6.7 This problem deals with the extension of the high temperature series expansion of the spin-1/2 Ising model, developed in Section 6.1, to include a field term.
(i) Show that the partition function of the spin-1/2 Ising model in a field can be written

$$\mathcal{Z} = \cosh^{\mathcal{B}} \beta J \cosh^{N} \beta H \sum_{\{s\}} \prod_{\langle ij \rangle} (1 + s_i s_j v) \prod_i (1 + s_i y)$$

where $v = \tanh \beta J$; $y = \tanh \beta H$, and $\mathcal{B}$ and N are the number of bonds and sites on the lattice respectively.
(ii) Show that the terms in $\mathcal{Z}$ can be represented by graphs on a lattice where each graph with l bonds and m odd vertices contributes a factor $2^N v^l y^m$, and hence that $\mathcal{Z}$ can be rewritten

$$\mathcal{Z} = \cosh^{\mathcal{B}} \beta J (2 \cosh \beta H)^N (1 + \mathcal{S}_0 + y^2 \mathcal{S}_2 + y^4 \mathcal{S}_4 + \ldots) \quad (6.21)$$

where $\mathcal{S}_m(v, N)$ is the contribution from all graphs with m odd vertices.
Taking the logarithm of eqn (6.21) and using the Linked Cluster Theorem[6] gives

$$-\beta \mathcal{F} = \mathcal{B} \ln \cosh \beta J + N \ln 2 \cosh \beta H + (\mathcal{S}'_0 + y^2 \mathcal{S}'_2 + y^4 \mathcal{S}'_4 + \ldots) \quad (6.22)$$

where the prime denotes that, for a given graph, only the part of the count proportional to N must be included.
(iii) Differentiate eqn (6.22) to show that the zero-field susceptibility is

$$\chi = \beta N + 2\beta \mathcal{S}'_2.$$

(iv) List the low order graphs with two odd vertices to show that the first few terms in the high temperature expansion of the zero-field susceptibility of the spin-1/2 Ising model on the square lattice are

$$\chi = \beta N + 2\beta(2v + 6v^2 + 18v^3 + \ldots).$$

[6] See eqn (6.7) for an explicit example of how this works.

7

Monte Carlo simulations

It could be argued that current physics research can be divided into three areas—theoretical, experimental, and computational. Numerical approaches, in which systems are mimicked as accurately as possible using a computer or in which computer models are set up to provide well-behaved experimental systems are increasingly providing a bridge between theory and experiment. The limitations on what can be done are set by the computational resources available.

A powerful numerical approach is the Monte Carlo method. It was introduced in 1953 at the dawn of the computer age and its range of applicability and accuracy have continued to increase with the development of more advanced computer technology. One of the simplest and most natural applications, which we shall focus on here, is to discrete spin models. However the technique is very widely used: to study continuous spin systems, fluids, polymers, disordered materials, and lattice gauge theories. Some examples are given at the end of this chapter.

7.1 Importance sampling

A common aim in statistical mechanics is to find the value of a thermodynamic variable, such as the energy or the magnetization, which is a weighted sum over all states in phase space

$$\langle A \rangle = \frac{\sum_{\{s\}} A e^{-\beta \mathcal{H}}}{\sum_{\{s\}} e^{-\beta \mathcal{H}}}. \quad (7.1)$$

For an Ising model on a lattice of N sites the sum is over 2^N configurations. This is a number which increases very quickly with N and a direct evaluation is feasible only for $N \lesssim 40$.

The first way one might try to get round this is to choose randomly a sample of the spin configurations, $\{s\}$, and, weighing them appropriately according to eqn (7.1), work out an estimate of the required average. This approach may be familiar as it is a standard technique used for the evaluation of integrals. However, it fails here because of the rapid variation of the Boltzmann factor, $e^{-\beta\mathcal{H}}$, with energy. Very few of the chosen configurations will be weighted by a sufficiently large factor to make a significant contribution to the average and a very unreliable estimate will result.

This problem occurs because only an extremely restricted part of configuration space is important in determining the averages. This we already know from statistical mechanics—the system spends the vast majority of its time in states with thermodynamic parameters within $O(1/\sqrt{N})$ of those describing thermodynamic equilibrium. Therefore it would seem sensible to restrict the sampling to these states. This is a technique known as importance sampling. But how to generate such a set of states? To try to find the probability distribution exactly the partition function would need to be calculated and this is tantamount to going back to the original problem of summing over an impossibly large number of states.

Luckily it turns out to be possible to generate a Markov chain of configurations (a sequence of states each of which depends only on the preceding one) which has the property that $\hat{A}_n$, the average of A over n successive states, converges to the thermodynamic average defined in eqn (7.1)

$$\hat{A}_n = \langle A \rangle + O(n^{-1/2}). \tag{7.2}$$

In the limit $n \to \infty$ each state is weighted by its Boltzmann factor, $e^{-\beta E}$. The disadvantage of this approach is that successive states of the Markov chain are highly correlated, which means that a much longer sequence of sample configurations is needed to achieve a given accuracy than if this were not the case.

The conditions on the transition probability between Markov states needed to achieve the result (7.2) are physically transparent. The transition probability must be normalized. It must be ergodic, that is all states must eventually be accessible. Finally, a sufficient condition is that it must obey detailed balance[1].

This does not specify the transition probability uniquely. The choice often used in Monte Carlo simulations is the Metropolis algorithm. A final state, $\{s\}_f$, is chosen from an initial state, $\{s\}_i$, by

[1]Parisi, G. (1988). *Statistical field theory*, p.346. (Addison-Wesley, Wokingham).

flipping one or more spins. The probability that the system is allowed to move from i to f is

$$\begin{aligned} P(\{s\}_i \to \{s\}_f) &= e^{-\beta(E_f - E_i)} && \text{if } E_f > E_i \\ &= 1 && \text{if } E_i \leq E_f \end{aligned} \tag{7.3}$$

where E_i and E_f are the energies of the initial and final states respectively.

There is a physically intuitive argument that shows that with this choice of transition probabilities the system tends asymptotically ($n \to \infty$) to a steady state in which the probability of a given configuration is $e^{-\beta E\{s\}}$. Consider m_r systems in a state $\{s\}_r$ and m_t in a state $\{s\}_t$ such that $E_t < E_r$. Using random numbers it is possible to construct a move such that the *a priori* probability of moving from state r to t is the same as that to move from t to r. (This is feasible but not always the case in realistic simulations.) Then, using eqns (7.3), the number of transitions from r to t and from t to r are

$$M_{r\to t} \propto m_r \tag{7.4}$$

$$M_{t\to r} \propto m_t e^{-\beta(E_r - E_t)} \tag{7.5}$$

respectively. The net number of transitions is

$$\Delta M_{r\to t} \propto \{m_r - m_t e^{-\beta(E_r - E_t)}\}. \tag{7.6}$$

The system will converge to a steady state where $\Delta M_{r\to t} = 0$ or

$$\frac{m_r}{m_t} = \frac{e^{-\beta E_r}}{e^{-\beta E_t}}. \tag{7.7}$$

7.2 Practical details

The steps involved in setting up a Monte Carlo simulation for a simple spin model are listed in the flow chart in Table 7.1. This is the basis of the program used to generate the spin configurations in Fig. 1.8. The procedure can be thought of in three parts. We concentrate in this section on the details of how to set up the program and return in the next to a fuller discussion of the problems inherent in the data analysis.

Setting up. The first task is to define a lattice of N sites, i, each of which is occupied by a spin, s_i. This needs to be done in such

Table 7.1. Flow diagram showing the steps in a Monte Carlo calculation of thermodynamic averages for a simple spin model. The system takes n_0 steps to reach equilibrium and the total number of steps is n_{max}

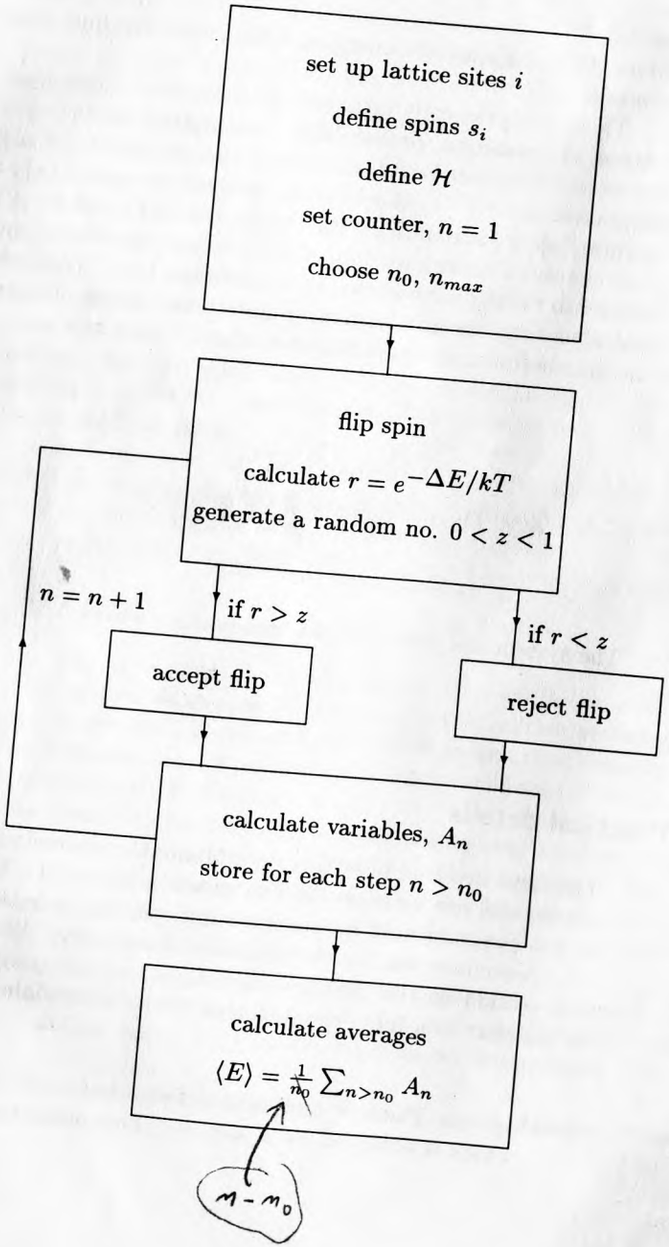

a way that a record is kept of the neighbours of each spin as its energy will be needed later. The parameters in the problem, such as the temperature and exchange interactions, should also be defined here.

Because N is necessarily finite thought must be given as to what to do with the spins on the boundaries of the system. These can either be left with fewer bonds than usual (free boundary conditions) or assumed to interact with the corresponding spin on the opposite face of the lattice (periodic boundary conditions). The latter option often gives the best results, but care must be taken that the system is not subject to false constraints. For example, simulations on a simple antiferromagnet with periodic boundary conditions can be expected to give inaccurate or spurious results if the length of the lattice is an odd number of spins.

Another consideration is the choice of initial values for the spins. Usually any choice will eventually lead to thermal equilibrium but it is helpful if this happens sooner rather than later. For a simple ferromagnet a ferromagnetically ordered state is likely to provide the most efficient initial configuration at low temperatures; at higher temperatures a random state provides the best starting point. We return to the problems of convergence to equilibrium and finite system size in Sections 7.3.1 and 7.3.3 respectively.

Generating the Markov chain. This is the heart of the program. It is summarized in the centre portion of Table 7.1. The steps are listed below

1. Select a spin, either randomly or sequentially. Calculate $r = e^{-\Delta E/kT}$ where $\Delta E = E_f - E_i$ is the change in energy associated with a possible spin flip (to a randomly chosen final state if the spin has more than two states).
2. Compare r to a random number $0 < z < 1$.
3. Flip the spin[2] if $r > z$.
4. Use the final configuration (whether the test spin was flipped or not) to generate the value of any thermodynamic quantity to be averaged. Store this value.

[2]It is not hard to convince oneself that this procedure reproduces the transition probability given by eqn (7.3): for $\Delta E < 0$, $r > 1$ and hence the spin is always flipped; for $\Delta E < 0$, the probability that $z < r$ is r and hence the spin is flipped with probability $r = e^{-\Delta E/kT}$.

It is important to be aware that any bias in the random number generator will introduce systematic errors into the results. The evidence is that the random number generators built into modern computers have sufficiently good statistics that the errors are insignificant compared to statistical errors. The question of how many configurations are needed to give satisfactory averages is discussed in Section 7.3.2.

Calculating the averages. Average the thermodynamic variables generated at each step of the Markov chain. Care must be taken not to include the initial states where the starting configuration still has an influence. The magnetization and energy are the easiest quantities to calculate as they are just sums over spins or products of spins.

7.3 Considerations in the data analysis

7.3.1 Influence of the starting configuration

During the first iterations of the Monte Carlo procedure the system is not in equilibrium, and hence these configurations cannot be included in the final averages. It can be hard to decide how many steps to exclude. One possibility is to perform several Monte Carlo runs with the same parameters but using different starting configurations. If the results agree to within statistical error it can be concluded that the influence of the starting configuration has been eliminated. A circumstance that can nullify this procedure, which has caused confusion in the past, is that a system can become stuck in a metastable state and feign true thermal equilibrium.

If the simulation is performed near the critical temperature, the additional problem of critical slowing down is encountered. Because of the increasing range of the correlations as criticality is approached, the time for relaxation to equilibrium τ diverges

$$\tau \sim \xi^z \tag{7.8}$$

with $z \sim 2$ for most models. In a finite system the divergence is suppressed; the smaller the system the quicker equilibrium can be achieved for a given temperature. However, at the same time finite-size corrections become more severe and a balance between these and equilibration times must be struck in the design of a Monte Carlo simulation.

An example of raw data from a Monte Carlo simulation showing the approach to equilibrium and the importance of excluding the initial

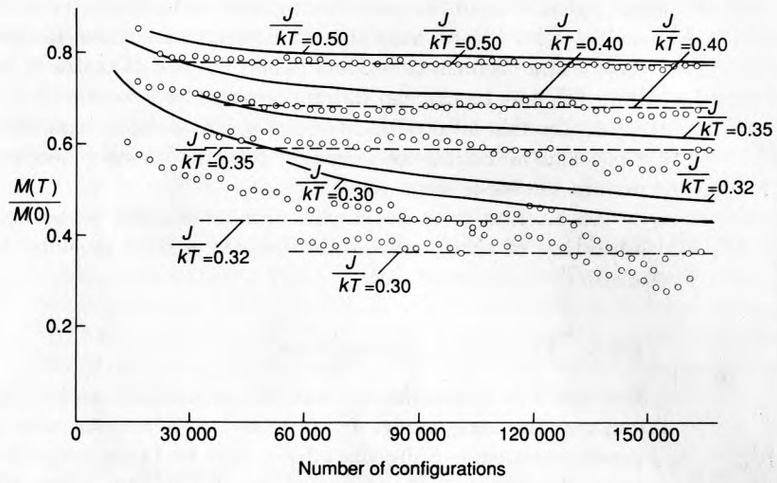

Fig. 7.1. Magnetization of an Ising ferromagnet on a cubic lattice of size $10\times10\times10$ with periodic boundary conditions plotted as a function of the number of Monte Carlo configurations for different temperatures. Open circles denote averages taken over the three preceding Monte Carlo steps per spin. Full curves give a running average if no initial configurations are excluded. The dashed lines are the final estimates of the magnetization where initial configurations have been excluded. After Binder, K. and Rauch, Z. (1969). *Zeitschrift für Physik*, **219**, 201.

configurations is shown in Fig. 7.1. Note that equilibrium is attained after a few Monte Carlo steps per spin for temperatures sufficiently far from the critical point but a slower relaxation and larger fluctuations are observed closer to T_c ($J/kT_c = 0.22$).

7.3.2 Statistical errors

To obtain reliable results for the equilibrium value of an observable, $\langle A \rangle$, the average must be taken over a time much longer than that over which the Monte Carlo states are correlated. This becomes more difficult near the critical point or if there are metastable states in the system. It can be shown that the deviation of $\hat{A}_n$ from $\langle A \rangle$ is normally distributed in the limit $n \to \infty$. Thus standard data analysis can be applied to determine the statistical error.

Dividing the equilibrium configurations into independent blocks

and calculating $\hat{A}$ for each block gives a set of essentially independent estimates of $\langle A \rangle$, the variance of which gives a value for the sampling error. The problem is to know when a block of states is long enough that different blocks can indeed be considered mutually independent. A test for this is to perform the analysis using several different block sizes. The blocks are long enough when the variance becomes independent of the block size.

An alternative is to average over several different runs. The disadvantage of this procedure is that the system must be equilibrated afresh for each set of data.

7.3.3 Finite-size corrections

Because it is impossible to simulate an infinite system—for a three-dimensional Ising model $N = (128)^3$ is a realistic size on modern supercomputers— finite-size effects must be taken into account. Away from the critical point, where the correlation length is small compared to the system size, this is usually not a major problem and the parameters of the simulation can be chosen so that the errors due to the limitations in the number of spins are small compared to statistical errors.

The problem becomes much more acute as a continuous phase transition is approached, because on a finite lattice the correlation length is prevented from becoming infinite. As a result any singularities associated with the phase transition are shifted and rounded. The best compromise is to obtain high quality data for lattices of different linear dimension L, and extrapolate to $L = \infty$.

7.4 Examples

7.4.1 The three-dimensional Ising model

A lot of effort has been put into Monte Carlo simulations for the three-dimensional Ising model. This is partly because of its suitability for fast algorithms and partly because of the intrinsic interest in obtaining precise values for the critical parameters. Special purpose machines, in which the time-consuming spin updating is carried out by a specially constructed processor, have been built in several places. These are orders of magnitude cheaper than supercomputers but can only carry out the specific task for which they were designed. They can achieve an accuracy comparable to that obtained by careful programming on the most powerful conventional computers.

As an example of the run times that can be achieved we give some

figures from a machine at Santa Barbara, USA[3]. This can update 10^7 spins each second. Lattice sizes of up to $N = 64 \times 64 \times 64$ were used and data for 10^7–10^8 Monte Carlo steps per spin collected for each size. At the critical temperature on the largest lattices the time to come to equilibrium was of the order of 7000 Monte Carlo steps per spin.

Using these data the result for the critical temperature was $K_c = J/kT_c = 0.221650(5)$ where the figure in brackets is the estimate of the error in the last digit. This agrees with, and is comparable in accuracy to, the best estimate from series expansions, $K_c = 0.221655(5)$[4]. The value obtained for the exponent ratio $\gamma/\nu = 1.98(2)$ is considerably less precise because of the problems of finite-size effects. Better results can be obtained using the Monte Carlo renormalization group, a technique which combines the strengths of the renormalization group and Monte Carlo simulations. This will be described in Chapter 9.

7.4.2 More complicated systems

Although Monte Carlo is particularly well suited to simulations of the Ising and other discrete spin models it was originally introduced in relation to fluids and has proved useful both here and in many other contexts. The most fundamental difference between simulations on different systems is in the choice of test configuration.

For example, for fluids, one possibility is to choose a molecule at random and allow it to move through a distance chosen at random between 0 and Δ in a random direction. The most accurate results are obtained if Δ is chosen so that approximately half the trials are accepted. Many different models have been considered in the literature, ranging from a gas of hard sphere molecules to attempts to incorporate realistic interatomic potentials. Common aims are to calculate the equation of state or the pair correlation function.

With today's computational power it is feasible to obtain realistic results for even more complicated systems. One example is solutions of polymers, long chain molecules, where Monte Carlo has been particularly useful in looking at properties which depend on the polymer topology rather than the details of the chemistry. Here the so-called reptation technique is one of the most efficient ways of generating suitable sequences of states. Starting from an arbitrary configuration the end of one of the chains is removed at random and added to the other

[3]Barber, M. N., Pearson, R. B., Toussaint, D., and Richardson, J. L. (1985). *Physical Review*, **B32**, 1720.

[4]Adler, J. (1983). *Journal of Physics A: Mathematical and General*, **16**, 3585.

end of the chain. As long as the self-avoidance of the chains is preserved the move is accepted according to the usual Metropolis criterion. Results have been obtained for such diverse problems as the changes in chain morphology as the temperature or solvent composition are varied, for the properties of chain molecules at surfaces, and for the dynamics of tangled polymers.

7.5 Problem

7.1 Write a Monte Carlo program to determine the temperature dependence of the energy and magnetization of a two-dimensional Ising model on a square lattice. Choose a lattice size appropriate to the power of the computer you are using. Useful illustrative results can be obtained using lattices of size as small as 6×6. Discuss

(i) the initial conditions used
(ii) the boundary conditions
(iii) the number of steps required to achieve thermodynamic equilibrium
(iv) error bars for the results at each temperature
(v) the effect of the finite system size.

8

The renormalization group

The approaches described so far in this book have given a broad p nomenological understanding of critical phenomena. However, altho a substantial framework of results and connections has been built we have, as yet, no explanations for the following:

1. Continuous phase transitions fall into universality classes char terized by a given value of the critical exponents.

2. For a given universality class there is an upper critical dimensi above which exponents take on mean-field values.

3. Relations between exponents, which follow as inequalities fro thermodynamics, hold as equalities.

4. Critical exponents take the same value as the transition tempe ature is approached from above or below.

5. Two-dimensional critical exponents often appear to be rationa fractions.

What is needed is a theory, based on the physics of what is hap pening at the critical point. We argued in Chapter 1 that the specia feature of criticality is that the correlation length is infinite and that the critical system is invariant on all length scales. The aim is to write down a (hopefully short, elegant, and comprehensible) mathematical theory which embodies this physics and explains all the observations listed above. A useful theory will also allow the calculation of critical exponents and transition temperatures, if not exactly, then within an accurate and well-controlled approximation scheme.