

Efficient Monte-Carlo Sampling of Ising Models

Lars Bogner

<https://github.uio.no/larsbog/FYS4150>

(Dated: November 19, 2025)

This work presents a numerical study of the two-dimensional Ising model using a custom implementation of Markov Chain Monte Carlo sampling with single-spin Metropolis updates. The model is simulated on square lattices with periodic boundary conditions, and the statistical properties of the system, i.e., energy, magnetization, heat capacity, and susceptibility, are extracted across a range of temperatures. Multiple independent Markov chains are executed in parallel to accelerate parameter sweeps, enabling efficient exploration of temperature-dependent behavior. Validation against analytical results for the $L = 2$ lattice demonstrates excellent agreement for energy and magnetization, confirming the correctness of the implementation. Detailed equilibration studies highlight the role of ergodicity and illustrate the contrasting dynamical behavior below and above the critical temperature. The simulations reproduce the characteristic signatures of the continuous phase transition at T_C , including the emergence of large correlated domains and the divergence of response functions. The methods and performance characteristics of the implementation provide a robust foundation for further investigations of critical phenomena in lattice spin systems.

CONTENTS

I. Introduction	1
II. Methods	1
A. Ising Model	1
1. Critical Temperature	2
2. Ergodic Property	3
B. Markov Chain Monte Carlo	3
C. Parallelization	3
D. Implementation Details	3
1. Tools and Usage of AI	4
III. Results and discussion	4
A. Validation against Analytical Results	4
B. Equilibration of the System	5
C. Phase Transition Behavior	5
1. Limit of Infinite Lattice Size	6
D. Performance Analysis	7
E. Statistical Properties of the Ising Model	7
IV. Conclusion	7
A. Analytical expectation values for the 2x2 lattice	8

cific heat, as functions of temperature. In order to efficiently sample the Ising model, we approximate the system state using Markov Chain Monte Carlo (MCMC) methods with rejection sampling. These are a class of algorithms that sample complex probability distributions by constructing a Markov chain with the desired distribution as its equilibrium distribution. In the context of the Ising model, these methods enable us to generate spin configurations representative of the system's behavior at a given temperature. Parallel computation can significantly speed up the sampling process, enabling the concurrent evaluation of multiple systems. This is particularly useful for large numbers of computations under different conditions, such as at various temperatures.

In [section II](#), the methods employed to implement MCMC sampling of the Ising model are delineated, including the specific algorithms employed and the parallelization techniques utilized. In [section III](#) the results of the aforementioned simulations are presented, including measurements of key statistical properties of the Ising model. The implications of these findings are discussed as well. Finally, in [section IV](#) we offer a summary of our work and put forward several suggestions for future research directions.

I. INTRODUCTION

The Ising model is a mathematical representation of ferromagnetism in statistical mechanics. This model consists of discrete variables, known as spins, which can take on one of two values, typically represented as +1 (up) or -1 (down). These spins are arranged on a lattice and interact with their nearest neighbors. The model is used to study phase transitions, such as the transition from a magnetized to a non-magnetized state, and has applications in physics, biology and the social sciences.

It allows us to measure the statistical properties of the system, such as magnetization, susceptibility and spe-

II. METHODS

A. Ising Model

In a two-dimensional system of size $L \times L$, each lattice site i contains a spin variable s_i that can take on values of either +1 (up) or -1 (down). The energy of a given configuration of spins is described by the sum over nearest-neighbor interactions:

$$E = -J \sum_{\langle i,j \rangle} s_i s_j. \quad (1)$$

The coupling constant, J , is the parameter that quan-

tifies the strength of the interaction between neighboring spins. The sum is taken over all pairs of nearest neighbors, defined as $\langle i, j \rangle$. In order to facilitate the process of summation without the occurrence of double counting, the total energy can be expressed in the following manner:

$$E = -\frac{J}{2} \sum_i \sum_{j \in \text{NN}(i)} s_i s_j, \quad (2)$$

where $\text{NN}(i)$ represents the set of nearest neighbors of site i . Depending on the configuration of the neighboring spins, the flipping of a single spin, s_k , will result in a change in the energy of the system by an amount

$$\Delta E = \Delta s_k \cdot J \sum_{j \in \text{NN}(k)} s_j, \quad (3)$$

where $\Delta s_k = s_k^{\text{new}} - s_k^{\text{old}} \in \{-2, 2\}$ is the change in the spin variable at site k . To substantiate this claim, it is essential to recognize that the energy contribution from site k prior to the flip is given by:

$$E_{\text{old}} = -J s_k^{\text{old}} \sum_{j \in \text{NN}(k)} s_j. \quad (4)$$

The discrepancy in the factor 2 in comparison to [Equation 2](#) emerges due to the fact that s_k is a constituent of the sum over nearest neighbors for each of its neighbors, resulting in double counting. Subsequent to the flipping of the spin at site k , the contribution to the energy from that site is updated, and after the flip, it becomes

$$E_{\text{new}} = -J s_k^{\text{new}} \sum_{j \in \text{NN}(k)} s_j. \quad (5)$$

The alteration in energy resulting from the flip is then

$$\Delta E = E_{\text{new}} - E_{\text{old}} = -J(s_k^{\text{new}} - s_k^{\text{old}}) \sum_{j \in \text{NN}(k)} s_j. \quad (6)$$

This expression demonstrates how the energy change depends exclusively on the local configuration of spins surrounding site k .

The total magnetization of the system M is derived from the summation of all spin variables

$$M = \sum_i s_i. \quad (7)$$

Both extensive quantities, defined as those that scale with the system size, can be normalized per spin to obtain intensive quantities

$$\epsilon = \frac{E}{N}, \quad m = \frac{M}{N}, \quad (8)$$

where $N = L^2$ is the total number of spins in the lattice.

The Ising model displays a specific heat capacity, C_V , and magnetic susceptibility, χ , both of which are defined as follows, once more normalized per spin:

$$\frac{C_V}{N} = \frac{N \left(\langle \epsilon^2 \rangle - \langle \epsilon \rangle^2 \right)}{k_B T^2}, \quad (9)$$

$$\frac{\chi}{N} = \frac{N \left(\langle m^2 \rangle - \langle m \rangle^2 \right)}{k_B T}, \quad (10)$$

where k_B is the Boltzmann constant and T is the temperature of the system. The notation $\langle \cdot \rangle$ denotes the expectation value, which can be estimated through averaging over sampled configurations. The average over a sufficiently large number of Monte Carlo samples will converge to the true expectation value. Given that expectation values are computed from sampled configurations, it is imperative to ensure that the system's state space is sampled correctly. In the Ising model, it is assumed that the configurations follow the Boltzmann distribution. This distribution dictates the probability of a configuration with energy E at temperature T as follows:

$$P(E) = \frac{e^{-E/(k_B T)}}{Z}, \quad (11)$$

with the partition function Z defined as

$$Z = \sum_{\text{all states}} e^{-E/(k_B T)}. \quad (12)$$

1. Critical Temperature

In two dimensions, the Ising model undergoes a continuous phase transition at the critical temperature T_C separating an ordered low-temperature phase from a disordered high-temperature phase. In the context of the model, spontaneous magnetization is exhibited when the temperature falls below the critical temperature, T_C . This phenomenon occurs because, even in the absence of an external magnetic field, the magnetization m remains non-zero. As the temperature approaches T_C from below, the magnetization decreases continuously, eventually vanishing according to a characteristic power law. Concurrently, fluctuations in energy and magnetization exhibit a marked increase, as evidenced by the divergent behavior of both the heat capacity [Equation 9](#), and the susceptibility [Equation 10](#). Precisely at T_C , the system becomes scale-invariant, and configurations manifest correlations that extend across the entire lattice. In numerical simulations, such phenomena manifest as strong critical slowing down, necessitating sufficiently long sampling times to resolve the sharply peaked response functions. The exact solution for the square lattice with nearest-neighbor interactions yields the following result: $k_B T_C / J = 2 / \ln(1 + \sqrt{2}) \approx 2.269$, providing a precise reference point for validating numerical results.

2. Ergodic Property

The implementation of Monte Carlo methods within the Ising model is contingent upon the assumption of ergodicity, which stipulates that the Markov chain generated by successive spin updates must possess the capacity to attain all physically relevant configurations within a sufficient time frame. In practice, this implies that single-spin flip dynamics must form a connected set over the configuration space, thereby avoiding traps in metastable regions. The Metropolis-Hastings update rule, which is employed for sampling purposes, ensures detailed balance by ensuring that the stationary distribution of the Markov chain corresponds to the Boltzmann distribution. Ergodicity guarantees the convergence of long-time averages of observables, such as ϵ and m , to ensemble averages. As the critical point is approached, ergodicity remains valid; however, the time required to decorrelate successive configurations grows significantly due to the emergence of large correlated domains. This phenomenon, often referred to as critical slowing down, necessitates meticulous care in determining autocorrelation times to obtain statistically independent samples and reliable estimates of thermodynamic quantities. When an ergodic Markov chain is considered, the equivalence between ensemble averages and long-time averages is guaranteed in the limit of a sufficient number of Monte Carlo steps. For any observable A , the ensemble expectation value $\langle A \rangle$ obtained from the Boltzmann distribution coincides with the time-averaged estimate $\hat{A}$ extracted from the simulation. It is imperative to note that the aforementioned equality is only valid once the chain has thoroughly explored configuration space. This underscores the significance of ergodicity and adequate sampling length as essential prerequisites for reliable numerical measurements. However, it does allow for the utilization of an adequately sampled time average in place of the expectation value. The time-averaged estimate is defined as follows:

$$\hat{A}(t) = \frac{1}{t} \int_0^t dt' A(t'). \quad (13)$$

B. Markov Chain Monte Carlo

To ensure the proper sampling of the Boltzmann distribution and to facilitate efficient exploration of the configuration space, we employ Markov Chain Monte Carlo (MCMC) methods. MCMC methods generate a sequence of configurations (states) where each configuration depends only on the previous one, forming a Markov chain. The fundamental objective is to establish a transition probability between states that satisfies detailed balance, thereby ensuring the system's convergence to the desired equilibrium distribution. In each Monte Carlo cycle, the objective is to flip each spin in the lattice precisely once. The acceptance of a proposed spin flip is determined us-

ing the Metropolis-Hastings algorithm, which involves calculating the energy change, denoted by ΔE , associated with the flip. In the event that $\Delta E \leq 0$ the flip is accepted without condition, as it results in a lower energy state. In the event that $\Delta E > 0$, the flip is accepted with a probability that is determined by the Boltzmann factors

$$P_{\text{accept}} = e^{-\Delta E/(k_B T)}. \quad (14)$$

This acceptance criterion ensures that the system can explore higher energy states, allowing it to escape local minima and sample the configuration space more effectively. The acceptance of the flip is contingent upon the random number r being drawn uniformly from the interval $[0, 1]$. More specifically, the flip is accepted if $r < P_{\text{accept}}$.

C. Parallelization

Parallelization of MCMC methods is challenging due to the sequential nature of Markov chains. However, it is possible to exploit the independence of multiple simulations at different temperatures or initial conditions to achieve parallelism. The implementation of multiple independent Markov chains operating concurrently, with each chain simulating the Ising model at an distinct temperature, has been demonstrated to result in a substantial acceleration of the overall sampling process. Each chain functions independently, facilitating straightforward parallelization across multiple processors. This approach is particularly beneficial when studying temperature-dependent properties of the Ising model, as it enables simultaneous exploration of the system's behavior across a range of temperatures.

D. Implementation Details

The simulations were executed using a custom C++ implementation of the two-dimensional Ising model with periodic boundary conditions. Each simulation instance constructs an $L \times L$ lattice with spins $s_{ij} \in \{\pm 1\}$. The lattice is initialized in one of two ways: either in a fully ordered configuration or with randomly assigned spins drawn from a uniform distribution. The generation of random numbers is achieved via a Mersenne Twister engine (`std::mt19937`), with each run being seeded individually. To ensure that periodic boundary conditions are met, the neighboring indices are computed using modulo arithmetic

$$\begin{aligned} \text{NN}(i, j) = \{ & ((i + 1) \bmod L, j), \\ & ((i - 1 + L) \bmod L, j), \\ & (i, (j + 1) \bmod L), \\ & (i, (j - 1 + L) \bmod L) \}. \end{aligned} \quad (15)$$

The spin dynamics follow the standard single-spin Metropolis algorithm. A Monte Carlo step involves the

selection of a lattice site uniformly at random, the computation of the local energy change ΔE associated with the flipping of that spin, and the acceptance of the flip with probability P_{accept} . In order to minimize computational expenditure, all conceivable Boltzmann factors (equation (14)) associated with the four-neighbor interaction are initially calculated and subsequently stored in a lookup table.

During the simulation, the instantaneous energy and absolute magnetization are evaluated after each attempted update. The values and their squares are then accumulated to compute ensemble averages and variances. Subsequently, the code derives the heat capacity and magnetic susceptibility from these moments, as outlined in equations (9) and (10). Optional time-resolved data-energy, magnetization, and other intermediate quantities, are stored at each Monte Carlo cycle when evolution recording is enabled.

Furthermore, to support large parameter sweeps, the code incorporates a scheduling layer that manages multiple independent simulations. The configuration of each run is delineated by a specific lattice size, a designated temperature, the number of Monte Carlo cycles employed, the initialization mode, and the RNG seed. A master seed generator ensures statistical independence across runs. The scheduler distributes the runs across a number of CPU threads specified by the user, executing them concurrently. Upon completion of a given run, an output file is generated, containing the averaged observables. When the relevant option is enabled, an additional file is also generated, which contains the full time evolution of the system.

1. Tools and Usage of AI

For concurrent execution, the implementation utilizes the C++ library `OpenMP`. The executables were compiled using `g++` (GCC) 15.2.1 20251022 (Red Hat 15.2.1-3) with optimization flag `-O3` enabled. All executions were carried out on an AMD Ryzen 7 PRO 4750U Processor, equipped with 16 GB of RAM, operating on the Fedora Linux 42 distribution. A large impact of IO operations on performance is not anticipated, given the utilization of an NVMe SSD for all read and write operations.

For the purposes of analysis, the Python 3.13 programming language was utilized in conjunction with the libraries `NumPy`, `Pandas`, `SciPy`, and `Matplotlib`. In the course of composing this report, the large language model `ChatGPT` was employed to proofread various sections and to propose enhancements in terms of phrasing and grammar. All figures were created manually, without the use of AI assistance. During the implementation of the code, the large language model provided by `GitHub Copilot` was used to improve coding speed, but not for generating large code blocks. A comprehensive review and testing of the code was conducted to ensure its correctness.

Table I. Comparison of ensemble properties with results obtained from analytical calculations in table III. Ensemble properties sampled from 10 simulations with $L = 2$, 10000 Monte Carlo cycles and random initialization.

Quantity	Ensemble Average	Relative Error
$\langle \epsilon \rangle$	-1.9958 ± 0.0014	$(0.0071 \pm 0.0692) \%$
$\langle m \rangle$	0.9987 ± 0.0005	$(0.0004 \pm 0.0463) \%$
$\langle \frac{C_V}{N} \rangle$	0.033 ± 0.011	$(3.4959 \pm 34.3238) \%$
$\langle \frac{\chi}{N} \rangle$	0.0038 ± 0.0014	$(4.2046 \pm 35.9493) \%$

III. RESULTS AND DISCUSSION

A. Validation against Analytical Results

In appendix A of the appendix, a summary of the expected behavior of a correctly sampled lattice of $L = 2$ is provided, in terms of its expectation values. Subsequent to the generation of said results, a comparison is formulated against a set of 10 simulations, which were executed for a duration of 10000 Monte Carlo cycles at a temperature of $k_B T = 1.0$. To assess the uncertainty associated with each ensemble mean, the standard deviation of the ensemble values is employed. The uncertainty in the relative error is derived through the application of Gaussian error propagation to the ensemble variances. The results are presented in table I and demonstrate a high degree of congruence with the anticipated analytical expectation values for the energy and magnetization of the system. Both quantities manifest relative deviations in the order of $10^{-3} \%$ to $10^{-4} \%$. The comparison indicates discrepancies of 3.4959% for the specific heat capacity of the system, accompanied by a substantial uncertainty range of $> 30 \%$. In a similar vein, the relative error of the susceptibility demonstrates a large deviation of $(4.2046 \pm 35.9493) \%$. The significant uncertainty hinders the rise of any tension between the theoretical prediction and the simulated quantity. Consequently, the observed alignment is deemed to be highly satisfactory in accordance with theoretical expectations, thereby validating the implementation of the model.

A comparatively extended simulation time frame was selected for the lattice size. However, a convergence analysis reveals that a minimum of 1000 steps are necessary to achieve results that are equivalent in quality. In order to execute this analysis, the deviation between time averages $\hat{\epsilon}$, $\hat{m}$ and theoretical predictions $\langle \epsilon \rangle$, $\langle |m| \rangle$ as a function of the Monte Carlo cycle for each of the simulations. The results are displayed in figure 1. Following a brief period of equilibration, the presence of artifacts becomes evident within the range of 10^{-2} to 10^{-3} , with an upper boundary delineating the maximum deviation. This observation signifies the stability of the system within its equilibrium state.




Figure 1. Evolution of simulated 2x2 lattice structures over the duration of 10 000 Monte Carlo cycles. Comparison of ensemble averages $\hat{\epsilon}$, $|\hat{m}|$ and theoretical predictions $\langle \epsilon \rangle$, $\langle |m| \rangle$. We use the values from [table III](#) as the prediction values.

B. Equilibration of the System

In order to provide further validation of the working principle of our implementation, it is possible to evaluate the equilibration/thermalization process. The system evolution is recorded by initializing the system from either a fully ordered state, where all the spins are parallel, or a fully disordered state, where the spin directions are chosen from an entirely uncorrelated uniform distribution. At a constant temperature, a specific degree of ordering is anticipated to arise. Below the critical temperature of the system, the energetic advantage of ordering is sufficient to cause nearly complete sorting of the system. At temperatures above the critical temperature, a disruption in the microscopic ordering of spins becomes possible, resulting in the loss of magnetization at the macroscopic level of the system. The assessment of disruptions to the system's optimal order necessitates the utilization of the normalized energy of the system. An optimal ordering of the system will result in minimal energy, while any disruption to this ordered structure will result in an increase in energy. A steady normalized energy indicates successful equilibration, i.e., the energy is in a steady state relative to the system's temperature. In order to illustrate the discrepancies occurring on very small time scales in comparison to the behavior observed on the long time scale, the following quantities are graphed: ϵ is the current normalized energy of the system at a specific time step and the averaged normalized energy $\hat{\epsilon}$:

$$\hat{\epsilon}(t) = \frac{1}{t} \int_0^t dt' \epsilon(t'). \quad (16)$$

For a system of $L = 20$, a total of 2000 Monte Carlo cycles are simulated and the normalized energies ϵ , $\hat{\epsilon}$ are recorded. This procedure is repeated for two temperatures: $k_B T = 1.0$ and $k_B T = 2.4$ and the two extremes




Figure 2. Evolution of the normalized current energy ϵ (slightly colored line) and the normalized time-averaged energy $\hat{\epsilon}$ (fully colored lines) for systems of lattice size $L = 20$.

of ordering, i.e., fully ordered and fully randomized initial configurations.

The results of this study are shown in [figure 2](#). The observed behavior is consistent with the anticipated outcomes. Specifically, it has been determined that the system approaches a state of equilibrium energy for both temperatures after approximately 1000 Markov Chain Monte Carlo cycles. When the fully unordered state is employed as the initial condition, a higher-than-equilibrium energy state is achieved, accompanied by a rapid decay towards the thermalized state. The temperature of $k_B T = 1$ results in complete ordering of the system, irrespective of the initial configuration. For lower temperatures, fewer stochastic deviations from equilibrium are observed, indicating that the system remains predominantly ordered even on short time scales. For the higher temperature point that was sampled, there are noticeable deviations from the time-averaged energy $\hat{\epsilon}$. On the long-time-scale average, the system nevertheless approaches an equilibrium state with respect to the energy. In summary, the model's behavior aligns with the theoretical framework.

C. Phase Transition Behavior

In order to observe the transition between the two phases of the two-dimensional Ising model, a simulation was conducted in which randomly initialized systems were observed at different temperature points for 50 000 MCMC cycles. A range of temperatures, from $k_B T = 1.0$ to 2.4 is scanned at an offset between simulations of $k_B \Delta T = 1 \times 10^{-3}$. In order to accelerate the simulation, no output is recorded during the simulation, and the systems parameters are extracted solely at the end of the simulation. This procedure is carried out for four lattice sizes, namely for $L \in \{40, 60, 80, 100\}$. The resulting system parameters are then plotted as a




Figure 3. Calculated normalized energy $\langle \epsilon \rangle$, magnetization $\langle |m| \rangle$, specific heat capacity $\langle \frac{C_V}{N} \rangle$ and susceptibility $\langle \frac{\chi}{N} \rangle$ as a function of system temperature for different lattice sizes. Systems were evolved for 50×10^3 MCMC cycles.

function of the set temperature and the lattice size. A plot containing the observables can be found in [figure 3](#). The average energy of the system increases linearly with the temperature, irrespective of the system's size. For temperatures above the critical temperature, the system exhibits minimal spread, with larger systems demonstrating increased energies. In the context of absolute magnetization, a decline is observed above the critical temperature threshold. This phenomenon can be attributed to the disruption of the global magnetization, which results in the manifestation of averaging effects within the lattice. The magnitude of the drop is more pronounced with increasing lattice sizes, as the averaging effect is statistically more significant. For temperatures below the critical temperature, a decrease in average absolute magnetization with an increase in temperature is observed, as the added dynamic in the system reduces perfect ordering. The specific heat capacity and the susceptibility demonstrate a diverging tendency around the critical temperature. A comparative analysis of the temperature-dependent specific heat capacity reveals a uniform trend across all systems, with the exception of a shift in the point of divergence. This variation in the critical temperature is attributed to the influence of the lattice size on the system's behavior. As the lattice size increases, the degree of divergence concomitantly rises.

1. Limit of Infinite Lattice Size

The diverging behavior is used in a second step to approximate the point of the critical temperature for each lattice size. For both the specific heat capacity and the




Figure 4. Fit to the specific heat capacity $\langle \frac{C_V}{N} \rangle$ and susceptibility $\langle \frac{\chi}{N} \rangle$ as a function of system temperature for different lattice sizes. Fit to [equation \(17\)](#).

susceptibility data a fit to the function

$$f(T) = Ae^{-\frac{(T-\mu)^2}{2\sigma^2}} + C \quad (17)$$

is performed. As μ denotes the maximum point of the Gaussian function, it is employed as an estimator for the critical temperature. The results of all fits are summarized in [figure 4](#). The fit is performed using the function `optimize.curve_fit` from the `scipy` package, with the minimization of squared residuals. Subsequently, the values obtained for the critical temperature in dependence of the lattice size are employed to approximate the critical temperature of an infinitely sized lattice. To elaborate, the critical temperature is plotted against the lattice size, and a secondary least squares minimization is performed on the relation.

$$T_C(L) = T_C(\infty) + aL^{-1}. \quad (18)$$

It is evident that two distinct values are obtained for the critical temperature of infinite lattices. Specifically, the first value is determined to be 2.284 and is derived from the data points extracted from the susceptibility. The second value, 2.271, is derived from the measurements on the specific heat capacity. The fit to obtain those results is shown in [figure 5](#). The data points obtained via the specific heat capacity demonstrate a high degree of agreement with the anticipated result, $k_B T_C = 2.269$. The discrepancy between the measurements obtained from the susceptibility suggests the necessity for a more sophisticated function to ascertain the critical temperatures from the simulations, i.e., improving [equation \(17\)](#). In order to achieve more precise measurements, it is necessary to increase the number of samples taken in the region of the critical temperature. Due to the scarcity of computing resources, this step has not been implemented in the present study.




Figure 5. Fit of the values obtained from figure 4 to the relation of the system size dependent critical temperature.

D. Performance Analysis

The performance of the simulation was evaluated through benchmarking in both a single-threaded and a 12-thread multithreaded configuration using the system time utility. For each of the executions, a total of 500 identical systems of lattice size $L = 100$ were simulated for 1000 Monte Carlo cycles. Large numbers of identical systems are deliberately selected to minimize the impact of a subset of threads that complete execution before others. This process enhances the overall run time without exhausting all available resources. In order to facilitate the uninterrupted execution of background processes, a configuration of 12 threads was selected on a system with a total of 16 threads. The single-threaded execution required a wall time of 217.634s, while the multithreaded run completed in 29.580 s. This corresponds to a speedup of

$$S = \frac{217.634}{29.580} \approx 7.36, \quad (19)$$

which is substantial, though below the ideal linear scaling expected for perfectly parallel work on 12 threads. The resulting parallel efficiency is therefore $\mathcal{E} = S/12 \approx 0.61$.

The reported CPU times provide additional insight into the system's performance. The single-threaded version accumulated 215.592s of user time, consistent with full utilization of one core. In contrast, the multithreaded run consumed 339.713s of user time over a wall interval of 29.580s. This corresponds to an average of roughly 11.5 concurrently active hardware threads. This finding suggests that the parallel implementation utilizes the majority of available cores during the majority of the runtime.

One potential explanation for the observed limited parallel efficiency is overhead due to concurrent processing, including task delegation, I/O, and an augmented thermal load on the CPU, which can result in thermal throttling. Given that the CPU utilized for the evaluation is




Figure 6. Normalized counts per discrete energy value of simulations using $L = 20$. Total of 50 simulation of 2000 cycles.

part of a notebook computer, the final point is particularly salient, given the prevalence of inadequate cooling capacity in compact designs. Subsequent studies are required to ascertain the impact of all factors on the reduced parallel efficiency.

E. Statistical Properties of the Ising Model

The evaluation of the probability distribution of the normalized energy is achieved through a comprehensive analysis of the distribution of epsilon for systems of $L = 20$ at the two temperatures previously specified in section III B. For each set of initial conditions, 25 identical simulations were executed, each with a different seed, and the resulting trajectories were collected. The results are then displayed in a normalized histogram, using the counts of each unique energy value present in the dataset. The approximate probability density function is displayed in figure 6.

IV. CONCLUSION

We have shown, that our implementation of the two-dimensional Ising Model without the presence of a magnetic field behaves, as expected from theoretical predictions. Especially showing relative errors in the order of $10^{-3} \%$ to $10^{-4} \%$ in quantities that are linear in the spin quantum numbers. Furthermore the statistical phenomena of phase change at the critical temperature has been proven. In future work it may be possible to improve the accuracy of the phase change behavior given more computing resources. We additionally propose a parallelization approach for calculating the system evolution for multiple systems concurrently. We obtain a parallel efficiency of $\mathcal{E} = 0.61$ with the potential of further speed up under ideal conditions.

Table II. Possible configurations of a 2x2 lattice with periodic boundary conditions. Assuming $|s_i| = 1$ and properties of each configuration.

$\sum_i \delta_{s_i,1}$	Degeneracy	Energy E	Magnetization M
0	1	-8J	-4
1	4	0	-2
2	4	0	0
2	2	8J	0
3	4	0	2
4	1	-8J	4

Appendix A: Analytical expectation values for the 2x2 lattice

In [table II](#), a comprehensive list of all possible configurations of a 2×2 lattice with periodic boundary conditions is provided, along with their respective degeneracies, energies, and magnetizations. It is possible to compute the partition function, Z , by utilizing the aforementioned information

$$Z = 2e^{8J\beta} + 2e^{-8J\beta} + 12, \quad (\text{A1})$$

where $\beta = 1/(k_B T)$. From the partition function, we can derive the expectation values for the energy per spin state as

$$\begin{aligned} \langle \epsilon \rangle &= -\frac{J}{ZN} \sum_i E_i \cdot \text{deg}_i \cdot e^{-\beta E_i} \\ &= \frac{6}{Z} (e^{-8\beta} - e^{8\beta}), \end{aligned} \quad (\text{A2})$$

with $N = 4$. In general an expectation value of the ensemble of a general quantity A can be expressed as

$$\langle A \rangle = \frac{1}{ZN} \sum_i A_i \cdot \text{deg}_i \cdot e^{-\beta E_i}. \quad (\text{A3})$$

The following quantities are furthermore obtained analytically:

$$\langle \epsilon^2 \rangle = \frac{72}{Z} (e^{-8\beta} + e^{8\beta}) \quad (\text{A4})$$

$$\langle |m| \rangle = \frac{2}{Z} (2 + e^{8\beta}) \quad (\text{A5})$$

$$\langle m^2 \rangle = \frac{8}{Z} (1 + e^{8\beta}). \quad (\text{A6})$$

Using those expression it is also straightforward to obtain $\langle \frac{C_V}{N} \rangle$ and $\langle \frac{\chi}{N} \rangle$.

For the numerical evaluation at $k_B T = J = 1$ of all expectation values, `Python` is used. The expectation values shown in [table III](#) are obtained from the numerical analysis.

Table III. Numerical expectation values for the intensive thermodynamic properties of the 2x2 lattice.

Quantity	Numerical Expectation Value
$\langle \epsilon \rangle$	-1.996
$\langle m \rangle$	0.999
$\left\langle \frac{C_V}{N} \right\rangle$	11.984
$\left\langle \frac{\chi}{N} \right\rangle$	2.996