorithms that includes population annealing and an

equilibrium version of simulated annealing. We have implemented algorithms in this class in the microcanonical ensemble and applied them to the 20-state Potts model, which displays a strong first-order transition. In agreement with previous work, we find that simulating a thermal first-order transition can be effectively carried out in the microcanonical ensemble yielding high precision results. In these applications of microcanonical annealing, the purpose of annealing is to collect data at all energies in the coexistence region to obtain canonical averages at the phase transition temperature using reweighting.

Our study of the evaporation transition showed a power-law distribution of sizes of disordered droplets. This is a surprising finding and contradicts results for the evaporation/condensation transitions in the Ising model as the magnetization is varied at fixed temperature [5]. This power-law behavior deserves further study.

We compared the performance of equilibrium simulated annealing, population annealing, and a hybrid algorithm that interpolates between the two. For the system sizes and number of sweeps considered here, equilibrium simulated annealing was found to be the most efficient algorithm followed by population annealing. The hybrid algorithm performed worst except for the largest size ($L = 60$) where it slightly outperformed population annealing, though for this size neither algorithm equilibrated the system with the allotted number of Monte Carlo sweeps. For future work, it would be interesting to compare the performance of multi-canonical methods to the annealing methods introduced here.

We conjecture that SA outperforms PA and HA because it takes advantage of the exponential convergence to equilibrium of Markov chain Monte Carlo in contrast to the $1/R$ convergence to equilibrium of sequential Monte Carlo. Note that for $L = 40$ the magnetization integrated autocorrelation time τ is bounded by $\tau \leq 2000$ while the number of sweeps per annealing step in this region is 7.6×10^6 so there are more than $10^3 \tau$ sweeps at every energy and, assuming the rate of convergence to equilibrium is approximately τ , systematic errors will be extremely small. By contrast, in PA the number of sweeps per annealing step carried out in the transition region is much less than τ and equilibration is achieved primarily by resampling, which converges only as $1/R$. The hybrid algorithm suffers from both an inadequate population size for effective resampling and too few Monte Carlo sweeps.

A toy model illustrates the above intuition. Suppose that the distance from equilibrium, Δ behaves as $\Delta = (1/R)e^{-t/\tau}$ where R is the population size, t the number of MCMC sweeps for each replica at each temperature and τ is the autocorrelation time of the MCMC, which is here assumed to be constant. We allot each algorithm the same total number of sweeps per annealing step so that the scaled number sweeps, c , defined as $c = Rt/\tau$ is held fixed. Note that c is the number of autocorrelation times for each annealing step in simulated annealing. In terms of R and c we seek to minimize $(1/R)e^{-c/R}$ for fixed c subject to the constraints that the population size is at least one, $R \geq 1$, and that at least one MCMC sweep is performed on each replica at each annealing step, $t \geq 1$. The latter constraint is required to ensure that some exploration and decorrelation is carried out at each annealing step. In

terms of population size, this constraint can be written as $R \leq c\tau$. It is straightforward to show that there are two minima for Δ and they occur at the two endpoints, SA ($R = 1$), and PA ($R = c\tau$). At the SA endpoint $\Delta = e^{-c}$ while at the PA endpoint $\Delta \approx 1/c\tau$. In addition there is a maximum at $R = c$ where $\Delta = (1/ec)$. While this toy model cannot be expected to provide quantitative results, it is expected to yield qualitative comparisons between algorithms.

The first conclusion from the toy model is that if the autocorrelation time is sufficiently short that t/τ can be made large, then it is better to use SA. A second conclusion is that it does not pay to compromise between many MCMC sweeps and a large population. In our case c , at least in the difficult coexistence region, is approximately 10^3 so that $\Delta_{SA} \ll \Delta_{PA}$ and we expect that SA is the preferred algorithm. The hybrid algorithm with $R = 10^2$ is not far from the worst choice, $R = c$ while PA is closer to the second optimum at $R = c\tau$ so we expect that for our choices of parameters HA would perform worse than PA.

We can conclude from both the simulation results and the simple toy model that SA is the best of the microcanonical annealing methods for high precision studies of the transition in large- q Potts models and, likely, other systems with thermal first-order transitions and simple symmetry breaking. However, for systems that display a rough free-energy landscape without obvious symmetries, we expect PA or HA to outperform SA. In each run of SA for the q -state Potts model only one of the q low temperature phases is discovered because the barriers between phases are too high. This deficiency does not affect any of the observables we measured because of the symmetry between the q low temperature phases. The situation is different in the case of multiple inequivalent phases or a glassy phase with a rough free-energy landscape. To apply SA to a system with a free-energy landscape with multiple inequivalent minima separated by high barriers it is necessary to combine many runs of SA using weighted averaging. Good statistics requires a large number of runs so that each important minimum is found in many runs. We conjecture that it is better to do a single run of HA or PA with the same annealing schedule as for SA and with $R = M$ replicas where M is the number of runs used for SA. For HA or PA, the resampling step distributes the work assigned to each minimum according to the weight of that minimum so that the more important minima are more accurately sampled and minima that are irrelevant at the target low temperature are removed from the population. Evidence supporting this conjecture can be found in Ref. [13], where it was shown that PA with population size R is far more likely to find ground states of Ising spin glasses than $M = R$ runs of SA with the same annealing schedule.

The microcanonical annealing algorithms discussed here may also be useful for sampling equilibrium systems with rough free-energy landscapes and solving hard combinatorial optimization problems in situations where behavior similar to first-order transitions occurs. For example, spin glasses display temperature chaos [27–29], which is the phenomena that typical states of the system changes discontinuously with temperature. It would be interesting to see whether microcanonical population annealing performs better than canonical population annealing for systems with temperature chaos.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation (Grant No. DMR-1507506). We thank Chris

Amey, Wolfhard Janke, Helmut Katzgraber, Roman Kotecký, Gili Rosenberg, Lev Shchur, and Martin Weigel for useful discussions.

-
- [1] B. Berg and T. Neuhaus, Multicanonical Ensemble: A New Approach to Simulate First-Order Phase Transitions, *Phys. Rev. Lett.* **68**, 9 (1992).
- [2] F. Wang and D. P. Landau, Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States, *Phys. Rev. Lett.* **86**, 2050 (2001).
- [3] T. Nogawa, N. Ito, and H. Watanabe, Evaporation-condensation transition of the two-dimensional Potts model in the microcanonical ensemble, *Phys. Rev. E* **84**, 061107 (2011).
- [4] V. Martin-Mayor, Microcanonical Approach to the Simulation of First-Order Phase Transitions, *Phys. Rev. Lett.* **98**, 137207 (2007).
- [5] M. Biskup, L. Chayes, and R. Kotecký, On the formation/dissolution of equilibrium droplets, *Europhys. Lett.* **60**, 21 (2002).
- [6] P. Schierz, J. Zierenberg, and W. Janke, First-order phase transitions in the real microcanonical ensemble, *Phys. Rev. E* **94**, 021301 (2016).
- [7] K. Leung and R. K. P. Zia, Geometrically induced transitions between equilibrium crystal shapes, *J. Phys. A: Math. Gen.* **23**, 4593 (1990).
- [8] C. D. Gelatt S. Kirkpatrick and M. P. Vecchi, Optimization by simulated annealing, *Science* **220**, 671 (1983).
- [9] K. Hukushima and Y. Iba, Population annealing and its application to a spin glass, in *The Monte Carlo Method In The Physical Sciences: Celebrating the 50th Anniversary of the Metropolis Algorithm*, Vol. 690, edited by J. E. Gubernatis (AIP, Melville, NY, 2003), pp. 200–206.
- [10] J. Machta, Population annealing with weighted averages: A Monte Carlo method for rough free-energy landscapes, *Phys. Rev. E* **82**, 026704 (2010).
- [11] J. Callahan and J. Machta, Population annealing simulations of a binary hard-sphere mixture, *Phys. Rev. E* **95**, 063315 (2017).
- [12] P. Kar, W. Nadler, and U. H. E. Hansmann, Microcanonical replica exchange molecular dynamics simulation of proteins, *Phys. Rev. E* **80**, 056703 (2009).
- [13] W. Wang, J. Machta, and H. G. Katzgraber, Comparing Monte Carlo methods for finding ground states of Ising spin glasses: Population annealing, simulated annealing, and parallel tempering, *Phys. Rev. E* **92**, 013303 (2015).
- [14] W. Wang, J. Machta, and H. G. Katzgraber, Population annealing: Theory and application in spin glasses, *Phys. Rev. E* **92**, 063307 (2015).
- [15] C. Amey and J. Machta, Analysis and optimization of population annealing, *Phys. Rev. E* **97**, 033301 (2018).
- [16] L. Barash, J. Marshall, M. Weigel, and I. Hen, Estimating the density of states of frustrated spin systems, *New J. Phys.* **21**, 073065 (2019).
- [17] A. Doucet, N. de Freitas, and N. Gordon (eds.), *Sequential Monte Carlo Methods in Practice* (Springer-Verlag, New York, 2001).
- [18] W. Wang, J. Machta, H. Munoz-Bauza, and H. G. Katzgraber, Number of thermodynamic states in the three-dimensional Edwards-Anderson spin glass, *Phys. Rev. B* **96**, 184417 (2017).
- [19] L. Y. Barash, M. Weigel, L. N. Shchur, and W. Janke, Exploring first-order phase transitions with population annealing, *Eur. Phys. J. Spec. Top* **226**, 595 (2017).
- [20] S. Martiniani, J. D. Stevenson, D. J. Wales, and D. Frenkel, Superposition Enhanced Nested Sampling, *Phys. Rev. X* **4**, 031034 (2014).
- [21] M. Creutz, Microcanonical Monte Carlo Simulation, *Phys. Rev. Lett.* **50**, 1411 (1983).
- [22] F. Y. Wu, The Potts model, *Rev. Mod. Phys.* **54**, 235 (1982).
- [23] G. Ossola and A. D. Sokal, Dynamic critical behavior of the Swendsen-Wang algorithm for the three-dimensional Ising model, *Nucl. Phys. B* **691**, 259 (2004).
- [24] L. Y. Barash, M. Weigel, M. Borovský, W. Janke, and L. N. Shchur, GPU accelerated population annealing algorithm, *Comput. Phys. Commun.* **220**, 341 (2017).
- [25] W. Janke and S. Kappler, Monte Carlo study of pure-phase cumulants of 2d q -state Potts models, *J. Phys. I (France)* **7**, 663 (1997).
- [26] R. J. Baxter, Potts model at the critical temperature, *J. Phys. C: Solid State Phys.* **6**, L445 (1973).
- [27] S. R. McKay, A. N. Berker, and S. Kirkpatrick, Spin-Glass behavior in Frustrated Ising Models with Chaotic Renormalization-Group Trajectories, *Phys. Rev. Lett.* **48**, 767 (1982).
- [28] A. J. Bray and M. A. Moore, Chaotic Nature of the Spin-Glass Phase, *Phys. Rev. Lett.* **58**, 57 (1987).
- [29] W. Wang, J. Machta, and H. G. Katzgraber, Chaos in spin glasses revealed through thermal boundary conditions, *Phys. Rev. B* **92**, 094410 (2015).