Convergence of Monte Carlo simulations to equilibrium

Onuttom Narayan* and A. P. Young†
Department of Physics, University of California, Santa Cruz, California 95064
(Received 19 July 2000; published 16 July 2001)

We give two direct, elementary proofs that a Monte Carlo simulation converges to equilibrium provided that appropriate conditions are satisfied. The first proof requires detailed balance while the second is quite general.

DOI: 10.1103/PhysRevE.64.021104 PACS number(s): 05.10.Ln, 75.40.Mg

Monte Carlo simulations are widely used in statistical physics. If the algorithm satisfies the detailed balance condition, it is easy to show that the desired distribution is a *stationary* distribution, i.e., if the system is, by some means, put into the desired distribution, it will subsequently stay in this distribution. It is harder, but nonetheless crucial, to also show that, starting from a general distribution, the algorithm will *converge* to the desired distribution. Although this can be proved without too much difficulty for a system with a finite number of states [1], this proof is unfamiliar to most physicists. The argument in the original paper of Metropolis *et al.* [2] makes convergence plausible but is not a proof [3]. Most physics texts on Monte Carlo methods either do not give a proof of convergence [4] or refer the reader [5,6] to rather abstract derivations in the mathematics literature [7], or rely on the Frobenius-Perron theorem [6] which is unfamiliar to most physicists.

In this paper we present two proofs of convergence which are self-contained and use only elementary methods. We feel that it is useful to present these derivations here because (i) it is not widely known in the physics community that it is not difficult to prove convergence, at least for systems with a finite number of states, and (ii) our proofs are (to the best of our knowledge) different from and as simple as existing proofs in the mathematical literature. Even the proof of Ref. [1], which is of comparable simplicity, is hard to understand physically; this is discussed more fully near the end of the paper.

The first proof relies on the Monte Carlo algorithm satisfying the condition of *detailed balance*. Although this is true in essentially all Monte Carlo simulations, it is not strictly required for the algorithm to converge to the equilibrium distribution. In the second half of this paper, we shall also present a more general proof which relaxes this condition, and which is very different from the proof assuming detailed balance.

Throughout this paper, we shall assume that the system being considered has a finite number of states. Most systems in physics can be approximated as such by discretizing any continuous variables sufficiently finely; for instance, in molecular simulations, it should be permissible to limit the phase space for any particle to a sufficiently large region, and then discretize it in small intervals. A rigorous proof of convergence for systems with an infinite number of states is much more complicated [11].

*Present address: Physics Department, Indian Institute of Science, Bangalore 560012, India.
†Corresponding author.

In the simplest case, the desired distribution that the Monte Carlo method seeks to simulate is the Boltzmann distribution at some temperature. However, in dealing with glassy systems, it is sometimes more efficient to simulate a different distribution, as in the multicanonical ensemble [8], the $1/k$ ensemble [9], and parallel tempering [10]. The results presented in this paper are valid whenever the desired distribution is a stationary distribution of the algorithm, and detailed balance—or, more generally, ergodicity—is satisfied. This includes the cases mentioned above.

The essential ingredients of the Monte Carlo method in statistical physics are the (non-negative) ''transition rates,'' $w_{l\to m}$, defined to be the probability that, given the system is in state l at ''time'' t, then it will be in state $m(\neq l)$ at time $t+1$. We define time to be incremented by one every Monte Carlo move (not sweep) and assume initially that all moves are equivalent, so the $w_{l\to m}$ do not depend on time. An example would be flipping a single spin chosen at random. The important case of sequential updating will be discussed later.

The probability that the system is in state l at time t is defined to be $P_l(t)$. The evolution of these probabilities is governed by the ''master equation,''

$$P_l(t+1)-P_l(t)=\sum_{m\neq l}\left[P_m(t)\,w_{m\to l}-P_l(t)\,w_{l\to m}\right]. \tag{1}$$

The first term on the right-hand side describes transitions into state l from m (which therefore increases P_l and so has a plus sign) while the second term describes transitions out of state l, which decreases $P(l)$. Note that only terms with $m\neq l$ contribute. We can also define $w_{l\to l}$ to be the probability that the system stays in state l, i.e., $w_{l\to l}=1-\Sigma_{m\neq l}w_{l\to m}$, or equivalently,

$$\sum_m w_{l\to m}=1. \tag{2}$$

Equation (2) implies that the master equation can be written

$$P_l(t+1)=\sum_m P_m(t)\,w_{m\to l}, \tag{3}$$

where the term $m=l$ is now included.

A necessary condition for the method to work is that the desired distribution, $P^{\rm eq}$, is *stationary*, i.e., if $P_l(t)=P_l^{\rm eq}$ for all l then $P_l(t+1)=P_l^{\rm eq}$. This means that the right-hand side of Eq. (1) must vanish for $P=P^{\rm eq}$. The condition of *detailed*

balance consists of the assumption that *each* term on the right-hand side of Eq. (1) separately vanishes for $P=P^{\rm eq}$, i.e.,

$$P_l^{\rm eq} w_{l\to m} = P_m^{\rm eq} w_{m\to l}. \tag{4}$$

In the first part of this paper, we shall assume that Eq. (4) is satisfied.

We start with the following quantity, which is a measure of the deviation from equilibrium,

$$G = \sum_l \frac{1}{P_l^{\rm eq}} (P_l - P_l^{\rm eq})^2 = \sum_l \left(\frac{P_l^2}{P_l^{\rm eq}} - 1 \right) \tag{5}$$

evaluated at time t, where the last expression follows because P and $P^{\rm eq}$ are normalized.

At time $t+1$ we indicate (for compactness of notation) the probabilities by P_l' and the corresponding value of G by G'. We will show that G monotonically decreases, i.e.,

$$\Delta G \equiv G' - G \leqslant 0, \tag{6}$$

where the equality only holds if G and G' both vanish, so the system is in equilibrium. This shows that the system will eventually approach arbitrarily close to the equilibrium distribution.

Using Eqs. (3) and (5), ΔG can be written as

$$\Delta G = \sum_{l,m,n} \left[w_{m\to l} w_{n\to l} \frac{P_m P_n}{P_l^{\rm eq}} \right] - \sum_l \frac{P_l^2}{P_l^{\rm eq}}. \tag{7}$$

In the first term on the right-hand side of Eq. (7) we use the detailed balance condition, Eq. (4), to replace $w_{m\to l}$ by $w_{l\to m} P_l^{\rm eq}/P_m^{\rm eq}$, and in the second term we can use Eq. (2) to insert a factor of $\Sigma_m w_{l\to m}$ (and interchange the indices l and m). This gives

$$\Delta G = \sum_{l,m,n} \left[w_{l\to m} w_{l\to n} P_l^{\rm eq} \frac{P_m P_n}{P_m^{\rm eq} P_n^{\rm eq}} \right] - \sum_{l,m} w_{m\to l} \frac{P_m^2}{P_m^{\rm eq}}. \tag{8}$$

Applying the detailed balance relation again and incorporating a factor of $\Sigma_n w_{l\to n}$, the last term in the above equation can be written as

$$- \sum_{l,m,n} w_{l\to m} w_{l\to n} P_l^{\rm eq} \left(\frac{P_m}{P_m^{\rm eq}} \right)^2. \tag{9}$$

Taking the half the sum of this and the same expression with m replaced by n, we finally get

$$\Delta G = -\frac{1}{2} \sum_{l,m,n} w_{l\to m} w_{l\to n} P_l^{\rm eq} \left(\frac{P_m}{P_m^{\rm eq}} - \frac{P_n}{P_n^{\rm eq}} \right)^2, \tag{10}$$

where terms with $m=l$ and $n=l$ are included.

Equation (10) is the main result of this part of the paper. It shows that ΔG is definitely negative unless, for every state l, all states which can be reached from l in a single move—equivalently, with detailed balance, all states from which l can be reached in a single move—have probabilities proportional to the equilibrium probabilities. The most natural scenario is that *all* states satisfy this with the same proportionality constant (which must be unity), i.e., the system is in equilibrium. However, ΔG also vanishes if $P_m/P_m^{\rm eq}$ assumes different values for states which have no common one-step descendants. Hence, to achieve full equilibrium, the algorithm must also be ergodic, i.e., starting from a given state, after a sufficiently long time there is non-zero probability for the system to be in any state. The condition of ergodicity is sufficient to ensure that even if ΔG is accidentally zero at some time step, it must decrease later, since any two states must have common descendants after several time steps. If in addition $w_{l\to l} \neq 0$ for all l, which is usually true, the one-step descendants of a set of states must include the set itself, so that it is not possible to break up all the states of the system into subsets with no common one-step descendants across two subsets. Thus if this condition is satisfied, ΔG cannot be zero (without the system being in equilibrium) at *any* time step.

We will distinguish between a process which is ergodic and one which satisfies the lesser condition of being "irreducible." In the latter, the system will eventually sample all states starting from a given initial state [6], but, *at a fixed later time*, the probability for some of the states is zero. A familiar example which is irreducible but not ergodic is the Ising model at infinite temperature simulated using Metropolis updating, for which the probability to flip is unity in this limit. Clearly after an odd time, the number of flipped spins must be odd and vice-versa. For such a non-ergodic system, it is possible for $P_m/P_m^{\rm eq}$ and $P_n/P_n^{\rm eq}$ to be different for states which have no common descendant at any fixed later time (for the Ising model example given here, states which differ by an odd number of spin flips).

Note that Eq. (10) does not give an estimate for *how fast* equilibrium is reached.

For random updating considered so far, the probability of making a transition is the same for every move, i.e., writing Eq. (1) as

$$P_l(t+1) = \sum_m \Gamma_{lm} P_m(t), \tag{11}$$

then Γ, the transition matrix (related to w by $\Gamma_{lm} = w_{m\to l}$), is the same for each "time" t. However, for sequential updating, the transition matrix depends on which site is being updated, so, for a complete sweep, we have

$$\Gamma = \Gamma^{(1)} \Gamma^{(2)} \cdots \Gamma^{(N)}, \tag{12}$$

where $\Gamma^{(i)}$ is the transition matrix for updating spin i. Although the $\Gamma^{(i)}$ individually satisfy the detailed balance condition, the transition matrix for the whole sweep, Γ, does *not*, [6] because the probability of the reverse transition, $m\to l$ say, for a whole sweep, is related to the probability of transition $l\to m$ in the desired way only if the spins are updated in the reverse order. Despite overall lack of detailed balance, convergence to the equilibrium distribution is still obtained

for sequential updating because G decreases at *each* step, as long as each of the transition probabilities, $\Gamma^{(i)}$, satisfies the detailed balance condition.

More generally, detailed balance is not necessary, see Ref. [6] and references therein. In the rest of this paper, we give a derivation of the necessary and sufficient conditions for convergence to $P^{\rm eq}$:

(i) the algorithm has $P^{\rm eq}$ as a stationary distribution,

(ii) for any pair of states (i,j), there exists some T_{ij} and some state k_{ij} such that at the T_{ij}'th time step the probabilities to have reached k from i and j are both non-zero.

We shall see that (ii) [with (i)] implies ergodicity.

We first prove that conditions (i) and (ii) are sufficient. Note that Eq. (1) is linear in the probabilities P_l, so that if we define $\delta P_l = P_l - P_l^{\rm eq}$, then [with condition (i)] δP_l also satisfies Eq. (1), with $\Sigma_l \delta P_l = 0$.

It is convenient to use a measure of the deviation from equilibrium that is different from Eq. (5):

$$L = \sum_l |\delta P_l|. \tag{13}$$

Like G, the quantity L is positive unless the distribution has converged to $P^{\rm eq}$. We denote by $W_{l\to m}(t)$ the probability that the system, starting out in state l, reaches state m after t time steps. This is obtained by ''iterating'' the transition rates $\{w\}$. Then

$$L(t) = \sum_m \left| \sum_l W_{l\to m}(t)\, \delta P_l(0) \right| \leqslant \sum_{m,l} W_{l\to m}(t) |\delta P_l(0)|. \tag{14}$$

Using the result $\Sigma_m W_{l\to m}(t) = 1$, we see that $L(t) \leqslant L(0)$. We now compare the two sides of the inequality in Eq. (14), by choosing a specific value of m and carrying out the l-summation. If all states l for which $W_{l\to m} \neq 0$ have the same sign for $\delta P_l(0)$, it is clear that the l-summation on both sides are equal. On the other hand, if some of these states have $\delta P(0) > 0$ and others have $\delta P(0) < 0$, the l-summation on the left-hand side has both positive and negative terms, and must be less than the corresponding sum on the right-hand side. Thus so long as there is at least one state m which receives ''contributions'' from two states (i,j) with opposite $\delta P(0)$, i.e., $\delta P_i(0)$ and $\delta P_j(0)$ have opposite signs and $W_{i\to m} \neq 0$ and $W_{j\to m} \neq 0$, we see that $L(t) < L(0)$. Condition (ii) ensures that for *any* (i,j) this is the case for $t = T_{ij}$. Thus $L(t)$ stays constant until $t = \min[T_{ij}]$, where the minimum is taken over all (ij) for which δP_i and δP_j have the opposite sign, and then decreases at that time step. The time t can be reinitialized to zero at this point, and the whole argument repeated again. Note that nothing in the argument requires the transition rates $w_{l\to m}$ to be time independent, so long as conditions (i) and (ii) are always satisfied. Also, as with the first approach above, no estimate has been obtained for *how fast* $L(t)$ approaches zero.

It may seem surprising that $L(t)$ need not decrease at every time step, whereas Eq. (10) shows that G must. If we start out with δP positive and negative on two states that are well separated from each other (in the sense that many time steps are required before the two states have common descendants), it is clear that $L(t)$ does not change at first. However, G decreases because the positive and negative δP's both get ''smeared out'' over several states. Thus different measures can depict the approach to equilibrium at different rates. Since any physical observable O is given by $\langle O \rangle = \Sigma_l O_l P_l$, the error $\langle \delta O \rangle \leqslant \Sigma_l |O_l \delta P_l| \leqslant L \max_l[|O_l|]$, so that L is a conservative indicator of how observables approach equilibrium.

We have proved the sufficiency of conditions (i) and (ii); the necessity can be easily demonstrated. The necessity of (i) is obvious. If (ii) is violated, starting with an initial condition $\delta P_i(0) = -\delta P_j(0)$ and all other δP_l's equal to zero, the positive and negative regions for δP stay separate for all time and cannot neutralize each other.

Condition (ii) [with (i)] is equivalent to the (seemingly stronger) condition of ergodicity. If the system were not ergodic, then starting with $P_l(0) = \delta_{il}$ would lead to a $P(t)$ with $P_j(t) = 0$ for some j, since all states would not be accessible from the state i at time t. Therefore $P(t)$ could not have converged to $P^{\rm eq}$. Since we have seen that (i) and (ii) imply convergence, they must imply ergodicity.

The proof of convergence in Ref. [1] is comparable in simplicity to the proof given above. However, it applies the *transpose* Γ^T of the transition matrix to the initial state of the system (represented as a column vector), and shows that under iteration the initial state evolves to a column vector whose entries are all identical. Although it is possible [1] to deduce the desired properties of Γ from this, there is no physical interpretation of evolution under Γ^T, which does not even conserve probability. In contrast, the proof given above shows how fluctuations from the desired distribution are ''mixed'' by time evolution, with positive fluctuations annihilating negative ones. It is also physically clear why ergodicity is essential for this annihilation process to proceed to completion: Positive and negative fluctuations are then assured of encountering each other under time evolution.

Although we have considered the conditions for convergence to $P^{\rm eq}$, in practice Monte Carlo simulations are not carried out using such an ensemble average, but by simulating a single system and taking a time average of many measurements. If we restrict ourselves to the case when the transition matrix Γ is time-independent, this defines a homogeneous Markov chain [12]. It can then be shown that if the Markov chain is irreducible and has $P^{\rm eq}$ as a stationary distribution, the time average $\langle P(t) \rangle_t$ converges to $P^{\rm eq}$. We indicate how to prove this result here: since $L(t) \leqslant L(0)$ in general, none of the eigenvalues of Γ can have modulus greater than unity [13]. Irreducibility ensures [14] that the eigenvalue 1 has a unique eigenvector, $P^{\rm eq}$. The time averaging removes contributions from oscillatory eigenvalues of Γ, of the form $e^{i\theta}$ $(\theta \neq 0)$. By comparison, the stronger condition of ergodicity (ii) [with (i)] is enough to rule out oscillatory eigenvalues, since no time averaging is needed for convergence to $P^{\rm eq}$. It is not clear how this proof that irreducibility is sufficient would generalize to time-dependent transition rates, since the evolution of P is then not directly related to the eigenvalues of $\Gamma(t)$.

We thank M. E. J. Newman for useful correspondence.

[1] J. G. Kemeny and J. L. Snell, *Finite Markov Chains* (Van Nostrand, Princeton, 1960).

[2] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).

[3] Metropolis *et al.* [2] show that if P_l/P_m is greater than the ratio of equilibrium probabilities then the net number of transitions $l \to m$ is negative so P_l/P_m decreases, and vice-versa. This shows that the algorithm *may* converge but does not rule out the possibility that it could overcompensate for the deviation from equilibrium leading to oscillatory or divergent behavior.

[4] See, e.g., M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, Oxford, 1999); K. Binder and D. S. Stauffer, in *Applications of the Monte Carlo Method in Statistical Physics*, 2nd. ed., edited by K. Binder, Topics in Current Physics Vol. 36 (Springer, Berlin, 1987).

[5] See, e.g., K. Binder, in *Monte Carlo Methods in Statistical Physics*, edited by K. Binder, Topics in Current Physics, Vol. 7 (Springer, Berlin, 1987), p. 1; J. M. Hammersley and D. C. Handscomb, *Monte Carlo Methods* (Methuen and Co., London, 1964).

[6] A. D. Sokal, in Quantum Fields on the Computer, edited by M. Creutz, 1992; Monte Carlo Methods in Statistical Mechanics: Foundations and New Applications, Cours de Troisiéme Cycle de la Physique en Suisse Romande (Lausanne, 1989).

[7] J. L. Doob, *Stochastic Processes* (Wiley, New York, 1953).

[8] B. A. Berg and T. Celik, Phys. Rev. Lett. **69**, 2292 (1992).

[9] B. Hesselbo and R. B. Stichcombe, Phys. Rev. Lett. **74**, 2151 (1995).

[10] K. Hukushima and K. Nemoto, J. Phys. Soc. Jpn. **65**, 1604 (1996).

[11] K. L. Chung, *Markov Chains with Stationary Transition Probabilities*, 2nd ed. (Springer, New York, 1967).

[12] This includes the possibility of cycling periodically through a set of transition matrices, since the composite matrix defined in Eq. (12) is time-independent, although it need not satisfy detailed balance even if its constituents do.

[13] If the vectors of Γ do not span the vector space, let $\{V_i\}$ be the eigenvectors with eigenvalues $\{v_i\}$. Then it is possible to construct a C_1 in the complement of the eigenspace, such that $\Gamma C_1 = \lambda C_1 + \Sigma_i d_i V_i$. By taking a suitable linear combination of C_1 with the V_i's, one can find a C_1' such that $\Gamma C_1' = \lambda C_1 + \Sigma_i d_i' V_i$, where $d_i' \neq 0$ *only* if $\lambda = v_i$. Therefore $\Gamma^N C_1' = \lambda^N C_1' + N\lambda^{N-1}\Sigma_i d_i' V_i$. The non-increasing property of $L(t)$ then implies $|\lambda| < 1$. The process of constructing C_i's can be iterated to show that the entire complement of the eigenspace of Γ is annihilated by $\lim_{N\to\infty}\Gamma^N$.

[14] Otherwise, a suitable linear combination of the several eigenvectors would yield an eigenvector P' whose elements would not all have the same sign; since $\Gamma^N P' = P'$ for all N, by the discussion after Eq. (14) this would require Γ^N not to connect the positive and negative elements of P' for any N, which would contradict the assumption that Γ is irreducible.