Simulations: The Ising Model

Asher Preska Steinberg,* Michael Kosowsky, and Seth Fraden
Physics Department, Brandeis University, Waltham, MA 02453

(Advanced Physics Lab)
(Dated: May 5, 2013)

The goal of this experiment was to create Monte Carlo simulations of the 1D and 2D Ising model. To accomplish this the Metropolis algorithm was implemented in MATLAB. The dependence of magnetization on temperature with and without an external field was calculated, as well as the dependence of the energy, specific heat, and magnetic susceptibility on temperature. The results of the 2D simulation were compared to the Onsager solution.

I. INTRODUCTION

As technology has improved in the 20^{th} and 21^{st} centuries, simulations have become widely used to make predictions about complex systems in a breadth of different fields, ranging from sports and games to science and engineering. For physical scientists, some knowledge of simulations and computing has become essential. Not only can simulations help elucidate processes that are difficult to run experiments on in the lab, but they can save precious experiment time and money. The focus of this project was creating simulations of the Ising model.

A. Properties of Magnetic Materials

This experiment dealt with two types of magnetic materials; *ferromagnets* and *paramagnets*. From quantum mechanics we know that electrons have there own intrinsic angular momentum, or *spin*. Each electron has a *magnetic dipole moment* ($\vec{\mu}$):

$$\vec{\mu} = \gamma \vec{S} \tag{1}$$

. Where γ is the gyromagnetic ratio and $\vec{S}$ is the spin angular momentum of the electron [1]. The magnetic properties these materials exhibit are due to the magnetic moments of electrons. When these materials are placed in a magnetic field ($\vec{B}$), the magnetic field exerts a torque on each of these magnetic dipoles, $\vec{\mu} \times \vec{B}$, which lines up the dipoles parallel to the field. Each pair of electrons in an atom have opposing spins, so when all the electrons in each atom are paired, the effect of the torque is nullified. But a material with atoms that have odd numbers of electrons, or unpaired electrons, is placed in an external magnetic field the magnetic moments of these unpaired electrons are rotated so they align parallel with the field [1, 2]. This alignment is not perfect due to thermal effects. Materials in which the magnetic moments align in the presence of magnetic fields but randomize when that field is taken away are known as *paramagnetic*.

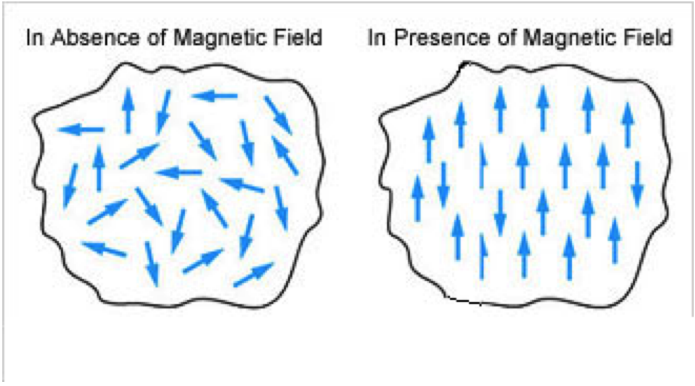

FIG. 1: Illustration of alignment of paramagnetic material in a magnetic field [3].

In *ferromagnetic* materials such as iron (hence the name **ferro**-magnetism), it is favorable due to quantum mechanical effects for the magnetic dipoles to point in the same directions as their neighbors. These materials form what are called *magnetic domains* in which all of the spins in a certain region of the material aligned. These materials have a large number of domains that are oriented in different directions, so on average the magnetic moments of each of these domains cancel out and the net magnetization of the material is zero. At low enough temperatures, magnetic domains can be aligned by an external magnetic field (see figure 3). The external field exerts a torque on each dipole, but since the magnetic dipoles tend to align with their neighbors, they resist. At the domain boundaries there are neighboring spins that are lined up in different directions, and the torque will start to re-orient the spins that are not aligned with the magnetic field so that domains parallel to the field will grow. With a strong enough field, the dipoles in the entire chunk of material will be oriented in the direction of the field [2]. In ferromagnetic materials this process is not reversible, and the material has a non-zero magnetization even when the field is removed [4].

For ferromagnetic substances there exists a temperature known as the *critical temperature* or *Curie point* (T_c) at which the material is no longer ferromagnetic (see figure 2). At low temperatures, the interactions between the spins cause them to align. Above the Curie tem-

* asherps@gmail.com

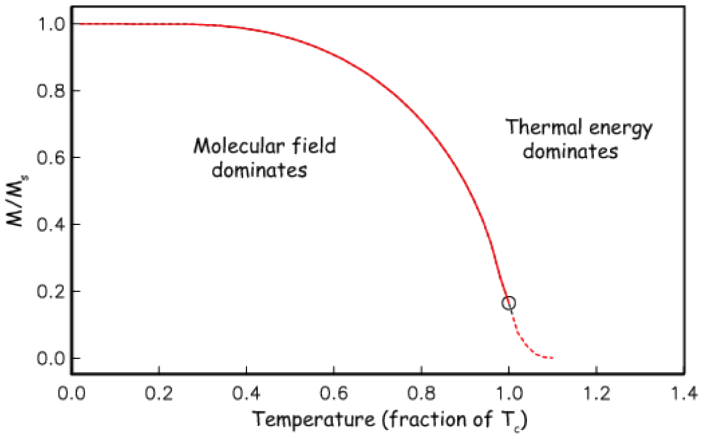

FIG. 2: Example of what the phase transition of a ferromagnetic material at the Curie point looks like. Vertical axis is magnetization (M/M_s), horizontal is temperature (T/T_c) [5].

perature, the material undergoes a *phase transition* and becomes paramagnetic, with all the magnetic moments orienting randomly due to thermal effects.

B. The Ising Model

The simplest system that exhibits a phase transition is the *Ising model*. Though in this report the Ising model will be used to model the phase transition of ferromagnetic materials, this model is broadly applicable. Many papers are published each year applying the Ising model to problems in social behavior, neural networks, and other topics. The model was first proposed by Wilhem Lenz, who gave the problem to his pupil Ernst Ising. Ising solved the model exactly in one dimension, but was disappointed to see that there was no phase transition. About 20 years later, Lars Onsager solved the Ising model exactly in two dimensions in the absence of an external magnetic field. The two dimensional model has a phase transition [4].

In the Ising model, the total energy of the system for a lattice with N spins is given as:

$$E = -J \sum_{i,j=nn(i)}^{N} s_i s_j - H \sum_{i=1}^{N} s_i \tag{2}$$

The first term represents the spin-spin interaction between a spin and its nearest neighbors. In the Ising model each spin (s_i) is either up or down. In this report, the exchange constant (J) is greater than zero and the external magnetic field is in the up direction. The second term in eq. 2 represents the energy of a magnetic dipole in a magnetic field, where μ has been incorporated into H. H will be referred to as the magnetic field [4].

C. Phase Transitions

In the last section it was mentioned that there is no phase transition in the one dimensional Ising model. Why is this? Imagine a simple one dimensional system of seven spins without an external magnetic field. The ground state of the system is when all the spins are aligned and the energy is $E = -7J$. Now imagine flipping some of those spins, creating two domains separated by a domain wall (see figure 4). In this new configuration the energy of the system is $E = -5J$. So the energy cost of creating one domain wall is $2J$. In a one dimensional grid of length N, there are $N-1$ different sites this domain wall can be placed, so the change in entropy associated with creating one domain wall is $\Delta S = -kTln(N-1)$. Therefore the free energy cost of creating one domain wall is:

$$\Delta F = 2J - kTln(N-1) \tag{3}$$

In one dimension, when the temperature is greater than zero and as $N \to \infty$ the creation of a domain wall always lowers the free energy of the system. Since it is favorable to create domain walls when $T > 0$, more and more domains are created until the system is completely random. Therefore in one dimension, the system is paramagnetic at all temperatures and there is no phase transition [4].

In two dimensions, for an $L \times L$ grid the energy cost of creating a domain wall is $2JL$ (see figure 5). In an $L \times L$ grid the domain wall can be placed at L different columns, and the entropy is on the order of lnL. The free energy cost of creating one domain wall in two dimensions is approximately:

$$\Delta F \approx 2JL - kTlnL \tag{4}$$

In this scenario $\Delta F > 0$ as $L \to \infty$. This means it is not favorable to create domain walls in two dimensions, and the spins will remain aligned due to the interactions between spins. As mentioned previously, above the critical temperature thermal effects dominate, leading to the system becoming disordered. There is therefore a phase transition in two dimensions.

D. The Metropolis Algorithm

The expectation value of an observable A is given by the equation:

$$\langle A \rangle = \frac{\sum_i A_i e^{-\beta E_i}}{\sum_i e^{-\beta E_i}} \tag{5}$$

Where A_i is the value of the observable for an individual state (i) [7]. Using equation 5 to calculate observables is not the most computationally efficient method. For example, if one had a simple 10×10 grid of spins, there would be 2^{100} states to sum over. Another method is to use the *Metropolis algorithm* which is a type of *Monte*

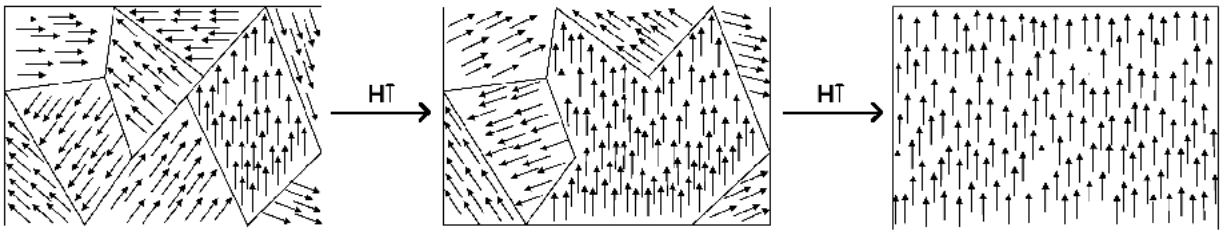

FIG. 3: Alignment of domains when a ferromagnetic material is exposed to an external magnetic field. On the left is an example of unaligned domains. As the magnetic field, H, is turned on, domains aligned in the direction of the field grow. The spins will remain aligned even after the magnetic field is removed [6].

FIG. 4: The creation of a domain wall in one dimension. Domain wall is indicated in red [4].

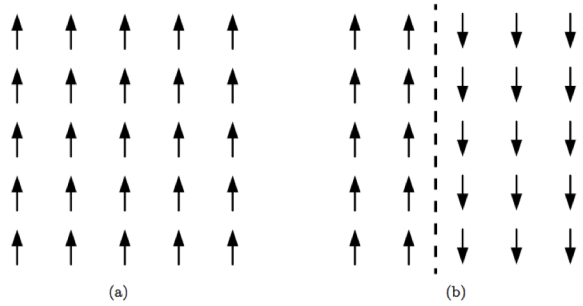

FIG. 5: The creation of a domain wall in two dimensions. (a) shows the system with all spins aligned, (b) shows the system after the creation of one domain wall. Domain wall is indicated by dashed line [4].

Carlo simulation. Monte Carlo simulations are a class of algorithms that utilize repeated random sampling to find results [8].

The Metropolis algorithm works as follows. First, a lattice or grid of spins is created. Next, a spin in the grid is chosen at random and the change in energy (ΔE) due to flipping the spin is calculated based on its interactions with its nearest neighbors (the two nearest neighbors in 1D, four nearest in 2D) and the magnetic field (see equation 2). If $\Delta E < 0$ the spin is always flipped. If $\Delta E > 0$ then the spin is flipped with a probability (p) determined by the Boltzmann factor, $p = e^{-\beta \Delta E}$. Following this, a very large number of other random spins are sampled, then the system is sampled even more to account for the time it takes the system to equilibrate [4].

To get $\langle E \rangle$ and $\langle E^2 \rangle$ of the system the energies of each spin were added up and divided by the total number of spins. To obtain the average magnetization ,$\langle M \rangle$, we set $\mu = 1$ and simply summed all the spins (which had values of $s_i = \pm 1$) then divided by the total number of spins. Using these quantities, the specific heat (C_v) and magnetic susceptibility (χ) in a constant magnetic field can be obtained in terms of the variance of the energy and magnetization [4, 7, 9, 10]:

$$C_v = \frac{\partial \langle E \rangle}{\partial T}$$

$$= -\frac{\beta}{T} \frac{\partial \langle E \rangle}{\partial \beta}$$

$$= \frac{\beta}{T} \frac{\partial^2 \ln Z}{\partial \beta^2}$$

$$= \frac{\beta}{T} \frac{\partial}{\partial \beta} \left(\frac{1}{Z} \frac{\partial Z}{\partial \beta} \right)$$

$$= \frac{\beta}{T} \left(\frac{1}{Z} \frac{\partial^2 Z}{\partial \beta^2} - \frac{1}{Z^2} \left(\frac{\partial Z}{\partial \beta} \right)^2 \right)$$

$$C_v = \frac{\beta}{T} (\langle E^2 \rangle - \langle E \rangle^2) \tag{6}$$

Similarly for magnetic susceptibility:

$$\chi = \frac{\partial \langle M \rangle}{\partial H} = \beta (\langle M^2 \rangle - \langle M \rangle^2) \tag{7}$$

II. EXPERIMENTAL

A. Basics of Code

The code that we used to carry out these simulations in two-dimensions is shown in the supplemental materials section. The programming language that was chosen for this experiment was MATLAB. Natural units of $J/k = 1$

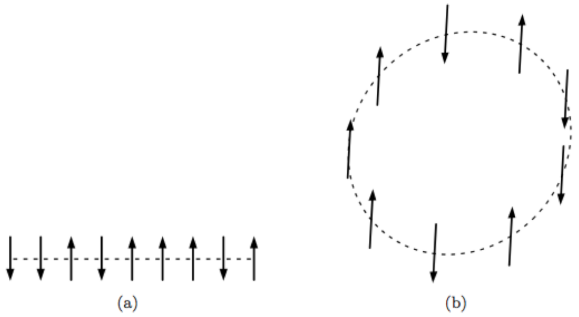

FIG. 6: Examples of boundary conditions. **(a)** "Free ends". The spins on either end are missing one neighbor. **(b)** "Periodic" or "torroidal" boundary conditions. Now the one dimensional array in (a) has been set up so the spin on the left end of (a) is neighbors with the spin on the right end [4].

were used. Two types of boundary conditions were considered for this project. The first was *free ends*, in which the spins on the edges of the lattice are missing a neighbor (see figure 6). The second was *periodic* or *torroidal* boundary conditions, in which spins on one edge of the lattice are not only neighbors with the spins directly next to them, but also with the spins that make up the other edge of the lattice. For this experiment, periodic boundary conditions were chosen. Simulations were run in both one dimension and two dimensions. In the experiments, two different initial conditions were used. The first was starting the simulations where all the spins were aligned randomly. The second, was where all the spins were initially aligned, as if a magnetic field had been used to align all the spins before the start of the experiment.

In the code, MATLAB would randomly chose a point in the lattice and determine the change in energy of flipping a spin as $\Delta E = -2E$ since $s = \pm 1$. As described in section I D, if $\Delta E < 0$, the spin was always flipped. If $\Delta E > 0$, MATLAB would compare the probability determined by the boltzmann factor (p) to a random number between 0 and 1. If the random number was less than p the spin was flipped. Generally, the grid would be randomly sampled 50 $\times$ grid size times, after which the magnetization was measured and the grid would be sampled another 50 $\times$ grid size times and the magnetization was measured again. If the difference between the two magnetization values was less than 5%, the system was considered equilibrated. If not, the system was sampled until the variations in magnetization were under 5%.

III. RESULTS

A. One Dimensional Ising Model

The results for the one dimensional Ising model with $H = 0$ are shown in figures 7-10. In all of the trials shown, the initial state of the system at each temperature was the aligned state (all spins were aligned in same direction). For the plot of the average energy against temperature in figure 7, the basic trend was as expected. As the temperature increased, the energy of the system slowly began to increase. This was because as the temperature was increased, thermal effects become more and more important and the system wants to go to the most entropically favorable state. So even though though the energy due to the spin-spin interactions is increasing, it is still more favorable to the system to go to the most disordered state.

Looking at the plot of the average magnetization against temperature (see figure 8) the results are basically what one would expect, with a few exceptions. At temperatures from around $T = 1 - 10$ the magnetization of the system is very close to zero. This makes sense because the 1D Ising model is not supposed to have a phase transition, and the material is supposed to be paramagnetic (see argument in section I C). At temperatures lower than $T \approx 1$ though, the magnetization is $M \approx 1$, so there appears to be a phase transition somewhere around here. Since it is well known from Ernst Ising's solution that there should be no phase transition in the 1D Ising model, what is causing this?

It turns out that at very low temperatures the Metropolis algorithm fails to give a very physical representation of the system. The way the Metropolis algorithm determines whether or not to flip a spin when $\Delta E > 0$ is by comparing a random number to the Boltzmann factor. But as $T \to 0$ and we have:

$$\lim_{T\to 0} p = \lim_{T\to 0} e^{-\Delta E/kT} = 0 \tag{8}$$

So at temperatures close to zero, the Metropolis algorithm will only flip spins so they align, but will never flip them out of alignment. This is why for very low temperatures the magnetization is always $\langle M \rangle = 1$ and also for extremely low temperatures the energy is always $E = -1$, which is the energy of the system when all the spins are aligned (see figure 7). The same problem appears when the Metropolis algorithm is used to calculate the magnetization and energy against temperature in the 2D case, but it is less apparent because all the spins are supposed to remain aligned at low temperatures in the 2D case.

In figure 11 the magnetization vs temperature was compared when $H = 0$ and $H = 1$. The results were as expected. With $H = 0$, for most of the temperatures (excluding very low temperatures, as explained) the magnetization is close to 0, since more and more domain walls are being created until the system is completely random. With $H = 1$, the spins remain aligned for a larger range of temperatures. This is because alignment with the external magnetic field lowers the energy of the spins, and thermal effects do not make the system completely random until the temperature is really increased.

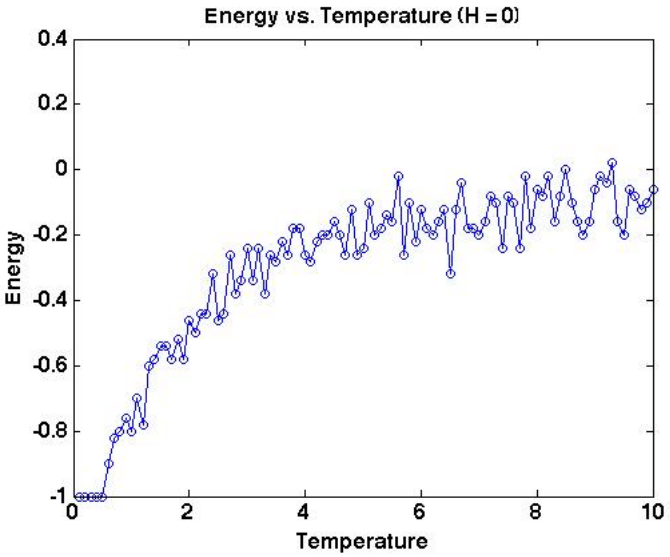

FIG. 7: Energy vs temperature for 1D Ising model with $H = 0$. The grid length was 200 spins, and $H = 0$. At each temperature the system started with all the spins aligned.

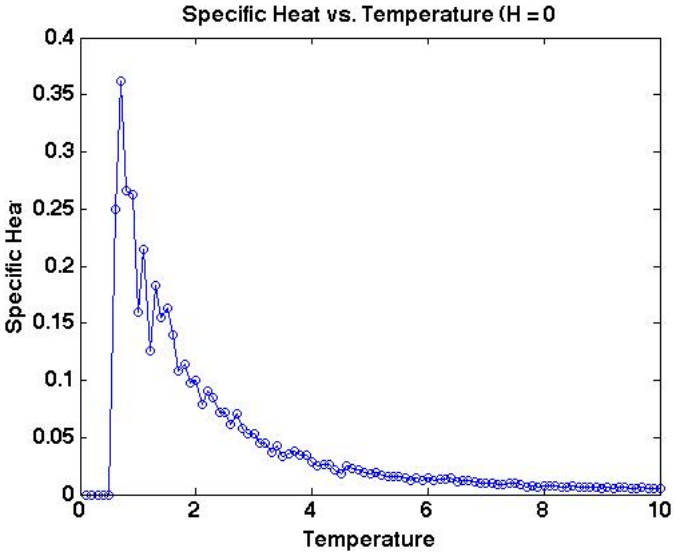

FIG. 9: Specific heat vs temperature for 1D Ising model with $H = 0$. The grid length was 200 spins, and $H = 0$. At each temperature the system started with all the spins aligned.

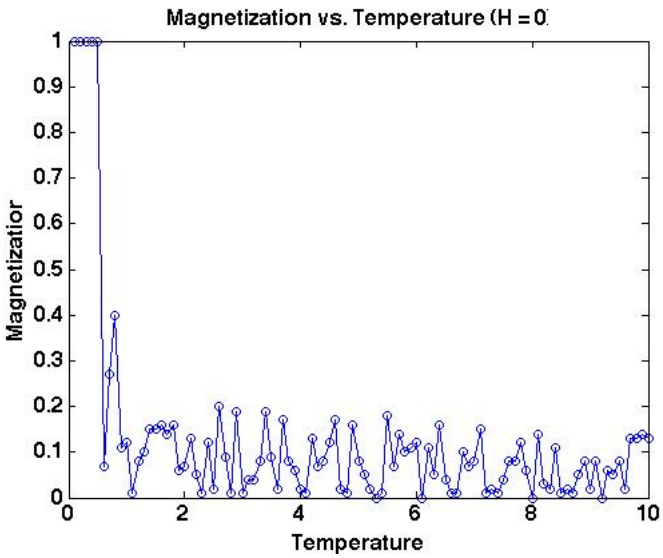

FIG. 8: Magnetization vs temperature for 1D Ising model with $H = 0$. The grid length was 200 spins, and $H = 0$. At each temperature the system started with all the spins aligned.

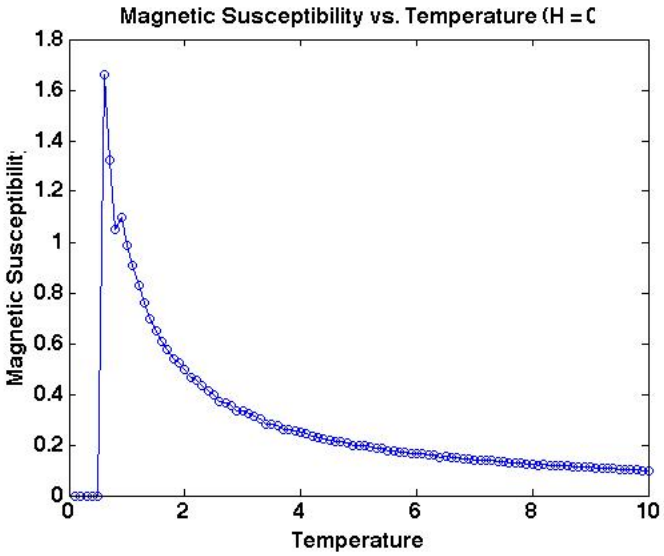

FIG. 10: Magnetic Susceptibility vs temperature for 1D Ising model with $H = 0$. The grid length was 200 spins, and $H = 0$. At each temperature the system started with all the spins aligned.

B. Two Dimensional Ising Model

The results for the two dimensional Ising model with $H = 0$ are shown in figures 12-15. The Onsager solution gives the critical temperature (T_c) as [4]:

$$\frac{KT_c}{J} = \frac{2}{ln(1+\sqrt{2})} \approx 2.269 \tag{9}$$

In our results there is a phase transition from ferromagnetic to paramagnetic somewhere around this temperature. In figure 12, the plot of energy vs temperature, we see that the energy due to the spin-spin interactions is at a minimum since all the spins are aligned, but somewhere around $T = 2-3$ there is a sharp rise in energy as the material switches from ferromagnetic to paramagnetic and the spins are oriented randomly, raising the energy due to the spin-spin interactions. In the plot of magnetization

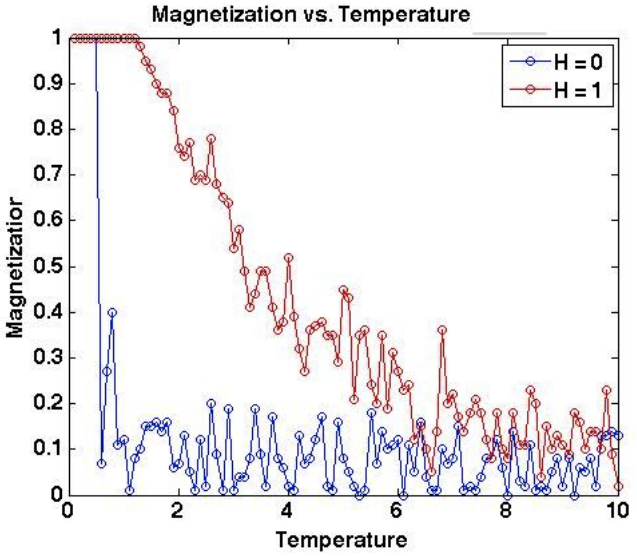

FIG. 11: Magnetization vs temperature for 1D Ising model with $H = 0$ and $H = 1$. The grid length was 200 spins. At each temperature the system started with all the spins aligned.

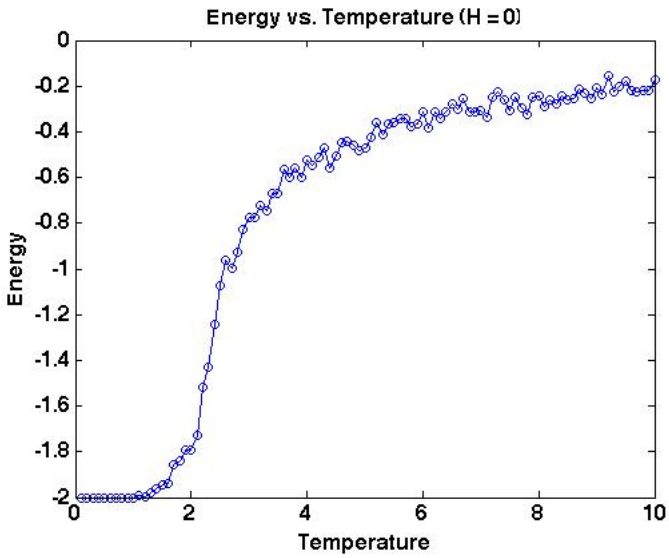

FIG. 12: Energy vs temperature for 2D Ising model with $H = 0$. The grid size was 50×50, and $H = 0$. At each temperature the system started with all the spins aligned.

against temperature (fig. 13) a phase transition is also observed. At temperatures below $T = 2-3$ the spins are aligned and $M \approx 1$. For temperatures above $T = 2-3$ the spins are randomly oriented and $M \approx 0$. The plots of specific heat and magnetic susceptibility against temperature (figures 14 and 15 respectively) show divergent behavior at the critical temperature as expected. The Onsager solution for magnetization is given as [11]:

$$M = \begin{cases} \left(1 - \left[\sinh \log\left(1+\sqrt{2}\right)\frac{T_c}{T}\right]^{-4}\right)^{1/8} & T < T_c \\ 0 & T > T_c \end{cases}$$

To compare our results to the Onsager solution, the results for magnetization from our simulation were plotted along with the above equation in figure 16.

In figure 17 the magnetization vs temperature was compared when $H = 0$ and 1. With $H = 1$, instead of seeing a sharp transition at the critical temperature as compared to $H = 0$, the magnetization instead gradually decrease to $\langle M \rangle = 0$. This is because like in the one dimensional case, the energy of each spin is lowered when they align with the external magnetic field, so the order of the system is increased and there is not a sharp transition from an aligned state to the random orientation of spins. As the temperature is increased to $T = 8-10$, thermal effects dominate, and the net magnetization is close to zero.

To further examine the effects of applying an external magnetic field to the system, a trial was run where the temperature of the system was fixed at $T = 3.5$ while the magnetic field was increased from $H = 0$ to 10 (figure 18). Without any external magnetic field, the spins are all randomly oriented at $T = 3.5$ (see figure 13) because $T > T_c$ and thermal effects have randomized the system. In figure 18, as the magnitude of the external magnetic field is increased, the system goes from a state where the spins are oriented randomly to a state where the system is completely aligned with the magnetic field. This is because with a strong enough external magnetic field, the aligning torque exerted by the magnetic field (discussed in section I A) is so strong that it overcomes thermal effects and lines up the majority of the magnetic dipoles in the direction of the magnetic field.

C. Equilibration of System

As mentioned in section II, the code had a mechanism built in that would keep on sampling until the system was considered equilibrated at the temperature, or when the magnetization was changing less than 5%. It was important to sample the system sufficiently, because as the system was sampled more and more, each individual spin was more likely to be in its most probable microstate. As more and more spins went to their most probable microstate, the overall system went to its most probable macrostate or the state the system will be in once its equilibrated. As the system approaches this state, the fluctuations in the expectation values of thermodynamic quantities begins to decrease, and the decrease in these fluctuations was used as an indicator that the system had reached equilibrium. An example of how the code checked if the fluctuations had decreased is shown in figure 19.

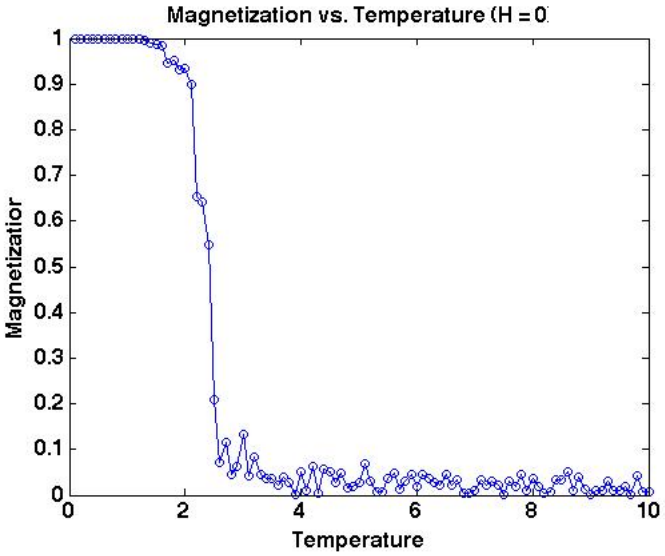

FIG. 13: Magnetization vs temperature for 2D Ising model with $H = 0$. The grid size was 50×50, and $H = 0$. At each temperature the system started with all the spins aligned.

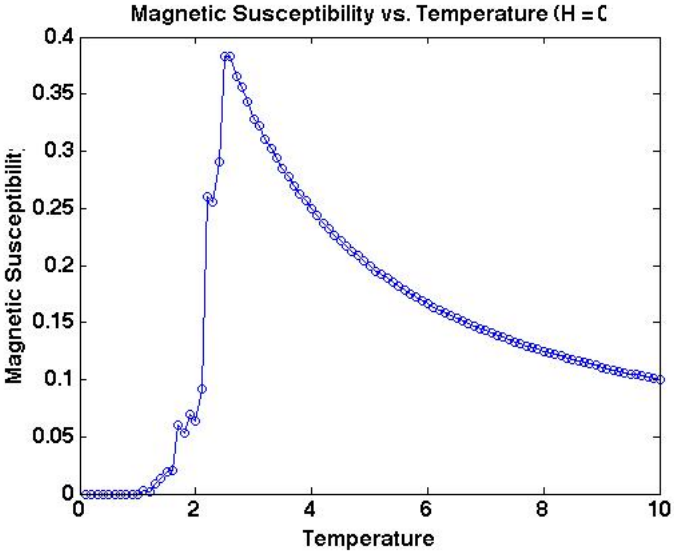

FIG. 15: Magnetic Susceptibility vs temperature for 2D Ising model with $H = 0$. The grid size was 50×50, and $H = 0$. At each temperature the system started with all the spins aligned.

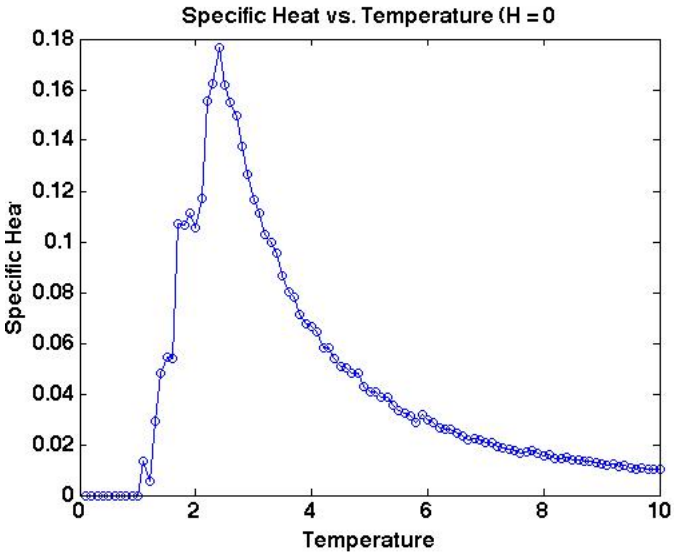

FIG. 14: Specific heat vs temperature for 2D Ising model with $H = 0$. The grid size was 50×50, and $H = 0$. At each temperature the system started with all the spins aligned.

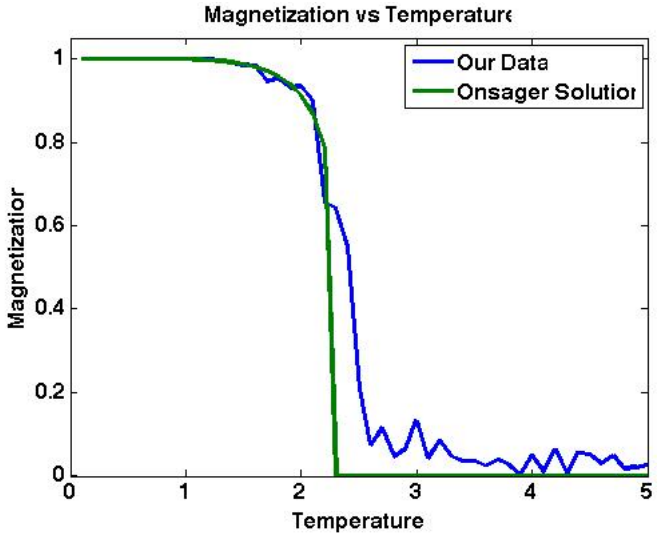

FIG. 16: Magnetization vs temperature for 2D Ising model with $H = 0$. The Onsager solution is plotted in green, the results of our simulation in blue. For our simulation, the grid size was 50×50, and $H = 0$. At each temperature the system started with all the spins aligned.

D. System Size

One of the other constraints of these simulations was the system size, or the number of spins that were in the system. In a real system there are 6.022×10^{23} atoms per mol of material, each of these atoms having electrons with magnetic moments. Unfortunately we simply do not have the computational power to make estimates on a system of that scale. When using random sampling to determine the values of different thermodynamic quantities, the smaller the system size, the more the system will be susceptible to random noise. As the system size is increased, the random noise affects the system less and less, and only large changes (not noise) will really have an effect on the system. To try to determine the effects of changing our system size, a trial was run in which the magnetization and energy were measured at $T = 6$ for

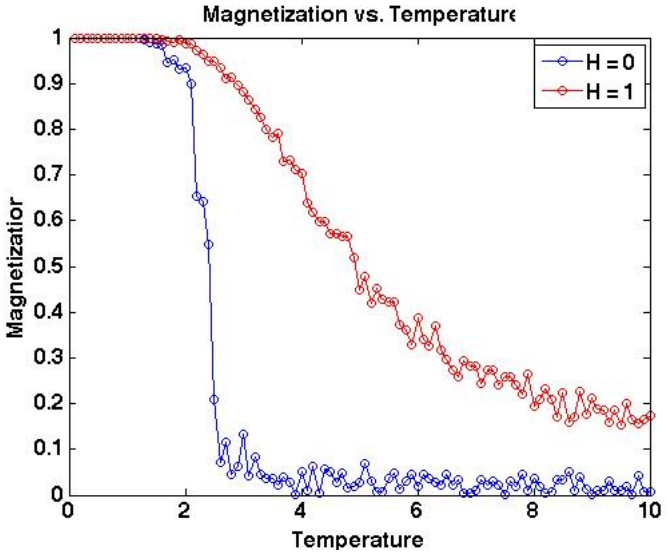

FIG. 17: Magnetization vs temperature for 2D Ising model with $H = 0$ and 1. The grid size was 50×50. At each temperature the system started with all the spins aligned.

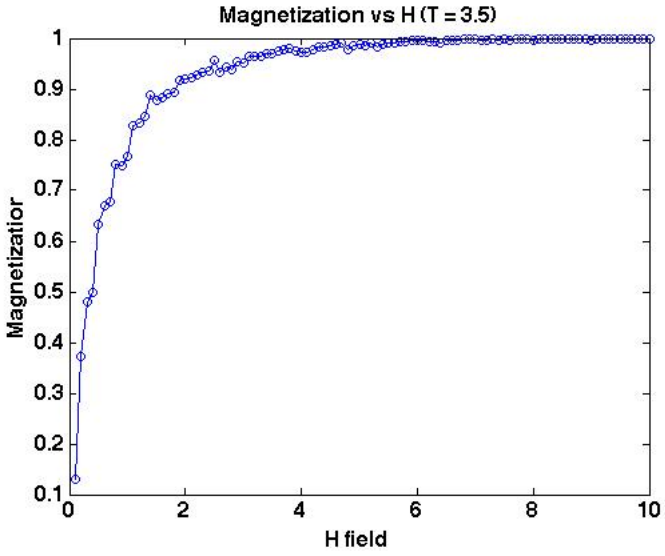

FIG. 18: Magnetization vs magnetic field for two dimensional Ising model with $T = 3.5$.The grid size was 50×50, and $H = 0$. The system started will all the spins aligned.

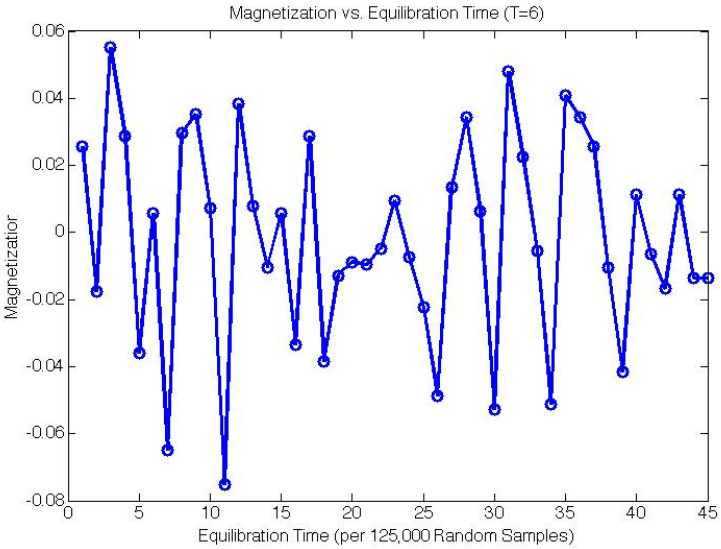

FIG. 19: Equilibration of a 50×50 grid of spins. The code checked if the system was equilibrated every 125,000 samples by determining if the magnetization of the system had changed by less than 5%. It can be seen that after the 40^{th} check the systems fluctuations are getting smaller and then are finally less than 5%.

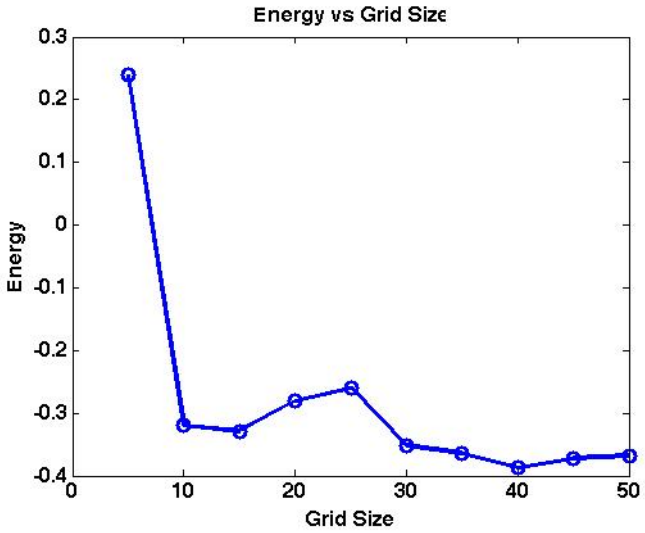

FIG. 20: Energy at $T = 6$ calculated for systems of different sizes. The sizes ranged from 5×5 to 50×50. The system was started with all spins aligned.

systems of different sizes (see figures 20 and 21). In both figures the fluctuations in the measured values of observables appear to decrease some as the size of the grid is increased. Since the values of energy and magnetization are still varying some once the size of the grid that was used in this experiment is reached, it means that at this size the results are still subject to some randomness.

IV. IMPROVEMENTS

There were many improvements that could have been made to this experiment. The first lesson that was quickly learned was to run simulations on a small system to attempt to debug the code before running simulations on very large grids. At the suggestion of Professor Fraden, some of our trials were visualized in MATLAB.

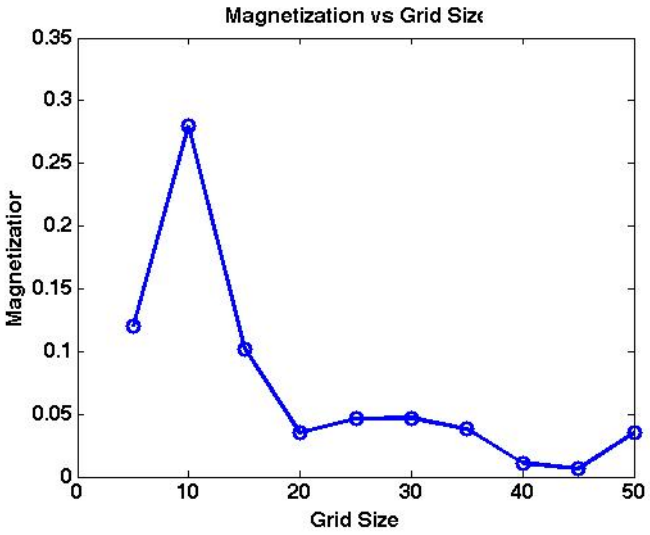

FIG. 21: Magnetization at $T = 6$ calculated for systems of different sizes. The sizes ranged from 5×5 to 50×50. The system was started with all spins aligned.

This was extremely useful in trying to see if the code was creating a good model of the system. Realizing these things sooner could have cut down on the time that was spent trying to understand if the metropolis algorithm was being implemented correctly in the code.

If total simulation time was not a concern another two aspects of the experiment that could have been optimized are the way that the code checks if the system has equilibrated and the size of the system. Instead of checking the system's magnetization every 125,000 samples (for a 50×50 grid) and comparing the magnetization at the current checkpoint to the previous one and seeing if the difference was less than 5%, the deviations in the magnetization between multiple checkpoints could have been compared. This would have allowed us to see if the fluctuations were really decreasing, or it was simply random chance that the magnetization between two checkpoints changed very little. Increasing the size of the system that the simulation was run on would have led to better agreement with the Onsager solution, as explained in section III D. Both of these improvements would have meant an increase in simulation time for each run.

To improve the time it took to simulate the system, a different programming language could have been chosen for the project. A group of graduate students achieved much faster run times using *Perl*.

V. CONCLUSION

In this experiment Monte Carlo simulations were created of the 1D and 2D Ising model. To do this the Metropolis algorithm was implemented in MATLAB. The dependence of the energy, magnetization, specific heat, and magnetic susceptibility of the system on temperature were calculated for the 1D and 2D Ising model. The results of the simulation for magnetization against temperature with and without an external magnetic field were compared. Some simple tests were done to try to understand the effects of increasing the amount of sampling and the system size in Monte Carlo simulations. The experiment was a success in that we learned some of the basics of how to run simulations and achieved results that were comparable to the Onsager solution.

VI. CODE

Please see the attached file, 2DIsingModel.m, for a sample code that was used in this project.

ACKNOWLEDGMENTS

I would like to thank my lab partner, Michael Kosowsky, for all his help and insight. We would both like to thank Andreas Rauch, Sathish Akella and especially Professor Seth Fraden for all their extremely useful help and advice, as well as their patience.

[1] D. J. Griffiths, *Introduction to Quantum Mechanics* (Pearson Education, Inc., 2005).
[2] D. J. Griffiths, *Introduction to Electrodynamics* (Pearson Education, Inc., 2008).
[3] Magnetism, http://electrical.blogspot.com/2012/06/types-of-magnetism.html (2009).
[4] H. Gould and J. Tobochnik, *Thermal and Statistical Physics* (Princeton University Press, 2009).
[5] UCSD, http://magician.ucsd.edu/Essentials/WebBook76x.png (2013).
[6] Wikipedia, https://en.wikipedia.org/wiki/Magnetic_domain (2013).
[7] K. A. Dill and S. Bromberg, *Molecular Driving Forces: Statistical Thermodynamics in Biology, Chemistry, Physics, and Nanoscience* (Garland Science, Taylor & Francis Group, LLC, 2011).
[8] Wikipedia, http://en.wikipedia.org/wiki/Monte_Carlo_method (2013).
[9] K. Zengel, http://fraden.brandeis.edu/courses/phys39/simulations/Simulation.html (2013).
[10] L. Larrimore, http://fraden.brandeis.edu/courses/phys39/simulations/Simulation.html (2013).
[11] Wikipedia, http://en.wikipedia.org/wiki/Ising_model (2013).