Monte-Carlo simulation of small 2D Ising lattice with Metropolis dynamics

Paul Secular*
Imperial College London
(Dated: 6th February 2015)

Results of a Monte-Carlo simulation of the nearest-neighbour Ising model are presented. A 10 by 10 Ising lattice with periodic boundary conditions was sampled using the Metropolis algorithm. Equilibrium properties of the system in zero external magnetic field showed close agreement to both Onsager's analytic solution and previous numerical studies of finite-size systems. The Metropolis dynamics in a non-zero external field were also investigated, revealing metastability and hysteresis phenomena characteristic of permanent magnets.

I. INTRODUCTION

The Ising model (or, more properly, the Lenz–Ising model [1]) is one of the simplest, yet one of the most important, systems in statistical physics [2]. It was introduced in 1920 by Lenz as a model of a ferromagnet and the analytic solution of the one-dimensional case was famously published by Ising in 1925 [1]. For further discussion of the physical motivation of the model and its limitations see [3]. In two dimensions and above, the thermodynamic limit of the model exhibits a second-order phase transition at a critical point (dividing paramagnetic and ferromagnetic states). In 1944 Onsager published the analytic solution of the two-dimensional model in the absence of an external field [4]. Very few other exact results are known, meaning the problem is usually attacked computationally with Monte-Carlo methods.

Although numerical studies of the Ising model are most commonly carried out on large lattices, with the results extrapolated to the infinite, thermodynamic limit, small systems and their finite-size behaviour may also be of some practical interest as they can be seen physically as corresponding to magnetic nanoparticles [5].

In the present work we use the Metropolis algorithm to investigate both the equilibrium properties and the nonequilibrium dynamics of a 10 by 10 square Ising lattice with periodic boundary conditions. The energy of this system is given by:

$$E = -\frac{1}{2}\sum_{i,j} J_{ij}\mathbf{S}_i \cdot \mathbf{S}_j - \mu \sum_k \mathbf{S}_k \cdot \mathbf{B} \tag{1}$$

and its magnetisation by:

$$\mathbf{M} = \frac{1}{N^2}\sum_k \mathbf{S}_k \tag{2}$$

For a homogeneous system, the nearest neighbour interaction energy is constant, so we have:

$$E = -\frac{1}{2}J\sum_{i,j} \mathbf{S}_i \cdot \mathbf{S}_j - \mu N^2 \mathbf{M} \cdot \mathbf{B} \tag{3}$$

If we measure energy in units of J and $\mathbf{B}$ in units of J/μ, we have the dimensionless form:

$$E = -\frac{1}{2}\sum_{i,j} \mathbf{S}_i \cdot \mathbf{S}_j - N^2 \mathbf{M} \cdot \mathbf{B} \tag{4}$$

For convenience, we also define a dimensionless inverse temperature:

$$\beta = \frac{J}{k_B T} \tag{5}$$

II. METHODOLOGY

A. The Metropolis algorithm

The Metropolis algorithm is "the most successful and influential" of all Monte-Carlo methods [6]. Conceptually, it gives the simplest way to correctly sample the Boltzmann distribution, which describes the canonical ensemble and hence the Ising model in the thermodynamic limit. Readers unfamiliar with the algorithm, may wish to review for example [7] or [8].

The Metropolis algorithm can be seen as a random walk through state space forming a Markov chain of sampled states. Thus subsequent samples are not independent but strongly correlated. This means it is important to average over large numbers of iterations in order to sample enough states far away from one another that they are effectively uncorrelated. The problem with this is that the algorithm dramatically slows down near the Ising model's critical point, making it unfeasible to get precise results close to the phase transition, especially for large systems [9]. In other words, it can take an impractical amount of time for the algorithm to reach equilibrium states. Because of this, cluster-flipping (as opposed to single spin-flipping) algorithms, such as the Wolff algorithm, are often employed, which are more efficient but more complex to implement [10].

* secular@physics.org

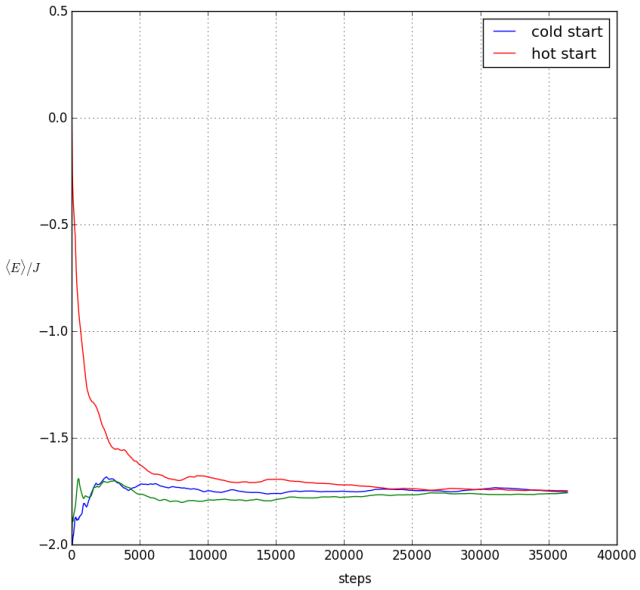

Figure 1. Mean energy per spin versus Monte-Carlo steps during equilibration of 10×10 Ising lattice at $\beta = 0.5$ and $B = 0.0$. Equilibrium was reached after 363 sweeps of the lattice (36300 steps). The green curve is the measurement system (started cold) whilst the blue and red curves are the comparison systems which start in a cold and hot state, respectively. Notice that the three curves converge.

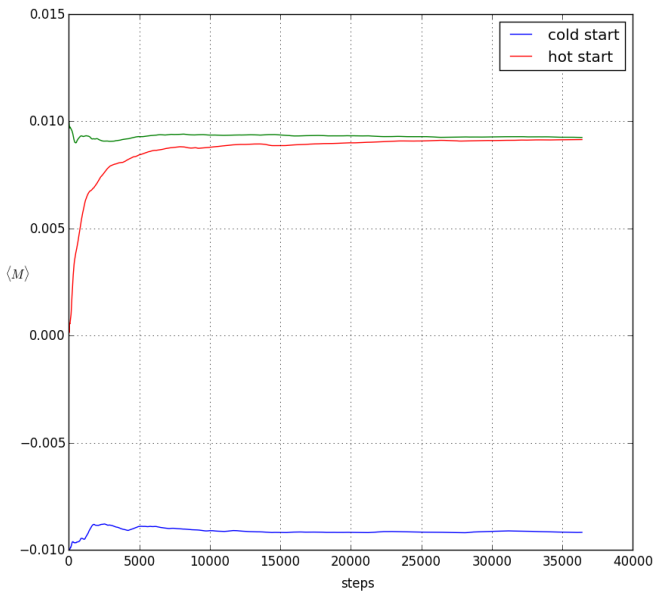

Figure 2. Mean magnetisation versus Monte-Carlo steps during equilibration of 10×10 Ising lattice at $\beta = 0.5$ and $B = 0.0$. Equilibrium was reached after 363 sweeps of the lattice (36300 steps). The green curve is the measurement system (started cold) whilst the blue and red curves are the comparison systems which start in a cold and hot state, respectively. Here the green and red curves converge to the positive magnetisation equilibrium state, whilst the blue curve settles in the negative magnetisation equilibrium state. These two equilibria exist because of the symmetry of the system in zero external field.

B. Pseudo-random number generation

At the heart of all Monte-Carlo methods is a pseudo-random number generator. In the present work, the NumPy numerical library for the Python language was used to generate pseudo-random numbers. NumPy uses the Mersenne Twister algorithm [11]: a popular algorithm first introduced by Matsumoto and Nishimura in [12]. The importance of testing even a good random number generator with a given model is demonstrated by Ferrenberg and Landau in [9]. Their conclusion that the Metropolis algorithm is less sensitive than the Wolff algorithm to correlations introduced by random number generators, justifies its choice in the present work.

C. Equilibration

The simulation code was designed to equilibrate a "measurement system" by performing Monte-Carlo steps until either an equilibration condition is met or the specified maximum number of lattice sweeps is reached. It then carries out a number of measurement sweeps providing the statistics needed to perform estimation of the system's thermodynamic properties. At the start of the simulation, an initial configuration is chosen. Two options are available: a "cold start" in which all spins are aligned in the same direction (chosen randomly) or a "hot start" in which each spin is independently oriented at random. The cold start corresponds to absolute zero and the hot start to an infinite temperature. As the temperature is varied step-by-step, the final equilibrium configuration of the measurement system is persisted and used as the initial state for the following temperature.

In order to determine when equilibrium is reached, two "comparison systems" are used: one initialised with a hot start and the other with a cold start. The code compares running averages of the energy per spin and the magnitude of the magnetisation of all three systems. A comparison of energy alone was found to be inadequate, since the systems could reach metastable states of very similar energy but noticeably different magnetisation. Once the values of the running averages have converged to within a small bound, the systems are considered to have reached equilibrium. The equilibration process is illustrated in figures 1 and 2.

After experimentation and visual inspection, a bound of 0.01 for both energy and magnetisation was found to give the best results. With larger values, the system was sometimes incorrectly classified as equilibrated. On the other hand, in zero external field, smaller values led to the system not being classified as equilibrated at all and left it flipping between the two equilibrium states (positive and negative magnetisation).

D. Performance

The simulation was coded in Python (the language taught to physics undergraduates at Imperial College London), which is an interpreted language. Because of this, the code ran slowly compared to similar codes reported in the recent statistical physics literature and the results achievable were thus restricted. This was an important factor in the decision to focus solely on results for a 10 by 10 lattice. Nevertheless, the code was optimised for speed using standard software engineering practices; in particular, a purely procedural paradigm using lists and arrays was implemented. Memoisation and local variables were also employed. While an object-oriented design would be preferable in terms of program structure and scalability, object creation and invocation add an overhead and thus affect performance.

III. RESULTS

A. Thermodynamic properties

The simulation was used to estimate the mean equilibrium energy and magnetisation of the 10 by 10 Ising lattice. It was found that averaging results over all three systems led to better statistics. With no external field present, systems sometimes fluctuated between the positive and negative magnetisation equilibria. Because of this spontaneous reversal in the magnetisation, its magnitude rather than its value was averaged (since the average of the magnetisation itself would be zero). This had the effect of offsetting the average magnetisation of the paramagnetic phase from zero to a small positive value (figure 4). The specific heat capacity and magnetic susceptibility were also computed by making use of the fluctuation-dissipation theorem which relates them to the variance in the energy and magnetisation respectively [2]. While the measurement sweeps of the lattice were carried out, the sample mean and variance were calculated iteratively using a numerically stable, one-pass method due to Welford [13].

Results for all four thermodynamic properties were computed for values of β between 0 and 1 at three different values of the external field: 0.0, 0.5 and 1.0 (figures 3-7). It was found that applying an external magnetic field has the effect of shifting the curves to the left implying that the phase transition occurs at a smaller value of β. In other words, when an external field is applied, the system doesn't need to be as cold for a phase transition to occur because this bias encourages the spins to more easily align in the same direction as that of the field itself. The peaks of the specific heat capacity and magnetic susceptibility curves also decrease with increasing external field; in the latter case, by one or two orders of magnitude (figures 6 and 7). These findings match those of previously published numerical studies of small Ising lattice systems (e.g. [14] and [8]).

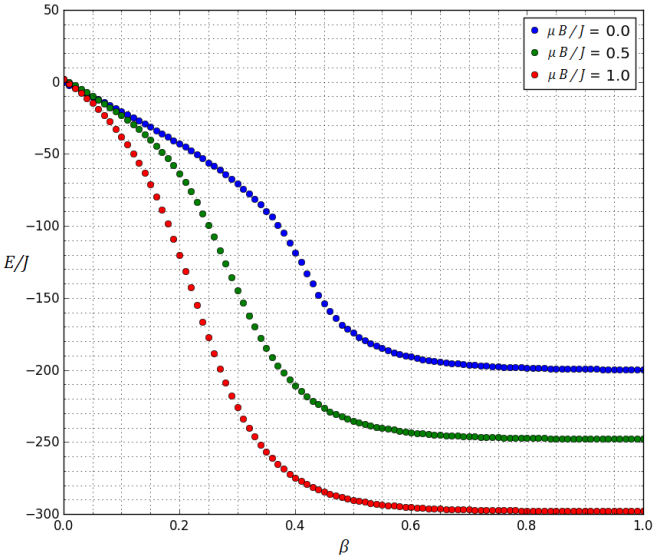

Figure 3. Mean equilibrium energy vs inverse temperature for 10×10 Ising model at three different values of external magnetic field. Data were calculated using 10000 measurement sweeps and the results were averaged over three systems with different initial configurations.

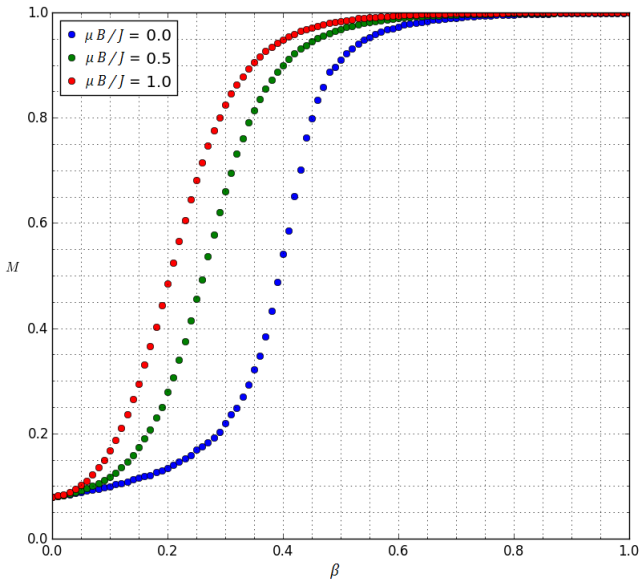

Figure 4. Mean of equilibrium magnetisation's absolute value vs inverse temperature for 10×10 Ising model at three different values of external magnetic field. Data were calculated using 10000 measurement sweeps and the results were averaged over three systems with different initial configurations.

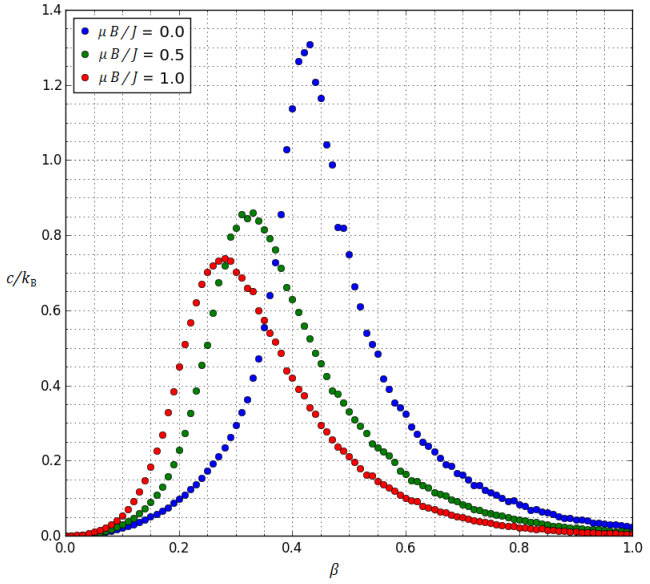

Figure 5. Mean equilibrium specific heat capacity vs inverse temperature for 10×10 Ising model at three different values of external magnetic field. Data were calculated using 10000 measurement sweeps and the results were averaged over three systems with different initial configurations. Notice how the maximum (peak) value decreases with increasing external field.

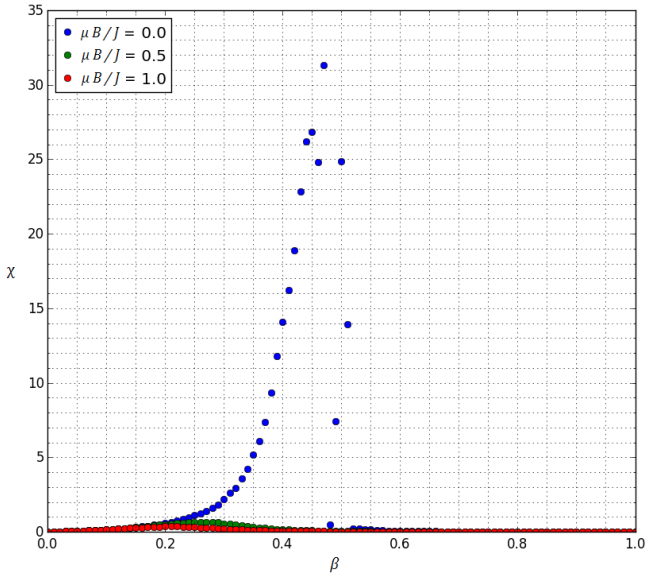

Figure 6. Mean equilibrium magnetic susceptibility vs inverse temperature for 10×10 Ising model at three different values of external magnetic field. Data were calculated using 10000 measurement sweeps and the results were averaged over three systems with different initial configurations. The maximum (peak) value decreases dramatically with increasing external field. At this scale the curve for zero external field is clear but the other two curves are barely visible.

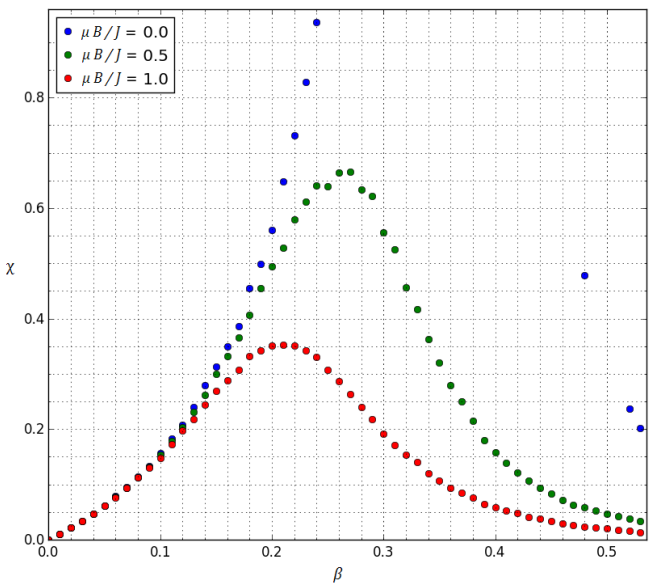

Figure 7. Mean equilibrium magnetic susceptibility vs inverse temperature for 10×10 Ising model at three different values of external magnetic field. Data were calculated using 10000 measurement sweeps and the results were averaged over three systems with different initial configurations. The maximum (peak) value decreases dramatically with increasing external field. At this scale the peak of the curve for zero external field lies far above the graph.

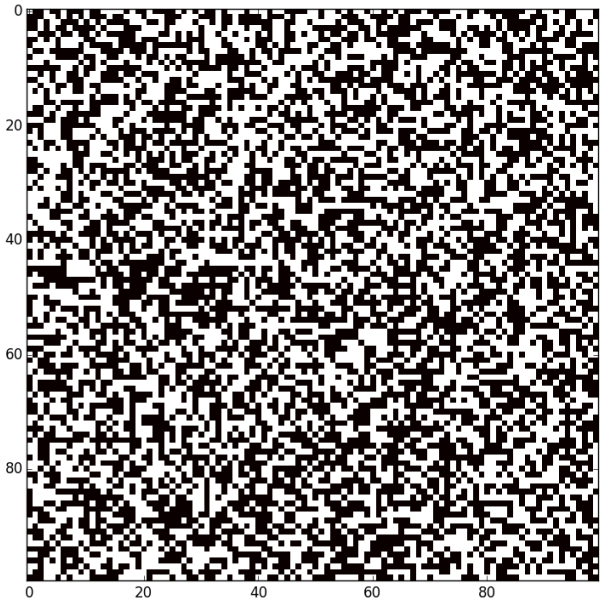

Figure 8. Equilibrium microstate of the 100×100 Ising lattice corresponding to a macrostate with inverse temperature of $\beta = 0$ (i.e. an infinite temperature) in zero external field. The spins appear to be oriented completely randomly with no obvious correlations.

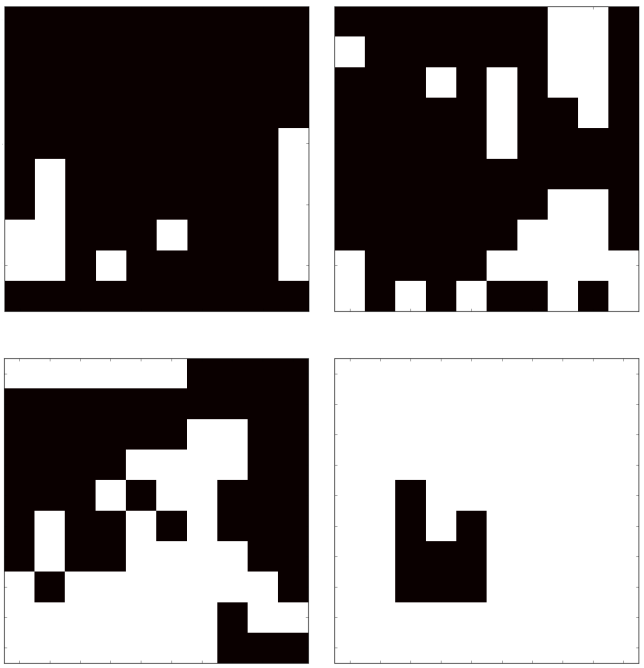

Figure 9. Four random microstates of the 10×10 Ising lattice with periodic boundary conditions corresponding to an equilibrium macrostate with inverse temperature of $\beta = 0.41$ and zero external field. In each of these states, it is clear that two large clusters (wrapped across the boundaries) have formed. This emergent structure is analogous to the domains in a ferromagnet.

To gain further insight into the behaviour of these thermodynamic quantities, the spatial pattern of spins in the lattice corresponding to equilibrium microstates in zero external field were examined. At an inverse temperature of $\beta = 0$, the system is infinitely hot and hence the spins are randomly oriented as expected. This randomness is more obvious in larger lattices (as illustrated in figure 8). When the inverse temperature is close to the critical point, the equilibrium microstates comprise spins distributed primarily in two randomly shaped clusters (figure 9). Finally, when β is increased to approximately 0.8 or above (corresponding to a very cold system), the equilibrium microstates show all spins aligned in the same direction, with just one or two occasional single spin-flip fluctuations.

These spatial patterns give an intuitive picture of the behaviour of the magnetisation, specific heat capacity and susceptibility as the temperature is varied. At the critical point, there is a bifurcation leading to two equally probable equilibrium macrostates (corresponding to positive or negative magnetisation). The system here is stuck between these two possibilities and fluctuates from one to the other. This explains why the susceptibility and heat capacity suddenly peak around this point. These quantities are known to diverge completely in the infinite Ising model, but of course our small, finite-size system doesn't allow that and instead these quantities reach a finite maximum [2]. Another property of the infinite model is that it becomes scale-independent at the critical point, with an infinite correlation length and hence clusters of aligned spins of all possible sizes [15]. In the finite lattice however, the maximum correlation length and cluster size are bounded by the dimensions of the system.

Equilibrium microstates were also examined for an external field of 0.5. Because of the small lattice size, it was difficult to tell if the non-zero magnetic field affected the pattern of spins. It seemed as though there may have been an additional domain or two formed around the (shifted) phase transition, but additional work will be required to clarify this (probably with larger lattice sizes).

B. Metropolis dynamics

When applied to the Ising model, the Metropolis algorithm can actually play two different roles. In the context of equilibrium thermodynamics, as discussed above, it can be used to sample the microstates of the system with the correct probability. However, it can also be seen as a way of supplying the Ising model with dynamics, allowing for the study of nonequilibrium, time-dependent statistics. When the Ising model is treated in this way, it is commonly referred to as the kinetic Ising model. The Metropolis dynamics have no direct physical motivation (and others, such as the closely-related Glauber dynamics [16], are also commonly used) but, importantly, they have been proven to give the correct equilibrium behaviour in the infinite Monte-Carlo time limit [17]. This is not necessarily true of other dynamics [18].

The nonequilibrium response of the Ising lattice at constant temperature to a linearly varying external field was investigated. At low temperatures (i.e. $\beta \gtrsim 0.5$), hysteresis was found to occur. Typical results are shown in figures 10 and 11. Increasing the number of equilibration sweeps reduced the area of the hysteresis loops. When the system was allowed to equilibrate fully at each value of the external field (modelling a quasi-static process) the hysteresis disappeared, as expected. At high temperatures (i.e. $\beta \lesssim 0.4$), no hysteresis was observed. These findings agree with those in the literature, for example [8].

When the system is in a non-zero external field, the symmetry between the two equilibrium states is broken and the state corresponding to a magnetisation in the same direction as the field becomes the single minimum of the system's free energy. However, the other state becomes a highly metastable local equilibrium state coexisting with the global equilibrium (ground state) [19]. The observed hysteresis is due to a first-order phase transition from the metastable state to the true equilibrium occurring at different values of the applied field [5]. Since metastability and hysteresis also occur in permanent magnets [20], these results show that the kinetic Ising model exhibits qualitatively correct behaviour.

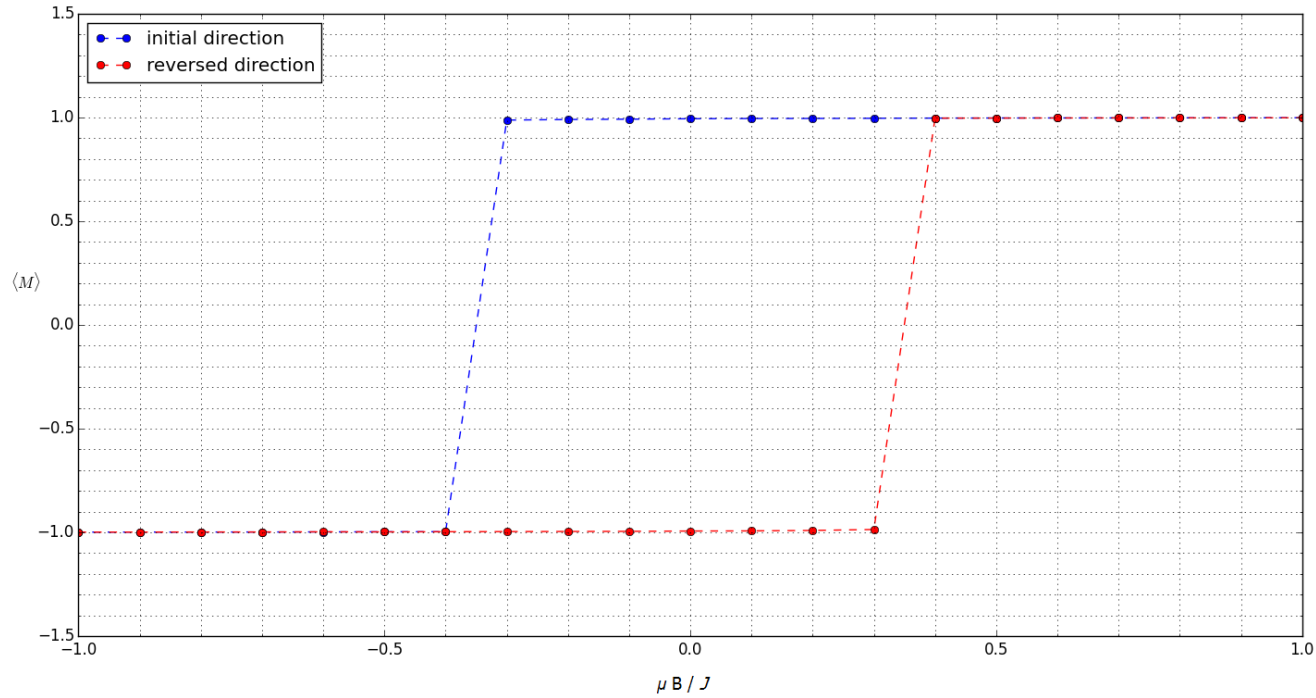

Figure 10. Plot of mean magnetisation (not its absolute value) against external magnetic field with $\beta = 0.75$. Notice the hysteresis loop formed by changing the external field linearly from 1.0 to -1.0 and back again.

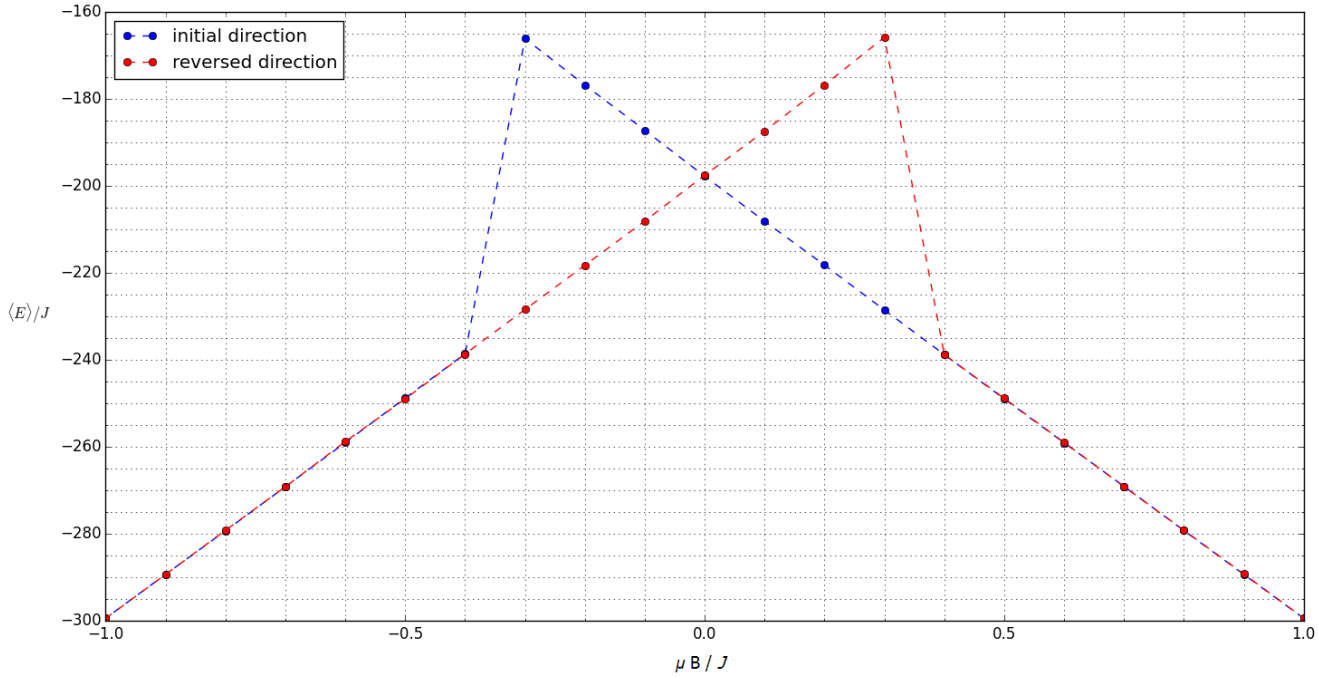

Figure 11. Plot of mean energy against external magnetic field with $\beta = 0.75$. The external field was varied linearly from 1.0 to -1.0 and back again, leading to a hysteresis effect.

IV. CONCLUSION

The Ising model is one of the most studied systems in statistical physics [10] and has even found application in other diverse areas of science [2]. This work has focused on a small 2D square Ising lattice with periodic boundary conditions, but the simulation code was developed in such a way that it can be applied to rectangular systems of any size. The results presented here are in excellent agreement with previous studies, suggesting that the algorithms are robust. The work naturally lends itself to three extensions:

1) The addition of error estimation, since neither statistical nor systematic (e.g. finite sampling time) errors were considered here in any detail. This would need to take account of the correlation time of the Metropolis sampling. The simplest approach would be to use the bootstrap or jackknife methods [10].

2) Investigation of the dependence of the results on lattice size and dimensionality (a 3D lattice would be more physical but computationally expensive).

3) Investigation of the dependence of the results on boundary conditions. It would be productive to try the more realistic, free edge boundary conditions in order to better model a magnetic nanoparticle as has been done previously by Cirillo and Lebowitz [5].

ACKNOWLEDGMENTS

The author would like to acknowledge helpful discussions with Ms. Catherine Smith and Prof. Kim Christensen.

[1] S. Brush, Reviews of Modern Physics **39**, 883 (1967).
[2] K. Christensen and N. R. Moloney, *Complexity and criticality* (Imperial College Press, London, 2005).
[3] A. M. Guenault, *Statistical physics* (Chapman & Hall, London, 1995), 2nd ed.
[4] L. Onsager, Physical Review **65**, 117 (1944).
[5] E. N. M. Cirillo and J. L. Lebowitz, Journal of Statistical Physics **90**, 211 (1998).
[6] I. Beichl and F. Sullivan, Computing in Science Engineering **2**, 65 (2000).
[7] K. Binder and D. P. Landau, *A guide to Monte Carlo simulations in statistical physics* (Cambridge University Press, Cambridge, 2005), 2nd ed.
[8] N. J. Giordano and H. Nakanishi, *Computational physics* (Pearson Prentice Hall, Upper Saddle River, NJ, 2006), 2nd ed.
[9] A. Ferrenberg, D. Landau, and Y. Wong, Physical Review Letters **69**, 3382 (1992).
[10] M. E. J. Newman and G. T. Barkema, *Monte Carlo methods in statistical physics* (Clarendon Press ; New York, Oxford, 1999).
[11] *numpy.random.RandomState - NumPy v1.9 manual* (2014), accessed: 21/12/2014, URL `http://docs.scipy.org/doc/numpy/reference/generated/numpy.random.RandomState.html`.
[12] M. Matsumoto and T. Nishimura, ACM Trans. Model. Comput. Simul. **8**, 3 (1998).
[13] B. P. Welford, Technometrics **4**, 419 (1962).
[14] D. Landau, Physical Review B **13**, 2997 (1976).
[15] J. P. Sethna, *Statistical mechanics : entropy, order parameters and complexity* (Oxford University Press, Oxford, 2006).
[16] R. J. Glauber, Journal of Mathematical Physics **4**, 294 (1963).
[17] F. Martinelli, E. Olivieri, and E. Scoppola, Journal of Statistical Physics **61**, 1105 (1990).
[18] P. L. Krapivsky, S. Redner, and E. Ben-Naim, *A Kinetic View of Statistical Physics* (Cambridge University Press, 2010).
[19] R. J. McCraw and L. S. Schulman, Journal of Statistical Physics **18**, 293 (1978).
[20] K.-H. Muller, G. Fuchs, and O. Gutfleisch, in *High Magnetic Fields: Science and Technology*, edited by F. Herlach and N. Miura (World Scientific, 2006).