Nuclear Physics B235[FS11] (1984) 417–434

MICROCANONICAL SIMULATION OF ISING SYSTEMS

Gyan BHANOT

The Institute for Advanced Study, Princeton, New Jersey 08540, USA

Michael CREUTZ

Brookhaven National Laboratory, Upton, NY 11973, USA

Herbert NEUBERGER

Physics Department, Rutgers University, New Brunswick, NJ 08903, USA

Received 1 December 1983

Numerical simulations of the microcanonical ensemble for Ising systems are described We explain how to write very fast algorithms for such simulations, relate correlations measured in the microcanonical ensemble to those in the canonical ensemble and discuss criteria for convergence and ergodicity.

1. Introduction

Over the last two decades considerable effort has been invested in the development and application of numerical simulation methods [1] to statistical systems. Traditionally, in applications dealing with spin or gauge systems, only stochastic methods were used. It has been suggested that it is also of interest to investigate deterministic evolutions [2, 3]. The purpose of this work is to present a study of pseudo-deterministic simulations of Ising systems. Our main aim is to find a method which is very fast even on general purpose computers. Our main concern is that, unlike the stochastic case [4], we know of no rigorous proof that establishes the desired long-time behavior of the process. We approach this question pragmatically, concentrating on an exactly soluble system, the 2-d Ising model. For this system we are able to check that, within experimental errors, the desired limiting behavior is indeed obtained.

The plan of the paper is as follows: in sect. 2 we describe the simulation algorithm and its implementation on the computer. In sect. 3 theoretical expectations are presented. The numerical results obtained in the experiment are given in sect. 4.

They are shown to agree with theory. The advantages of microcanonical versus canonical simulations in general are discussed in sect. 5.

An appendix has been added in which we compare the evolution of stochastic processes which create canonical ensembles to that of deterministic processes which are supposed to give rise to microcanonical ensembles. This clarifies why it is very hard to rigorously prove that the latter can be used for evaluating statistical averages. The rest of our paper shows that, in practice, microcanonical simulations can be safely used.

2. The algorithm and its implementation

The algorithm is an adaptation of the one presented in ref. [3]. We consider a 2-dimensional Ising system with energy

$$E_s = \sum_{\langle ij \rangle} (1 - s_i s_j), \tag{1}$$

where the sum runs over nearest neighbor spins s_i on the sites of an $L \times L$ lattice. Following ref. [3] imagine a demon with a sack of energy E_d. By decree, E_d is non-negative and bounded from above by E_d^{max}. Given an initial lattice distribution of spins with energy E_s and a demon with energy E_d, one generates a sequence of configurations with constant $E_d + E_s = E_T$. The algorithm that accomplishes this is the following:

The demon hops from site to site trying to flip spins. At each site, ΔE_s, the change in E_s on flipping the spin, is computed. If ΔE_s is positive, the demon must supply energy to flip the spin. He does so if his sack contains enough energy. Energy supplied to the lattice depletes the demon's own energy. If, on the other hand, ΔE_s is negative, the spin flip lowers the lattice energy. The energy is accepted by the demon (and the spin is flipped) if the demon has space in his bag to hold the extra energy. If ΔE_s is zero, the spin flip is always accepted. This procedure, repeated several times over the lattice spins, generates an ensemble of states at fixed E_T from which correlation functions involving the spins can be computed.

In our experiment we worked on a 64×64 square lattice. Since an Ising spin takes only two values, we maximally utilize storage by using one computer bit per spin [5]. Spin "up" corresponds to the bit being zero and spin "down" to it being one. Our codes were written on a VAX 11/780 which has a 32-bit word length. We store the 64×64 lattice in 128 words, indexed by an integer I, $I = 1, \ldots, 128$. A given row of the lattice is coded into two successive words in the list of 128 words. For example, the words indexed by $I = 1$ contains in the rightmost bit position the spin corresponding to the lower right corner of the lattice, in the next bit position, the third spin in the bottom row counting from right to left and so on. The important thing is

that each word contains spins that do not interact with each other via the action (eq. (1)) and so, may be updated simultaneously.

We use 32 demons, each of which is allowed four energy states labeled 00, 01, 10 and 11 in binary notation. These are stored in two 32-bit words D1 and D2. The ith bit in D1(D2) contains the coefficient of $2^0(2^1)$ in the energy of the ith demon. It is easy to see that the minimum non-zero change in a demon's energy corresponds to four units of lattice energy. The total energy E_T of the system of spins plus the demons is then given by

$$E_T = E_s + 4E_d. \tag{2}$$

Given the demon energies D1 and D2 and the products S · S1, S · S2, S · S3, S · S4 of the spin word S (which is being updated) with its four neighboring spin words S1, S2, S3 and S4, one must construct a boolean function B, which produces the final values SN, D1N and D2N for the spin and demon energy words using the algorithm described above. On a bit level, B has to reproduce table 1. In this table, the entries under "Initial spin state" are the four corresponding bits of the products S · Si, $i = 1,2,3,4$ (permutations of these bits yield the same result for SN, D1N, D2N). Under "Initial demon energy" are given the two bits of the energy of the demon corresponding to the spin being updated. In the table, Y means the spin flip is accepted and N that it is not. The numbers beside Y or N in each box denote the two bits of the final demon energies.

In table 2 we show a piece of Fortran code that produces table 1. All variables are of the integer type and XOR, OR, AND and NOT are standard boolean operations. The way the code works is as follows: first notice that in the cases where the spin flip is accepted in table 1, the final demon energy, DN, is given by DN = 2 + D − · NOT · (S · S1 + S · S2 + S · S3 + S · S4) where D is the initial demon energy. In the code, lines 7–10 add 2 to the demon energies and in lines 11–16, · NOT · (S · Si) is subtracted successively for each i. The cases where the spin-flip is not accepted in table 1 are kept track of in the bits of a separate word: "REJECT". The bits of

TABLE 1

The boolean function B that gives new demon energies and spin flips from the products S Si, ($i = 1,2,3,4$), and the initial demon energies

Initial spin state	Initial demon energy							
	00		01		10		11	
1111	Y	10	Y	11	N	10	N	11
1110	Y	01	Y	10	Y	11	N	11
1100	Y	00	Y	01	Y	10	Y	11
1000	N	00	Y	00	Y	01	Y	10
0000	N	00	N	01	Y	00	Y	01

Y and N stand for site spin flip and no-spin flip respectively

TABLE 2
Fortran code to produce table 1

```
c       Input: S,S1,S2,S3,S4,D1,D2
        I(1) = not and (S,S1)
        I(2) = not and (S,S2)
        I(3) = not and (S,S3)
        I(4) = not and (S,S4)
        IR = 0
        REJECT = 0
        D10 = D1
        D20 = D2
        D3 = D2
        D2 = not · D2
        DO 1 K = 1,4
        IR = or(IR, and (I(K), ·not or (D1,or (D2,D3))))
        D1 = xor(D1,I(K))
        D2 = xor(D2, and (I(K),D1))
        D3 = and (D3, not · and (D1,D2))
1       continue
        REJECT = or (D3, or (IR,REJECT))
        ACCEPT = not REJECT
        SN = xor(S,ACCEPT)
        D1N = or (and (D10,REJECT), and (D1,ACCEPT))
        D2N = or (and (D20,REJECT), and (D2,ACCEPT)
c       Output SN,D1N,D2N
```

REJECT are zero unless: (a) DN becomes negative (insufficient energy in the demon to flip the spin), or (b) DN $\geqslant 4$ (demon gains more energy from the flip than he can hold). The final spin and demon are constructed using REJECT to decide what changes are to be made to the bits of the initial spin and demons (lines 19–21).

To improve thermal contact between the demons we scrambled their energies after each pass through the lattice. This was done by cyclically permuting the bits of the demon energies randomly. Because of this, our simulations cannot be called strictly deterministic.

We should remark that it is not necessary in principle to have an upper bound on the demon energy. However, practically speaking, the demon energy should be restricted from above as tightly as possible because this results in the fastest algorithms. One should ensure of course that the bound is not so severe that it prevents the demons from flipping spins.

Using this logic in our program we achieved a speed of 6×10^5 spin-updates/sec on a VAX 11/780. Using a more sophisticated code and special tricks, Shapiro [8] has succeeded in gaining an extra factor of four in speed. This means a rate of about 2.5×10^6 spin-updates/sec. Therefore, on a VAX 11/780, one obtains a speed only a factor of 10 slower than special purpose machines built for the Ising model [6].

3. Theoretical predictions

Ideally, our numerical procedure should generate configurations of the spin system (denoted by C) and of the demon system (denoted by c) with equal probability and with constant E_T. The average over passes of any observable $\Theta(C, c)$ should approach the following ensemble average:

$$\overline{\overline{\Theta}}(E_T) = \frac{\Sigma_{C,c}\delta_{E_s(C)+4E_d(c),E_T}\Theta(C,c)}{\Sigma_{C,c}\delta_{E_s(C)+4E_d(c),E_T}}. \tag{3}$$

Let $g(E_d)$ and $G(E_s)$ be the number of states of the demons and the spin system at energy E_d and E_s respectively. Then

$$\sum_{C,c}\delta_{E_s(C)+4E_d(c),E_T} = \sum_{E_s,E_d} G(E_s)g(E_d)\,\delta_{E_s+4E_d,E_T}. \tag{4}$$

In the absence of demons, microcanonical averages for the 2-d Ising model are defined by

$$\overline{\Theta}(E) = \frac{\Sigma_C \delta_{E_s(C),E}\Theta(C)}{\Sigma_E G(E)}. \tag{5}$$

The known exact solution for the 2-d Ising model gives analytic expressions for canonical averages [7]:

$$\langle\Theta\rangle_\beta = \frac{\Sigma_E G(E)\,e^{-\beta E}\overline{\Theta}(E)}{\Sigma_E G(E)\,e^{-\beta E}}. \tag{6}$$

We therefore have to relate the averages of eq. (3) (measured) to those of eq. (6) (predicted). The average of eq. (5) is an intermediate step: For an observable which is independent of the demons we have

$$\overline{\overline{\Theta}}(E_T) = \frac{\Sigma_{E_d} g(E_d)\overline{\Theta}(E_T-4E_d)G(E_T-4E_d)}{\Sigma_{E_d} g(E_d)G(E_T-4E_d)}. \tag{7}$$

When the volume of the Ising system tends to infinity we expect that

$$G(E) \sim e^{VS(E/V)}\left(1+O\left(\frac{1}{V}\right)\right), \tag{8}$$

where $e = E/V$ is kept away from its lower and upper bounds. It is important to realize that even in the disordered phase (i.e. in the presence of a mass gap) the

power corrections in eq. (8) are nonvanishing. For most physical systems, given a $\beta > 0$, the sums over E in eq. (6) are dominated by a single energy given by

$$\beta = \frac{\mathrm{d}S}{\mathrm{d}e}. \tag{9}$$

β is conjugate to e. Eq. (9) defines implicitly a function $e(\beta)$. As $V \to \infty$ we have

$$\langle \Theta \rangle_\beta \underset{V\to\infty}{\to} \bar{\Theta}(Ve(\beta)). \tag{10}$$

In particular

$$\frac{1}{V}\langle E_s \rangle_\beta \underset{V\to\infty}{\to} e(\beta). \tag{11}$$

The function $e(\beta)$ is therefore known from the exact solution. It can be extracted from the experiment as follows: We use in eq. (3) a Θ which describes the demon energy distribution:

$$\Theta = P_\varepsilon = \delta_{E_d,\varepsilon} = \delta_{E_s, E_T - 4\varepsilon},$$

$$\bar{\bar{P}}_\varepsilon(E_T) = \frac{g(\varepsilon)\Sigma_{E_s} G(E_T - 4\varepsilon)\,\delta_{E_s+4\varepsilon,\,E_T}}{\Sigma_{E_s E_d} G(E_T - 4E_d) g(E_d)\,\delta_{E_s+4E_d,\,E_T}}. \tag{12}$$

Hence,

$$\bar{\bar{P}}_\varepsilon(E_T) \underset{V\to\infty}{\sim} N_\beta g(\varepsilon)\, \mathrm{e}^{-4\beta_T\varepsilon - \gamma\varepsilon^2}, \tag{13}$$

where N_β fixes the normalization and $e(\beta_T) = E_T/V$. γ is given by $8/VC$, where $C = \mathrm{d}e/\mathrm{d}\beta$ (see eq. (11)) and γ vanishes as $V \to \infty$. (At the phase transition point even the product γV vanishes when $V \to \infty$.) To be able to extract $e(\beta)$ from the measured $\bar{\bar{P}}_\varepsilon(E_T)$ one needs $g(\varepsilon)$. It is easy to see that in our case

$$g(\varepsilon) = \sum_{k=0}^{\min[32,\,\varepsilon/2]} \binom{32}{\varepsilon - 2k}\binom{32}{k}. \tag{14}$$

One can also look at various moments of $\bar{\bar{P}}_\varepsilon(E_T)$. For example

$$\bar{\bar{E}}_d = \sum_\varepsilon \varepsilon \bar{\bar{P}}_\varepsilon(E_T) \underset{V\to\infty}{\to} 32\left[\frac{1}{1+\mathrm{e}^{4\beta_T}} + \frac{2}{1+\mathrm{e}^{8\beta_T}}\right]. \tag{15}$$

Ignoring the small γ term in eq. (13) one can show that, for an arbitrary $g(\varepsilon)$ the best estimate for β_T is obtained by measuring $\bar{\bar{E}}_d$ and using eq. (15).

At infinite volume the averages in eqs. (3), (5) and (6) are equal for $\beta = \beta_T$. However, the finite volume corrections are power-like in (3) and (5). It turns out that

our accuracy is such that the $O(1/V)$ corrections are needed. We therefore need an expression for Θ in powers of $1/V$ where the coefficients are given in terms of $\langle\Theta_i\rangle$, the latter being exactly computable. Standard manipulations give, up to order $1/V^2$

$$\langle\Theta\rangle_{\beta_T} = \overline{\Theta}\big(\langle E_s\rangle_{\beta_T}\big) + \frac{1}{2}\frac{\partial^2\overline{\Theta}}{\partial E^2}\bigg|_{E=E_T}\langle\big(E_s - \langle E_s\rangle_{\beta_T}\big)^2\rangle\beta_T\,,$$
$$\overline{\overline{\Theta}}(E_T) = \overline{\Theta}\big(\overline{E}_s\big)\,. \tag{16}$$

It is therefore useful to define β^* by

$$\langle E_s\rangle_{\beta^*(E_T)} = \overline{\overline{E}}_s(E_T)\,. \tag{17}$$

β^* and β_T differ to order $1/V$. From eq. (16) we now obtain

$$\langle\Theta\rangle_{\beta^*(E_T)} = \overline{\overline{\Theta}}(E_T) - \frac{1}{2V}\frac{d}{d\beta}\left[\frac{d\langle\Theta\rangle_\beta/d\beta}{de/d\beta}\right]_{\beta=\beta_T} + O\left(\frac{1}{V^2}\right). \tag{18}$$

Eq. (18) can be used only after eq. (17) has been applied in order to obtain $\beta^*(E_T)$. It is the dependence of various observables on E_s that will be checked against theory to order $1/V$. The natural control parameter in our ensemble is the nearest neighbor spin-spin correlation.

4. Numerical results

We did one very long run at a fixed value of E_T ($E_T = 4096$). 10000 initial lattice sweeps were discarded and 40000 more were kept to compute averages. We obtained,

$$\overline{\overline{e}} = 0.50421 \pm 0.00002\,,$$
$$\overline{\overline{e}}_{diag} = 0.37629 \pm 0.00008\,,$$
$$\overline{\overline{e}}_{nnn} = 0.30038 \pm 0.00011\,. \tag{19}$$

Here, $\overline{\overline{e}}_{diag}$ and $\overline{\overline{e}}_{nnn}$ are expectation values of the product of two spins along the shortest diagonal and on-axis-next-nearest-neighbors respectively. The average demon energy gave via eq. (15),

$$\beta = 0.381 \pm 0.001\,. \tag{20}$$

This value lies in the disordered phase. All the error bars quoted here are truly

statistical. They were measured by bunching our data until the correct dependence on the bunch size $1/\sqrt{N_{\text{bunch}}}$ was seen.

These correlation functions can be computed exactly for the Ising model. First, we find (see eq. (17)),

$$\beta^* = 0.380383 \pm 0.000009, \tag{20a}$$

which agrees with eq. (20).

The exact results [7] for the other correlations at this temperature are

$$\langle e_{\text{diag}} \rangle_{\beta^*} = 0.376460 \pm 0.000025,$$

$$\langle e_{\text{nnn}} \rangle_{\beta^*} = 0.300722 \pm 0.000025. \tag{21}$$

Eqs. (19) and (21) do not agree. It is here that we use the effects of finite volume corrections computed before. Using eq. (18), we calculate

$$\langle e_{\text{diag}} \rangle_{\beta^*} = \bar{\bar{e}}_{\text{d}} + 0.00013 + \text{O}\left(\frac{1}{V^2}\right),$$

$$\langle e_{\text{nnn}} \rangle_{\beta^*} = \bar{\bar{e}}_{\text{nnn}} + 0.00032 + \text{O}\left(\frac{1}{V^2}\right), \tag{22}$$

and the discrepancy is accounted for.

We also compared the observed demon energy distribution with the prediction of eqs. (13–14). This is shown in fig. 1 which is a plot of E_{d} versus $\log[g(E_{\text{d}})/P(E_{\text{d}})]$. The slope of this curve yields 4β. Moreover, the coefficient γ in eq. (13) may also be estimated from this figure. We find $\gamma \lesssim 3 \times 10^{-4}$ which justifies neglecting it.

A different simulation was done next at various values of E_{T} discarding 200 sweeps and averaging over 1000. Figs. 2 and 3 show $\bar{\bar{e}}$ and $\bar{\bar{e}}_{\text{diag}}$ as functions of β (which was computed from the average demon energy). The open circles correspond to an initial start of a disordered lattice and the dark points to an initially ordered lattice. The solid curve is the exact canonical result.

While individually, figs. 2 and 3 show rather substantial deviations from the exact result, the errors in both curves are highly correlated between the two graphs. This suggests that the source of the discrepancy is in the estimate of β. Indeed, if we plot $\bar{\bar{e}}_{\text{diag}}$ versus β^* (see fig. 4), the agreement with the exact result improves dramatically. The main observation is that most of the discrepancy between the exact answer and the data in figs. 2 and 3 comes from the $1/V$ corrections in β. In fig. 4 there are present of course the other $1/V$ effects (see eq. (18)) that we have discussed. However, these are too small to be seen on this plot.

In fig. 5 we plot the quantity $\langle e_{\text{diag}} \rangle_{\beta^*}^{\text{exact}} - \bar{\bar{e}}_{\text{diag}} - (1/V \text{ correction})$. This represents the intrinsic statistical error of our microcanonical ensemble. In fig. 6 we plot the

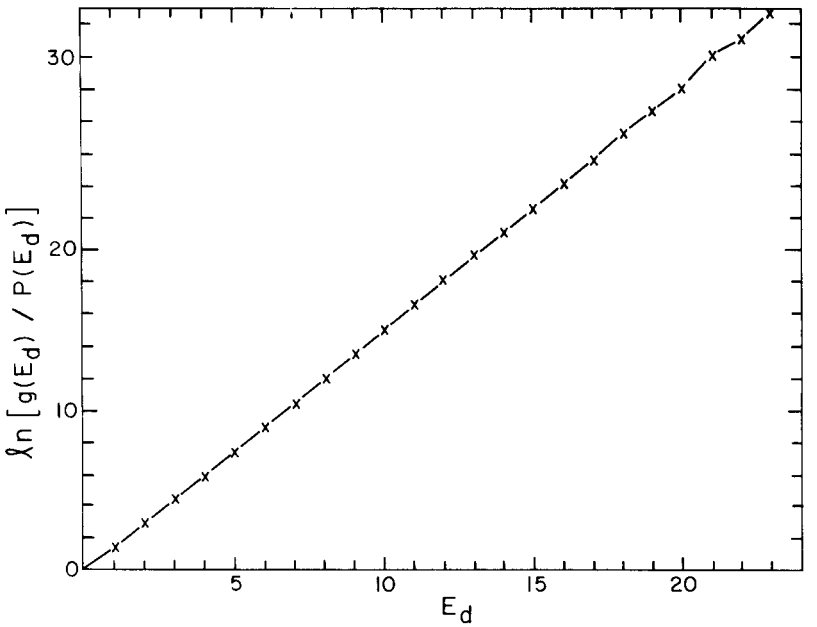

Fig 1 Plot of $\log[g(E_d)/P(E_d)]$ versus E_d for a run at $E_T = 4096$

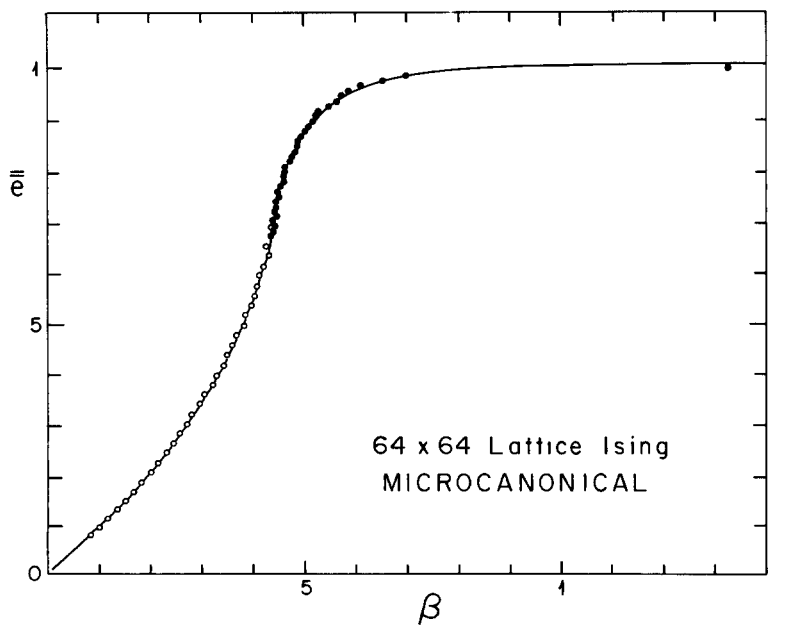

Fig 2 Nearest-neighbour spin correlation function $\bar{e}$ versus β determined from the demon energy The solid line is the exact result in the canonical ensemble

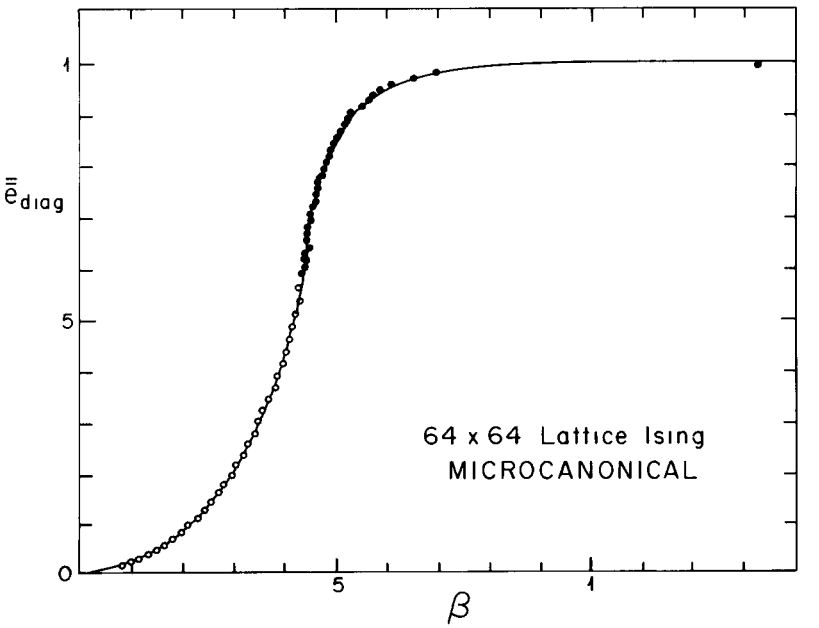

Fig 3 $\bar{\bar{e}}_{diag}$ versus β

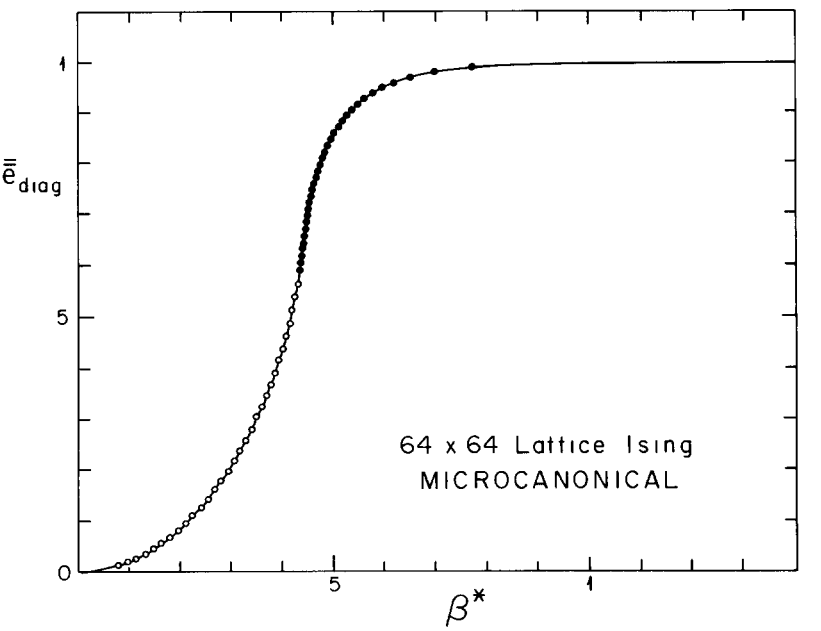

Fig 4 The data of fig 3 plotted against β^* (see eq (17))

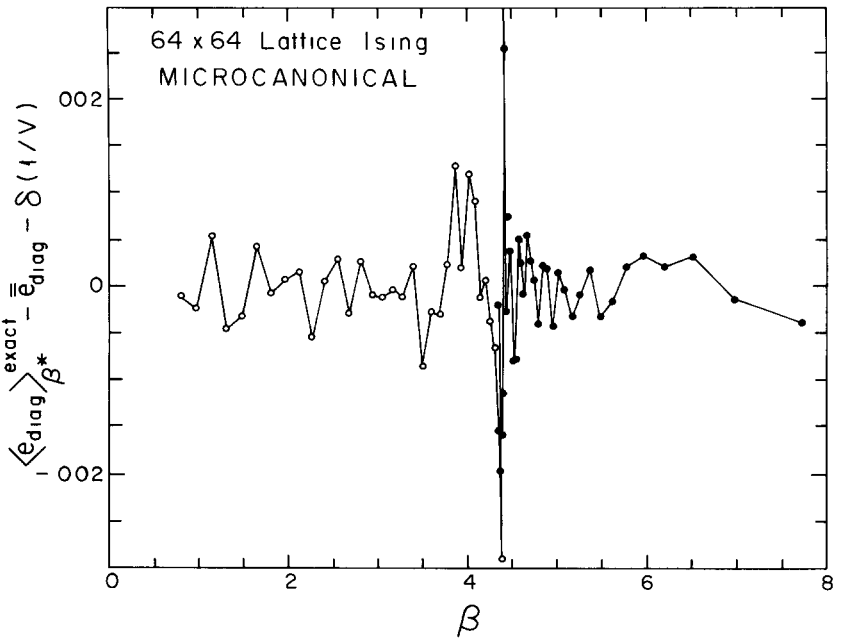

Fig 5 Statistical errors in the microcanonical ensemble for $\bar{\bar{e}}_{diag}$

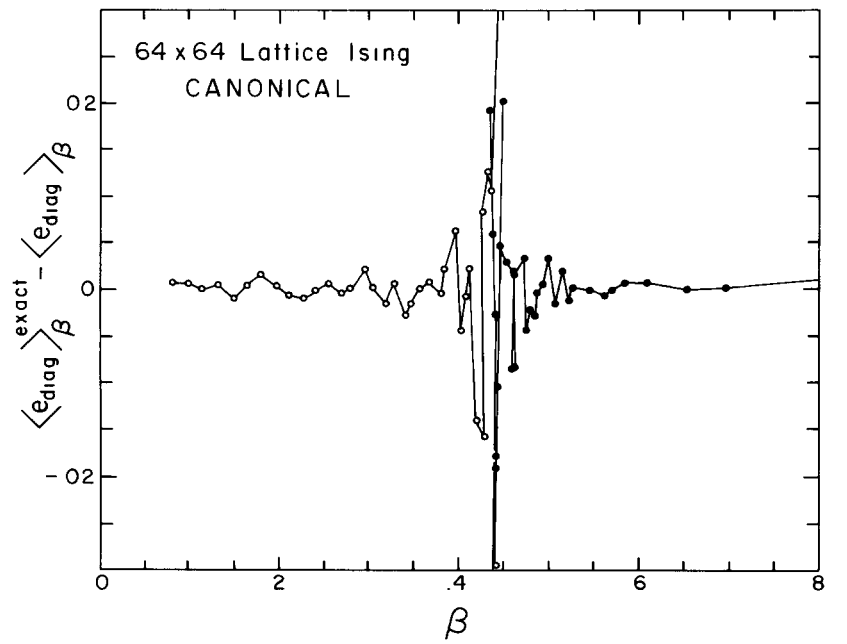

Fig. 6 Statistical errors in the canonical ensemble for $\bar{\bar{e}}_{diag}$ using the same amount of CPU time as in fig 5

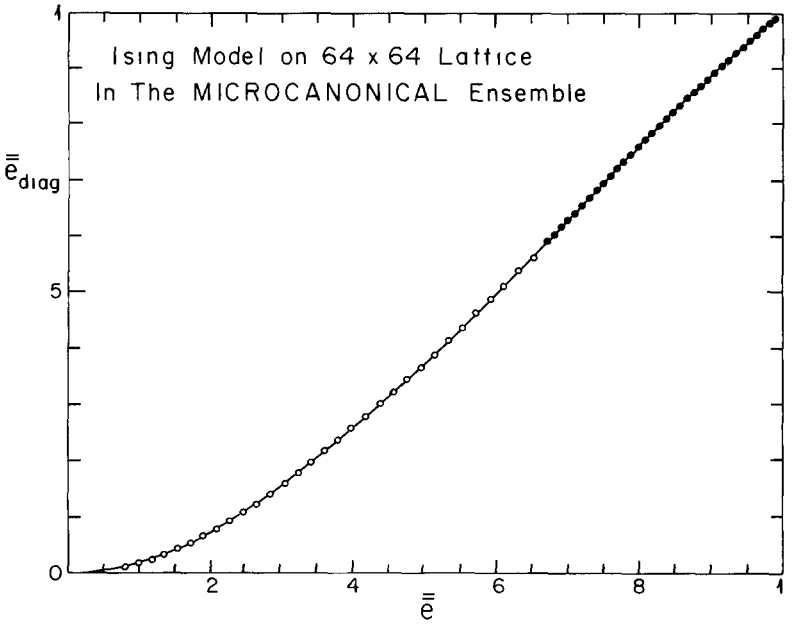

Fig 7 $\bar{\bar{e}}_{diag}$ versus $\bar{\bar{e}}$ The solid line is the exact (canonical) result This graph should be compared with figs 2, 3 and 4

corresponding quantity (viz. $\langle e_{diag}\rangle_\beta^{exact} - \langle e_{diag}\rangle_\beta$) in the canonical ensemble. This data was taken by running a metropolis algorithm at the same values of β as in fig. 2 with hot and cold lattice starts and using the same amount of computer time. A comparison of the scales on the ordinates of figs. 5 and 6 shows that the statistical error in the microcanonical runs, for the same amount of computer time is much smaller than for the canonical. Therefore, the difference in speed of the two programs does reflect itself in the accuracy of the results.

Finally, in fig. 7, we plot $\bar{\bar{e}}_{diag}$ versus $\bar{\bar{e}}$ along with the exact curve in the canonical ensemble. The agreement is very good, as expected.

5. Discussion

Even so called stochastic processes become deterministic when implemented on a computer, because only pseudo-random number generators are used. Therefore, strictly speaking, the distinction between our simulation (even in the absence of the random elements of the procedure) and standard Monte Carlo experiments is not as fundamental as it might seem. Nevertheless, the two methods are conceptually different. The most significant aspect of this difference is the type of control parameters one can use. In standard simulations, one can control couplings (gener-

alized temperatures). In the microcanonical simulation, one controls directly some of the observables (for practical reasons they have to be reasonably local). In our case we hold the internal energy fixed.

This feature of microcanonical simulation might be advantageous because physical quantities are expected to depend in a smoother way on the internal energy than on the temperature. For example, in the 4-dimensional SU(2) lattice gauge theory, the violent dependence of the plaquette average on the coupling β in the crossover region seems to be a reflection of a pair of singularities in the complex β plane. When the plaquette variable is used as the independent variable to parametrize other observables, the effect of these singularities seems to be neutralized [9].

One may also speculate that it would be easier to use microcanonical simulations for numerical studies of truncated renormalization group recursion relations. Instead of using a set of couplings to obtain matching between "elementary" correlations and their "block spin" counterparts, one would directly control the value of at least a subset of the elementary correlations.

Microcanonical simulations are not expected to be slower than stochastic ones. For simple systems like the Ising models, where the computation of random numbers takes up a significant fraction of computer time, microcanonical simulations can be made faster by the construction of relatively simple boolean functions. Thus, ordinary sequential machines can be used in an effectively multiprocessing mode. For this, as we have seen, one uses multiple spin packing in a computer word, with simultaneous updating of spins in the word.

One disadvantage of microcanonical simulation is the absence of a mathematical proof for ergodic or mixing behavior. A more detailed discussion of this is given in the appendix. However, a small amount of randomness can always be introduced at no significant cost in speed. In our case, we used a random shift of the demons for this purpose. This probably neutralizes any undesirable consequences of possible violations of ergodicity. In our simulations, we have checked for "bad" initial states where the system gets trapped into subsets of the constant energy surface. We found that this only happens for very special configurations. Our high accuracy run indicates that one is probably quite safe from such states as long as some indeterminancy is present in the system. Therefore, for Ising systems, our method is a faster way to obtain numerical results.

The algorithm we use is easily generalizable to higher dimensions. On a 3-dimensional lattice of size $120 \times 128 \times 128$, we have programs which run at the rate of 24.0 million spin-updates/sec on a CDC-7600.

Discussions with H. Hamber, O. Martin, J. Shapiro and S. Wolfram are gratefully acknowledged. The work of G.B. and H.N. was supported by the DOE under grant no. DE-AC02-76ER02220 and that of M.C. was supported by the DOE under grant no. DE-AC02-76DH00016.

Appendix

In this appendix, we compare the convergence criteria for stochastic and deterministic processes. We will first review the basic features of stochastic methods usually employed in numerical simulation and then compare them with deterministic ones. Since we deal with Ising systems in a finite volume, we have a finite space of states. This allows the use of discrete notation. Most of what we have to say generalizes to infinite spaces.

In a standard Monte Carlo study, one simulates a Markov process on the space of states $\{C\}$. The process is defined by the transition probabilities $p(C_1 \to C_2)$. These $p(C_1 \to C_2\}$ build a stochastic matrix P.

$$P_{C_1C_2} = p(C_2 \to C_1). \tag{A.1}$$

It is then required that, for a given function on $\{C\}$ (e.g. the action $E(C)$) and a non-negative initial vector Q (with components $Q(C)$), we should have

$$\lim_{n\to\infty} P^n Q = N e^{E(C)}. \tag{A.2}$$

For a given Q, the limit in eq. (A.2) exists iff the matrix P has no eigenvalues $\exp(i\varphi)$ with $\varphi \neq 2\pi n$ [4]. The limit will be Q independent iff the unit element is a non-degenerate eigenvalue of P. The fact that $\Sigma_C p(C_1 \to C) = 1$ ensures that unity is an eigenvalue. Frobenius' theorem [10] ensures that all other eigenvalues μ of P satisfy $|\mu| \leqslant 1$.

Given $E(C)$, the problem is to find $p(C_1 \to C_2)$ such that P satisfies the requirements of the existence and uniqueness of the limit as given in eq. (A.2).

A possible solution is obtained by requiring detailed balance.

$$e^{E(C_1)} p(C_1 \to C_2) = e^{E(C_2)} p(C_2 \to C_1). \tag{A.3}$$

For the particular choice $Q(C) = e^{E(C)}$, one has

$$\begin{aligned}(PQ)(C) &= \sum_{C_1} p(C_1 \to C)\, e^{E(C_1)} \\ &= \sum_{C_1} p(C \to C_1)\, e^{E(C)} = Q(C).\end{aligned} \tag{A.4}$$

The only additional requirements are: (i) that there be no complex eigenvalue with modulus unity except the unit element and (ii) that the eigenvalue unity be nondegenerate.

Define,

$$X(C_1, C_2) = p(C_1 \to C_2)\, e^{(E(C_2) - E(C_1))/2}. \tag{A.5}$$

Clearly, $X(C_1, C_2) = X(C_2, C_1)$. Therefore, P is similar to a symmetric matrix and has only real eigenvalues. The eigenvalue -1 is still possible and, by Frobenius' theorem [10], will occur iff a relabeling of states brings P into the form,

$$P = \begin{pmatrix} 0 & P_1 \\ P_2 & 0 \end{pmatrix}, \tag{A.6}$$

where, dim P_1 = dim P_2 because of the symmetry of X. To avoid this, we must pick $p(C_1 \to C)$ in eq. (A.3) such that the space of states cannot be broken into two disjoint parts with the only possible transitions being those which take a state from one part to the other.

The other requirement $p(C_1 \to C_2)$ has to satisfy is that the eigenvalue unity be simple. This is ensured by making P non-decomposable; i.e. no relabeling of states can bring P to the form,

$$P = \begin{pmatrix} P_1 & 0 \\ P_2 & P_3 \end{pmatrix}. \tag{A.7}$$

This means that there does not exist a splitting of the space of states into two disjoint, non-empty sets with transitions from one particular set into the other disallowed.

An obvious way to meet the above requirements is to have no forbidden transitions. This is the situation in the "heat bath" simulations [1]. In the metropolis simulation, where an updating table is used, it must be ensured that an arbitrary state C_f can be reached in a finite number of steps from any initial state C_i.

In summary, once detailed balance is satisfied, one has only to avoid two undesirable possibilities: (1) that there exists an initial state from which the evolution does not cover the entire space of states and (2) that there are undamped oscillations between two subsets of states*.

Instead of a transition probability $p(C_1 \to C_2)$, the deterministic methods use a fixed mapping φ. One looks then for density functions defined over the space of states that are invariant under this mapping. Such a density function defines an ensemble. This is the analogue of the limiting distribution $e^{E(C)}$ in the previous discussion. The most familiar case is when the space of states is the phase space of a closed physical system and the mapping φ is the one-parameter group of diffeomorphisms representing hamiltonian flow.

The two features one would like φ to have are that it be measure preserving (Liouville's theorem) and that it conserves the energy. The invariant density then is the microcanonical ensemble. When the number of degrees of freedom are large,

* For a generalization to gauge theories, where the space is continuous but compact, see the appendix of ref [11]

microcanonical averages can be expanded in powers of the inverse volume with expansion coefficients given in terms of the canonical averages. The main concern of ergodic theory is to what extent the microcanonical ensemble represents average properties of the long time behavior of arbitrary motions of the system [12].

In ergodic theory, one deals with a space of states M over which some density μ and a volume preserving mapping φ is defined.

Let f be a function on M (an observable). The time average of f is

$$f^*(x) = \lim_{n\to\infty} \frac{1}{n} \sum_{m=0}^{n-1} f(\varphi^m(x)), \tag{A.8}$$

while its ensemble average is,

$$\langle f \rangle = \int d\mu(x) f(x). \tag{A.9}$$

The simplest question is: does $f^*(x)$ exist and to what extent is it related to some ensemble average? One can show [12] that if φ is measure preserving, $f^*(x)$ exists, it is invariant, $(f^*(\varphi(x)) = f^*(x))$ and

$$\langle f^* \rangle = \langle f \rangle. \tag{A.10}$$

More precisely, $f^*(x)$ exists for any x except perhaps a subset of M of measure zero.

The first problem of ergodic theory is the dependence of f^* on the initial condition x. A system is ergodic if $f^*(x)$ is independent of x. A system is ergodic iff every invariant subset of M is either of zero volume or else its volume equals that of M. Therefore, in the presence of conservation laws, one has to restrict M to the subspace where the conserved quantities are given prescribed values. In hamiltonian systems, this always means at least a restriction to surfaces of constant energy.

Thus, ergodicity allows one to use time averages to evaluate ensemble averages. This is not really sufficient to make the simulation practical. The requirement of the stochastic process (eq. (A.2)) was much stronger. It required a long time limit which is independent of initial conditions. We would be satisfied with a condition which is a little weaker; viz. an analogue of eq. (A.2) but with expectation values:

$$\lim_{n\to\infty} \int d\mu(x) f(\varphi^n(x)) g(x) = \langle f \rangle \langle g \rangle. \tag{A.11}$$

One can think of g as defining some initial distribution (analogue to Q in eq. (A.2)), properly normalized ($\langle g \rangle = 1$) and of f as being our observable.

The above property is called mixing and implies ergodicity. Eq. (A.2) was true iff the matrix P of transition probabilities satisfied certain spectral properties. We shall see later that an analogous spectral requirement arises here also. The main difference

will be that whereas in the previous case, it was very easy to choose P to satisfy the spectral conditions, this is very hard to do in our case. There are very few examples of dynamical systems which have been proved to be mixing. This is the source of our concern.

The space H of all complex-valued functions on M is a Hilbert space with inner product given by,

$$(f, g) = \int \mathrm{d}\mu(x) f^*(x) g(x). \tag{A.12}$$

φ induces a unitary operator U on H:

$$(Uf)(x) = f(\varphi(x)). \tag{A.13}$$

Our system is ergodic iff unity is a simple eigenvalue of U. For P we knew that all eigenvalues are inside the unit circle. For U we know that they are on it. If the system is ergodic and mixing, the single discrete eigenvalue is unity. It is then true that for any f and g,

$$\lim_{n \to \infty} (U^n f, g) = (f, 1)(1, g) \tag{A.14}$$

and this is equivalent to mixing.

Thus, in analogy with the previous case, we have to get rid of the oscillations induced by eigenvalues that are pure phases. For a system with a finite number of states this is rigorously impossible because the reversibility of φ implies Poincaré recurrences. It is only when the volume of the system tends to infinity that mixing might hold exactly.

A sufficient condition [12] for mixing is that as $V \to \infty$, there exist a basis $\{f_\alpha^J\}_{\alpha \in I, J \in Z}$ of H such that,

$$U f_\alpha^J = f_\alpha^{J+1}. \tag{A.15}$$

In practice, periodic recurrences can be avoided by introducing a small amount of randomness in an otherwise deterministic process. For example, one can sample the sites of the lattice randomly instead of in a well defined sequence.

Let $\{C\} = S$ be the set of states of our Ising system. We also need an auxiliary space of states $\{D\} = d$. We consider henceforth $M = S \times d$. On M we define the measure μ which assigns equal weight to every configuration. We need now a mapping φ: $M \to M$ which preserves this measure. This means that φ must have an inverse. For practical purposes we want φ to be both simple and implementable by parallel processing. Therefore we want φ to be local. We look for a φ which is expressible in terms of "elementary" mappings acting on individual sites. Such an elementary mapping acts on one site of the Ising system and on the set d. We require

that the mapping preserve one quantity: the "energy" E_T, which is an observable on M. We also require E_T to be additive. The space d is the space of states of the "demons." One can easily generalize this to include any number of conserved quantities (one would always require additivity). The space d must be manageable; i.e. canonical or microcanonical averages on it can be computed analytically.

The demons serve two purposes: the first is that they make ergodic behavior more likely. Indeed, because of the simplicity of φ, surfaces of constant E_s might split into regions disjoint under φ (trapping because of impenetrable barriers). The existence of the demons allows limited violations of the conservation of E_s which helps to overcome "small" barriers. The second purpose that is served by the demons is that they allow the "measurement" of variables conjugate (in the thermodynamic sense) to the conserved quantities. This is because, for small demon systems, the Ising system acts as a macroscopic reservoir. Even when the space d is large, one can still make it manageable (see ref. [2]) but there the addition of more conserved quantities (conjugate to more couplings) becomes a problem.

This discussion explains why it is difficult to rigorously establish ergodicity or mixing in deterministic microcanonical simulations. However, the injection of small amounts of randomness into the process probably makes it convergent in a sense which is likely to be even stronger than mixing. We consider the present paper as experimental proof of this claim.

References

[1] M Creutz, Quarks, gluons and lattices (Cambridge University press, to appear)
[2] D Callaway and A Rahman, Phys Rev Lett 49 (1982) 613,
D Callaway, Lattice gauge theory in the microcanonical ensemble, Argonne preprint, ANL-HEP-PR-83-04 (January 1983)
[3] M Creutz, Phys Rev Lett 50 (1983) 1411
[4] N G van Kampen, Stochastic processes in physics and chemistry (North-Holland, 1981)
[5] L Jacobs and C Rebbi, Brookhaven preprint BNL-27842 (1980)
[6] R B Pearson, J L Richardson and D Toussaint, ITP, Santa Barbara preprints, NSF-ITP-81-139 (January 1982) and NSF-ITP-82-98 (October 1983)
[7] B M McCoy and T T Wu, The two-dimensional Ising model (Harvard University Press, Cambridge, Mass 1973)
[8] J Shapiro, private communication
[9] M Falcioni, E Marinari, M L Paciello, G Parisi and B Taglienti, Phys Lett 108B (1982) 331
[10] F R Gantmacher, Applications of the theory of matrices, translated by J L Brenner (Interscience, New York, 1959)
[11] K Wilson, Monte Carlo calculations for the lattice gauge theory, Cornell University preprint, CLNS/80/442 (January 1980)
[12] V I Arnold, A Avez, Ergodic problems of classical mechanics (Benjamin, New York, 1968),
P R Halmos, Lectures on ergodic theory (Chelsea, New York, 1956)