

# PHYSICAL REVIEW LETTERS

VOLUME 74

20 MARCH 1995

NUMBER 12

## Monte Carlo Simulation and Global Optimization without Parameters

Bobby Hesselbo and R. B. Stinchcombe

*Theoretical Physics, University of Oxford, 1 Keble Road, Oxford, OX1 3NP, United Kingdom*

(Received 31 May 1994)

We propose a new ensemble for Monte Carlo simulations, in which each state is assigned a statistical weight  $1/k$ , where  $k$  is the number of states with smaller or equal energy. This ensemble has robust ergodicity properties and gives significant weight to the ground state, making it effective for hard optimization problems. It can be used to find free energies at all temperatures and picks up aspects of critical behavior (if present) without any parameter tuning. We test it on the traveling salesperson problem, the Edwards-Anderson spin glass, and the triangular antiferromagnet.

PACS numbers: 02.70.Lq, 02.50.Ng, 02.60.Pn, 64.60.Ak

The method of Monte Carlo simulation has proved very useful for studying the thermodynamic properties of model systems with moderately many degrees of freedom. The idea is to sample the system's phase space stochastically, using a computer to generate a series of random configurations. We take the phase space to consist of  $N$  discrete states (with label  $i$ ), though the method applies equally to continuous systems. Often only a tiny fraction of the phase space (the part at low energy) is relevant to the properties being studied, due to the strong variation of the Boltzmann weight  $\exp(-\beta E_i)$  in the canonical ensemble (CE). It is then helpful to sample in an ensemble (with relative weights  $w_i$  and absolute probabilities  $p_i = w_i / \sum_j w_j$ ), which is concentrated on this region of phase space. The Metropolis algorithm [1] samples directly in the CE, and is good at determining many physical properties (with the notable exception of the free energy). The price to be paid for this is that successive configurations are not independent (typically they have a single microscopic difference), but instead form a Markov chain with some equilibration time  $t_{\text{eq}}(w_i)$ .

We may distinguish two important characteristics of a Monte Carlo simulation: its ergodicity [measured by  $t_{\text{eq}}(w_i)$ ] and its pertinence [measured by  $N_s(w_i; I)$ , the average number of independent samples needed to obtain the information  $I$  that we seek]. We should choose  $w_i$  so as to minimize the total number of configurations that need to be generated, which is proportional to  $t_{\text{eq}}(w_i)N_s(w_i; I)$ . It is easy to specify an ensemble which

would yield the sought information if independent samples could be drawn from it, but an ensemble with too much weight at low energies may become fragmented into "pools" at the bottoms of "valleys" of the energy function, and so have a large equilibration time. For example, it is well known that at low temperatures the Metropolis algorithm can get stuck in ordered or glassy phases. Ergodicity may be improved by sampling instead in a nonphysical ensemble with a broad energy distribution, which allows the valleys to be connected by paths passing through higher energies [2-4]. A weight assignment leading to such a distribution cannot in general be written as an explicit function of energy alone; rather it is an algorithm's purpose to find this assignment, which then tells us about the density of states  $\rho(E)$ . This reversal (starting with the distribution and finding the weights) of the usual Monte Carlo process can be achieved using a series of normal simulations, adjusting the weight  $w_i$  after each run so that the resulting energy distribution  $\rho_{w_i}(E)$  converges to the desired one. Although one might need more samples from such a broad energy ensemble (BEE) than from a particular CE in order to find properties relating to that CE, it is possible for a single BEE simulation to provide information on properties over a range of temperatures. BEEs are also helpful for finding free energies [5] (since relative normalizations can be determined only for overlapping distributions [1,6]) and for sampling across regions of negative heat capacity in the vicinity of first-order phase transitions [7].

The energy distribution used by a BEE algorithm is a free parameter [8], and is often taken to be uniform [this was called the “multicanonical ensemble” (MCE) by Berg and Celik [2]]. It would, however, be natural to look for an optimal most general distribution, i.e., one with the best worst-case performance in terms of ergodicity and pertinence. We will consider only monotonically decreasing weight assignments  $w_i$ , implemented using the Metropolis scheme of accepting a transition  $i \rightarrow j$  with probability  $\min(w_j/w_i, 1)$ . Our proposal is to use the ensemble with weight

$$w_i = 1/k_i, \quad (1)$$

where  $k_i$  is the number of states with energies up to and including  $E_i$ . This ensemble has the property that  $\log(N)N_s$  independent samples from it convey as much information, concerning any property, as  $N_s$  independent samples from any rival ensemble [9] (the factor  $\log N$ , which is a measure of the ensemble’s worst-case pertinence, is smaller for this ensemble than for any other). In particular, of order  $\log N$  independent samples from this ensemble are sufficient both to find the ground state and to determine the normalization of the density of states. While the best worst-case ergodicity is probably obtained by sampling at infinite temperature, this is useless in terms of pertinence. We expect reasonable ergodicity for the  $1/k$  ensemble since if we require a rival ensemble to assign an equal probability to some state, then its transition rates from this state may exceed those in the  $1/k$  ensemble by a factor of at most  $\log N$ . In contrast, the equilibration time for uniform energy sampling may be made arbitrarily large by choosing a suitably unreasonably reparametrized Hamiltonian  $H' = f(H)$ , where  $f(H)$  is a monotonically increasing function (the  $1/k$  ensemble is invariant under such operations).

The  $1/k$  ensemble is equivalent to uniform entropy sampling [i.e.,  $\rho_{1/k}(E) \propto dS/dE \equiv 1/T(E)$ ] since for practical purposes the entropy  $S$  is given by  $S(E_i) \sim \log k_i$ . Like the CE, it has a sensible thermodynamic limit in that the relative weight of states with a single microscopic difference remains of order unity as  $M \rightarrow \infty$ , where  $M$  is the system size. However, whereas in the CE fluctuations in intensive quantities such as energy density typically go to zero like  $M^{-1/2}$ , in the  $1/k$  ensemble they are independent of  $M$ , with the result that the  $1/k$  ensemble is non-self-averaging even for simple systems such as the ferromagnetic Ising model. For example, if the physical system has a second-order phase transition at some temperature  $T_c$ , this will be reflected by a power law contribution to the spin-spin correlation function in the  $1/k$  ensemble [10], with a new exponent

$$G_{1/k}(r) - G_{1/k}(\infty) \sim r^{-(d-2+\eta)-(1-\alpha)/\nu}. \quad (2)$$

In spite of this, the correlation function (determined by spatial average) for a state drawn from the  $1/k$  ensemble is likely to be exponentially decaying, with a random

correlation length. To obtain (2), we first note that uniform sampling of the entropy leads to smooth (i.e., at least once differentiable) sampling of the energy, at least in systems or regimes where the heat capacity and temperature are strictly positive, since

$$\frac{d\rho_{1/k}(E)}{dE} \propto -\frac{1}{T^2 C}. \quad (3)$$

The  $1/k$  ensemble may be expressed as a linear combination of canonical ensembles (in the thermodynamic limit):

$$\begin{aligned} \rho_{1/k} &\xrightarrow{N \rightarrow \infty} \int p_{\text{CE}}(T(\bar{E})) \rho_{1/k}(\bar{E}) d\bar{E} \\ &\equiv \int p_{\text{CE}}(T) \rho_{1/k}(T) dT, \end{aligned} \quad (4)$$

where  $\bar{E}$  is the normalized energy and  $p$  represents any probability assigned in an ensemble, since relative fluctuations in the CE go to zero in the thermodynamic limit. Close to the critical energy [letting  $t = (T - T_c)/T$ ] we find

$$\rho_{1/k}(t) \sim t^{-\alpha}, \quad (5)$$

where  $\alpha$  is the critical exponent describing the divergence of the heat capacity. Under a real-space renormalization with scale factor  $b$ ,  $\rho_{1/k}(t)$  is carried by the flow ( $t \rightarrow b^{1/\nu} t$ , where  $\nu$  is the exponent describing the divergence of the correlation length) away from the fixed point, and so reduces by a factor  $b^{-(1-\alpha)/\nu}$ . Thus there is a contribution to  $G_{1/k}(r)$  which, as in the canonical critical ensemble, scales under renormalization group transformations, though with an extra factor of  $b^{-(1-\alpha)/\nu}$ . This reflection of critical properties (which normally require parameter tuning) in the  $1/k$  ensemble shows that it in some sense exhibits (by means of nontrivial probability distributions) possible behaviors of the system over all temperatures.

In principle,  $1/k$  sampling may be implemented by an algorithm whose only parameters [11] are the number of Monte Carlo steps to use at each stage of the convergence process (which should be enough for equilibration to have occurred, and might be determined by the algorithm). Specifically, we may represent the  $n$ th approximation  $\rho^n(E)$  to the density of states as a set of delta functions and use the recurrence

$$\rho^{n+1}(E) = \frac{\sum_{\text{samples}} \delta(E - E_i)/w_i^{n+1}}{\sum_{\text{samples}} 1/w_i^{n+1}}, \quad (6)$$

with

$$w_i^{n+1} = \begin{cases} \sigma^n(E)^{-1} & \text{if } E \geq E_{\min}^n, \\ \widetilde{\sigma}^n(E)^{-1} & \text{if } E < E_{\min}^n, \end{cases} \quad (7)$$

where  $\sigma^n(E) = \int_{-\infty}^{E_i} \rho^n(E) dE$  is the integrated density of states and  $\widetilde{\sigma}^n(E)$  is an extrapolation of  $\sigma^n(E)$  below the lowest sampled energy  $E_{\min}^n$ . Note that when sampling

a continuous space, one should use  $w_i^{n+1} = [\sigma^n(E) + \sigma_{\text{offset}}]^{-1}$  in order to make the sampled entropy range finite.  $\sigma^n(E)$  may be evaluated in order  $\log N_\delta^n$  steps, where  $N_\delta^n$  is the number of delta functions used to represent  $\rho^n$  (memory constraints may limit  $N_\delta^n$  so that some grouping procedure is needed for the delta functions). We have also found it useful to represent  $\rho^n(E)$  as a histogram and to compute and store the bias function  $w_i$  before each run. One requires only that the histogram is fine enough to resolve variation in  $\rho(E)$ . Uniform energy sampling, in contrast, needs a specific choice of histogram, which must be coarse enough to have good statistics. In  $1/k$  sampling, Eq. (7) automatically interpolates as finely as permitted by the data, short of curve fitting. However, curve fitting is helpful in determining  $\tilde{\sigma}^n(E)$ , since with each run the range of energies being sampled increases to cover energies where the predicted  $\sigma(E)$  used in (7) is not wrong by a large factor. The first run may use  $w_i^0 = \text{const}$ , which is likely to lead to progressively increasing equilibration times in the following runs as the sampled energy range extends further down.

The improved ergodicity of BEEs makes them attractive for use in hard optimization problems [3,4]. While their applicability may be similar to that of simulated annealing (see, e.g., [12]), their behavior differs in that they offer “open-ended” improvement, since they never commit to a particular valley, but continue to search for better solutions. They also dispense with the need for a cooling schedule, which is a crucial parameter for simulated annealing algorithms. Although a “cautious” BEE algorithm may spend most of its time visiting highly nonoptimal configurations, this could be offset by using parallel computation, such as one might anticipate being readily available in the future (equilibration time, on the other hand, is a basic constraint on an algorithm’s performance).

Our first test of  $1/k$  sampling is a 100-city traveling salesperson problem (see, e.g., [13]), with moves consisting in segment transport or reversal, following [14]. Figure 1 shows  $\sigma^{45}(E)$ , where each run was continued until  $2 \times 10^6$  transitions had been accepted, and (7) was used with the trivial extrapolation  $\tilde{\sigma}^n(E < E_{\min}^n) = \sigma^n(E_{\min}^n)/2$ .

An additional run was conducted starting with a randomized configuration but using the previously obtained density of states; Fig. 2 shows the length of the best-so-far tour as a function of the number of cost evaluations [which we consider to be more relevant than computer time since it is more characteristic of an algorithm and since some optimization problems, e.g., protein folding (which has been studied using the MCE [15]), may involve expensive cost calculations]. This should be regarded as an upper limit for the performance of  $1/k$  sampling in that not all of the 45 iterations of (6) could be eliminated by extrapolating the density of states, and as a lower limit in that no parallelism was used.

If we know  $N$ , then the absolute value of  $\rho_{\min}^n$  provides a useful measure of progress during global optimization,

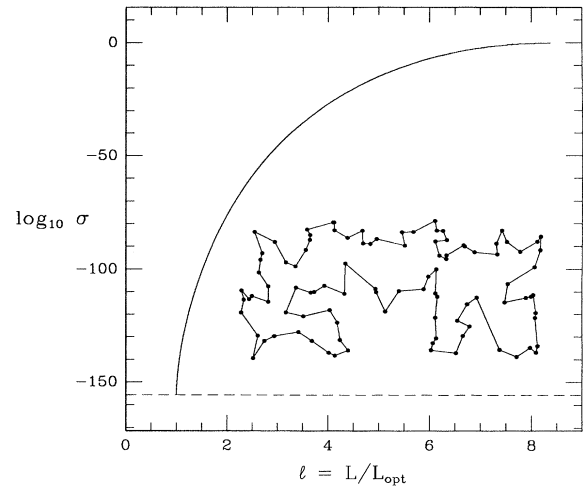


FIG. 1. Integrated density of states for the archived traveling salesperson problem “kroA100” [18,19]. Normalization is with respect to the established optimal tour length  $L_{\text{opt}}$ , as listed in the archive. The dashed line shows  $\rho_{\min} = 2/99$ . Inset: the optimal tour.

since  $N\rho_{\min}^n$  serves as an estimate for the number of states at or below the lowest energy sampled (assuming ergodicity). In this way, using runs of  $16 \times 10^6$  accepted transitions on the problem instance shown in Fig. 1, we obtained a ground state entropy of  $0.15 \pm 0.15$  bit, with a variance of 0.6 bit.

In order to compare  $1/k$  sampling with the multicanonical ensemble, we performed simulations of the Edwards-Anderson model with Ising spins  $s_i = \pm 1$  and

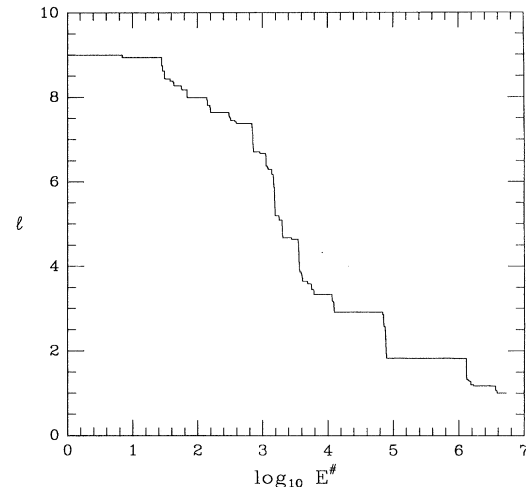


FIG. 2. Length of the best-so-far tour as a function of the number of cost evaluations  $E^\#$  for a particular run. Among 10 such runs, the number of cost evaluations required to find the optimal tour varied between  $\sim 2 \times 10^6$  and  $\sim 64 \times 10^6$ . The plateaus are due to excursions back up to nonoptimal configurations.

nearest-neighbor interactions  $J_{ij} = \pm 1$  (with  $\sum J_{ij} = 0$ ), on a  $12 \times 12$  square lattice with periodic boundary conditions. Figure 3 shows the energy visitation densities  $H(E)$  and the calculated entropy,  $s(E) = \log 2 + \log[\sigma(E)]/12^2$ , for one realization. For nine realizations we computed the ergodicity times in sweeps (MC steps per spin), following [2]. We found  $\tau_{1/k}^e$  to vary between 1199 and 19512, with median 2025, while  $\tau_{1/k}^e/\tau_{\text{MCE}}^e$  was more sharply peaked at  $0.69 \pm 0.04$ . The ground state entropies were  $s(E_0) = 0.080 \pm 0.019$  nat per spin.

The last application reported here is a simulation of a regular system with frustration, the triangular antiferromagnet, on a  $48 \times 48$  parallelogram with periodic boundary conditions. Using five runs of  $7.4 \times 10^5$  sweeps, we obtained a ground state entropy of 0.323 20, with a variance of 0.000 15, which is consistent with the exact bulk value [16] of  $(2/\pi) \int_0^{\pi/3} \log(2 \cos \omega) d\omega \approx 0.323 07$ .

These simulations show that  $1/k$  sampling has significant advantages over existing techniques. For the traveling salesperson problem it found the global optimum, its only parameter being the number of iterations to use. Lee and Choi [17] have obtained good results for large scale traveling salesperson problems using a “multicanonical annealing” algorithm which is based on the MCE, but constrained to a certain energy range which is then “annealed.” While this approach is less greedy than simulated annealing, we believe that ergodic algorithms will have a higher probability of finding the global optimum in the limit of many samples or of much parallelization.  $1/k$  sampling may, however, benefit from being truncated above some fixed energy, provided this isn’t so low as to compromise ergodicity. The results for the spin glass show that  $1/k$  sampling has faster equilibration and more weight for low-

lying states than the MCE, though it would be worthwhile to continue this comparison to larger systems. It would also be interesting to compare the variance of the ground state entropy results for the triangular antiferromagnet with that obtained by other methods.

$1/k$  sampling may also be useful for determining the functional form of a density of states, since it is completely impartial on account of its reparametrization invariance. Unfortunately, the equilibration times of BEE algorithms are rather long, going as  $M^2$  in the best case and as more than  $M^3$  for the Edwards-Anderson spin glass [2]. While BEE algorithms may be unnecessarily cautious for studying well-behaved systems when free energies are not required, the slower equilibration for the spin glass probably reflects the intrinsic difficulty of this problem. We suggest that  $1/k$  sampling may be especially useful for obtaining complete and reliable information on the properties of relatively small systems, since it has, among a large class of ensembles, the most general applicability in terms of the number of independent samples needed, combined with robust ergodicity properties and a minimum requirement for input from the user.

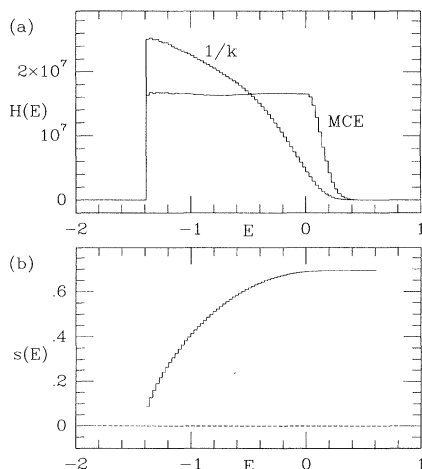


FIG. 3. Results for simulations using  $6.4 \times 10^6$  sweeps on one realization of the  $12 \times 12$  Edwards-Anderson spin glass. (a) Histogram  $H(E)$  of the energy visitation density in the  $1/k$  ensemble and in our implementation of the MCE following [2]. (b) The entropy per spin for the same system.

- [1] N. Metropolis, M. Rosenbluth, M. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
- [2] B. A. Berg and T. Celik, *Phys. Rev. Lett.* **69**, 2292 (1992); B. A. Berg, *Int. J. Mod. Phys.* **3**, 311 (1992).
- [3] E. Marinari and G. Parisi, *Europhys. Lett.* **19**, 451 (1992); A. P. Lyubartsev *et al.*, *J. Chem. Phys.* **96**, 1776 (1992).
- [4] B. A. Berg, *Nature (London)* **361**, 708 (1993).
- [5] G. M. Torrie and J. P. Valleau, *J. Comput. Phys.* **23**, 187 (1977). This appears to have been the first use of a BEE, and was called “umbrella sampling.”
- [6] J. P. Valleau and D. N. Card, *J. Chem. Phys.* **57**, 5457 (1972).
- [7] B. A. Berg and T. Neuhaus, *Phys. Rev. Lett.* **68**, 9 (1992).
- [8] J. Lee, *Phys. Rev. Lett.* **71**, 211 (1993); **71**, 2353 (1993).
- [9] This is because  $\log(N)N_s$  samples from the  $1/n$  ensemble may be processed (by deleting each sample with a suitable probability) to produce  $N_s$  samples from a rival ensemble  $w'_i$ . The  $1/n$  ensemble is optimal in this sense since we could want to know about properties arbitrarily close to some energy  $E_i$ , in which case  $(p'_i/p_i)N_s$  samples from  $w_i$  tell us the same as  $N_s$  samples from  $w'_i$ , and the probability assigned by the  $1/n$  ensemble is everywhere smaller by the same factor (viz  $\log N$ ) in comparison with the greatest probability that any decreasing distribution could assign.
- [10] Note that the connected part of the correlation function  $G_{1/n}(r)$  does not tend to zero at infinity since no symmetry is broken in the  $1/n$  ensemble—distant spins really are correlated. A function which does go to zero is  $G'_{1/n} = \langle m_i m_j \rangle - \langle m_i \text{sgn}(S) \rangle \langle m_j \text{sgn}(S) \rangle$ , in which the subtracted part may be interpreted as the value of  $\langle m_i m_j \rangle$  that would be predicted from separate observation of each spin together with the total spin  $S$ .
- [11] We regard the choice of allowed transitions as being part of the problem specification since it is problem

- specific. A good choice, which should as far as possible allow progressive energy reductions, is nevertheless very important.
- [12] S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi, *Science* **220**, 671 (1983).
- [13] Z. Michalewicz, in *Gentic Algorithms + Data Structures = Evolution Programs* (Springer-Verlag, Heidelberg, 1992), Chap. 10.
- [14] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, in *Numerical Recipes in C* (Cambridge Univ. Press, Cambridge, UK, 1992), 2nd ed., Chap. 10.9.
- [15] U.H.E. Hansmann and Y. Okamoto, *J. Comp. Chem.* **14**, 1333 (1993).
- [16] G.H. Wannier, *Phys. Rev.* **79**, 357 (1950).
- [17] J. Lee and M. Y. Choi, *Phys. Rev. E* **50**, R651 (1994).
- [18] G. Reinelt, *ORSA J. Comput.* **3**, 376 (1991).
- [19] The TSP library may be obtained by anonymous ftp to titan.cs.rice.edu as the file "tsplib.tar" in the directory "public."