

The Two-Dimensional Ising Model

An Introduction to Monte Carlo Statistical Mechanics

A. Jackson (November 1997)

Contents

1	Introduction	2
2	Theory	2
2.1	Statistical Physics	3
2.2	The Ising Model	4
2.3	The Phase Change Behaviour of the Ising Model	5
3	Method	6
3.1	Basics of the Monte Carlo Method	6
3.1.1	Sampling Techniques	6
3.1.2	The Metropolis Importance Sampling Algorithm	7
3.1.3	The Markov Chain	8
3.1.4	Initial State & Boundary Conditions	9
3.2	Finite Size Effects	9
3.3	Multi-Canonical Monte Carlo Simulation	10
3.4	Ensemble Averages and Error Estimation	11
4	Results	11
4.1	Results via Importance Sampling	11
4.2	Finite Size Effects	12
4.3	Results via Multi-Canonical Simulation	13
5	Discussion/Conclusions	16
6	Bibliography/References	16

1 Introduction

This report will outline the theory and findings of a simple Ising model simulation, and is intended to act as an introduction to the general concepts involved in applying Monte Carlo techniques to statistical physics. The two-dimensional Ising ferromagnet model is one of the simplest examples from this field, and yet its implementation allows many of the techniques of this kind of Monte Carlo simulation to be demonstrated. Also, the 2-D Ising model is one of the few statistical physics models that can be solved analytically, and so it provides a ‘gold standard’ against which to test numerical simulation techniques.

Firstly, the theory behind the Ising model will be examined, from the basic concepts of statistical physics to the specific implementation of the Ising ferromagnet. Also, the general behaviour expected from the Ising model (including the phase change behaviour) will be discussed.

Following the theory, the concepts of the Monte Carlo approach to statistical physics problems will be outlined, with the Ising model being used as an example. The basic methods of canonical Monte Carlo will be explained, along with the limitations of such simulations. Following this the concepts of multi-canonical Monte Carlo simulation will be introduced. The next step will be to outline the ways in which properties such as the magnetisation of a system can be extracted from the simulation, and how the errors in these results can be estimated.

Results will be presented for the canonical and multi-canonical approaches. This will lead to a discussion on the performance of the simulations, including notes on how the techniques might be improved. Finally, the wider applicability of this work will be examined, with reference to the recently developed lattice-switch MCMC technique.

2 Theory

Statistical physics is the study of systems with many degrees of freedom, the behaviour of which can be described in terms of the variation of a few macroscopic quantities. In order to determine the behaviour of such a system, we can choose to fix three of the macroscopic quantities, and then calculate how the other properties of the system behave. Each set of fixed macroscopic quantities forms an ensemble. For example, if the number of particles (N), the volume of the system (V) and the energy of the system (E) are all fixed, then we have the micro-canonical ensemble, which corresponds to an isolated system. The choice of ensemble can be made in order to make the problem as easy to solve as possible, as all ensembles are equivalent in the thermodynamic limit (i.e. as $N \rightarrow \infty$). This report mainly concerns the canonical ensemble, where N and V are fixed, along with the temperature T . This corresponds to a system in a heat bath, free to exchange energy with its surroundings.

2.1 Statistical Physics

We start with a system for which we can calculate the energy of any microscopic configuration it may be in. These configurations, called microstates, may be a particular array of spins on a mesh, or they may be defined by the locations of a set of particles and the interactions between them. While there may be very many possible microstates in a system with many degrees of freedom, the system can still be described in terms of a few macroscopic variables, each set of which defines a macrostate.

The basic assumption behind statistical physics is that the system is equally likely to be found in any of the possible microstates. However, the density of states (the degeneracy) means that the probability of occupying most of the possible macrostates will be small, and the system will be confined to a small range of macrostates for a given set of canonical conditions.

The number of microstates (Ω) for a particular macrostate (for the canonical ensemble, a particular set of values for N , V & T) can be used to form the Boltzmann entropy:

$$S(N, V, T) = k_B \ln \Omega(N, V, T) \quad (1)$$

Where k_B is Boltzmann's constant. This means that when a system is in equilibrium (i.e. the state of maximum entropy), the system is occupying the energy state which has the highest degeneracy. This expression allows the Boltzmann entropy to be related to the equations of thermodynamics, and expressions for quantities such as the temperature and free energy to be found, in terms of Ω . For the canonical ensemble, we consider a system that is free to exchange energy with a large heat bath. The degeneracy of the bath (Ω_B) determines the probability (P_i) of finding the smaller system in a particular microstate (i):

$$P_i = \frac{\Omega_B(E - E_i)}{\sum_j \Omega_B(E - E_j)} \quad (2)$$

Where E_i is the energy of the i th microstate of the sub-system, and E is the total energy of the whole system (i.e. including the heat bath). Via thermodynamic theory, equation (1) can be used to re-express the degeneracy in terms of the temperature of the system:

$$P_i = \frac{e^{-\beta E_i}}{\sum_j e^{-\beta E_j}} \quad (3)$$

This forms the Boltzmann distribution, the denominator of which is known as the partition function ($Z(\beta)$), and where β is the inverse temperature ($1/k_B T$). The mean energy of a system can then be found as simply:

$$\bar{E} = \sum_i E_i P_i = \frac{\sum_i E_i e^{-\beta E_i}}{Z} \quad (4)$$

While this is a quantum mechanical expression, the eigenstates E_i can easily be re-expressed in classical terms to form a continuous integration over all phase-space. It is these integrations, continuous or discrete, over configuration/phase space which make the numerical solution of statistical

physics problems difficult. While the general form of the behaviour can be found via simulation, it is generally not possible to find the partition function as this depends of the density of states of the system ($g(E)$). Replacing the sum over eigenstates with a sum over energy states we have:

$$Z = \sum_E g(E)e^{-\beta E} \quad (5)$$

Where the scale of the density of states is not know, but the form of it is. However, if we know how many microstates make up any particular energy state, then we can calibrate the scale of $g(E)$. This is useful because once the magnitude of the partition function is known, then we can calculate the Helmholtz free energy:

$$F = -k_B T \ln Z \quad (6)$$

The minimum value of which will correspond to the state the system will choose to occupy under a given set of conditions (N, V, and T). Although it is not generally possible to find the absolute values of the free energy (i.e. when Z is unknown), the relative free energy of two particular states can usually be found, and this indicates the relative stability of the two configurations. The advantage of the Ising model is that not only is the analytical solution known (in 2-D), but the number of microstates for the ground state of the system is also known. This means that the partition function and thus the relative free energy of the Ising system can be calculated. We can also find the specific heat of the system from:

$$C = \left(\frac{\delta \bar{E}}{\delta T} \right)_v \quad (7)$$

$$\bar{E} = -\frac{\delta \ln Z}{\delta \beta} \quad (8)$$

Which, it can be shown, lead to the expression:

$$C = \frac{1}{k_B T^2} \frac{\delta^2 \ln Z}{\delta \beta^2} = \frac{\sigma_E^2}{k_B T^2} \quad (9)$$

Where σ_E^2 is the variance in the energy distribution (i.e. the width of the energy peak).

2.2 The Ising Model

This simple model of a ferromagnet is based upon the concept of interacting spins on an unchanging lattice. The particular form of this model considered here is that of an infinite regular square lattice in two-dimensions only. The spin at each lattice site (i) interacts with its four nearest neighbour sites such that the overall hamiltonian for the system is:

$$H = -J \sum_{\langle i,j \rangle} S_i S_j \quad (10)$$

Where the sum over $\langle i, j \rangle$ represents a sum over the nearest neighbours of all the lattice sites, and where J is the nearest neighbour interaction energy (which can be set to one, as this just changes the scale of the units of energy in the system). This spin on each site is defined by the value $S_i = \pm 1$.

By setting up a system obeying this Hamiltonian, we can find the mean energy and the specific heat per spin for the Ising ferromagnet from equations (4) and (9) respectively:

$$\bar{E} = \frac{1}{N} \frac{\sum_i E_i e^{-\beta E_i}}{Z} \quad (11)$$

$$C = \frac{\sigma_E^2}{N k_B T^2} \quad (12)$$

Also, we can calculate the mean magnetisation from the following ensemble average:

$$\bar{M} = \frac{1}{N} \sum_i S_i \quad (13)$$

We are especially interested in the variation of these quantities with temperature, as this system will undergo a phase change, and so discontinuities should appear in its macroscopic properties at a particular value of T (T_c , corresponding to $\beta_c = 0.440686$). As mentioned above, the number of microstates for the ground state of the system is known, as the lowest energy configuration is either all up or all down (i.e. $g(E_{gs}) = 2$). In this way the absolute free energy of any state can be found.

It should be possible to compare the behaviour of \bar{E} , C , \bar{M} and also $Z(\beta)$ with the results in the literature to determine if the simulation is giving reasonable results.

2.3 The Phase Change Behaviour of the Ising Model

The balance between the thermal energy available ($k_B T$), and the interaction exchange energy (J) defines the behaviour of the Ising model, and a general picture of the system behaviour can be worked out from this information alone. At high temperature (i.e. at low β), where the thermal energy is much greater than the interaction energy, the interaction will be drowned out by the randomising effects of the heat bath, and so the pattern of spins on the mesh will be random and there will be no overall magnetisation. However, as the temperature is lowered, the correlations between the spins can begin to take effect. Lower temperatures allow clusters of like spins to form, the size of which defines the correlation length (ξ) of the system.

This correlation length increases as T decreases until, at the critical temperature (T_c), it becomes infinite, and the spins form a fractal pattern with as many up-spins as down-spins. As ξ diverges, so do the ergodic time (t_e) and the specific heat (C).

After this critical point, the system must make a decision as which spin direction is favoured, until at $T = 0$ all the spins are pointing either upwards or downwards. Both choices are energetically equivalent, and there is no way to pre-guess which one the system will choose. The breaking of a

symmetry, such having to choose the favoured direction for the spins, and the divergence of system parameters such as ξ , t_e and C are characteristic of phase changes in general.

3 Method

Whilst the problem we need to solve is neatly expressed in equation (4), the actual solution of such a summation is somewhat problematic due to the vast number of possible microstates which all but the simplest systems can explore. For example, if we take a relatively small number of spins for our Ising model simulation, say a 10×10 grid, then we have 100 degrees of freedom, and so the number of possible microstates is 2^{100} ($\sim 1.3 \times 10^{30}$ combinations). Going through all of these combinations will take far too long (even if we could calculate the energy of one billion combinations per second, which we cannot, the integration would still take 4.0×10^{13} years to complete).

A second option would be to perform some sort of quadrature integration procedure, where data is sampled from a grid of point in configuration space. The problem with this it that the Boltzmann factor is only significant in a tiny proportion of this space, and this procedure would usually miss the areas of interest. Furthermore, if the quadrature grid is fine enough to allow a reasonable integration, the time required for the calculation will become far too long for it to be useful. It is clear from this that a different approach is required.

3.1 Basics of the Monte Carlo Method

The Monte Carlo technique works by taking a random sample of points from phase space, and then using that data to find the overall behaviour of the system. This is analogous to the idea of an opinion poll. When trying to work out the opinion of the general public on a particular issue, holding a referendum on the subject is costly and time-consuming. Instead, we can ask a randomly sampled proportion of the population and use the average of their opinion as representative of the opinion of the whole country. This is much easier than asking everybody, but is not as reliable. We must ensure that the people asked are chosen truly at random (to avoid correlations between the opinions), and be aware of how accurate we can expect our results to be given that we are only asking a relatively small number of people. Issues such as these all have analogous issues for the implementation of the Monte Carlo technique.

3.1.1 Sampling Techniques

The most basic sampling technique consists of just randomly choosing points from anywhere within the configuration space of the system, and is called simple sampling. A large number of spin patterns are generated at random (for the whole mesh), and the average energy and magnetisation calculated from that data. However, this technique tends to suffer from exactly the same problems as the quadrature approach, often sampling from unimportant areas of phase space. The chances of a randomly created array of spins producing an all-up/all-down spin pattern is remote

($\sim 2^{-N}$), and a high-temperature random spin array is much more likely. The most common way to avoid this problem is by using importance sampling, which works by applying weights to the microstates. We can modify equation (4) to include these weights:

$$\bar{E} = \frac{\sum_i \frac{E_i}{N_i} e^{-\beta E_i}}{Z} \quad (14)$$

Where N_i is the probability of the system being in microstate i . Clearly, if we can modify the simulation procedure to make the probability of being in state i cancel out with the Boltzmann terms, i.e.

$$N_i = \frac{e^{-\beta E_i}}{Z} \quad (15)$$

Then equation (14) will simplify to give:

$$\bar{E} = \frac{1}{N} \sum_i E_i \quad (16)$$

One very useful and widely used technique for the implementation of this idea is the algorithm of Metropolis *et al.*

3.1.2 The Metropolis Importance Sampling Algorithm

The basic concept here is the idea of altering the probability of finding the system in microstate i by modifying the probability of the system moving to or from that state, relative to the other states in the system. These transition probabilities are based on the detailed balance condition, which states that the probability of moving from a state o to a state n must equal the probability of the opposite (n to o) transition occurring. Expressing this mathematically, we have:

$$N_o \pi(o \rightarrow n) = N_n \pi(n \rightarrow o) \quad (17)$$

Where N_i is the probability of being in macrostate i , and π represents the transition probability, such that:

$$\pi(o \rightarrow n) = \alpha(o \rightarrow n) \times acc(o \rightarrow n) \quad (18)$$

Where α represents the probability of generating the transformation ($o \rightarrow n$) out of all the possible transformations, and acc is the probability of accepting this transition. We can choose to make α constant, so all possible transitions have an equal probability of being generated, and then use equations (15), (17) and (18) to find the ratio of the acceptance probabilities:

$$\frac{acc(o \rightarrow n)}{acc(n \rightarrow o)} = \frac{N_n}{N_o} = e^{-\beta(E_n - E_o)} \quad (19)$$

While there are many ways to implement this balance condition, the technique used in the Metropolis algorithm is as follows. First generate the new trial microstate (n) from the current microstate

(o), ensuring that α is a constant. Then, accept this trial move with a probability given by:

$$acc(o \rightarrow n) = \frac{N_o}{N_n} \quad \text{when } N_n < N_o \quad (20)$$

$$= 1 \quad \text{when } N_n \geq N_o \quad (21)$$

The question that remains is, how the transition state n should be selected (when in the state o) in order to allow the algorithm to work effectively.

3.1.3 The Markov Chain

Generating a whole new random map of spins for each transition (as mentioned above) could work, but the implementation of the Metropolis algorithm in this regime would make little difference as the most likely spin patterns produced by this technique would be energetically equivalent random (high T) spin patterns. A more effective technique is to use a Markov chain. This is based on the concept of generating the new trial state n from the current state o by changing the current state only very slightly. In fact, the standard approach is to alter the state of a single degree of freedom in o to generate n . For the Ising model, this corresponds to flipping a single spin in the current microstate to create the trial microstate.

There are numerous ways of doing this, each of which has advantages and disadvantages. The most straightforward technique is the typewriter algorithm, which just scans across the mesh from left to right and from top to bottom, trial flipping each spin along the way. While very efficient in terms of processor time, this routine can suffer from problems due to the correlated way in which the trial moves are generated.

Although more computationally expensive, a less heavily correlated technique is to move randomly about the mesh, selecting the next trial spin flip position at random from all N positions. As this technique depends on the calculation of at least one random number, it will be slower than the simple typewriter technique.

An often-used form of the random spin selection is the random walk algorithm. This works by randomly choosing the next spin-flip by selecting one of the nearest neighbours to the previous spin-flip position. This will proceed to move the spin-trial position around the mesh, in such a way that corresponds well to the generation of clusters of like spins. This routine runs at approximately the same speed as the completely random selection routine, and is known to perform well, and so random walking shall be used in this simulation.

This implementation of the Markov chain algorithm leads to a new perspective on the Monte Carlo simulation technique. When the system is changing by only a small amount each time, it behaves like a time-dependent stochastic simulation (for example, a Molecular Dynamics simulation with an added random element to account for the presence of a heat bath). While the Markov chain corresponds to configurations at equilibrium, with no concept of time involved, it can be thought of in terms of a time-dependent simulation, and the dynamics of the system can be observed (up to a point, as the system has no concept of momentum). This leads to the idea of Monte Carlo time, the unit of which corresponds to each of the N degrees of freedom having been visited once (on

average) by the random walking algorithm. While the dynamical behaviour of the Ising model will not be studied here, the units of Monte Carlo ‘time’ steps (MCS) will be in use.

3.1.4 Initial State & Boundary Conditions

The primary parameter which defines the Ising model problem is the number of spins we will be considering (N). A square mesh is being used and side lengths $L = 10 - 20$ (i.e. $N = 100 - 400$) give reasonable results whilst not taking too long to calculate a reasonable number of Markov chain steps. However, we are trying to simulate the behaviour of an infinite lattice, and so periodic boundary conditions are used. The limitations of using such a small array of spins to simulate an infinite lattice via periodic boundary conditions are examined in section 3.2 below.

The final condition that we need to start the simulation is the initial condition of the spins on the mesh. There are two main possibilities, the ‘cold’ start and the ‘hot’ start. A cold start simply means that we start the simulation from the ground state, with all spins pointing in the same direction, whereas the hot start means we begin at the other extreme, with all spins arranged randomly. Given the system has enough time to equilibrate, the particular form of the initial conditions should not affect the final results, and so no comparison of these techniques has been included here. The hot start will be used throughout.

3.2 Finite Size Effects

One problem with simulations such as these is the approximation of an infinite mesh by a finite mesh subject to periodic boundary conditions. The problem is that a finite mesh is much more likely to become ordered than an infinite mesh, as this ordering of the spins will occur when the correlation length ξ becomes of the order of the side length L , as opposed to when $\xi \rightarrow \infty$. Thus the position of the energy peak (i.e. \bar{E}) will change depending on the size of the system (of course, the total energy will scale linearly with the size of the system, but the form of the mean energy per spin distribution will also change). Another expected effect due to the finite size of the model is the variation in the width of the energy peak. Away from the critical point, the variance of the peak will vary as $1/L^2$ and so the standard deviation as $1/L$. This results in the width of the peak being too great for smaller mesh sizes, and so the specific heat will be overestimated for low N systems. However, at the critical temperature, the mean energy per spin varies as $L^{-\frac{(1-\alpha)}{\nu}}$ and the variance as $L^{-\frac{\alpha}{\nu}}$, where α and ν are critical exponents.

The way in which these effects change the results must be understood in order to allow accurate comparison of simulation data with the analytical results from the system in the thermodynamic limit ($L \rightarrow \infty$). The effects can be examined by plotting the system parameter of interest (in this case, the energy of the system at the critical point, E_c , and the critical specific heat, $\propto \sigma_E^2$) against $1/L$ for a range of L . In this way, the value of the system parameter can be extrapolated back to $1/L = 0$, corresponding to an infinite system. For the special case of the 2-D Ising model, the critical exponents are such that both plots should be linear.

3.3 Multi-Canonical Monte Carlo Simulation

Another problem with importance sampling Monte Carlo, often more serious than finite-size effects, is that the simulations can easily become non-ergodic. For the Ising model, at temperatures less than or equal to the critical temperature the ergodic time will be sufficiently large for the simulation not to sample all configuration space. The magnetisation can be mostly upwards or mostly downwards, and as the probability of passing from mostly up to mostly down is very low, the system can become stuck in one or the other configuration space for most of the simulation. This leads to an unrealistically high probability of one of the two states, when each should be equally likely. Also, at the critical point the system is not likely to visit the ground state, and unless this region is well sampled, we cannot use our knowledge of $g(E_{gs})$ to find the critical partition function.

Multi-canonical Monte Carlo attempts to avoid this problem by applying weightings to the *macrostates* of the system, in much the same way as importance sampling applies weights to the microstates. This is achieved by breaking the distribution of a macroscopic variable up into a histogram consisting of $r = 1 \dots N_m$ bins. The probability of being in a particular microstate (eqn(3)) is modified such that:

$$P_i = \frac{1}{Z} e^{-\beta E_i + \eta_r} \quad i \in r \quad (22)$$

Where the microstate i falls within the r th macrostate bin. This modifies the detailed balance condition (19) such that:

$$\frac{N_n}{N_o} = e^{-\beta(E_n - E_o) + (\eta_{r(n)} - \eta_{r(o)})} \quad (23)$$

Where $\eta_{r(i)}$ refers to the weighting for the macrostate which contains the microstate i . Under this scheme, the sampling within the limits of each bin is exactly the same as before (i.e. Boltzmann importance sampling), but the sampling at the edges of each bin is modified to make the transition from the current bin to the neighbouring bin more or less likely than under importance sampling alone. Various schemes for the estimation and refinement of these weights have been suggested, but here the relatively simple *visited states; mean value estimator* scheme will be used. The weights are based on the data gathered in the course of sampling the macrostates, in the form of a histogram C_r consisting of a total of N_c visits to the macrostate bins (i.e. macrostates in the range of the histogram). Initially the weights ($\eta_r^{(0)}$) are all set to zero, and so the first run (after equilibration) consists of normal importance sampling. After this the following iteration is used to modify the weights:

$$\eta_r^{(n+1)} = \eta_r^{(n)} - \ln[C_r^{(n)} + 1] + k \quad (24)$$

Where k is an arbitrary constant fixed by the convention $\min_r(\eta) = 0$, thus ensuring all weights are positive. After each run, this iteration attempts to flatten the distribution of microstates, as the histogram data $C_r^{(n)}$ will tend to the value $1/N_m$ for each bin as $\eta_r^{n+1} \rightarrow \eta_r^{(n)}$. Once sufficient iterations have been performed, a single long run can be made with the final weights, and the weights can then be removed from the collected data:

$$e^{-\beta E_i} = e^{-\beta E_i + \eta_r} \times e^{-\eta_r} \quad (25)$$

The resulting distribution can then be analysed in the same way as the results from a straightforward importance sampling simulation.

3.4 Ensemble Averages and Error Estimation

While extracting the ensemble averages from a Monte-Carlo simulation is easy, we must be able to estimate the errors in the results in order to determine how reliable the simulation is. The variance in the ensemble averages is not a suitable error indicator, as this can correspond to a real physical quantity (for example, the variance in the energy of the Ising model corresponds to the heat capacity of the system, see equation (9)). This being the case, a different error estimation technique must be used.

While there are other techniques, the simple and reliable block error technique has been used here. Given we know the approximate correlation time of the system, we can split a long run of data up into N_b equal sections, each being longer than the correlation time. These blocks of data can then be treated as N_b independent sets of experimental results, and the ensemble average of any particular macroscopic quantity in each block (A_b) can be used to estimate the error in that quantity by taking the mean from each block and finding the variance of the block values (as they should form a gaussian distribution):

$$\sigma^2(\{A\}) = \frac{\sigma^2(\{A_b\})}{N_b} \quad (26)$$

Now, as the number of observations we make grows larger, the estimated value of the error in our results becomes smaller, as we would expect.

4 Results

4.1 Results via Importance Sampling

All the simulations sections 4.1 and 4.2 were carried out using a 10×10 system that had been given 1×10^4 MCS to equilibrate. After this, the histograms and block averages were calculated using 4 blocks of 1×10^4 MCS of simulation time.

Figure 1 shows the general form of the variation in the distributions of the energy and magnetisation as a function of temperature. Above the critical temperature ($\beta = 0.25$), the energy is high and the magnetisation forms a gaussian distribution about $M = 0$. At the critical point ($\beta = 0.4406868$), the energy distribution has widened and lowered, and the magnetisation has split into the two possibilities ($-M$ and $+M$). Below the critical temperature (in this case at $\beta = 0.5$), the energy has lowered further (beginning to include the ground state more often), and the magnetisation has moved nearer to ± 1.0 . Also, the right-hand peak is much taller than the left-hand peak. This is because the ergodic time of the system is now quite large, and the probability of moving from the $-M$ to the $+M$ side (or vice versa) is low. This means that during the relatively short simulation time, the system has mostly occupied the $-M$ region.

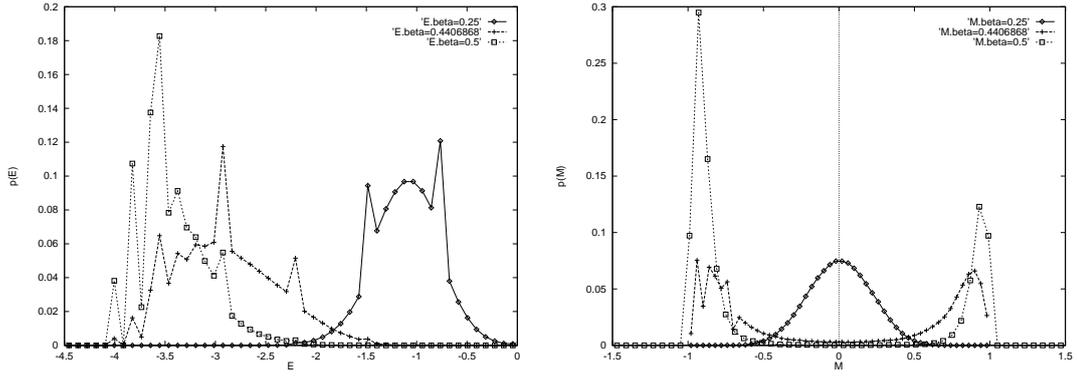


Figure 1: Variation of the energy and magnetisation distributions for temperatures above, below and at the critical point.

The overall trends are shown in figure 2. In the thermodynamic limit, the energy plot would contain a sharp step indicating the phase change and the specific heat would diverge to infinity. However, due to the finite size of the spin array, these discontinuities have been smoothed out.

While the poor resolution of the magnetisation plot obscures the details of the behaviour in the $T \leq T_c$ range, the plot is consistent with the expected results in that the average magnetisation per spin falls from 1.0 at $T = 0$ to 0.0 at $T = T_c$. In fact the results show that the magnetisation is not quite zero at T_c , and this is because the code looks at the absolute value of the magnetisation, and so the average value of the gaussian distribution about $M = 0$ does not always produce an ensemble average of $M = 0$, and in fact depends on the width of the distribution.

Error bars are shown in all three plots, but are very small compared to the scale of the results.

4.2 Finite Size Effects

The variation of the energy and the specific heat as functions of inverse system size are shown in figure 3 below. The data points appear to approximate the expected linear form reasonably well, with the least squares fit routine used giving the following results for the gradients (m_E , m_C) and the intercepts (c_E , c_C) from the set of coordinates and error bars for each graph:

$$m_E = -2.3394 \pm 0.4279(18.3\%)$$

$$c_E = -2.2585 \pm 0.0274(1.2\%)$$

$$m_C = 1.8122 \pm 0.1283(7.1\%)$$

$$c_C = 0.1450 \pm 0.0082(5.7\%)$$

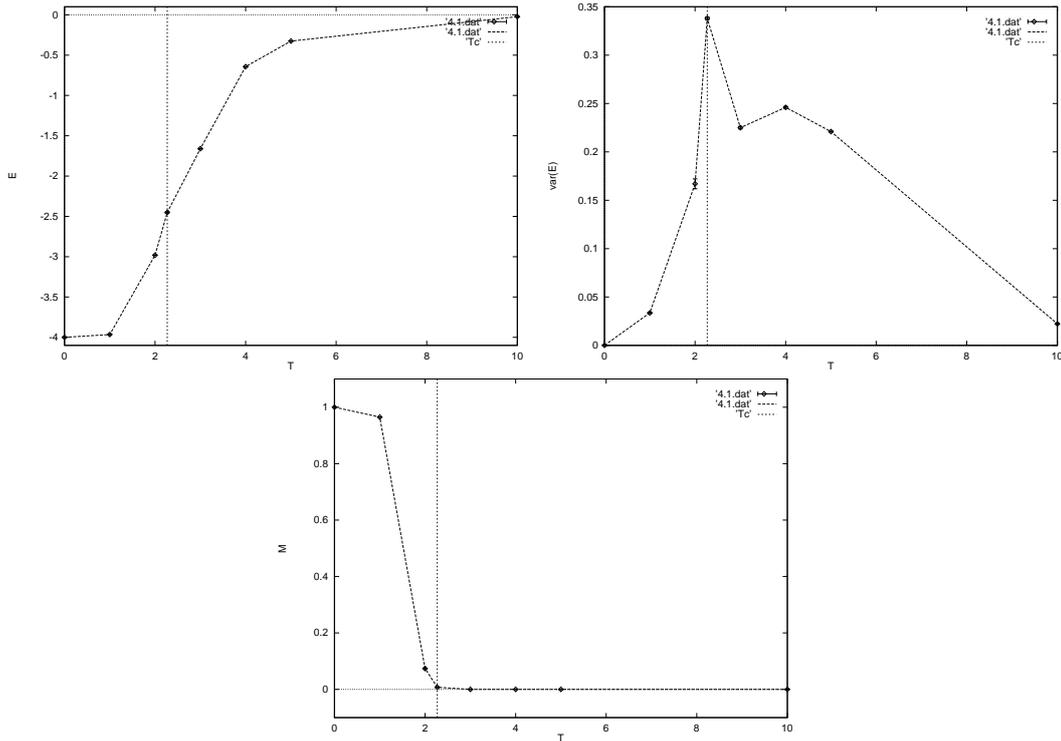


Figure 2: Variation of the average energy, the variance in the energy (proportional to the specific heat) and the magnetisation as a function of temperature.

Although the low number of data points used leads to quite high error in the least-squares fit results, the intercept of the specific heat (at 0.1450) corresponds favourably to the result quoted by Ferdinand and Fisher ^[1] of 0.1381.

4.3 Results via Multi-Canonical Simulation

The results in this section were collected from four iteration runs of the multi-canonical weighting code, each run being 3×10^4 MCS in length, for a 15×15 Ising model. A histogram consisting of 60 bins over the range -4.5 to 0.0 was used. The number of bins was chosen to be small enough to avoid the possibility of there being empty bins throughout the distribution after each run, and also large enough to give reasonable resolution of the form of the results.

Figure 4 shows the resulting energy distribution for the four multi-canonical runs. The initial unweighted run has the same form as before, but successive iterations of the weights quickly modifies the distribution, with the fourth run giving an almost flat distribution over the histogram range.

The way in which the iteration procedure modifies the weights is illustrated in figure 5. It is clear

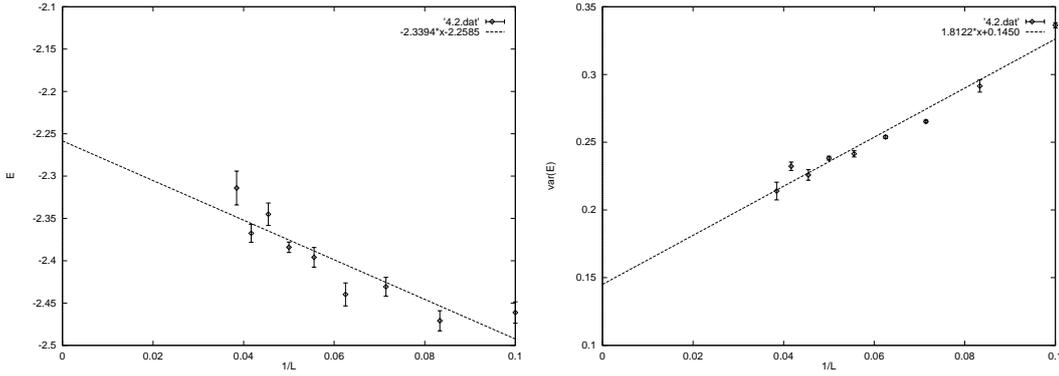


Figure 3: Variation of the average energy, and the variance in the energy as a function of inverse system size ($1/L$).

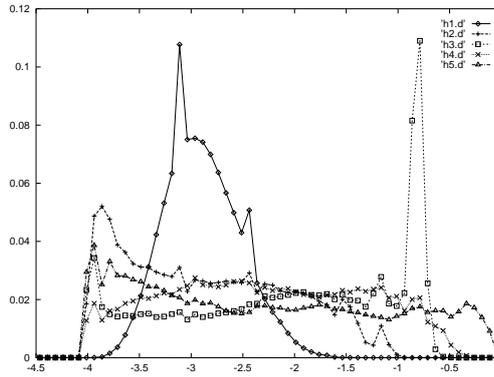


Figure 4: Distribution of energy states from four successive multi-canonical weight iterations.

that once a energy range is well sampled, then the weights change little upon iteration, and the weighting reaches out from the initially well sampled region to the initially unsampled region as the iteration proceeds.

The results shown in figures 4 & 5 both compare well with the visited states; mean value estimator results of Smith and Bruce [5].

The form of the results for each iteration following the removal of the weights is shown in figure 6 below. Successive iterations can be seen to allow the edges of the energy peak to be more accurately sampled. This includes the improved sampling from the ground state.

When calculating the partition function, we have to assume that the number of visits to the ground state is equal to the number of visits to the ground state bin. The validity of this assumption depends on the resolution of the histogram, as too few histogram bins will mean far too many different states are being counted as the ground state. The suitability of the chosen number of bins used in these results was only established by finding how many bins could be efficiently filled by the relatively short histogram collection runs I was using. Despite this, the calculation has been

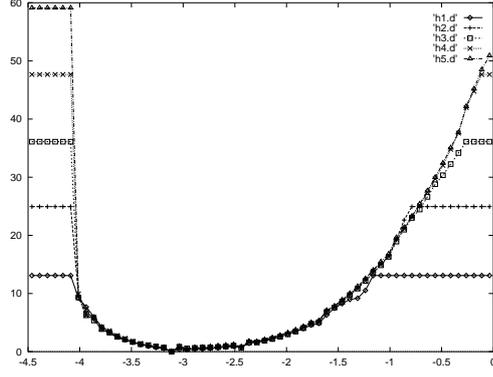


Figure 5: Distribution of multi-canonical weights during four successive iterations.

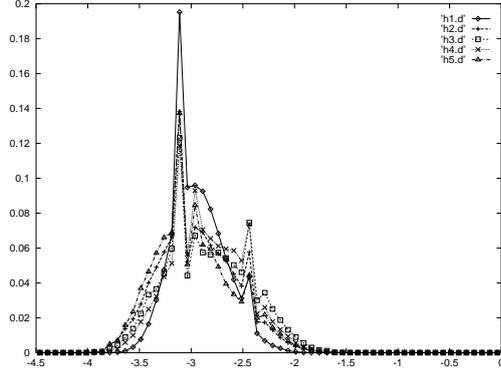


Figure 6: Distribution of energy states following removal of the multi-canonical weights for all four iterations.

found to yield reasonable results.

At the critical temperature the ratio of the probabilities of occupying the ground state and the mean energy state was found to be 2.867×10^{-4} . Thus the ground state will only be visited by the Ising model at the critical temperature once for every ~ 3500 MCS, which explains why the ground state is poorly sampled under a Boltzmann importance sampling scheme.

This ratio gives the probability of being in the ground state at the critical temperature, and can be substituted into equation (3) once it has been re-expressed in terms of a sum over energy and the density of states ($g(E)$, c.f. equation(5)). Then the critical partition function can be found for this system from the knowledge that $g(E_{gs}) = 2$, $E_{gs} = -4N$ and $\beta_c = 0.4406868$. This allows the critical free energy to be calculated:

$$F_{crit} = -920.083$$

Which, given how coarse the simulation histogram was, and the short period used for the histogram collection runs, compares favourably with the published large crystal critical free energy result (as calculated from equation (109c) of Onsager ^[7]) of $F_{crit} = -1006.245$.

5 Discussion/Conclusions

While the quantitative results of the simulations are not in strong agreement with the published results, the qualitative behaviour of the simulations does agree with that expected from the literature. The inaccuracy in the numerical results can be attributed to the relatively short data gathering runs used ($\sim 10^4$ MCS as opposed to $\sim 10^7$ MCS in the literature). Given these facts, it is reasonable to assume that I have successfully implemented the Ising model simulation in both the canonical and multi-canonical Monte Carlo schemes. However, it is worth noting that the code has not been fully optimised for speed and that improvements could be made by, for example, using look-up tables for the energy change probabilities to avoid carrying out so many calculations of the exponential of the transition energy at the very core of the code.

The advantage of the multi-canonical scheme over Boltzmann importance sampling has been demonstrated in the case of the critical free energy calculation (although I have not given any proof that the ground state really is poorly sampled under Boltzmann importance sampling). The multi-canonical scheme can prove to be of even greater value when the simulation involves sampling from two or more different regions of configuration space (corresponding to two or more regions of the order parameter distribution) where the probability of moving from one region to another is very small under normal sampling conditions.

One application of this technique is the lattice-switch Monte Carlo method of Bruce, Wilding & Ackland [8], where the multi-canonical weightings are used to allow a hard-sphere solid system to explore an area of configuration space where a lattice switch between its two possible solid state structures (*fcc* and *hcp*) can be performed. In this way, the system can explore both possibilities and so determine the relative free energies of the two states (as well as the absolute free energies). Their results were found to be in agreement with those of thermodynamic integration, but with significantly greater accuracy. The next step will be to extend this technique to the Lennard-Jones solid, which can also occupy more than one solid state form, to help establish a clearer picture of its solid state phase diagram.

6 Bibliography/References

References

- [1] A E Ferdinand & M E Fisher. *Phys. Rev.* **185**, pp 832 (1969).
- [2] K Binder & D W Heermann, *Monte Carlo Simulation in Statistical Physics*, 2nd Edition. Springer-Verlag (1992).
- [3] D Frenkel & B Smit, *Understanding Molecular Simulation*. Academic Press (1996) .
- [4] F Mandl, *Statistical Physics*, 2nd Edition. Wiley (1988).
- [5] G R Smith & A D Bruce, *A study of the multi-canonical Monte Carlo method*. *J. Phys. A.* **28**, pp 6623 (1995).

- [6] Y M Yeomans, *Statistical Physics of Phase Transitions*. Oxford University Press (1992).
- [7] L Onsager. Phys. Rev. **65**, pp 117 (1944).
- [8] A D Bruce, N B Wilding & G J Ackland, *Free energies of crystalline solids: a lattice-switch Monte Carlo method*. Phys. Rev. Lett. **79**, pp 3002 (1997).