

Optimal schedules for annealing algorithms

Amin Barzegar,¹ Firas Hamze,¹ Christopher Amey,² and Jonathan Machta^{3,4}

¹*Microsoft Quantum, Microsoft, Redmond, Washington 98052, USA*

²*Department of Physics, Brandeis University, Waltham, Massachusetts 02453, USA*

³*Department of Physics, University of Massachusetts, Amherst, Massachusetts 01003, USA*

⁴*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501, USA*



(Received 11 March 2024; accepted 16 May 2024; published 6 June 2024)

Annealing algorithms such as simulated annealing and population annealing are widely used both for sampling the Gibbs distribution and solving optimization problems (i.e., finding ground states). For both statistical mechanics and optimization, additional parameters beyond temperature are often needed such as chemical potentials, external fields, or Lagrange multipliers enforcing constraints. In this paper we derive a formalism for optimal annealing schedules in multidimensional parameter spaces using methods from nonequilibrium statistical mechanics. The results are closely related to work on optimal control of thermodynamic systems [Sivak and Crooks, *Phys. Rev. Lett.* **108**, 190602 (2012)]. Within the formalism, we compare the efficiency of population annealing and multiple weighted runs of simulated annealing (“annealed importance sampling”) and discuss the effects of nonergodicity on both algorithms. Theoretical results are supported by numerical simulations of spin glasses.

DOI: [10.1103/PhysRevE.109.065301](https://doi.org/10.1103/PhysRevE.109.065301)

I. INTRODUCTION

Simulated annealing [1] and related Monte Carlo algorithms such as annealed importance sampling [2], population annealing [3,4], parallel tempering [5,6], simulated tempering [7], the Wang-Landau algorithm [8], and nested sampling [9] play an important role both in equilibrium sampling and optimization and have been widely applied and studied in physics, mathematics, computer science, and operations research.

Simulated annealing (SA) consists of a Markov Chain Monte Carlo algorithm (MCMC) that samples the Gibbs distribution together with an annealing schedule. The temperature parameter of the MCMC is adjusted during the run of the algorithm according to the annealing schedule which starts at some high temperature where the MCMC mixes rapidly to some low target temperature where mixing is slow. In optimization applications, the target temperature is chosen so that there is a significant probability of finding an optimal or near optimal solution (ground state). The hope, born out by both theory and numerical experiments, is that for finding optima it is much more efficient to anneal to a low temperature than to simply run the MCMC at the low temperature.

Simulated annealing was introduced to find optima (ground states) but it can be extended [2,10] to be used as a Gibbs sampler. In sampling mode, it is known as annealed importance sampling (AIS) and consists of multiple independent runs of SA. Gibbs averages are obtained as weighted averages over these runs. Population annealing (PA) consists of many parallel, interacting runs of SA. At each step in the anneal, a birth-death process, also known as resampling, removes some of these runs and copies others to keep the population near the equilibrium distribution and PA natively functions as a Gibbs sampler. A similar idea is implemented in diffusion Monte Carlo [11] and substochastic Monte Carlo [12] where

the population represents a wave function and resampling implements time evolution.

In this paper, we explicitly discuss the optimization of annealing schedules for population annealing and annealed importance sampling although the ideas are expected to carry over, at least qualitatively, to other annealing algorithms. Our formalism specifies annealing schedules that, for a given amount of computational work, minimize systematic errors in sampling the Gibbs distribution. The same schedules, with end points at sufficiently low temperatures, are expected to work well for optimization.

A great deal of work has been done to determine optimal annealing schedules for SA. It can be proved [13,14] that SA will converge to the global optimum if the temperature schedule decreases to zero logarithmically slowly but these results are not useful in practical application. A variety of ad hoc schedules have been proposed and studied, see, for example, Refs. [15,16] and references therein. The proposed annealing schedules that are most closely related to the present work are based on constant “thermodynamic speed” [17,18]. Less effort has been devoted to optimizing annealing schedules for PA with several proposals discussed in Refs. [19–21].

Annealing algorithms are most frequently deployed with a single control parameter, usually inverse temperature, however, there are many applications where multiple control parameters are useful. In statistical physics, examples include sampling Gibbs distributions with external fields or chemical potentials. In optimization applications, parameters are employed to enforce constraints via penalty terms. During the anneal these penalties are typically increased so that constraints are initially soft but fully enforced at the end of the anneal. In the present work we develop a formalism for optimized annealing schedules for multiple control parameters.

The paper is organized as follows. In Sec. II we review population annealing and develop the theoretical ideas for optimal annealing paths in the context of multiple control parameters and smooth annealing paths. The new theoretical material is mostly contained in Secs. IIB and IIE. In Sec. IIB we introduce continuous time annealing and, in conjunction with the Appendix, describe a new resampling method, pairwise residual resampling, that can be carried out quasi-continuously without introducing significant errors.

The theory of optimal annealing paths is developed in Sec. IIE and is formally similar to the theory of optimal control of thermodynamic systems introduced by Sivak and Crooks [22] and their collaborators [23]. Following Sivak and Crooks [22] we observe that the optimal annealing schedules, at least near equilibrium, are geodesics relative to a metric defined by the so-called “friction tensor.” The components of the friction tensor are integrals of time correlation functions of the forces conjugate to the control parameters. In the familiar case of a single control parameter, usually inverse temperature, the conjugate force is the energy. It is perhaps not surprising that optimizing annealing algorithms and optimizing the control of thermodynamics systems are closely related since annealing algorithms were motivated by thermodynamics annealing and are computer instantiations of controlled thermodynamic systems.

The friction tensor formalism assumes the dynamics, here a Markov chain, is ergodic and near equilibrium throughout the annealing schedule. However, in many realistic situations the annealing path passes through a region where ergodicity is broken, for example, at a first-order transition or in the low temperature phase of spin glasses. In Sec. IIE 2 we extend the formalism to a simple case of ergodicity breaking and explore how transitions into nonergodic regions of parameter space contribute systematic errors in population annealing. Within a nonergodic region the global correlation times appearing in the standard friction tensor diverge and must be replaced by an average over correlation times within each isolated local minimum of the free energy, which motivates the introduction of a restricted friction tensor in Sec. IIE 3.

In Sec. III we extend the ideas developed for population annealing to annealed importance sampling and discuss when one or the other algorithm is expected to be superior. In Sec. IV we present results from population annealing simulations for finding spin-glass ground states using various annealing schedules. These results show that the restricted friction tensor schedule is superior to two of the commonly used schedules for population annealing. The paper closes with a discussion in Sec. V.

II. POPULATION ANNEALING

Population Annealing is a sequential Monte Carlo method [24] that sequentially transforms or “anneals” a population of R replicas of the system from an easy to sample region of parameter space to a difficult to sample region of parameter space. During the anneal, two processes are used to carry out the transformation of the population from one distribution to another. The first is the action of an MCMC algorithm whose invariant measure is the Gibbs distribution at the current value of the parameters. The second step consists of resampling the

population, described in detail below. Simulated Annealing can be understood as a special case of PA where the population has one member ($R = 1$) and, therefore, no resampling occurs.

Annealing is carried out following an annealing schedule, which is a trajectory in the space of parameters of the Gibbs distribution. In many applications there is a single parameter, the inverse temperature β , which controls the average energy of the Gibbs distribution. In this work, we consider annealing in multidimensional parameter spaces.

We write the generalized Gibbs distribution, $\pi(\sigma, \lambda)$ for the probability of microstate σ given the vector of parameters $\lambda = \{\lambda_i\}$ in the form

$$\pi(\sigma, \lambda) = \exp[F(\lambda) - E(\sigma, \lambda)], \quad (1)$$

where $E(\sigma, \lambda)$ is the dimensionless energy as a function of system configuration σ and $F(\lambda)$ is the dimensionless free energy defined as

$$F(\lambda) = -\log \sum_{\sigma} \exp[-E(\sigma, \lambda)]. \quad (2)$$

We use the notation $\langle A \rangle_{\lambda}$ to indicate an average of observable A over the generalized Gibbs distribution with control parameter vector λ .

In many cases, the parameters appear linearly in the energy function. For example, consider an Ising model with coupling strength J in an external field h . The dimensionless energy function is given by

$$E(\{S_i\}, \{\beta J, \beta h\}) = -\beta J \sum_{\langle i,j \rangle} S_i S_j - \beta h \sum_i S_i, \quad (3)$$

where the spins S_i take values ± 1 and the summation is over the edges of the graph of coupled spins. The two control parameters are $\lambda_1 = \beta J$ and $\lambda_2 = \beta h$.

The annealing schedule, $\lambda(t)$, is parameterized by time t , which ranges over the interval 0 to T . Time is measured in units of elementary moves of the MCMC method. For example, in the case of the Ising model with the Metropolis algorithm as the MCMC a unit of time corresponds to attempting a single spin flip. In standard implementations of PA and SA time progresses in discrete steps but, to simplify the analysis, here we consider time to be a real parameter and assume that elementary moves of the MCMC occur as a Poisson process, as in the case of Glauber dynamics. An elementary MCMC move at time t is governed by parameters $\lambda(t)$. In the case of PA, time is considered to run in parallel for each member of the population so the annealing time variable is independent of the population size R .

A. Resampling and free energy

For PA, values of the control parameters are taken to be stepwise constant functions of t with jumps at regular intervals separated by step time dt . Resampling is taken to occur instantaneously at the time of the steps in $\lambda(t)$ and is a birth-death process in the population of replicas. Let t be a resampling time and t^+ be a time immediately after the parameter change at time t (more mathematically, t^+ is shorthand for taking a one-sided limit). Let $w_r(t)$ be the reweighting factor for replica r between the Gibbs distribution with parameters $\lambda(t)$

and $\lambda(t^+)$,

$$w_r(t) = \frac{1}{Q(t)} e^{-E[\sigma_r(t), \lambda(t^+)] + E[\sigma_r(t), \lambda(t)]}, \quad (4)$$

where $\sigma_r(t)$ is the state of replica r , R is the population size, and $Q(t)$ is a normalization factor used to control the population size,

$$Q(t) = \frac{1}{R} \sum_{r=1}^R e^{-E[\sigma_r(t), \lambda(t^+)] + E[\sigma_r(t), \lambda(t)]}. \quad (5)$$

The reweighting factors determine the (joint) distribution of the copy numbers n_r , the number of copies in the population at time t^+ of replica r at time t . The copy numbers are nonnegative integers whose expectations satisfy

$$\mathbb{E}(n_r) = w_r. \quad (6)$$

Note that $n_r = 0$ means replica r is removed from the population and $n_r = 2$ means a copy of the replica is added to the population. Equation (6) ensures that if the population at time t is a good sample from the Gibbs distribution with parameters $\lambda(t)$ then, after resampling, the population at time t^+ will be a good sample from the Gibbs distribution with parameters $\lambda(t^+)$.

There are many ways to implement resampling that satisfy Eq. (6) with either fixed population size or variable population size [25,26]. Population annealing is typically implemented with nearest integer resampling [26,27]. In nearest integer resampling n_r is independently chosen from either the floor or ceiling of w_r such that $\text{Prob}(n_r = \lfloor w_r \rfloor) = [R/R(t)](\lceil w_r \rceil - w_r)$, and $\text{Prob}(n_r = \lceil w_r \rceil) = [R/R(t)](w_r - \lfloor w_r \rfloor)$, where $R(t)$ is the population size at time t before the resampling event and the factor $[R/R(t)]$ is employed so that on each step the population is pushed toward the target size, R . The formalism developed below assumes resampling with a fixed population size, which is discussed in detail in the Appendix.

At any time during the anneal, an estimator for observable A can be obtained from the population average, \tilde{A} ,

$$\tilde{A}(t) = \frac{1}{R} \sum_r A_r(t), \quad (7)$$

where $A_r(t)$ is the value of A in replica r at time t .

The resampling normalization, $Q(t)$ [Eq. (5)] provides an estimator of the free energy \tilde{F} according to

$$\tilde{F}[\lambda(t_2)] = F[\lambda(t_1)] - \sum_{t_1 < t \leq t_2} \log Q(t), \quad (8)$$

where the sum is over resampling times in the interval $(t_1, t_2]$. This relation follows from the Jarzynski equality [10,28] and is also derived in the context of population annealing in Ref. [4].

B. Continuous time annealing

Although practical applications of annealing algorithms employ piecewise constant schedules, here we will develop a formalism for continuous annealing schedules. We can achieve arbitrarily good approximations to a smooth $\lambda(t)$ by successively subdividing the time step, dt between resampling in a piecewise constant schedule. Henceforth, in this work we assume the schedule is a smooth function of time. It is

not obvious that resampling is well-behaved when the time, dt between resamplings becomes small. In the Appendix we introduce a new resampling scheme, which we call pairwise residual resampling (PRR), that is well-suited to doing frequent resampling. Pairwise residual resampling shares features with nearest integer resampling, systematic resampling and residual resampling [26] but is designed to work only in the situation that the weights w_r are very close to unity as is the case for sufficiently small time steps. Other resampling schemes such as multinomial and Poisson resampling are not well behaved in the $dt \rightarrow 0$ limit.

For small time steps, the resampling normalization $Q(t)$ can be expanded to leading order in dt as

$$Q(t) = 1 - \frac{1}{R} \sum_{r=1}^R E(\sigma_r(t), \lambda(t^+)) - E(\sigma_r(t), \lambda(t)) \quad (9)$$

$$\rightarrow 1 - \frac{dt}{R} \sum_{r=1}^R \dot{\lambda}_i \frac{\partial E_r(\lambda)}{\partial \lambda_i}, \quad (10)$$

where E_r is shorthand for the dimensionless energy of replica r and summation over repeated indices is assumed. The expected number of copies of replica r reduces to

$$w_r \rightarrow 1 - dt \dot{\lambda}_i (X_i^r - \tilde{X}_i). \quad (11)$$

Here X_i^r is the i th conjugate force of the r th replica,

$$X_i^r = -\frac{\partial E_r(\lambda)}{\partial \lambda_i}, \quad (12)$$

and \tilde{X}_i is the population average of X_i^r [see Eq. (7)]. In the case of the Ising model in a field, Eq. (3), $X_1 = \sum_{\langle i,j \rangle} S_i S_j$ is the dimensionless coupling energy and $X_2 = \sum_i S_i$ is the dimensionless magnetization.

From Eqs. (8), (10), and (12) in the $dt \rightarrow 0$ limit, the free-energy estimator becomes

$$\tilde{F}[\lambda(t_2)] = F[\lambda(t_1)] + \int_{t_1}^{t_2} dt' \dot{\lambda}_i \tilde{X}_i(t). \quad (13)$$

The integral in Eq. (13) represents the generalized work associated with the anneal from t_1 to t_2 .

We conclude this section introducing some notation used in the rest of the paper. The average of a quantity A over the Gibbs distribution with parameters λ is written as $\langle A \rangle_\lambda$. If we carry out many runs of PA with schedule $\Lambda = \{\lambda(t) : 0 \leq t \leq T\}$, then the average over runs of the estimator \tilde{A} from each run is written as $\langle \tilde{A}(t) \rangle_\Lambda$. Note $\tilde{A}(t)$ may depend on the annealing path up to time t . However, if the anneal is sufficiently slow so that the population remains close to equilibrium, then $\langle \tilde{A}(t) \rangle_\Lambda \approx \langle A \rangle_{\lambda(t)}$. For example, the PA average for the free energy is

$$\langle \tilde{F}[\lambda(t_2)] \rangle_\Lambda = F[\lambda(t_1)] + \int_{t_1}^{t_2} dt' \dot{\lambda}_i \langle \tilde{X}_i(t) \rangle_\Lambda, \quad (14)$$

and for an infinitely slow anneal we recover the standard result for thermodynamic integration,

$$F[\lambda(t_2)] = F[\lambda(t_1)] + \int_{t_1}^{t_2} dt' \dot{\lambda}_i \langle X_i(t) \rangle_{\lambda(t)}. \quad (15)$$

C. Culling fraction

An important metric associated with resampling is the culling fraction ϵ , the expected number of replicas removed from the population during resampling. For both nearest integer resampling and pairwise residual resampling (see the Appendix), it is straightforward to see [21] that the culling fraction for a single resampling step is

$$\epsilon = \frac{1}{R} \sum_{\{r|w_r < 1\}} (1 - w_r). \quad (16)$$

Note that the complement, $1 - \epsilon$, can also be interpreted as the overlap between the Gibbs distributions before and after a nearest integer resampling step. Two quite effective proposed annealing schedules with a single control parameter for PA are a fixed culling fraction schedule [21] and, equivalently for nearest integer resampling, a fixed overlap schedule [19]. In Sec. IV B we compare the friction tensor schedules proposed below to fixed culling fraction schedules.

In the $dt \rightarrow 0$ limit all w 's are near one, so the culling fraction for a single resampling step is infinitesimal. The cumulative culling fraction $\epsilon(t_1, t_2)$ over segment (t_1, t_2) of the annealing path can be obtained by expanding w_r to leading order in dt and integrating over the annealing path segment. Furthermore, since the population average of w_r is one (or very nearly one for nearest integer resampling), the sum in Eq. (16) can be equivalently expressed as being over replicas with $w_r > 1$ except that the sign of the summand is changed. Using Eq. (11) we obtain an expression for the cumulative culling fraction along an annealing segment,

$$\epsilon(t_1, t_2) = \frac{1}{2R} \int_{t_1}^{t_2} dt \left\langle \sum_r |\dot{\lambda}_i (X_i^r - \tilde{X}_i)| \right\rangle_{\Lambda}. \quad (17)$$

If the population remains close to equilibrium, then the nonequilibrium average $\langle \cdots \rangle_{\Lambda}$ can be replaced by an equilibrium average $\langle \cdots \rangle_{\lambda(t)}$,

$$\epsilon(t_1, t_2) = \frac{1}{2} \int_{t_1}^{t_2} dt' \langle |\dot{\lambda}_i \delta X_i| \rangle_{\lambda(t)}. \quad (18)$$

In this expression, $\delta X = X - \langle X \rangle_{\lambda(t)}$.

If furthermore, the Gibbs distribution is a multivariate Gaussian, as is often a good approximation, then we have

$$\epsilon(t_1, t_2) = \int_{t_1}^{t_2} dt' \sqrt{\frac{\dot{\lambda}_i \sigma_{ij}^2 \dot{\lambda}_j}{2\pi}}, \quad (19)$$

where $\sigma_{ij} = \langle \delta X_i \delta X_j \rangle_{\lambda(t)}$ is the covariance matrix. In the simple case that the only control parameter is inverse temperature the integrand reduces to $\beta \sigma_e / \sqrt{2\pi}$ where σ_e is the variance of the dimensionless energy.

D. Weighted averaging and systematic errors

A simple average over multiple independent runs of PA reduces statistical but not systematic errors. Multiple independent runs of PA can be combined using weighted averaging to improve both statistical and systematic errors in measuring either the free energy or an observable [27]. Let \tilde{A}_m be the estimator of observable A from run m of PA and let \tilde{F}_m be the estimator of the dimensionless free energy in run m of

PA. The weighted average over M runs, \bar{A} is given by

$$\bar{A} = \frac{1}{M} \sum_{m=1}^M \tilde{A}_m \exp(-\tilde{F}_m + \bar{F}), \quad (20)$$

where \bar{F} is the weighted average free energy,

$$\bar{F} = -\log \frac{1}{M} \sum_{m=1}^M \exp(-\tilde{F}_m). \quad (21)$$

It is assumed in these expressions that the population size R and the annealing schedule in all runs are the same. For fixed R , in the limit of $M \rightarrow \infty$, the weighted free energy and all observables converge to their exact Gibbs value. In optimization applications, weighted averaging corresponds simply to choosing the best solution from all runs.

A single run of PA with finite population size will suffer from systematic errors. The magnitude of these errors can be obtained by comparing ordinary averaging over runs, which does not suppress systematic errors to weighted averaging, which becomes exact for large M . The central limit theorem suggests that for large R , the estimators for observables and the free energy should become Gaussian and then the systematic error is related to the (co)variances according to Ref. [27],

$$\Delta \tilde{A} = \text{cov}(\tilde{A}, \tilde{F}) \quad (22)$$

and

$$\Delta \tilde{F} = \text{var}(\tilde{F})/2, \quad (23)$$

where $\Delta \tilde{A}$ is the difference between the exact value of the observable and the expected result from a single PA run,

$$\Delta \tilde{A} = \langle \tilde{A} \rangle_{\Lambda} - \langle A \rangle_{\lambda}, \quad (24)$$

and $\langle \cdots \rangle_{\Lambda}$ is an (ordinary) average over runs of PA (or SA) with a schedule Λ that terminates at control parameter values λ . The analogous definition of $\Delta \tilde{F}$ is

$$\Delta \tilde{F} = \langle \tilde{F} \rangle_{\Lambda} - F(\lambda), \quad (25)$$

where $F(\lambda)$ is the free energy associated with the Gibbs distribution at control parameter values λ .

Inspired by the central limit theorem, we assume that for large R both $\text{cov}(\tilde{A}, \tilde{F})$ and $\text{var}(\tilde{F})$ scale as $1/R$ and we can then rewrite the systematic errors as

$$\Delta \tilde{F} = \rho_f / 2R \quad (26)$$

and

$$\Delta \tilde{A} = \rho_f C(A) / R, \quad (27)$$

where

$$\rho_f = \lim_{R \rightarrow \infty} R \text{var}(\tilde{F}) \quad (28)$$

and

$$C(A) = \lim_{R \rightarrow \infty} \frac{\text{cov}(\tilde{A}, \tilde{F})}{\text{var}(\tilde{F})}. \quad (29)$$

The quantity ρ_f , introduced in the context of PA in Ref. [27], sets the scale of systematic errors in free energy measurements. Furthermore, if the coefficient $C(A)$ is approximately

independent of the annealing schedule, then ρ_f also sets the scale of systematic errors in the observable A . Since A is calculated from the final population and not from the whole trajectory, we expect the $C(A)$ is only weakly dependent on the annealing path.

Thus, ρ_f can be loosely interpreted as the population size required to reach equilibrium. To minimize the computational work of doing equilibrium sampling with PA the annealing schedule should be chosen to minimize ρ_f . Much of the remainder of the paper is concerned with obtaining tractable expressions for ρ_f that are amenable to minimization.

While ρ_f is related to systematic errors, other measures have been proposed to estimate statistical errors in PA. One of these, ρ_t (see Refs. [21,26,27]) is based on the distribution of family sizes. In this context a family is a set of replicas descended from a single replica in the initial population and ρ_t is the sum of the squares of the family sizes divided by the population size. Thus, at the beginning of the simulation, $\rho_t = 1$ and is bounded by $\rho_t \leq R$. As shown in Ref. [21], for the Ising spin glass ρ_t and ρ_f typically differ by a constant that is small compared to ρ_f itself. Because of this close connection and the ease of measuring ρ_t , we use this measure to characterize the hardness of a spin-glass samples studied in Sec. IV B. A second measure of statistical errors is the effective population size introduced in Ref. [29], which is bounded above by R/ρ_t .

E. Friction tensor

In this section we obtain a formal expression for ρ_f and show that in the near equilibrium approximation, it can be expressed in terms of the friction tensor introduced in Ref. [22]. The variance of the free energy estimator is obtained by squaring Eq. (13) and subtracting mean values so that Eq. (28) becomes

$$\rho_f = \lim_{R \rightarrow \infty} R \int_0^T dt' \int_0^T dt'' \dot{\lambda}_i(t') \langle \delta \tilde{X}_i(t') \delta \tilde{X}_j(t'') \rangle_{\Lambda} \dot{\lambda}_j(t''), \quad (30)$$

where $\delta A = A - \langle A \rangle_{\Lambda}$. This expression is a time integral over a matrix of nonequilibrium correlation functions of the conjugate forces. The dynamics in the nonequilibrium average includes both the dynamics of the MCMC and of resampling, both of which promote the decay of correlations. In the more general setting of sequential Monte Carlo algorithms, similar expressions for the variance of the free energy estimator have been developed [24,30].

To gain greater insights into the meaning of Eq. (30) and to obtain a more tractable expression we make two related approximations:

- (1) The population remains close to an equilibrium sample during the anneal.
- (2) The control parameters vary slowly compared to the time scale for the decay of correlations.

Given these approximations we set $t' = t''$ in the time argument of the parameter velocities and we replace the nonequilibrium average by an equilibrium average at the current values of the control parameters. The resulting expression

takes the form of a time integral over a bilinear form,

$$\rho_f = \int_0^T dt \dot{\lambda}_i(t) \zeta_{ij}(\lambda(t)) \dot{\lambda}_j(t), \quad (31)$$

where $\zeta_{ij}(\lambda(t))$ is the “friction tensor,”

$$\zeta_{ij}(\lambda(t)) = 2 \int_0^\infty dt' \langle \delta X_i(0) \delta X_j(t') \rangle_{\lambda(t)}. \quad (32)$$

The assumption that correlations decay quickly justifies taking the limit to infinity in the integral defining the friction tensor. The friction tensor can be further decomposed into the outer product of the matrix of equilibrium correlation times and equal time covariances,

$$\zeta_{ij} = 2 \tau_{ij} \sigma_{ij}^2, \quad (33)$$

where $\sigma_{ij} = \langle \delta X_i \delta X_j \rangle_{\lambda(t)}$ and

$$\tau_{ij} = \int_0^\infty dt' \frac{\langle \delta X_i(0) \delta X_j(t') \rangle_{\lambda(t)}}{\langle \delta X_i \delta X_j \rangle_{\lambda(t)}}. \quad (34)$$

In the friction tensor approximation, the requirements for a good annealing path are intuitively plausible: Regions of parameter space with large values of variances or autocorrelation times should be avoided or, if they cannot be avoided, traversed slowly.

From this point onward there are two directions to explore. The first is to obtain a more detailed understanding of the implications of Eq. (31). Fortunately, this equation is formally identical and also conceptually related to the expression obtained by Sivak and Crooks [22,23] for minimum excess work paths in nonequilibrium thermodynamics so we can take over their analysis with minimal modification.

The second, more difficult task is to understand the role of resampling. In going from Eq. (30) to Eq. (31) resampling has been explicitly removed from the expression since the dynamics in the definition of the friction tensor is at a single temperature and resampling only plays a role when the temperature changes. It appears we have thrown away the baby with the bath water.

1. Optimal annealing paths for population annealing

The friction tensor formulation of ρ_f in Eqs. (31) and (32) is formally identical to the expression found by Sivak and Crooks [22] for minimum excess thermodynamic work, which allows us to adopt their analysis directly. They showed that the friction tensor defines a pseudometric space whose geodesics are minimum dissipation paths in a space of thermodynamic parameters. In one dimension, the geodesic simply determines the speed along the path and, when translated into an annealing schedule in the single parameter of inverse temperature, β the result is ([22], Eq. 12)

$$\dot{\beta} \propto \zeta^{-1/2}, \quad (35)$$

where, in this case, $\zeta = 2\tau_e C/\beta^2$, $C = \beta^2 \sigma_e^2$, is the heat capacity, σ_e is the standard deviation of the energy, and τ_e is the integrated autocorrelation time of the energy. This annealing schedule has the property that the rate of increase in ρ_f is constant along the path.

The annealing schedule (35) can be rewritten as $\dot{\beta} \propto \beta/\sqrt{\tau_e C}$, which makes clear the similarity to the constant ther-

modynamic speed schedule for SA proposed in Refs. [17,18], which takes the form $\dot{\beta} \propto \beta/\tau\sqrt{C}$. Here τ is the “apparent relaxation time” of the energy. Assuming that “apparent relaxation time” is interpreted as the energy integrated autocorrelation time the only difference between these schedules is an extra factor of $\sqrt{\tau_e}$ in the denominator.

Rotskoff and Crooks [31] show how to compute geodesics in the two-dimensional space of coupling energy and magnetic field for the Ising model and these ideas could be applied to determining annealing paths in the temperature-field plane.

2. Ergodicity breaking and the role of resampling

The friction tensor formalism developed at the end of Sec. II E does not take into account resampling since it uses Gibbs averages taken at a single value of the control parameters. In this section we discuss the role of resampling. In regions of parameter space where correlation times are short, the role of resampling is relatively limited and difficult to quantify. On the one hand, resampling helps keep the population close to the current Gibbs distribution so that the local equilibrium approximation in going from the full expression, Eq. (30) to the friction tensor formulation, Eq. (31) is more accurate. On the other hand, resampling increases the variance of the free energy and covariance of the free energy with observables because it is a stochastic process. It is not clear which of these effects dominates for a given problem so that it may be that AIS or simply using the MCMC at the target values of the control parameters is more efficient.

Resampling becomes important in regions of parameter space where correlation times become large so that, by itself, the MCMC cannot reach equilibrium. Here resampling may be required to maintain equilibrium as the control parameters change. While a full analysis of the combined effects of the resampling and the MCMC is beyond the scope of this work, we will investigate a simple and relatively generic situation to show how the friction tensor formulation can be modified to include both nonergodicity of the MCMC and the contributions of resampling to ρ_f .

In the simplified model, there are two disconnected regions in configuration space below an ergodicity breaking transition. Within each region the MCMC is ergodic but transitions between regions do not occur on the time scales of the simulation. We assume that within each of the two regions, the MCMC is able to achieve local equilibrium in the region on a reasonable time scale and it is these times that will enter into the modified friction tensor formulation. To further simplify the problem, suppose that inverse temperature β is the only control parameter, energy E is its conjugate field and the only relevant time scales are the energy autocorrelation times in the two regions. Finally, we assume that the difference between the entropies and energies of the two regions are independent of β . Let $\mathcal{E} > 0$ be the average energy difference. The entropy difference can be parameterized by a transition temperature, β_c where the two regions are equally probable in the Gibbs distribution. The conclusions we obtain for this simple two region model should apply, at least qualitatively, to situations where configuration space fractures into multiple disconnected regions with more complex free energy landscapes.

The simple two-region model is discussed in Sec. IV of Ref. [32]. The key result is that the probability, $p(\beta)$ of being in the lower energy (and lower entropy) regions obeys a Fermi distribution,

$$p(\beta) = \frac{1}{1 + e^{-(\beta - \beta_c)\mathcal{E}}}, \quad (36)$$

and that the entropy difference is $\beta_c\mathcal{E}$. Suppose that the MCMC becomes effectively nonergodic at β_n . The interesting and computationally difficult situation is when nonergodicity occurs before the transition, $\beta_n < \beta_c$, while the target temperature is after the transition, $\beta(T) > \beta_c$ so that the low-energy region dominates the distribution at the target temperature but is improbable at the ergodicity transition. As the anneal proceeds beyond β_n , resampling increases the fraction of the population in the low-energy region but stochasticity in the number of replicas in the low-energy region at β_n is not eliminated until well beyond β_c when nearly all the population is in the low-energy region.

To evaluate the contribution to ρ_f from resampling in the two-region model, suppose, without loss of generality, the energy in low-energy region is zero, and let η be the fraction of replicas in this zero energy region. Since the annealing schedule is in β the only conjugate field is the energy and $\tilde{X} = (1 - \eta)\mathcal{E}$.

If we ignore the stochasticity of resampling, then from Eq. (4), we have the following differential equation for the population in the low-energy region as a function of β ,

$$\frac{d\eta}{d\beta} = \mathcal{E}\eta(1 - \eta), \quad (37)$$

whose solution with initial condition $\eta = \eta(\beta_n)$ is

$$\eta(\beta) = \frac{\eta(\beta_n)e^{(\beta - \beta_n)\mathcal{E}}}{1 + \eta(\beta_n)(e^{(\beta - \beta_n)\mathcal{E}} - 1)}. \quad (38)$$

Since we are only concerned with resampling, we use β as the independent variable instead of time in calculating the double integrals defining ρ_f [Eq. (30)] and obtain, for an anneal from the ergodicity transition to infinite β ,

$$\rho_f = \lim_{R \rightarrow \infty} \mathcal{E}^2 R \int_{\beta_n}^{\infty} d\beta' \int_{\beta_n}^{\infty} d\beta'' \langle \delta\eta(\beta') \delta\eta(\beta'') \rangle_{\Lambda}, \quad (39)$$

where $\langle \dots \rangle_{\Lambda}$ is here an average over Gaussian initial conditions. At the ergodicity transition the population size in the zero energy region, $R\eta(\beta_n)$, is a Poisson random variable with mean and variance $Rp(\beta_n)$. For $Rp(\beta_n) \gg 1$ the distribution of $\eta(\beta_n)$ approaches a sharply peaked Gaussian distribution with mean $p(\beta_n)$ and variance $p(\beta_n)/R$.

To evaluate Eq. (39) we switch the order of averaging and integration. The only integral to be done is

$$g(x) = \int_0^{\infty} dz \frac{xe^{z\mathcal{E}}}{1 + x(e^{z\mathcal{E}} - 1)} - 1 \quad (40)$$

$$= \log(x)/\mathcal{E}. \quad (41)$$

The first term in the integrand is transcribed from Eq. (38) and the one is subtracted to make the integral convergent but does not affect the final result for ρ_f .

In terms of g , the expression for ρ_f takes the form

$$\rho_f = \lim_{R \rightarrow \infty} \mathcal{E}^2 R [\langle g(\eta(\beta_n))^2 \rangle_\Lambda - \langle g(\eta(\beta_n)) \rangle_\Lambda^2]. \quad (42)$$

To carry out these Gaussian averages in the large R limit, both $g(x)$ and $g(x)^2$ are expanded in a Taylor series around the mean $\eta(\beta_n) = p(\beta_n)$ and terms to second order in $[\eta(\beta_n) - p(\beta_n)]$ are kept. The final result is simple and intuitive,

$$\rho_f = 1/p(\beta_n). \quad (43)$$

The intuition behind this result is that once the population size exceeds $1/p(\beta_n)$ it becomes likely that the low-energy state is in the population and, via resampling, can grow to take over the population. Beyond this point, systematic errors are suppressed proportional $1/Rp(\beta_n)$.

It should be noted that a subtle error has been made in going from the exact Poisson distribution for the population in the low-energy region to the approximate Gaussian distribution. Specifically, with probability $e^{-Rp(\beta_n)}$, there are *no* members of the population in the low-energy region when the system becomes nonergodic. In this situation, the algorithm completely fails. Thus, it is more correct to say that ρ_f as calculated above controls systematic errors when the failure probability is small, $e^{-Rp(\beta_n)} \ll 1$.

Although the above calculation applies exactly to the simple two-region model, the qualitative result should hold if there are multiple disconnected competing regions, if there are first-order-like transition preceded by ergodicity breaking among several regions in configuration space along an annealing trajectory in several dimensions. Specifically, upon crossing each first-order-like transition one expects a contribution to ρ_f that behaves as $1/p$ where p is the probability of being in the lower energy region when ergodicity is broken.

Thus far we have assumed the MCMC is capable of equilibrating replicas within each disconnected region but we have not quantified the contribution to ρ_f associated with the MCMC due to correlation times and equal time covariance. To estimate these contributions it is necessary to modify the friction tensor formulation because resampling takes care of redistributing replicas between disconnected regions and the MCMC only decorrelates replicas within these regions. In the next section we introduce a restricted friction tensor for use in nonergodic regions.

3. Restricted friction tensor

Suppose there are k disconnected regions in configuration space and that as a function of time along the annealing path, the equilibrium probability of being in each region α is $p_\alpha[\lambda(t)]$. We assume that resampling distributes the population $\eta_\alpha(t) \approx p_\alpha[\lambda(t)]$, where η_α is the fraction of the population in region α . If we ignore the contribution of resampling to ρ_f and proceed as before from Eq. (30) to Eqs. (31) and (32), then we obtain, for the restricted friction tensor ζ_{ij} ,

$$\zeta_{ij}^r(\lambda) = 2 \int_0^\infty dt' \sum_\alpha p_\alpha(\lambda) \langle \delta_\alpha X_i(0) \delta_\alpha X_j(t') | \alpha \rangle_\lambda, \quad (44)$$

where $\langle \cdots | \alpha \rangle_\lambda$ is a conditional Gibbs average over the (dynamically determined) regions labeled by α and $\delta_\alpha X = X - \langle X | \alpha \rangle_\lambda$. Note that if one region dominates the average then

this expression reduces to the unrestricted friction tensor expression. When there are several important regions of configuration space, the restricted friction tensor is an average over the friction tensors associated with each region.

In general, the conditional averages associated with the restricted friction tensor are not known, however, we can estimate the restricted friction tensor from the correlation functions appearing in the full friction tensor if there is a clear separation of time scales between the local relaxation time within regions and global relaxation time. We have from the law of total covariance:

$$\begin{aligned} \langle \delta X_i(0) \delta X_j(t') \rangle &= \sum_\alpha p_\alpha \langle \delta_\alpha X_i(0) \delta_\alpha X_j(t') | \alpha \rangle \\ &+ \sum_\alpha p_\alpha [\langle X_i \rangle - p_\alpha \langle X_i | \alpha \rangle] \\ &\times [\langle X_j \rangle - p_\alpha \langle X_j | \alpha \rangle]. \end{aligned} \quad (45)$$

Note the second term on the right-hand side (RHS) of Eq. (45) does not decay, but if we subtract off the linearly increasing contributions to the integral defining the unrestricted friction tensor we obtain the restricted friction tensor. To measure the restricted friction tensor, we can do a linear fit of the time integral whose limit is the full friction tensor,

$$2 \int_0^x dt' \langle \delta X_i(0) \delta X_j(t') \rangle_\lambda \rightarrow a_{ij}(\lambda)x + \zeta_{ij}^r(\lambda), \quad (46)$$

where $a_{ij}(\lambda)$ the second term Eq. (45) is independent of x . We can write the restricted friction tensor in the same form as Eq. (32),

$$\zeta_{ij}^r = 2\tau_{ij}^r \sigma_{ij}^2, \quad (47)$$

where the restricted integrated correlation time matrix τ_{ij}^r is defined by this equation.

III. ANNEALED IMPORTANCE SAMPLING

Simulated annealing (PA with $R = 1$) can also be used with weighted averaging; this application of SA is known as *annealed importance sampling* (AIS) [2,10]. Annealing importance sampling permits simulated annealing to be used not only as an optimizer but also as a Gibbs sampler. In what follows, we will compare a single run of PA with population R to a run of AIS consisting of the weighted average over R^* runs of SA to better understand the pros and cons of resampling. In an abuse of notation, for AIS we use \tilde{F} and \tilde{A} to represent the estimator resulting a single run of AIS of size R^* . Multiple runs of AIS with the same R^* can be additionally combined using Eqs. (20) and (21). With this notation, combining Eqs. (13) and (12) with $R = 1$ and Eq. (21) with $M = R^*$, we have

$$\tilde{F}[\lambda(t_2)] = F[\lambda(t_1)] - \log \frac{1}{R^*} \sum_{r=1}^{R^*} \exp \left(- \int_{t_1}^{t_2} dt' \dot{\lambda}_i X_i^r \right), \quad (48)$$

where

$$X_i^r = - \frac{\partial E_r(\lambda)}{\partial \lambda_i}. \quad (49)$$

In AIS, observables are obtained from the analog of Eq. (20),

$$\tilde{A} = \frac{1}{R^*} \sum_{r=1}^{R^*} A_r \exp \left(- \int_{t_1}^{t_2} dt' \dot{\lambda}_i X_i^r + \tilde{F} \right), \quad (50)$$

where A_r is the value of the observable in run r of SA. The variance of \tilde{F} for AIS now involves the *exponential* of the integral along the annealing path. If the variance of the exponential is itself small, then the formula for ρ_f for AIS will reduce to the PA formula except that two-time correlators are with respect to the dynamics of the MCMC alone and include no resampling.

A. Role of reweighting in annealed importance sampling

We can analyze the role of reweighting in AIS in the context of the two-region model with the same assumptions as used in the previous sections. The sum in Eq. (48) contains two contributions, one from runs of SA that find the low-energy region and one from runs that do not so that

$$\tilde{F} = F_0 - \log(\eta(\beta_n) + (1 - \eta(\beta_n))e^{-(\beta - \beta_n)\mathcal{E}}). \quad (51)$$

Thus, ρ_f for AIS in the limit of large β is simply

$$\rho_f = \lim_{R^* \rightarrow \infty} R^* \text{var}(\log(\eta(\beta_n))), \quad (52)$$

which gives exactly the same result as for PA, $\rho_f = 1/p(\beta_n)$.

For the simple model of ergodicity breaking preceding a first-order-like transition, we see that AIS and PA perform equally well. However, for PA we have ignored the stochasticity associated with resampling so, in principle, AIS may be superior to PA. This conclusion conflicts with simulations that show that PA with population size R is far superior to R runs of SA for finding ground states of the three-dimensional Ising spin glass (Edward-Anderson model) [33]. We believe this difference is related to the fact that spin glasses and perhaps many other systems undergo a sequence of several first-order-like transitions as the temperature is lowered. This phenomenon is known as temperature chaos [34,35] in the spin-glass literature.

Furthermore, let us suppose the landscape of these successive transitions is nested in the sense that it is easier for a replica that has already passed through one of these transitions to pass through the next transition. In this situation PA will be superior because resampling brings the population into the low-energy region of the first transition so that more replicas are exploring the subspace in which the next transition is likely to be found. The simulation results of Ref. [33] suggest that the landscape is indeed nested.

IV. SIMULATIONS

In this section, we describe the results of a computational study comparing the performance of several one-dimensional annealing schedules for population annealing applied to the Ising spin glass. To compare the performance of these schedules we measured both ρ_f at low temperature and the probability of finding the ground state.

TABLE I. Definitions of the annealing schedules investigated in this paper. The schedules are defined by the inverse temperature step size $\delta\beta$ and sweeps per temperature step δs as a function of β . Note that the values given in this table are up to some normalization factors, which are used to fix the total number of temperature steps and sweeps. Here, c_0 is the uniform schedule, ϵ_0 is the fixed culling schedule, and $\zeta_1, \zeta_2, \zeta_3$ are the three types of friction tensor schedules. ϵ , ζ , and τ_e are the culling fraction, the friction tensor, and the integrated autocorrelation time of the energy, respectively.

Name	$\delta\beta$	δs
c_0	1	1
ϵ_0	$1/\epsilon$	1
ζ_1	$\tau_e/\sqrt{\zeta}$	τ_e
ζ_2	$1/\epsilon$	$\sqrt{\zeta}/\epsilon$
ζ_3	1	$\sqrt{\zeta}$

A. Numerical methods

We carried out PA simulations of 100 instances of the three-dimensional Ising spin glass with Gaussian disorder on a cubic lattice of size $10 \times 10 \times 10$ with periodic boundary conditions. We simulated 100 disorder samples using 5 annealing schedules with a target population size of $R = 10^5$ and nearest integer resampling. These disorder samples had been previously studied in Ref. [21]. For each annealing schedule, we carried out two separate sets of simulations. In the first set we allotted a total of 1000 Monte Carlo sweeps to each simulation and in the second set 3000 sweeps. In both cases, the elementary move was a Metropolis spin flip on a randomly chosen spin and one sweep consisted of 1000 elementary moves. For each annealing schedule, the sweeps were divided among 300 temperature steps starting at $\beta = 0$ and ending at $\beta = 5$. For each schedule and each sweep number, we carried out 20 independent runs and estimated ρ_f at $\beta = 5$ from the variance of these runs. For the probability of finding the ground state for a given disorder realization, we used the fraction of the 20 runs that found the same ground state energy as found in Ref. [21].

We studied three friction tensor schedules satisfying $\dot{\beta} = 1/\sqrt{\zeta}$ [see Sec. II E 1 and Eq. (35)], the commonly used fixed culling fraction schedule satisfying $\dot{\beta} = 1/\epsilon$ (see Sec. II C), and a constant β -step schedule, $\dot{\beta} = \text{const}$. The different annealing schedules are distinguished by the inverse temperature step size $\delta\beta$ and sweeps per temperature step δs as a function of β . In our continuous time formulation, it is only the ratio $\delta s/\delta\beta$ that is relevant to computing ρ_f so that there are many ways to implement schedules for a given $\dot{\beta}$. In the case of the friction tensor schedules, we choose three options that seem natural: (1) The number of sweeps per temperature step is constant, (2) the number of sweeps per temperature step yields a fixed culling fraction ϵ in each resampling step, and (3) the number sweeps is proportional to the energy integrated autocorrelation time, τ_e . If our simulations are close enough to the continuous time ($dt \rightarrow 0$) limit, then there is expected to be little difference in the performance of these three schedules.

The five schedules and their designations are given in Table I. In all cases, the constants of proportionality are fixed

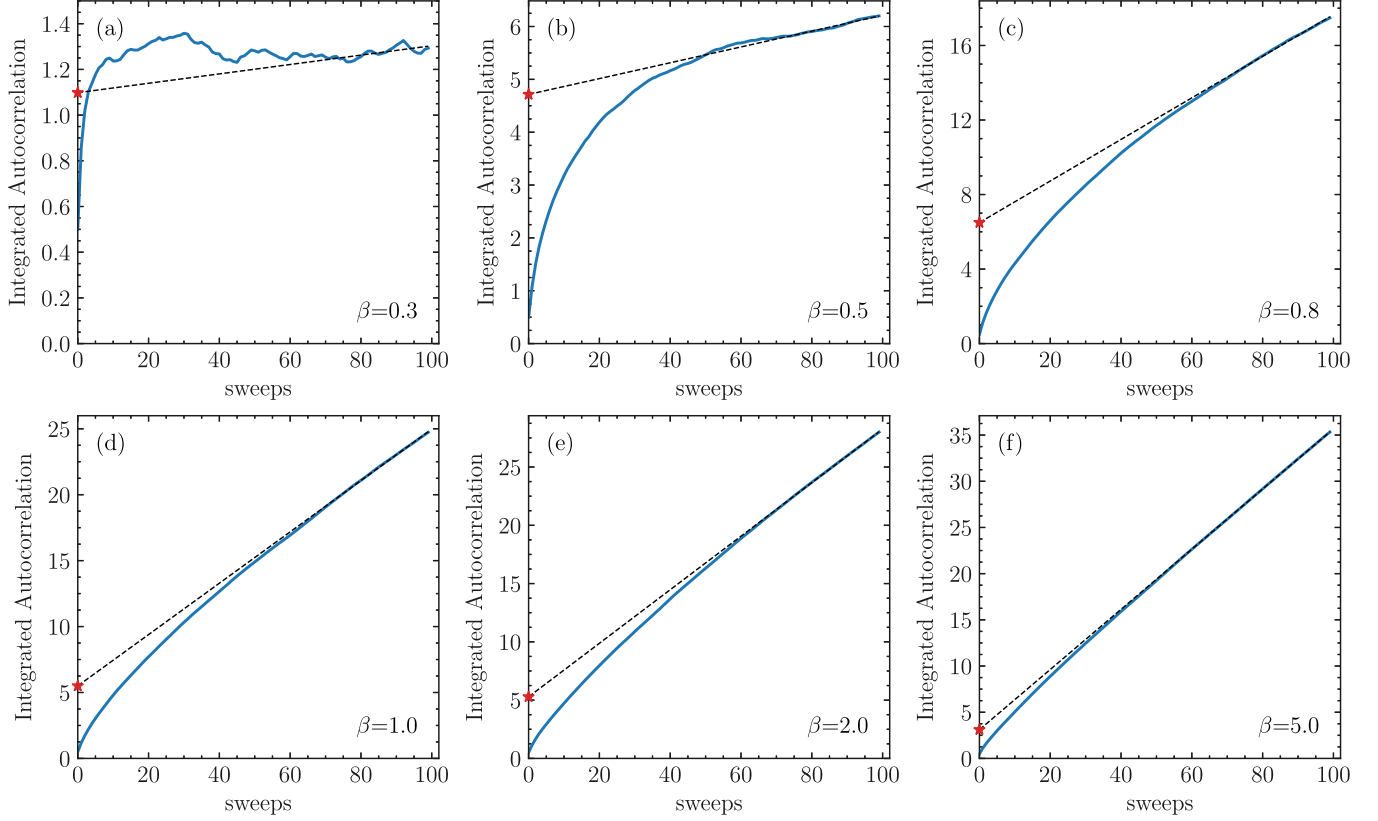


FIG. 1. Extraction of the energy autocorrelation time τ_e by linear fits to the integral of the energy autocorrelation function defined in Eq. (34). Panels (a–f) show examples of this procedure at different inverse temperatures β . We use a long run with 100 sweeps per temperature step for one realization in each bin. A linear curve is then fit to the data after 70 sweeps and the intercept (shown by the red stars) is interpreted as the autocorrelation time.

by the total number of temperature steps (300) and the total number of sweeps, either 1000 or 3000.

The 100 spin-glass samples used in our study were selected from a set of 5000 samples analyzed in Ref. [21]. As noted in Ref. [21], the distribution of hardness for PA, as measured by ρ_f or ρ_t is very broad and described by a log inverse Gaussian distribution. The 100 samples used in our study were selected to span the whole range of hardness encountered in the randomly generated 5000 Gaussian instances. Specifically, the 5000 instances were divided into 10 bins equally spaced on a logarithmic scale ranging from the easiest to the hardest instances. The 10 bins contain different numbers of instances but are of equal logarithmic width. Ten instances were then randomly selected from each of the ten bins so that our study set of 100 instances are roughly uniformly spread over the range of (log) hardness that would be encountered in a much larger sample of instances. The value of ρ_t measured in Ref. [21] is taken as the true measure of the hardness of each instance. As observed in Refs. [21,27], ρ_f and ρ_t differ by a relatively small constant for spin glasses and can be used interchangeably as a measure of computational hardness. We note that resources devoted to the simulations in Ref. [21] were much larger than used here and that we are reasonably confident that ρ_t has been measured in the asymptotic regime and the exact ground state was found.

1. Measuring the integrated autocorrelation time

Relaxation times become very long in spin glasses so that we use the restricted friction tensor formalism developed for nonergodic systems in Sec. II E 3, which, in any case reduces to the usual definition in the high temperature regime where equilibration is fast. For simplicity of notation we refer to the restricted energy autocorrelation time, defined in Eq. (47), as τ_e . To obtain τ_e we measured the autocorrelation function in preliminary runs with 100 sweeps per temperature step for one disorder realization for each of the 10 bins. The sum of the energy autocorrelation function from 0 to t sweeps, minus $1/2$, was fit to a linear function of t in the interval 70 and 100 sweeps. The intercept of the linear fit is identified as the restricted integrated energy autocorrelation time, τ_e . The value of τ_e used in the annealing schedule for each bin was obtained by averaging the value obtained from 20 runs on the exemplar of that bin.

Figure 1 shows the sum of autocorrelation function minus $1/2$ for several values of β for a disorder instance in bin 3 and a single long run. The red stars mark the location of the intercept, which is the value of τ_e . For the highest temperature shown ($\beta = 0.3$), the integrated autocorrelation function is noisy but appears to flatten indicating that the full autocorrelation time has been measured. However, for all larger values of β the integrated autocorrelation function is nearly linear

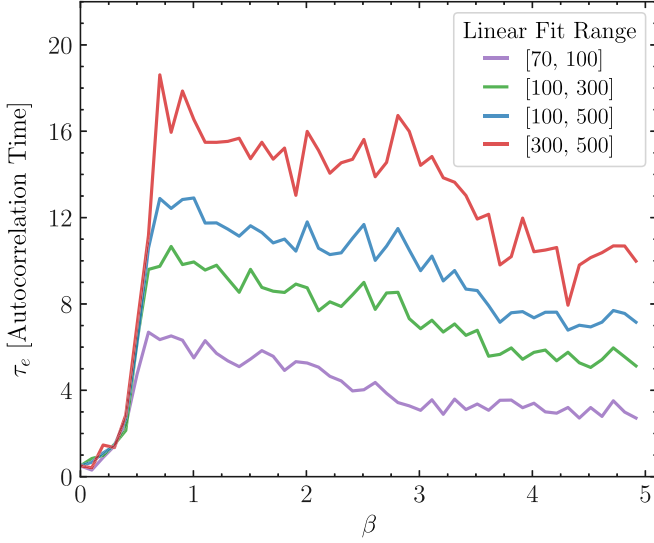


FIG. 2. Autocorrelation time τ_e versus inverse temperature β extracted by fitting a line to different sections of the integrated autocorrelation function as illustrated in Fig. 1. The integrated autocorrelation function was measured over 500 sweeps in total for a select disorder realization in bin 3. The noise in the curves is due to the fact that the data is taken from one single run here. For the performance comparison simulations, we used the bottom range [70, 100] (bottom curve) to define τ_e .

with a positive slope in the range 70 to 100 sweeps indicating that Metropolis dynamics are nonergodic on these time scales but that a restricted integrated autocorrelation time can be measured.

The lowest (purple) curve Fig. 2 shows τ_e as a function of β from the same run as shown in Fig. 1. We see that τ_e rises sharply as β increases and reaches a maximum somewhat before the critical temperature. Thereafter, in the nonergodic low-temperature regime, τ_e slowly decreases. The roughness of the curves is primarily because of the fact that this data is from a single run.

The simple assumption in Sec. II E 3 of two widely separated time scales does not hold for spin glasses where there are many time scales. This fact manifests as an ambiguity in measuring τ_e depending on the range over which the linear fit is carried out. Unfortunately, τ_e , which is the intercept of the linear fit, is very sensitive to the choice of the fit window. To demonstrate this point, we measured τ_e using different fit windows in the same long run as described above. The results are shown in Fig. 2. As the window is moved toward longer times, there are two main effects: (1) the sharply rising curve for small β extends closer to the critical point and reaches a higher value, and (2) τ_e in the large β phase is larger. However, in the low temperature region, the different curves for τ_e are similar except for a scale factor. Thus, the primary effect moving the fitting window to larger times would be to sharpen the peak in the $\delta s / \delta \beta$ curve and to move computational work from the high temperature to the low temperature phase. Given a limited total number of sweeps, it is not clear how to choose the best fitting window for measuring the restricted autocorrelation time.




FIG. 3. Number of sweeps per unit β as a function of inverse temperature β for the three types of schedules: friction tensor (red, ζ), fixed culling fraction (blue, ϵ_0), and constant β -step (green, c_0). The total number of sweeps, i.e., the area under the curves is 3000 in all three cases. The c_0 schedule distributes the sweeps uniformly in temperature, whereas the ϵ_0 and ζ schedules put more work at high temperatures and close to the critical point, respectively.

In addition to τ_e , the variance of the energy, σ_e^2 is an input to the friction tensor. Our estimator of this variance is the geometric mean of population variance immediately after a resampling step and immediately before the next resampling step. The temperature step size δs for the fixed culling fraction schedule was obtained by direct measurement of δs in the same simulations used to obtain τ_e and σ_e^2 , which were carried out using the adaptive fixed culling fraction schedule described in Ref. [21].

2. Annealing schedules

Figure 3 shows $\delta s / \delta \beta$ for the three types of schedule. A single curve labeled ζ can be used for all three friction tensor schedules since they all have the same functional form for $\Delta s / \Delta \beta$. The constant culling fraction schedule and the friction tensor schedules are similar in doing very little computational work at large β and instead relying on resampling. The friction tensor schedule differs from the fixed culling schedule in doing more sweeps in the region near the spin-glass phase transition (at $\beta = 1.07$, see Ref. [36]), where the autocorrelation time becomes large. The idea that more sweeps are required near a phase transition or where the autocorrelation time becomes large has been discussed in other studies of annealing and tempering algorithms [18,20,21,37,38]. Figure 3 shows the fixed culling and friction tensor schedules for bin 3 but in practice the schedules for all bins were roughly the same and it would have been sufficient to choose a single schedule for all bins.

The culling fraction is roughly proportional to the standard deviation of the energy [see Eq. (19)], which explains why it decreases monotonically with β . The friction tensor is proportional to the product of the variance of the energy and

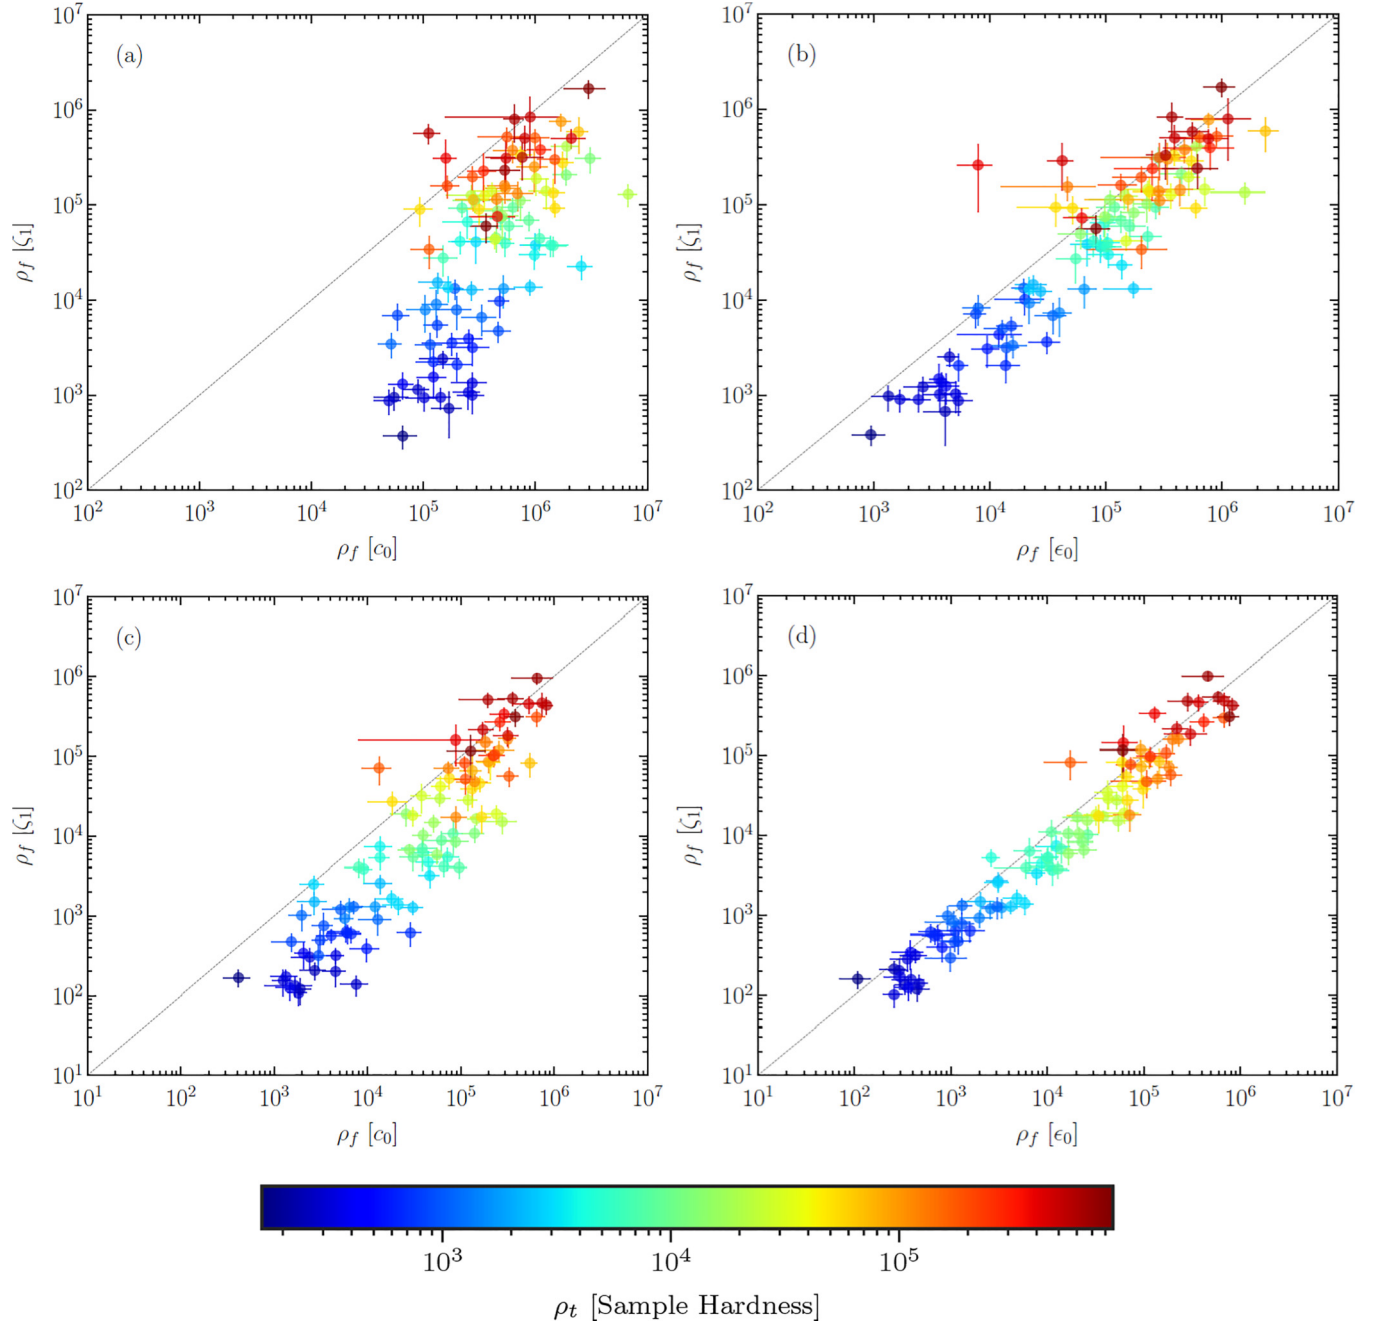


FIG. 4. Scatter plot comparison of ρ_f for the friction tensor schedules ζ_1 versus the flat schedule c_0 and constant culling schedule ϵ_0 . Each dot shows one sample with the vertical position of the dot representing ρ_f for the ζ_1 friction tensor schedule and the horizontal position representing the competing schedule, either constant β -step, c_0 or constant culling fraction ϵ . The color of the dot is the value of ρ_t , given in the color bar below the panels, for that sample obtained from long simulations [21] and represents the hardness of the corresponding instance. Easy samples are in the lower left, while the hard samples are concentrated in the upper right of each panel. Panels (a) and (b) are measured in 1000-sweep simulations, whereas (c) and (d) in 3000-sweep simulations. The value of ρ_f is not meaningful when it is greater than 10^5 so the comparison should only be trusted for the easier samples. The friction tensor schedule outperforms the other two schedules although the difference is smaller as more sweeps are used.

the restricted integrated autocorrelation time, which leads to a sharply peaked functional form for $\delta s / \delta \beta$ near the critical point.

B. Performance of annealing schedules

Figure 4 presents ρ_f values for the 100 disorder realizations comparing the three types of schedules. Each panel is

a scatter plot comparing a friction tensor schedule to one of the two other types of schedules and the coordinates of each dot represent the values of ρ_f for a single disorder realization and the two schedules labeled on the axes. The upper panels are results from 1000 sweep simulations and the lower panels form 3000 sweep simulations. The easiest samples are found in the lower left and the hardest in the upper right of each

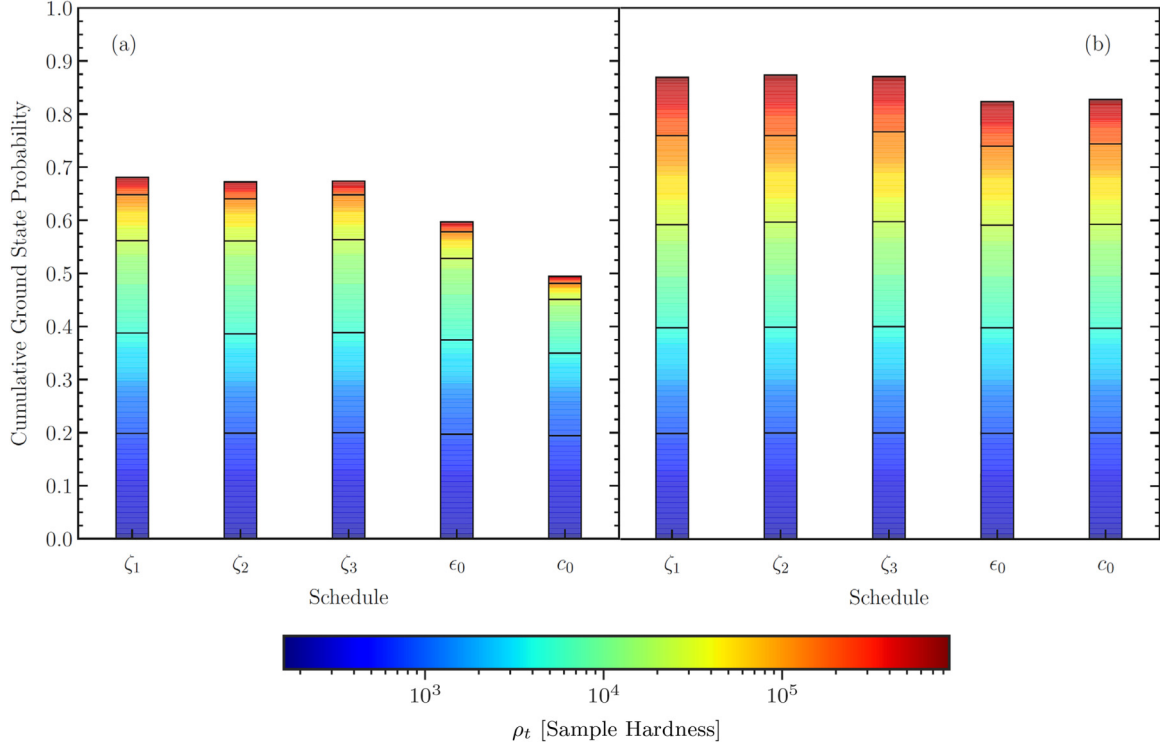


FIG. 5. Cumulative probability of finding the ground state for the five annealing schedules. The height of each bar represents the overall probability of finding the ground state for the 100 disorder samples. Each bar is divided by horizontal lines into 5 blocks each containing 20 disorder realizations whose hardness is represented by the color bar with the easiest problems in the bottom block and hardest problems in the top block. Panel (a) shows results from 1000 sweep simulations and panel (b) 3000 sweep simulations. The ground state is almost always found using all schedules for the easier problems but for the hardest problems the ground state is rarely found.

panel. We have shown only the ζ_1 friction tensor schedule. Similar scatter plots comparing the three friction tensor schedules are not shown because the ρ_f values are statistically indistinguishable.

When ρ_f is comparable to or larger than R the system is not well-equilibrated and the value of ρ_f is likely to be underestimated, both because the population size is too small and the number of independent runs to measure it is too small. For these reasons, ρ_f is best suited for comparing the performance of different schedules for the easier samples. We see in panel (a) that ρ_f for the constant β -step schedule is always near or above 10^5 suggesting it is unable to well-equilibrate any of the samples with 1000 sweeps. However the constant culling and friction tensor schedules do succeed in equilibrating most of the easier samples with 1000 sweeps. For the hardest samples, none of the schedules reach equilibrium even with 3000 sweeps. For the harder samples, we cannot draw meaningful comparisons between schedules based on the ρ_f measurements.

It is clear from Fig. 4 that the friction tensor and fixed culling schedules are far superior to the fixed β schedule for the easier samples. From Figs. 4(b) and 4(d) we see that the friction tensor schedule outperforms the fixed culling schedule though by a relatively small margin. Finally, it is interesting to observe that the differences between the schedules are most pronounced for the 1000 sweep simulations where computational resources are more limited.

We next turn to the probability of finding the ground state. This metric is better suited to comparing the performance of

the schedules for the harder samples in the study set. These samples are typically not well-equilibrated at the lowest temperature, nonetheless, the ground state is sometimes found.

Figure 5 is a bar graph showing the cumulative probability of finding ground states for all 100 samples for each of the three types of schedules. If an algorithm always found the ground state for every sample, then the height of the bar would be one. The easiest samples are at the bottom of the bar and the height of the bar at a given color level is the cumulative probability of finding ground states easier than the hardness represented by that color. The horizontal lines in each bar separate groups of two hardness bins. Figure 5(a) presents results for 1000 sweeps and Fig. 5(b) for 3000 sweeps.

First note the ground states are found with very high probability by all schedules for the two easiest bins (i.e., the height of the first horizontal line is near 0.2, the value expected if the ground state is found for all 20 samples in all 20 runs for that sample). However, for the hardest 20 samples, represented by the small rectangles at the top of each bar, the probability of finding ground states is much less than 0.2. As was the case in comparing the schedules by ρ_f the friction tensor schedule performs best and the constant β -step schedule worst. All algorithms improve in going from 1000 to 3000 sweeps and the difference between the algorithms diminishes. The most dramatic differences are found for the probability of finding the ground state for the 20 hardest samples using 1000 sweeps. The average probability per sample of finding the ground state is 0.16 for the friction tensor schedule, 0.09 for the fixed culling schedule, and 0.06 for the constant β -step schedule.

V. SUMMARY AND DISCUSSION

We have developed a formalism for quantifying the performance of annealing schedules that is applicable to population annealing and annealed importance sampling. This formalism applies to annealing schedules that are sufficiently slow that the system remains near equilibrium throughout the process. In this setting the performance metric for annealing algorithms, ρ_f is formally equivalent to the friction tensor metric introduced by Sivak and Crooks [22] for optimal control of thermodynamic systems. In the usual case of annealing in a single parameter (inverse temperature) an explicit equation determines the rate of change of the parameter during the anneal. The result is that the velocity (typically Monte Carlo sweeps per unit time) is proportional to the inverse of the square root of the friction tensor. The friction tensor itself is the product of the variance of the energy and the energy autocorrelation time.

In the interesting case of annealing in a multidimensional parameter space the optimal annealing path is a geodesic in the metric defined by the friction tensor [22]. Identifying geodesics or approximate geodesics in realistic multiparameter annealing situations is a topic for future work. Multidimensional annealing schedules could play an important role in constrained optimization where annealing parameters can be used to control the penalties associated with constraints.

In addition to near-equilibrium annealing, we also considered the common situation of ergodicity breaking associated with first-order transitions or temperature chaos. In these situations, the friction tensor formalism must be supplemented by a term associated with traversing the ergodicity-breaking region. In a first-order transition, there are two (or more) local minima in the free energy landscape and the transition occurs where the global minimum changes from a single “high temperature” local minimum to one (or more) “low temperature” minima. At some point in parameter space before the transition, ergodicity breaking generally occurs. In the context of a simple two-state model, we calculated that contribution to ρ_f is equal to the inverse of the probability of being in the low-temperature minima when ergodicity breaking occurs. In addition, we proposed a modified version of the correlation time matrix to include the presence of broken ergodicity. Overall, optimizing a multidimensional annealing path involves minimizing the sum of the friction tensor contributions from the continuous path segments and the contribution of the ergodicity-breaking transitions. It is left for future work to study how to deploy these ideas to find good annealing paths for a realistic multiparameter optimization problem.

Although the formalism applies to smooth annealing schedules that remain near equilibrium except at sharp ergodicity-breaking transitions, we expect it to be useful in more realistic settings. To explore the practical validity of these ideas, we carried out numerical simulations on the three-dimensional Ising spin glass using population annealing with inverse temperature as the sole annealing parameter. We found that an annealing schedule based on the friction tensor (modified for ergodicity breaking) outperformed schedules based on either a fixed culling fraction or constant steps in inverse temperature. The constant β -step schedule performed

very poorly related to the other two schedules because it allocates too much work to the Metropolis algorithm at low temperatures where both the variance of the energy and the modified autocorrelation time are small. However, the fixed culling fraction schedule, which has been employed in previous population annealing studies, performs almost as well as the friction tensor schedules. These two schedules turn out to be closely related. The difference is that the friction tensor includes a factor of the energy autocorrelation time and the culling fraction does not. The good performance of the fixed culling fraction schedule results perhaps from the small variation of the energy autocorrelation time when modified for ergodicity breaking. Given the simplicity of determining the culling fraction on the fly, a fixed culling fraction schedule is useful in practice. It may be that a fixed culling fraction schedule will also prove useful in the context of multiparameter annealing.

ACKNOWLEDGMENTS

We thank Arnaud Doucet, Helmut Katzgraber, Brad Lackey, and Martin Weigel for useful discussions.

APPENDIX: PAIRWISE RESIDUAL RESAMPLING

To carry out pairwise residual sampling (PRR), the set of replicas is divided into two subsets, S^+ and S^- according to whether the replica weight is greater or less than one, $S^+ = \{r | w_r > 1\}$ and $S^- = \{r | w_r < 1\}$. (It may be the case that some replicas have $w = 1$. These replicas are untouched by PRR.) We define the total positive and negative *residual* weights W^+ and W^- associated with S^+ and S^- , respectively, as

$$W^\pm = \sum_{r \in S^\pm} |w_r - 1|. \quad (\text{A1})$$

We make the assumption that the total residual weight is less than one ($W^+ + W^- < 1$), which will hold when dt is sufficiently small. Individual residual weights are then also less than one and since the sum of all w 's is R we must have that $W^+ = W^- \equiv W$. Note that the expected value of W/R is the culling fraction [see Eq. (16)].

Because the positive and negative residual weights are equal, we are allowed to resample in pairs, making an extra copy of a replica in S^+ and culling a replica in S^- so that the population size stays fixed. Furthermore, in the limit $W < 1 \ll R$, no changes to the population are made with probability $1 - W/R$ or one pair is copied/culled with probability W/R . Resampling involving more than one pair is suppressed by an extra factor of W/R and can be ignored in the large R limit. In the case that the culling/copying event is chosen, the culled replica r_- is chosen from S^- with probability $(1 - w_{r_-})/W$ and the copied replica r_+ chosen from S^+ with probability $(w_{r_+} - 1)/W$. It is straightforward to verify that requirement for the copy numbers, Eq. (6), is satisfied and that, like nearest integer resampling, for every replica in S^+ , n_r is either 1 or 2 while for every replica in S^- , n_r is either 0 or 1.

We now show that PRR is well-defined for arbitrarily small dt and that the overall contribution of resampling to the

variance of observables is negligible. Consider the variance of the population-averaged conjugate forces, \tilde{X}_k . These variances, through the Schwarz inequality, bound the contributions to the growth of ρ_f . Our aim is to show that the contribution to the growth of these variances due to resampling is negligible in the small dt and large R limit.

In a single resampling step, \tilde{X}_k is modified according to

$$\tilde{X}_k(t^+) = \frac{1}{R} \sum_r n_r X_k^r(t), \quad (\text{A2})$$

which can be rewritten as

$$\tilde{X}_k(t^+) = \tilde{X}_k(t) + \frac{1}{R} \sum_r (n_r - 1) X_k^r(t). \quad (\text{A3})$$

These equations can now be used to analyze the evolution of $\text{var}(\tilde{X}_k(t))$ in the absence of MCMC dynamics.

Using the Law of Total Variance we have that

$$\text{var}(\tilde{X}_k(t^+)) = \text{var}(\mathbb{E}(\tilde{X}_k(t^+) | \{\mathbf{X}^r(t)\})) + \langle \text{var}(\tilde{X}_k(t^+) | \{\mathbf{X}^r(t)\}) \rangle_\Lambda, \quad (\text{A4})$$

where the conditional expectation and conditional variance are carried out holding the population at time t fixed but allowing randomness in the choice of copy numbers, $\{n_r\}$ given the weights $\{w_r\}$. The unconditioned variances are with respect to the full nonequilibrium distribution induced by PA. Using Eq. (A2) and the fact that the expectation of n_r is w_r , the first term on the RHS of Eq. (A4) can be written as

$$\text{var}(\mathbb{E}(\tilde{X}_k(t^+) | \{\mathbf{X}^r(t)\})) = \frac{1}{R^2} \text{var}\left(\sum_r w_r X_k^r\right), \quad (\text{A5})$$

where in this equation and henceforth in the Appendix, a variable without a time argument will assumed to be at the time immediately before the resampling event t .

In the small dt limit we can use Eq. (11) and expand to leading order in dt to obtain

$$\text{var}(\mathbb{E}(\tilde{X}_k(t^+) | \{\mathbf{X}^r\})) = \frac{1}{R^2} \text{var}\left(\sum_r (1 - dt \dot{\lambda}_i (X_i^r - \tilde{X}_i)) X_k^r\right). \quad (\text{A6})$$

If the population is close to equilibrium and uncorrelated, then we can obtain the approximate expression,

$$\begin{aligned} \text{var}(\mathbb{E}(\tilde{X}_k(t^+) | \{\mathbf{X}^r(t)\})) \\ = \text{var}(\tilde{X}_k) + \frac{2}{R} dt \dot{\lambda}_i \text{Cov}(\delta X_i X_k, X_k), \end{aligned} \quad (\text{A7})$$

where the covariance is respect to the Gibbs distribution at parameter values $\lambda(t)$. We will show below that the second term in the decomposed variance Eq. (A4) is of negligible magnitude compared to $\text{var}(\tilde{X}_k)$. This enables us to rewrite the above relation as the differential equation

$$\frac{d\text{var}(\tilde{X}_k)}{dt} = \frac{2}{R} \dot{\lambda}_i \text{Cov}(\delta X_i X_k, X_k), \quad (\text{A8})$$

which describes the change in variances due to the change in the parameters of the Gibbs distribution. This equation is not exact because of correlation in the population but it reveals the order in R dependence. The $1/R$ behavior of $\text{var}(\tilde{X}_k(t))$ is expected since \tilde{X}_k is an average over a population of size R .

The second term on the RHS of Eq. (A4) contains the contribution to $\text{var}(\tilde{X}_k)$ arising from resampling. Our goal is to show that this term can be integrated in the $dt \rightarrow 0$ limit and is higher order in $1/R$. From Eq. (A3) and noting that $\tilde{X}_k(t)$ is fixed by the conditioning we have

$$\langle \text{var}(\tilde{X}_k(t^+) | \{\mathbf{X}^r\}) \rangle_\Lambda = \frac{1}{R^2} \left\langle \text{var}\left(\sum_r (n_r - 1) X_k^r | \{\mathbf{X}^r\}\right) \right\rangle_\Lambda. \quad (\text{A9})$$

The argument of the variance on the RHS of this equation is a random variable that is equal to zero with probability $1 - W/R$ and equal to $(X_k^{r^+} - X_k^{r-})$ with probability W/R . Note that the argument of the variance involves at most two terms in the sum over the population. The culling probability [see Eq. (16)] is an infinitesimal given by

$$W/R = \frac{dt}{2R} \sum_r |\dot{\lambda}_i (X_i^r - \tilde{X}_i)|, \quad (\text{A10})$$

and near equilibrium can be approximated by $dt \langle |\dot{\lambda}_i \delta X_i| \rangle_{\lambda(t)}$.

These observation lead to the conclusion that the resampling contribution to the increase in variance is order dt so it can be integrated and is well-defined in the $dt \rightarrow 0$ limit. Furthermore, the contribution from resampling is order $1/R^2$, whereas $\text{var}(\tilde{X}_k)$ is order $1/R$ so the resampling contribution is negligible in the large R limit.

-
- [1] S. Kirkpatrick Jr., C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, *Science* **220**, 671 (1983).
 [2] R. M. Neal, Annealed importance sampling, *Stat. Comput.* **11**, 125 (2001).
 [3] 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*, edited by James E. Gubernatis, Vol. 690 (AIP, College Park, MD, 2003), pp. 200–206.

- [4] J. Machta, Population annealing with weighted averages: A Monte Carlo method for rough free-energy landscapes, *Phys. Rev. E* **82**, 026704 (2010).
 [5] C. Geyer, Monte Carlo maximum likelihood for dependent data, in *Computing Science and Statistics: 23rd Symposium on the Interface*, edited by E. M. Keramidas (Interface Foundation, Fairfax Station, VA, 1991), p. 156.
 [6] K. Hukushima and K. Nemoto, Exchange Monte Carlo method and application to spin-glass simulations, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).

- [7] E. Marinari and G. Parisi, Simulated tempering: A new Monte Carlo scheme, *Europhys. Lett.* **19**, 451 (1992).
- [8] F. Wang and D. P. Landau, Determining the density of states for classical statistical models: A random walk algorithm to produce a flat histogram, *Phys. Rev. E* **64**, 056101 (2001).
- [9] J. Skilling, Nested sampling for general Bayesian computation, *Bayesian Anal.* **1**, 833 (2006).
- [10] C. Jarzynski, Equilibrium free-energy differences from nonequilibrium measurements: A master-equation approach, *Phys. Rev. E* **56**, 5018 (1997).
- [11] P. J. Reynolds, J. Tobochnik, and H. Gould, Diffusion quantum Monte Carlo, *Comput. Phys.* **4**, 662 (1990).
- [12] M. Jarret and B. Lackey, Substochastic Monte Carlo algorithms, [arXiv:1704.09014](https://arxiv.org/abs/1704.09014).
- [13] S. Geman and D. Geman, Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images, *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-6**, 721 (1984).
- [14] B. Hajek, Cooling schedules for optimal annealing, *Math. Oper. Res.* **13**, 311 (1988).
- [15] H. Cohn and M. Fielding, Simulated annealing: Searching for an optimal temperature schedule, *SIAM J. Optim.* **9**, 779 (1999).
- [16] A. G. Nikolaev and S. H. Jacobson, Simulated annealing, in *Handbook of Metaheuristics* (Springer, Berlin, 2010), pp. 1–39.
- [17] P. Salamon, J. D. Nulton, J. R. Harland, J. Pedersen, G. Ruppreiner, and L. Liao, Simulated annealing with constant thermodynamic speed, *Comput. Phys. Commun.* **49**, 423 (1988).
- [18] B. Andresen and J. M. Gordon, Constant thermodynamic speed for minimizing entropy production in thermodynamic processes and simulated annealing, *Phys. Rev. E* **50**, 4346 (1994).
- [19] L. Y. Barash, M. Weigel, M. Borovský, W. Janke, and L. N. Shchur, GPU accelerated population annealing algorithm, *Comput. Phys. Commun.* **220**, 341 (2017).
- [20] A. Barzegar, C. Pattison, W. Wang, and H. G. Katzgraber, Optimization of population annealing Monte Carlo for large-scale spin-glass simulations, *Phys. Rev. E* **98**, 053308 (2018).
- [21] C. Amey and J. Machta, Analysis and optimization of population annealing, *Phys. Rev. E* **97**, 033301 (2018).
- [22] D. A. Sivak and G. E. Crooks, Thermodynamic metrics and optimal paths, *Phys. Rev. Lett.* **108**, 190602 (2012).
- [23] P. R. Zulkowski, D. A. Sivak, G. E. Crooks, and M. R. DeWeese, Geometry of thermodynamic control, *Phys. Rev. E* **86**, 041148 (2012).
- [24] P. Del Moral, A. Doucet, and A. Jasra, Sequential Monte Carlo samplers, *J. R. Stat. Soc. B* **68**, 411 (2006).
- [25] R. Douc and O. Cappe, Comparison of resampling schemes for particle filtering, in *Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis (ISPA)* (IEEE, Piscataway, NJ, 2005), pp. 64–69.
- [26] D. Gessert, W. Janke, and M. Weigel, Resampling schemes in population annealing: Numerical and theoretical results, *Phys. Rev. E* **108**, 065309 (2023).
- [27] W. Wang, J. Machta, and H. Katzgraber, Population annealing: Theory and application in spin glasses, *Phys. Rev. E* **92**, 063307 (2015).
- [28] C. Jarzynski, Nonequilibrium equality for free energy differences, *Phys. Rev. Lett.* **78**, 2690 (1997).
- [29] M. Weigel, L. Barash, L. Shchur, and W. Janke, Understanding population annealing Monte Carlo simulations, *Phys. Rev. E* **103**, 053301 (2021).
- [30] N. Chopin, Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference, *Ann. Stat.* **32**, 2385 (2004).
- [31] G. M. Rotskoff and G. E. Crooks, Optimal control in nonequilibrium systems: Dynamic Riemannian geometry of the Ising model, *Phys. Rev. E* **92**, 060102(R) (2015).
- [32] J. Machta, Strengths and weaknesses of parallel tempering, *Phys. Rev. E* **80**, 056706 (2009).
- [33] 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).
- [34] A. J. Bray and M. A. Moore, Chaotic nature of the spin-glass phase, *Phys. Rev. Lett.* **58**, 57 (1987).
- [35] W. Wang, J. Machta, and H. G. Katzgraber, Chaos in spin glasses revealed through thermal boundary conditions, *Phys. Rev. B* **92**, 094410 (2015).
- [36] H. G. Katzgraber, M. Körner, and A. P. Young, Detailed study of universality in three-dimensional Ising spin glasses, *Phys. Rev. B* **73**, 224432 (2006).
- [37] H. G. Katzgraber, S. Trebst, D. A. Huse, and M. Troyer, Feedback-optimized parallel tempering Monte Carlo, *J. Stat. Mech.* (2006) P03018.
- [38] E. Bittner, A. Nußbaumer, and W. Janke, Make life simple: Unleash the full power of the parallel tempering algorithm, *Phys. Rev. Lett.* **101**, 130603 (2008).