

SYMMETRY BREAKING AND CLOCK MODEL INTERPOLATION IN 2D CLASSICAL
 $O(2)$ SPIN SYSTEMS

By

Leon Hostetler

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Physics—Doctor of Philosophy
Computational Mathematics, Science and Engineering—Dual Major

2023

ABSTRACT

The field of two-dimensional classical spin systems has been studied for many decades using a variety of analytical and numerical methods. This field of study contains many interesting models including the $O(2)$ model, which has an infinite-order Berezinskii-Kosterlitz-Thouless (BKT) transition and the class of q -state clock models which have second-order phase transitions when $q = 2, 3, 4$ but BKT transitions when $q \geq 5$. Motivated by attempts to quantum simulate lattice models with continuous Abelian symmetries using discrete approximations, we study an extended- $O(2)$ model that differs from the ordinary $O(2)$ model by the addition of an explicit symmetry breaking term. Its coupling allows us to smoothly interpolate between the $O(2)$ model (zero coupling) and a q -state clock model (infinite coupling). In the latter case, a q -state clock model can also be defined for non-integer values of q . Thus, such a limit can also be considered as an analytic continuation of an ordinary q -state clock model to non-integer q . In the infinite coupling limit, the extended- $O(2)$ model can be simplified, and so we start by establishing the phase diagram in that case. Using Monte Carlo and tensor methods, we show that for non-integer q , there is a second-order phase transition at low temperature and a crossover at high temperature. Next we establish the phase diagram at finite values of the coupling again using both Monte Carlo and tensor methods. We show that for non-integer q , the second-order phase transition at low temperature and crossover at high temperature persist to finite coupling. For integer $q = 2, 3, 4$, there is a second-order phase transition at infinite coupling (i.e. the clock models). At intermediate coupling, there are second-order phase transitions, but the critical exponents vary with the coupling. At small coupling, the second-order phase transitions may turn into BKT transitions.

ACKNOWLEDGMENTS

My time as a Ph.D. student at MSU has come to a close, and I am grateful that I have had this opportunity to do something that I truly love—to learn about and to study the basic laws of the universe—as my full-time job. The years spent at MSU were at times difficult, but they were also the most fulfilling years of my life so far.

A thesis like this, and the five years of research that led to it, is not the product of a single person. I have many people to thank for helping me. First and foremost, I'd like to thank my supervisor Alexei Bazavov who guided me throughout my Ph.D. education and without whom this research would not have happened. Second, I'd like to thank Yannick Meurice who went above and beyond coauthor and PI. Next, I want to thank the other members of my Ph.D. committee—Dean Lee, Huey-Wen Lin, Mohammad Maghrebi, and Andreas von Manteuffel. In addition to benefiting from their guidance on my research, I have had the privilege of taking and/or auditing courses from many of these exceptional educators.

A deep thanks to the many collaborators and coauthors that I have worked with including many of the above names but also Judah Unmuth-Yockey, Giovanni Pederiva, Andrea Shindler, and Brandon Henke. Without you, I would not be where I am today. A particularly deep thanks to two specific collaborators—Jin Zhang and Ryo Sakai—whose work with tensor methods was especially valuable to my research and plays a prominent role in this thesis.

A special thanks to Johannes H. Weber, Yannis Trimis, Thomas (TC) Chuna, Daniel Hoying, Yash Mandlecha, and Alejandro Salas for all the interesting and thought-provoking discussions. I also want to acknowledge the members of the QuLAT Collaboration for the excellent questions and general discussions.

This research was a project of the QuLAT Collaboration, which studies the foundations of quantum computing for gauge theories and quantum gravity. Monte Carlo results presented in this dissertation were obtained using computational resources and services provided by the Institute for Cyber-Enabled Research (ICER) at Michigan State University (MSU). Tensor renormalization group results were computed on Syracuse University HTC Campus Grid and at the National Energy

Research Scientific Computing Center (NERSC) a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 using NERSC awards HEP-ERCAP0020659 and HEP-ERCAP0023235. The research presented in this dissertation was supported in part by NSF award ACI-1341006 and supported in part by the Quantum Information Science Enabled Discovery (QuantISED) program of the U.S. Department of Energy (DOE) under Awards No. DE-SC0010113 and No. DE-SC0019139. The author was personally supported during the first and final years of his Ph.D. by Michigan State University via a University Distinguished Fellowship (UDF).

Finally, I would like to thank my parents. Although they inhabit a different world, they are curious about my work and always there for support. Last, but certainly not least, a loving thanks to my wife Esther for her support and for her patience with me as I put in long hours on esoteric work.

–Leon Hostetler, Aug. 17, 2023

TABLE OF CONTENTS

LIST OF ABBREVIATIONS	vi
LIST OF SYMBOLS	vii
CHAPTER 1 INTRODUCTION	1
CHAPTER 2 CLASSICAL SPIN SYSTEMS AND CRITICAL PHENOMENA ON THE LATTICE	4
2.1 Ising Model	4
2.2 Potts Models	14
2.3 $O(n)$ Models	21
2.4 Clock Models	28
CHAPTER 3 MCMC SIMULATIONS AND THEIR ANALYSIS	38
3.1 Markov Chain Monte Carlo	38
3.2 Updating Algorithms	40
3.3 Autocorrelations	51
3.4 Jackknife Error Analysis	56
3.5 Reweighting	60
3.6 Curve Fitting	62
3.7 Finite Size Scaling	67
3.8 MCMC in Practice	70
CHAPTER 4 THE EXTENDED- $O(2)$ MODEL	76
4.1 Introduction	76
4.2 The $h_q \rightarrow \infty$ Phase Diagram	80
4.3 The Extended- $O(2)$ Model at Finite- h_q	116
4.4 Future Questions	164
CHAPTER 5 QUANTUM FIELD THEORIES ON THE LATTICE	167
5.1 Introduction	167
5.2 QCD in the Continuum	167
5.3 Lattice Regularization	171
5.4 MCMC	175
5.5 Quantum Simulation of Lattice Quantum Field Theories	183
CHAPTER 6 SUMMARY	191
BIBLIOGRAPHY	194
APPENDIX A ADDITIONAL DETAILS ON CLASSICAL SPIN SYSTEMS	209
APPENDIX B TECHNICAL DETAILS OF SIMULATING THE EXTENDED- $O(2)$ MODEL	218

LIST OF ABBREVIATIONS

BKT	Berezinskii-Kosterlitz-Thouless
BMHA	Biased Metropolis Heatbath Algorithm
CDF	Cumulative distribution function
HMC	Hybrid Monte Carlo
HOTRG	Higher-order tensor renormalization group
LQFT	lattice quantum field theory
LQCD	lattice quantum chromodynamics
MCMC	Markov chain Monte Carlo
MC	Monte Carlo
NISQ	Noisy intermediate scale quantum
PDF	Probability density function
QCD	Quantum chromodynamics
QFT	Quantum field theory
RG	Renormalization group
RNG	Random number generator
SVD	Singular value decomposition
TRG	Tensor renormalization group
XY	I.e. the XY model also called the $O(2)$ model in two dimensions

LIST OF SYMBOLS

$\langle \dots \rangle$	Ensemble average taken over many lattice configurations
α	Specific heat critical exponent
β	Inverse temperature or magnetic critical exponent
β_c	Inverse temperature at critical point
C	Specific heat
χ, χ_M	Magnetic susceptibility
χ^2	Chi-square function or statistic
D	Number of dimensions
δ	Magnetic critical exponent
E	Internal energy
η	Correlation function critical exponent
$F(p)$	Structure factor
γ	Susceptibility critical exponent
H	Hamiltonian
$h, \vec{h}$	External magnetic field
h_q	Symmetry breaking strength in Extended-O(2) model
J, J_i, J_{ij}	Coupling strength
k_B	Boltzmann constant
L	The length of the lattice along one/each dimension
$M, \vec{M}$	Magnetization
N	Number of lattice sites (typically) i.e. the volume of the system
ν	Correlation length critical exponent
$O(n)$	The orthogonal group in dimension n
q	Number of “clock” states in clock models

Q	A goodness-of-fit measure
$S_i, \vec{S}_i$	Spin variable at lattice site i
σ	Standard deviation
T	Temperature
τ_{int}	Integrated autocorrelation time
θ_i, φ_i	Angle of spin vector at site i
U_M	Binder cumulant
ξ	Correlation length
Z	Partition function

CHAPTER 1

INTRODUCTION

Quantum chromodynamics (QCD), a key component of the Standard Model of particle physics, is the theory of the strong interaction that occurs between color-charged particles like quarks and gluons. The expectation value of observables is expressed in terms of infinite-dimensional path integrals. Evaluating these path integrals is not easy. One approach is that of perturbative QCD where the path integrals are expanded in powers of the coupling constant. On the other hand, lattice QCD is an *ab initio* nonperturbative approach where the theory is regularized on a lattice and a Wick rotation is performed so that the path integrals can be interpreted as classical partition functions. This allows us to apply the extensive machinery of classical statistical mechanics to QCD. In particular, equilibrium expectation values can be estimated by Monte Carlo simulation.

Lattice QCD using Markov Chain Monte Carlo (MCMC) methods is our most powerful tool for extracting predictions from quantum chromodynamics (QCD) in the non-perturbative regime, however, there are some serious limitations and challenges. The distribution weight for the gauge fields is proportional to the determinant of a very large matrix resulting from integrating over the fermionic degrees of freedom, which makes even equilibrium calculations difficult. The challenges become even more severe when we want to study non-equilibrium QCD or QCD at non-zero baryon density. In both cases, we run into sign problems which make our current approaches impractical. Quantum simulation may offer a way forward.

At its simplest, the idea of quantum simulation is to investigate quantum field theories by studying the behavior of simpler analogues in a laboratory setting. In a *digital* simulation, the Hamiltonian is mapped to a quantum machine, such as a universal gate-based quantum computer, which is time-evolved stroboscopically in discrete steps. This approach seems promising, however, the circuit depth of these algorithms and the lack of error correction available on current hardware, means this is more of a long-term approach. If we want useful results in the near-term, it seems that analog simulation may be the best path forward. In an *analog* simulation, the Hamiltonian of the field theory is mapped to a simpler quantum system, such as a lattice of cold atoms, which is

allowed to evolve continuously in real-time. Such real-time evolution would completely bypass the sign problem that occurs when we try to simulate real-time physics using the traditional MCMC approach.

Whether one pursues analog or digital quantum simulation, we will need to start with toy models. One example is the Schwinger model, which is quantum electrodynamics (QED) in 1+1 dimensions. The Schwinger model exhibits some features of QCD such as confinement, however, it is much easier to study since it is an Abelian model in 1+1 dimensions. Even simpler is the Abelian-Higgs model (i.e. scalar QED), which in 1+1 dimensions is the Schwinger model with the electron replaced by a complex scalar field. It has been shown that in principle the Abelian-Higgs model can be mapped to a Rydberg ladder [1–4]. In this proposed approach, the Abelian-Higgs Hamiltonian is mapped to an angular momentum Hamiltonian, which in turn is truncated to a finite number (e.g. 5) of angular momentum states, which in principle can be simulated on a Rydberg ladder. In the limit of infinite self-coupling and zero gauge coupling, the Abelian-Higgs model reduces to the classical $O(2)$ spin model.

In this dissertation, we study classical $O(2)$ -like spin systems in two dimensions. There are a number of reasons why these kinds of models are interesting to those working towards the quantum simulation of quantum field theories. The first is that the classical $O(2)$ spin model in two dimensions is a particular limit of the Abelian-Higgs model, which itself is a toy model for QCD. The implementation of such $O(2)$ -like models on an analog simulator may therefore be a first step toward the simulation of more complicated models like the Abelian-Higgs model or the Schwinger model. With the $O(2)$ model, we can add a symmetry-breaking term that allows us to break the $O(2)$ symmetry down to $\mathbb{Z}_q$ in a controlled way. This allows us to study the role of symmetry in such systems, and it allows us to study $\mathbb{Z}_q$ approximations of $U(1)$ or $O(2)$ symmetry. This is relevant for the “field digitization” of gauge theories. If we want to put gauge fields on a quantum simulator, whether analog or digital, we need some way to handle the infinite Hilbert spaces of the gauge fields with only a finite set of qubits. One way to handle this is to approximate the gauge fields, by representing them with a discrete subgroup. This is the idea of field digitization. For example,

for a $U(1)$ gauge theory such as Abelian-Higgs or the Schwinger model, one could approximate the gauge group by $\mathbb{Z}_q$. With the $O(2)$ model extended by this symmetry breaking term, we can continuously interpolate between different values of q .

With this dissertation, we have tried to present the material in a self-contained manner. This meant starting with an introduction to classical spin systems and critical phenomena on the lattice. This is the given in Chapter 2. In Chapter 3, we introduce MCMC simulations and their statistical analysis. The goal was for the first two chapters to be a pedagogical introduction to these topics. Then in Chapter 4, we present the bulk of the author's research. This is the chapter on the Extended- $O(2)$ model. In Chapter 5, we come full circle back to QCD and the quantum simulation of QCD. In Chapter 6 we summarize our findings.

CHAPTER 2

CLASSICAL SPIN SYSTEMS AND CRITICAL PHENOMENA ON THE LATTICE

2.1 Ising Model

2.1.1 Introduction

The Ising model was developed to model ferromagnetism. Despite its simplicity, in two and higher dimensions it possesses features that appear also in the quantum field theories which describe the Standard Model of particle physics. These mathematical models begin with the discretization of space or spacetime as a lattice. The fields of the physical theory are defined on the lattice sites and/or on the links between the sites, and the fundamental mechanics of the theory are encoded in the way in which the field variables at neighboring sites and/or links interact. Monte Carlo simulation of such lattice models provides a framework for studying these theories even in regimes where other methods fail. Techniques originally derived to study Ising-like systems are used today to solve problems in quantum field theory (QFT). As such, understanding these classical lattice spin systems and the tools used to study them can be a useful first step to understanding lattice quantum field theories.

In the Ising model, we define discrete “atomic spins” on a lattice of N sites with periodic boundary conditions. The Hamiltonian energy function, which gives the energy of a particular lattice configuration or microstate, is

$$H = - \sum_{\langle i,j \rangle} J_{ij} S_i S_j - \mu \sum_{j=1}^N h_j S_j, \quad (2.1)$$

where the first sum is over nearest neighbor pairs, $S_i = \pm 1$ is the spin-variable sitting at the lattice site i , J_{ij} is the interaction strength between the spin variables at sites i and j , μ is the magnetic moment, and h_j is the interaction strength between the spin variable at site j and an external magnetic field. Typically, the interaction strength for site variables is taken to be the same between all nearest neighbor pairs, i.e. $J_{ij} = J$, and the external magnetic field is taken to be constant,

$h_j = h$. The Ising Hamiltonian becomes

$$H = -J \sum_{\langle i,j \rangle} S_i S_j - h \sum_{j=1}^N S_j, \quad (2.2)$$

where μ has been absorbed in the new h . The spin-variables in this model take the values $+1$ and -1 . Alternatively, one can think of the spin-variables as unit vectors pointing in some reference direction e.g. $+\hat{y}$ or the opposite direction $-\hat{y}$. Then the magnetic field with strength h biases the spin variables toward the positive (if $h > 0$) or negative direction (if $h < 0$).

A mechanical system seeks to minimize its free energy. More precisely, at constant temperature T , the Helmholtz free energy $F = E - TS$ of a statistical mechanical system is minimized at equilibrium. This involves the interplay of two quantities—the internal energy E and the entropy S . For the Ising model, the states with the lowest internal energy are those in which the spins are aligned with each other, so states in which most or all spins are pointing in the same direction dominate when the temperature is low. On the other hand, when the temperature is high, the free energy is minimized when the entropy is large. High entropy states correspond to those with a lot of disorder, i.e., the spins are pointing in random directions. As the temperature of the system is increased, at some point it transitions from a mostly ordered phase to a mostly disordered phase. This transition can be thought of as governed by the competition between the propensity for energy minimization (which is ultimately due to conservation of energy) and the propensity for maximization of entropy (which is ultimately due to the statistical fact that the system is more likely to find itself in a more common configuration). The two-dimensional Ising model is of great theoretical interest because it is the simplest system in statistical mechanics to show a phase transition. It is also one of the few classical spin models whose exact solution is known in one and two dimensions. Furthermore, exact results can be computed on two-dimensional lattices. See Appendix A.1.

At the microscopic level, the behavior of the individual spin variables is governed by a *local* Hamiltonian—each spin-variable is affected only by its immediate neighbors. Nevertheless, at the macroscopic level, the system shows sudden and dramatic changes in thermodynamic quantities—e.g. the specific heat shows an infinite discontinuity in an otherwise smooth function—as the

temperature of the system crosses the critical temperature associated with the phase transition. Evidently, even well-behaved Hamiltonians governing the microscopic behavior can lead to non-analytic behavior in the thermodynamic limit.

The partition function of the Ising model is

$$Z = \sum_i e^{-\beta H_i}, \quad (2.3)$$

where the sum is over all possible lattice configurations, and

$$\beta = \frac{1}{T}, \quad (2.4)$$

is the inverse temperature in units where the Boltzmann constant is $k_B = 1$. The Helmholtz free energy is

$$F = -\frac{1}{\beta} \ln Z = E - \frac{1}{\beta} S. \quad (2.5)$$

2.1.2 Thermodynamics and Observables

The results in this section were produced using Markov chain Monte Carlo (MCMC) with a heatbath updating algorithm. MCMC methods and their statistical analysis will be introduced in Chapter 3. The MCMC code that was used to produce the Ising results in this section was adapted from publicly available Fortran code developed by Bernd Berg [5].

When working with lattice systems, it is often instructive to look at equilibrium lattice configurations at different temperatures. In two dimensions, this is particularly helpful, as lattice configurations can easily be visualized as images. For example, one can use `imshow` in `matplotlib` to display a lattice of numbers as a lattice of colors. In Figure 2.1, we show examples of equilibrium Ising configurations at different temperatures. These results were obtained after performing MCMC simulations on 400×400 lattices for different values of β . One can immediately see how the level of order changes for equilibrium states at different β . At high temperatures (i.e. small β), the equilibrium configurations are highly disordered. At low temperatures (large β), the equilibrium configurations are highly ordered i.e. they are almost fully magnetized in one direction. The images in Figure 2.1 are from independent simulations performed at different fixed β . They do




Figure 2.1 Example Ising configurations at different β (inverse temperature) from MCMC simulations performed on 400×400 lattices. At $\beta = 0$, the system is at infinite temperature and completely disordered/random. As β is increased, the system cools and increasingly large domains develop. Near $\beta = 0.4407$, the system undergoes a second-order phase transition in the infinite-volume limit. At larger values of β , the system is “frozen” with a single domain having taken over the entire lattice. Note that these images are obtained from independent MCMC simulations at fixed β .

not illustrate how a single lattice would vary as β is changed. Notice the appearance of magnetic “domains” where multiple regions of different magnetizations coexist.

The internal energy is

$$E = -\frac{\partial}{\partial \beta} \ln Z = \langle H \rangle, \quad (2.6)$$

where Z is the partition function. For a set of lattice configurations, the internal energy is calculated by computing H for each configuration and then taking the average. To normalize the energy across different lattice sizes, one typically reports the *energy density* E/N i.e. the energy per lattice site.

The specific heat is defined as

$$C = -\frac{\beta^2}{N} \frac{\partial E}{\partial \beta} = \frac{\beta^2}{N} (\langle E^2 \rangle - \langle E \rangle^2), \quad (2.7)$$

where, as usual, N is the volume or the number of sites in the lattice. The quantity on the right¹ can be computed on a sample of equilibrium lattice configurations by computing E and E^2 on each configuration and then averaging over the sample to get $\langle E \rangle$ and $\langle E^2 \rangle$. Note that the variance of a random variable X can be written as $\text{Var}[X] = \langle X^2 \rangle - \langle X \rangle^2$. Thus, the specific heat is proportional to the variance or the fluctuations of the internal energy E .

The magnetization is formally defined as the derivative of the log of the partition function with respect to the ordering field

$$M' = \frac{1}{\beta} \frac{\partial}{\partial h} \ln Z = \left\langle \sum_{i=1}^N S_i \right\rangle, \quad (2.8)$$

where the sum is over the sites of the lattice and the angle brackets indicate an average over an ensemble of lattices. Typically, one reports the magnetization density M/N i.e. the magnetization per lattice site. For a completely disordered lattice, $M/N = 0$ in the $N \rightarrow \infty$ limit. For a completely ordered lattice, $M/N = -1$ or $M/N = +1$. Since the lattice should be equally likely to magnetize in the $+1$ or the -1 direction, these will cancel out in the limit of infinite statistics. Therefore, it is generally more useful to use a proxy magnetization such as

$$M = \left\langle \left| \sum_{i=1}^N S_i \right| \right\rangle. \quad (2.9)$$

On a two-dimensional lattice with no external magnetic field, M is a nonzero quantity for $\beta > \beta_c$. This is the phenomenon of spontaneous magnetization. For the two-dimensional Ising model in infinite volume, the exact value is known to be

$$M(\beta) = \left[1 - \frac{1}{\sinh^4(2\beta)} \right]^{1/8}. \quad (2.10)$$

This formula was presented without proof by Onsager during a pair of conferences in 1948 [6] and proved by Yang in 1952 [7].

¹It is easy to show that $\partial E / \partial \beta = -\langle E^2 \rangle + \langle E \rangle^2$ by starting with the definition $\langle E \rangle = \sum_i E_i e^{-\beta E_i} / \sum_i e^{-\beta E_i}$, differentiating both sides, and then simplifying using $\langle O \rangle = \sum_i O_i e^{-\beta E_i} / \sum_i e^{-\beta E_i}$ for a general observable O .

The magnetic susceptibility, i.e. the fluctuations of the magnetization, can be formally defined as

$$\chi_{M'} = \frac{1}{\beta N} \frac{\partial M'}{\partial h} = \frac{1}{N} \left(\langle M'^2 \rangle - \langle M' \rangle^2 \right). \quad (2.11)$$

We use the proxy magnetization Eq. (2.9) instead of Eq. (2.8) and hence use

$$\chi_M = \frac{1}{N} \left(\langle M^2 \rangle - \langle M \rangle^2 \right). \quad (2.12)$$

In Figure 2.2, we show the internal energy density, specific heat, magnetization per site, and the magnetic susceptibility all versus β for lattice volumes going from 4×4 up to 64×64 . The Ising model has a second-order phase transition at $\beta_c = \ln(1 + \sqrt{2})/2 \approx 0.440687$ in the infinite-volume limit, and each of these thermodynamic functions shows a particular kind of response near the critical point. The specific heat and magnetic susceptibility, for example, diverge with increasing lattice size near the critical point.

2.1.3 Symmetry Considerations

In the Ising model with no external field, there is a global $\mathbb{Z}_2$ symmetry—that is, one can transform the spin variables $S \rightarrow -S$ without changing the energy function Eq. (2.2). This symmetry is spontaneously broken when at low energy, the system chooses either the configuration with most spins oriented up or the configuration with most spins oriented down. These two configurations have the same energy, however, there is a very large energy barrier between them, and it is very unlikely for the overall magnetization to flip orientations in the low energy sector. See Figure 2.3.

Since the two low-energy possibilities have the same energy and spin-flip probabilities, one can say that there are two “thermodynamically equivalent” low-energy sectors. A Monte Carlo simulation at very low energy (i.e. large β) may spend all of its time “frozen” in one of these sectors. The likelihood of “tunneling” to the other sector is very low since almost all spins would have to flip directions. In such a case, the Monte Carlo simulation is stuck in one sector and is not sampling the entire space. In the Ising model, this is not a problem since the two sectors are thermodynamically equivalent, and all measurements will be the same whether it’s stuck in




Figure 2.2 Here we show several thermodynamical functions versus β for the Ising model on several different lattice sizes. From top to bottom, we have the energy density, specific heat, magnetization per site, and the magnetic susceptibility. Notice how these functions behave near the critical point $\beta_c = \ln(1 + \sqrt{2})/2 \approx 0.4407$ where there is a second-order phase transition.




Figure 2.3 In the Ising model with no external field, there is a global $\mathbb{Z}_2$ symmetry—one can transform the spin variables $S \rightarrow -S$ without changing the energy function Eq. (2.2). This symmetry is spontaneously broken when at low energy, the system chooses either the configuration with most spins oriented up or the configuration with most spins oriented down. These two configurations have the same energy, however, there is a very large energy barrier between them, and it is very unlikely for the overall magnetization to flip orientations in the low energy sector.

one sector or the other. In a later chapter, we will see examples of models with explicitly-broken symmetries which result in thermodynamically *distinct* low-energy sectors, which cause a lot of trouble for the MCMC approach.

2.1.4 Topological Defects: Domain Walls

In ferromagnetism, a magnetic domain is a region in which the magnetization is in a uniform direction. When cooled below the Curie temperature, the magnetization of a ferromagnetic material spontaneously divides into many magnetic domains. An analogous phenomenon occurs in the Ising model. Near the critical point, many neighboring spins align with each other forming “domains” of aligned spins. This can be seen very clearly in panels (d) and (e) of Figure 2.1. The boundaries between domains of oppositely-oriented spins are called “domain walls”.

In the Ising model, magnetic domains are a consequence of the competition between energy-minimization and entropy maximization. Consider the Ising model in one dimension for some finite β . When all N spins are aligned

$$\cdots + \cdots$$

then the internal energy E is at its absolute minimum, but the entropy is zero. In contrast, consider the same model but now with a domain wall created by flipping all spins to the right of some arbitrary point

$$\cdots + + + + + + + \mid - - - - - - - - - - - \cdots$$

The internal energy is increased by 2, but the entropy is increased by $k_B \ln(N - 1)$ since there are $N - 1$ places to put the wall. Thus, the free energy changes by an amount $\Delta F = 2 - \ln(N - 1)/\beta$. Evidently, for finite β and for $N \rightarrow \infty$, the creation of a domain wall lowers the free energy. This implies that there cannot be spontaneous magnetization in the one-dimensional Ising model, and thus no phase transition at finite temperature.

In the two-dimensional Ising model, an argument from domain walls can be used to show the existence of spontaneous magnetization (and therefore a phase transition) [8–10]. Also, the low-temperature expansion of the two-dimensional Ising model can be expressed in terms of domain walls. See for example [11].

2.1.5 Second-order Phase Transitions

In the old Ehrenfest classification, phase transitions are labeled by the lowest derivative of the free energy that is discontinuous at the transition. In this scheme, a “first-order” transition exhibits a discontinuity in the first derivative of the free energy with respect to some thermodynamic variable. A “second-order” phase transition is continuous in the first derivative but discontinuous in the second derivative. The internal energy, Eq. (2.6), and the specific heat, Eq. (2.7), are the first and second derivatives, respectively, of the free energy with respect to β . The magnetization, Eq. (2.9), and magnetic susceptibility, Eq. (2.12), are formally the first and second derivatives of the free energy with respect to the external field h . In the modern classification, phase transitions are also divided into “first” and “second” order phase transitions, which mostly coincide with the Ehrenfest classification.

The Ising model has a second-order phase transition at the critical point $\beta_c = \ln(1 + \sqrt{2})/2 \approx 0.4407$. At the phase transition, the internal energy and magnetization are continuous, but the specific heat and magnetic susceptibility diverge. On a finite lattice, true divergences cannot occur,

Table 2.1 Here we show (for reference) the critical exponents of the Ising model in two dimensions. The final column gives the exact infinite-volume critical point.

α	β	γ	δ	η	ν	β_c
0	1/8	7/4	15	1/4	1	$\frac{1}{2} \ln(1 + \sqrt{2})$

however as seen in Figure 2.2, we do see peaks in the specific heat and magnetic susceptibility, which grow with increasing lattice size. Here we see that the specific heat diverges “logarithmically” with volume, i.e. the peak height increases by a constant value each time the lattice size is doubled. The magnetic susceptibility diverges faster.

Near a phase transition, the thermodynamic functions exhibit power law divergences described by a set of critical exponents: $\alpha, \gamma, \nu, \beta, \delta, \eta$. For example, if one defines the reduced temperature $\tau = (T - T_c)/T_c = (\beta_c - \beta)/\beta$, then the specific heat near the phase transition diverges as $C \propto \tau^{-\alpha}$ with some fixed value for the critical exponent² α . The magnetic susceptibility diverges as $\chi \propto \tau^{-\gamma}$, the correlation length diverges as $\xi \propto \tau^{-\nu}$, the ordering field and the magnetization are related as $h \propto M^\delta$, and the spin-spin correlator for spins separated by a distance r diverges as $\langle S_0 S_r \rangle \propto r^{-d+2-\eta}$. To list them nicely in one place, we repeat them here:

$$C \propto \tau^{-\alpha} \quad (2.13)$$

$$\chi \propto \tau^{-\gamma} \quad (2.14)$$

$$\xi \propto \tau^{-\nu} \quad (2.15)$$

$$M \propto \tau^\beta \quad (2.16)$$

$$h \propto M^\delta \quad (2.17)$$

$$\langle S_0 S_r \rangle \propto r^{-d+2-\eta} \quad (2.18)$$

A fixed set of critical exponents defines a universality class. For example, in Table 2.1, we list the exponents of the 2D Ising universality class.

²More completely, one can define two exponents α and α' to allow for different behavior on either side of the critical point.

2.2 Potts Models

2.2.1 Introduction

The Potts model is a generalization of the Ising model. Whereas each Ising spin variable takes on the two possible values ± 1 , in the Potts model, the spin variable takes on one of q possible values. In terms of notation, we follow mostly [5]. The Potts model is also called the “standard Potts model” or sometimes the “Ashkin-Teller-Potts model”. The “vector” or “planar” Potts model, on the other hand, refers to what is more commonly called the “clock model”, which we will look at in Section 2.4.

As with the Ising model, we consider a d -dimensional lattice with N sites and periodic boundary conditions. For the q -state Potts model, the i th lattice site contains a spin-variable q_i , which can take on the possible values $q_i = 0, 1, \dots, q-1$. Each spin interacts only with its nearest neighbors. The Hamiltonian energy function is

$$H = -2 \sum_{\langle i,j \rangle} J_{ij}(q_i, q_j) \delta_{q_i q_j} + \frac{2dN}{q} - \frac{2}{\beta} h \sum_{i=1}^N \delta_{0q_i}, \quad (2.19)$$

where the first sum is over nearest neighbor pairs of lattice sites and $J_{ij}(q_i, q_j)$ parameterizes the interaction between adjacent lattice sites i and j . In principle, J_{ij} could be different for each pair of adjacent sites and thus would be an array. Typically, $J_{ij} = J$ is a constant, and a common choice is $J = 1$. The second term is a convenient normalization term. The third term accounts for the possibility of an external magnetic field. With this choice of the magnetic field term, the external field interacts only with site variables in the state $q_i = 0$. When $J = 1$ and $q = 2$, the Potts model is equivalent to the Ising model.

For the rest of this section, we will assume $J = 1$ and $h = 0$, in which case the Hamiltonian function simplifies to

$$H = -2 \sum_{\langle i,j \rangle} \delta_{q_i q_j} + \frac{2dN}{q}. \quad (2.20)$$

The canonical partition function takes the familiar form

$$Z = \sum_i e^{-\beta H_i}, \quad (2.21)$$

where the sum is over lattice configurations.




Figure 2.4 Example q -state Potts model configurations at different q from MCMC simulations performed on 400×400 lattices near their critical points.

2.2.2 Thermodynamics and Observables

The results in this section were produced using MCMC with a heatbath updating algorithm. The MCMC code was adapted from publicly available Fortran code developed by Bernd Berg [5].

In Figure 2.4 we show example equilibrium configurations for the q -state Potts model at different q from MCMC simulations performed on 400×400 lattices near their critical points.

As usual, the internal energy³ is defined as

$$E = -\frac{\partial}{\partial \beta} \ln Z = \langle H \rangle, \quad (2.22)$$

At $\beta = 0$, the system is completely disordered and all possible values of q_i are equally likely. In that case, the expected contribution to the energy of site i is $1/q$, and so the sum in Eq. (2.20) is

³Sometimes it is more convenient to consider the “action” $S = \langle \sum_{\langle i,j \rangle} \delta_{q_i q_j} \rangle$ since for the Potts model, this quantity is a non-negative integer. For a given configuration, the action takes the possible values $S = 0, 1, 2, \dots, dN$. The largest possible action, $S_{max} = dN$, occurs when all spins have the same value.

canceled out by the normalization term $2dN/q$, resulting in

$$\left. \frac{E}{N} \right|_{\beta=0} = 0. \quad (2.23)$$

The minimum energy density, which occurs when $\beta \rightarrow \infty$ and all spins are aligned, is easily calculated to be

$$\left. \frac{E}{N} \right|_{\beta \rightarrow \infty} = 2d \left(\frac{1}{q} - 1 \right). \quad (2.24)$$

For the specific heat, we take the usual definition

$$C = -\frac{\beta^2}{N} \frac{\partial E}{\partial \beta} = \frac{\beta^2}{N} \left(\langle E^2 \rangle - \langle E \rangle^2 \right). \quad (2.25)$$

In two dimensions, the Potts model has a second order phase transition for $q \leq 4$ and a first order phase transition for $q \geq 5$ [12]. The infinite-volume critical point is [5]

$$\beta_c = \frac{1}{T_c} = \frac{1}{2} \ln(1 + \sqrt{q}), \quad q = 2, 3, \dots \quad (2.26)$$

Formally, the magnetization is the derivative of the log of the partition function with respect to the ordering field

$$M' = \frac{\partial}{\partial h} \ln Z = \left\langle 2 \sum_{i=1}^N \delta_{0q_i} \right\rangle. \quad (2.27)$$

However, in MCMC simulations without an ordering field, this quantity does not signal critical behavior, so we turn to proxy definitions of the magnetization. Typically, one wants the “magnetization” to be a measure of the amount of order in the lattice. For a system with a phase transition, it should be some order parameter whose derivative shows a divergence at the critical point. We start by defining the “occupation numbers”

$$N_{q_0} = \sum_{i=1}^N \delta_{q_i, q_0}, \quad q_0 = 0, 1, \dots, q-1. \quad (2.28)$$

That is, N_{q_0} gives the number of lattice sites with value q_0 . These numbers satisfy $N_0 + N_1 + \dots + N_{q-1} = N$. One approach, as in Ref. [5], is to define the Potts magnetization density with respect to q_0 as

$$M_{q_0} = \frac{1}{N} \langle N_{q_0} \rangle. \quad (2.29)$$

where the angle brackets as usual represent the ensemble average taken over many lattices at the same value of β . This magnetization can be understood as the average fraction of spin variables with spin value $q = q_0$. For example, in a 3-state Potts model with all spins in the state $q = 1$, we would have $M_1 = 1$ and $M_0 = M_2 = 0$. For the 2-state Potts model ($q_i \in \{0, 1\}$), the conventional definition of the Ising magnetization, Eq. (2.9), is the difference $m = \langle |N_0 - N_1| \rangle$. For example, in a 2-state Potts model with 75% of the spins in the state $q = 1$, the Potts magnetizations are $M_1 = 0.75$ and $M_0 = 0.25$, and the Ising magnetization density is $m/N = |0.25 - 0.75| = 0.50$. In the infinite-time limit, the expectation of the Potts magnetization, when there is no external field, is

$$M_{q_0} = \frac{1}{q}, \quad (2.30)$$

independent of β . Thus, the Potts magnetization does not signal critical behavior and cannot be used as an order parameter. However, the Potts magnetization squared

$$M_{q_0}^2 \equiv q \left\langle \frac{N_{q_0}^2}{N^2} \right\rangle, \quad (2.31)$$

does signal critical behavior. According to [5], a finite size scaling analysis would reveal a singularity at β_c in $M_{q_0}^2$ for first-order transitions and in the derivative of $M_{q_0}^2$ for second-order transitions.

In this work, we follow others (e.g. [13, 14]) in defining magnetization in the q -state Potts model. Now we are interested in the largest occupation number $N^* \equiv \max(N_0, N_1, \dots, N_{q-1})$ on a per-configuration basis. For example, if a particular 10×10 configuration of the 3-state clock model has thirty-three sites with $q_i = 0$, twenty-one sites with $q_i = 1$, and the remaining forty-six sites with $q_i = 2$, then $N^* = 46$ for that configuration. We then define the magnetization for the q -state Potts model as

$$M = \frac{q \langle N^* \rangle - 1}{q - 1}, \quad (2.32)$$

which is normalized such that the magnetization per site is

$$\left. \frac{M}{N} \right|_{\beta=0} = 0, \quad \left. \frac{M}{N} \right|_{\beta \rightarrow \infty} = 1. \quad (2.33)$$

For the magnetic susceptibility, we take the usual definition

$$\chi_M = \frac{1}{N} \left(\langle M^2 \rangle - \langle M \rangle^2 \right). \quad (2.34)$$

In Figure 2.5, we look at these thermodynamic quantities as functions of β for different values of $2 \leq q \leq 10$. These results are from MCMC simulations performed on 32×32 lattices. The top panel in this figure shows the energy density. As shown in Eq. (2.23), the energy density is zero at $\beta = 0$. This reflects the normalization chosen such that a fully disordered lattice corresponds to zero energy. The inflection points in the energy density curves occur near the critical points given by Eq. (2.26). As q increases, the critical point shifts toward larger β . Notice that for $q \geq 5$, the energy density seems to become discontinuous at the critical point. This is a reflection of the phase transition being first-order for $q \geq 5$. On an infinite lattice, there would be true discontinuities here. As β becomes large, the energy density asymptotically approaches the value given by Eq. (2.24). In the second panel, we plot the specific heat defined in Eq. (2.25). The specific heat shows a prominent peak near the critical point, and we see that this peak becomes taller and sharper at larger q . In the third panel, we show the magnetization defined in Eq. (2.32). In the disordered regime at small β , the magnetization is $M \sim 0$, whereas in the ordered regime at large β , $M \sim 1$. As with the energy, the magnetization also seems to become discontinuous at the critical point for $q \geq 5$. In the fourth and final panel, we plot the magnetic susceptibility defined in Eq. (2.34). This quantity shows peaks near the critical point similar to those seen in the specific heat.

2.2.3 First-order Phase Transitions

The q -state Potts model is known to have a second-order phase transition for $q = 2, 3, 4$ and a first-order phase transition for $q \geq 5$ [12]. In the infinite volume limit, this implies for $q = 2, 3, 4$, a continuous energy and magnetization but a discontinuous specific heat and magnetic susceptibility. For $q \geq 5$, it implies a discontinuous energy and magnetization and also of course discontinuous specific heat and susceptibility. Numerical studies are performed on finite lattices which cannot show true discontinuities, so it is at times tricky to discern the nature of a phase transition. Typically, numerical simulations have to be done on different size lattices and then one tries to extrapolate the




Figure 2.5 The energy density, specific heat, magnetization per site, and magnetic susceptibility of the q -state Potts model on a 32×32 lattice. These were obtained from low statistics MCMC simulations using a heatbath algorithm. Notice how the energy is continuous at $q = 2, 3, 4$, but becomes discontinuous at larger q . This is an illustration of the Potts model having a second-order phase transition at small q , but a first-order phase transition at larger q .




Figure 2.6 The specific heat (left) and magnetic susceptibility (right) of the 8-state Potts model both diverge sharply at the first-order phase transition near $\beta \approx 0.671227$. On a finite lattice, these divergences are finite, but indicate the presence of a phase transition in the infinite volume model. These examples were obtained from small-statistics MCMC simulations using a heatbath updating algorithm.




Figure 2.7 Histograms of the energy density at different values of β in the vicinity of a phase transition. In the left panel, histograms of the energy density near the second-order phase transition of the 2-state Potts model (i.e. the Ising model). In the right panel, histograms of the energy density near the first-order phase transition of the 5-state Potts model. Notice that the distribution becomes bimodal near a first-order critical point. These results were obtained on 100×100 lattices using an MCMC heatbath algorithm.

behavior to the infinite-volume limit. In Figure 2.6, we show the specific heat (left panel) and the magnetic susceptibility (right panel) for the 8-state Potts model on lattices of different size. These quantities are clearly diverging with volume.

Distinguishing between first and second order phase transitions is not always easy. The specific

heat and magnetic susceptibility diverge with volume in both cases. However, for first-order transitions, the energy and magnetization become discontinuous at the critical point. For second-order transitions, they do not. We see hints of these energy and magnetization discontinuities for $q \geq 5$ in Figure 2.5. For example, for $q = 10$, we see that as $\beta \rightarrow \beta_c$ from the left, the energy approaches $E/N \approx -1.5$, but as $\beta \rightarrow \beta_c$ from the right, the energy approaches $E/N \approx -3$. Histogramming the energy values from an MCMC simulation performed very close to the critical point β_c results in a bimodal distribution as seen in the right panel of Figure 2.7. In fact, one way to find the critical point β_c of a first-order phase transition is to vary β until the two peaks in such a histogram have the same height. A similar histogram of the energies near the critical point of a second-order phase transition, shows a single peak as seen in the left panel of Figure 2.7.

2.3 $O(n)$ Models

2.3.1 Introduction

We turn our attention now to our first spin system with continuous symmetry. The Hamiltonian energy function of the classical $O(n)$ model (aka n -vector model) on a lattice with N sites is

$$H = -J \sum_{\langle i,j \rangle} \vec{S}_i \cdot \vec{S}_j - \sum_{i=1}^N \vec{h} \cdot \vec{S}_i, \quad (2.35)$$

where the first sum is over nearest neighbor pairs, J gives the coupling strength, $\vec{S}_i$ is an n -component unit vector at the site i , and $\vec{h}$ is the external magnetic field.

We will focus on the case $n = 2$, which is commonly called the “XY model”. In this case, we can parameterize the spin variable using a single angle $\theta_i \in [0, 2\pi)$ as $\vec{S}_i \equiv (\cos \theta_i, \sin \theta_i)^T$. Assuming a constant and uniform magnetic field, we can choose $\vec{h} = h\hat{x}$. Then the Hamiltonian simplifies to

$$H = -J \sum_{\langle i,j \rangle} \cos(\theta_i - \theta_j) - h \sum_{i=1}^N \cos \theta_i. \quad (2.36)$$

For the rest of this section, we will assume $J = 1$ and no external field so that $h = 0$, in which case

$$H = - \sum_{\langle i,j \rangle} \cos(\theta_i - \theta_j). \quad (2.37)$$

Furthermore, we will focus on the $O(2)$ model in two dimensions. In the previous models we considered, there was a discrete symmetry, but here we have a model with continuous $O(2)/U(1)$ symmetry.

The XY model has been studied analytically [15], using Monte Carlo methods [16–23], and using tensor methods [24]. Unless otherwise noted, numerical results and plots presented in this section were obtained using a lattice codebase developed in-house.

The canonical partition function must now take into account that the degrees of freedom are continuous. Schematically, we can write

$$Z = \int dU e^{-\beta H(U)}, \quad (2.38)$$

where the integration is over lattice configurations U .

2.3.2 Thermodynamics and Observables

The results in this section were produced using a Biased Metropolis Heatbath Algorithm (BMHA) in an MCMC codebase developed by the author.

The internal energy and specific heat for the XY model are defined in the usual way.

We can define the magnetization vector for a given spin configuration as

$$\vec{M} = (M_x, M_y) = \left(\sum_{i=1}^N \cos \theta_i, \sum_{i=1}^N \sin \theta_i \right) = \sum_{i=1}^N \vec{S}_i, \quad (2.39)$$

where N is the number of lattice sites, and $\vec{S}_i$ is the spin vector at site i which makes an angle θ_i with some reference direction. With this definition, the ensemble average matches the formal definition of the magnetization as the derivative of the logarithm of the partition function

$$\langle \vec{M} \rangle = \frac{1}{\beta} \frac{\partial}{\partial \vec{h}} \ln Z = \left\langle \sum_{i=1}^N \vec{S}_i \right\rangle. \quad (2.40)$$

However, in MCMC simulations with zero external field (i.e. $\vec{h} = 0$) and no explicit symmetry breaking, the spontaneous magnetization is uninteresting: $\langle \vec{M} \rangle = 0$. So we use instead the magnitude of the vector magnetization

$$M = |\vec{M}| = \left| \sum_{i=1}^N \vec{S}_i \right| = \sqrt{\left(\sum_{i=1}^N \cos \theta_i \right)^2 + \left(\sum_{i=1}^N \sin \theta_i \right)^2}, \quad (2.41)$$

as a proxy magnetization. This serves as an order parameter to signal critical behavior. Other quantities that may be used as an order parameter include $\langle M \rangle$ and $\langle M_x \rangle$ or $\langle M_y \rangle$ or even $\langle M^2 \rangle = \langle M_x^2 \rangle + \langle M_y^2 \rangle$. In the thermodynamic limit, the Mermin-Wagner theorem requires that $\langle M \rangle = 0$, however, on any finite lattice, this quantity is nonzero at low temperature and may be used as an order parameter.

By the Mermin-Wagner theorem, a continuous symmetry cannot be spontaneously broken in a two-dimensional system with short-range interactions at any finite temperature. The XY model has a continuous $U(1)$ symmetry, and so this theorem implies that long-range order (i.e. magnetization) cannot occur. Physically, this can be understood as follows. In a low-dimensional system with a continuous symmetry such as the XY model, long-range order is more fragile and vulnerable to thermal fluctuations, and the long-range order that one might expect at low temperature ends up being destroyed by long wavelength spin-wave excitations [25]. On the other hand, the Mermin-Wagner theorem applies to the *infinite* XY model. Magnetization can and does occur at sufficiently low temperature in even the largest realizable finite-size implementation of the XY model. This effectively makes the true thermodynamic limit of the 2D XY model inaccessible [26–29].

Once the order parameter M is chosen, one defines the magnetic susceptibility in the usual manner.

In Figure 2.8, we show several thermodynamic quantities versus β for the XY model at different lattice sizes. In the top panel, we show the energy density, and in the second panel the specific heat. In the third panel, we show the magnetization per lattice site, and in the bottom panel, we show the magnetic susceptibility. Near the critical point, the magnetic susceptibility diverges with volume, but the specific heat does not. This is in contrast to the behavior of the specific heat near a first or second-order transition.

2.3.3 Topological Defects: Vortices

In the two-dimensional XY model, where the spin degrees-of-freedom are vectors that can take angles $\theta \in [0, 2\pi)$, we see a new kind of topological excitation—vortices, antivortices, and vortex-antivortex pairs.




Figure 2.8 Here we show several thermodynamic quantities versus β for the XY model at different lattice sizes. In the top panel, we show the energy density, and in the second panel the specific heat. In the third panel, we show the magnetization per lattice site, and in the bottom panel, we show the magnetic susceptibility. Near the critical point, the magnetic susceptibility diverges with volume, but the specific heat does not. This is in contrast to the behavior of the specific heat near a first or second-order transition.

We start by defining the winding number as in [25]

$$k = \frac{1}{2\pi} \sum_{\Gamma} \Delta\varphi_{ii'}, \quad (2.42)$$

where Γ is an oriented closed loop and

$$\Delta\varphi_{ii'} = \varphi_{i'} - \varphi_i \in [-\pi, \pi), \quad (2.43)$$

is the angular difference between spin $\varphi_{i'}$ and its loop neighbor φ_i . If an oriented closed loop has $k > 0$, it is a (positive) vortex of strength k , and if it has $k < 0$, it is a “negative vortex” or an “antivortex” of strength k . The most common value is $k = 0$. In practice, one can choose Γ to be an elementary plaquette measured in the counter-clockwise manner.

There are really two varieties of vortices that appear in these kinds of models. In a Monte Carlo simulation, “dynamical” vortices occur when a system is quenched from a very high temperature to a low temperature. The vortices and antivortices that appear during the equilibration period are large and prominent and make nice pictures as in Figure 2.9. In this figure we start with a system at infinite temperature and cool it very quickly to a cold temperature. At first there are a large number of tiny vortices and antivortices which quickly annihilate leaving a handful of vortex-antivortex pairs that last for a long time (in terms of Monte Carlo steps) before they too annihilate—leaving a quasi-ordered lattice. Dynamics are explored in [16, 30–32].

One can also study “equilibrium” vortices, by appropriately sampling equilibrium lattice configurations and identifying the vortices therein. In this case, one is not interested in the dynamical evolution of the vortex through Monte Carlo time. Rather, one may be interested in the density of vortices and their role in the BKT transition as one approaches the critical point β_c . Example equilibrium vortices are shown in Figure 2.10. Further discussion of vortices can be found in [18, 29, 33–35].

2.3.4 BKT Phase Transition

Typically, a phase transition in a classical spin system occurs between a disordered phase at high temperature and a broken symmetry phase with long-range order at low temperature. The XY model, which has continuous unbroken symmetry, cannot have long-range order. For this reason,




Figure 2.9 Here we show an example of the evolution of a 32×32 lattice in the XY model with periodic boundary conditions when it is started at infinite temperature and immediately quenched to nearly zero temperature. We start in panel (a) with random spins. This corresponds to an infinite temperature lattice. At this temperature, tiny vortices (black squares) and antivortices (red squares) proliferate. Vortex positions are identified by computing the winding number Eq. (2.42) of each plaquette. A winding number of $+1$ corresponds to a vortex, and a winding number of -1 corresponds to an antivortex. Panel (b) shows the same lattice after one BMHA sweep with the low temperature $\beta = 100$. Already many of the vortices and antivortices have annihilated and some order is starting to appear. Panel (c) shows the lattice after 10 BMHA sweeps with $\beta = 100$. Panels (d) and (e) show the lattice after 100 and 1000 BMHA sweeps. Now only two or three vortex-antivortex pairs remain. But they are large and annihilate slowly. By the 10000th BMHA sweeps (panel (f)), no vortex-antivortex pairs remain, and the lattice is in a state of quasi-long-range order.

such a model was believed to have no phase transition. However, the work of Berezinskii, Kosterlitz, and Thouless [15, 36, 37], showed that the XY model has a new kind of phase transition—one that separates a region with quasi-long-range order from one where correlations decay exponentially. The BKT transition is not accompanied with any symmetry-breaking, making it a rather special




Figure 2.10 A few examples of equilibrium vortices from independent Monte Carlo simulations of the XY model performed on 32×32 lattices with periodic boundary conditions. Vortex positions are identified by computing the winding number Eq. (2.42) of each plaquette. A winding number of $+1$ corresponds to a vortex, and a winding number of -1 corresponds to an antivortex. In panel (a), an equilibrium configuration is shown for $\beta = 0.8$, which is in the disordered phase. The vortex-antivortex pair density is high and many are unbound. Panel (b) comes from a simulation performed at $\beta = 1.0$, which is not far from the critical point of the XY model. Panel (c) shows an equilibrium configuration at $\beta = 1.3$, which is in the quasi-ordered phase. Here the vortex-antivortex pair density is small, and they are tightly bound.

case of a phase transition.

The phase transition in the XY model can be understood in terms of vortices [16, 38]. In the low temperature quasi-ordered phase, vortex-antivortex pairs are tightly bound by a logarithmic potential. As the temperature increases, they become unbound, and in the disordered phase, vortices and antivortices behave like a plasma of free charges. This is analogous to the phenomenon of confinement and deconfinement in QCD.

The low temperature phase (i.e. $\beta > \beta_c$) of the XY model goes by many names including “quasi-ordered phase”, “topological phase”, “critical phase”, “spin-wave phase”, “condensed BKT phase”, and “vortex-antivortex phase”. In this phase, the correlation length is infinite, and the spin-spin correlation function decays as a power of the distance. There is no true long-range order in this phase—thus the Mermin-Wagner theorem is satisfied—however, the system is quasi-ordered and on a finite lattice the system is magnetized in this phase. Vortices exist as vortex-antivortex pairs which are tightly bound by a logarithmic potential. Such tightly-bound vortex pairs are unimportant since their effect is short-ranged. The important large-scale fluctuations in this phase




Figure 2.11 Here we show the allowed “clock” angles for several different values of q . In the q -state clock model, the spin-variables sitting at lattice sites can only take the allowed values $\theta = \frac{2\pi k}{q}$ with $k = 0, 1, \dots, q - 1$.

are spin waves.

At the critical point $\beta_c = 1.11995(6)$ [39], there is an “infinite-order” BKT transition. As the temperature is increased, bound vortex-antivortex pairs dissociate at the critical temperature—leaving free vortices and antivortices.

The high temperature phase (i.e. $\beta < \beta_c$) goes by several names including “disordered phase” and “paramagnetic phase”. In this phase, the spin-spin correlation function decays exponentially with distance and free vortices and antivortices behave like a plasma of free charges.

The BKT transition can be distinguished from first and second-order transitions by the behavior of the specific heat. The specific heat typically diverges at first-order (see Figure 2.6) and second-order (see Figure 2.2) phase transitions but not at a BKT transition (see Figure 2.8). However, this check may not be sufficient. If the critical exponent α is negative, then the specific heat plateaus also for second-order transitions. See for example, [40]. For more on distinguishing BKT from first- and second-order transitions, see [40, 41].

2.4 Clock Models

2.4.1 Introduction

We now look at the q -state clock model, which is also called the “planar Potts model”, the “vector Potts model”, or the “ $\mathbb{Z}_q$ model”. The Hamiltonian energy function is the familiar

$$H = -J \sum_{\langle i,j \rangle} \vec{S}_i \cdot \vec{S}_j - \sum_{i=1}^N \vec{h} \cdot \vec{S}_i, \quad (2.44)$$

where the first sum is over nearest neighbor pairs, J gives the coupling strength, $\vec{S}_i$ is a 2-component unit vector at the site i , the second sum is over the N lattice sites, and $\vec{h}$ is the external magnetic field. This is precisely, the energy function of the $O(2)$ or XY model, however, the spins are now restricted to take the “clock” angles

$$\theta_i = \frac{2\pi k}{q}, \quad k = 0, 1, \dots, q-1. \quad (2.45)$$

See Figure 2.11. We parameterize the spin variables as $\vec{S}_i \equiv (\cos \theta_i, \sin \theta_i)^T$. Assuming a constant and uniform magnetic field, we can choose $\vec{h} = h\hat{x}$. Then the Hamiltonian simplifies to

$$H = -J \sum_{\langle i,j \rangle} \cos(\theta_i - \theta_j) - h \sum_{i=1}^N \cos \theta_i. \quad (2.46)$$

For the rest of this section, we will assume $J = 1$ and no external field so that $h = 0$, in which case

$$H = - \sum_{\langle i,j \rangle} \cos(\theta_i - \theta_j). \quad (2.47)$$

The q -state clock model is a rich model with connections to several other spin systems. For $q = 2$, the clock angles are $\theta_i \in \{0, \pi\}$, and so

$$\cos(\theta_i - \theta_j) = \begin{cases} 1 & \text{if } \theta_i = \theta_j \\ -1 & \text{if } \theta_i \neq \theta_j \end{cases}. \quad (2.48)$$

This is equivalent to the Ising model. For $q = 3$, the clock angles are $\theta_i \in \{0, 2\pi/3, 4\pi/3\}$, and so

$$\cos(\theta_i - \theta_j) = \begin{cases} 1 & \text{if } \theta_i = \theta_j \\ -\frac{1}{2} & \text{if } \theta_i \neq \theta_j \end{cases}. \quad (2.49)$$

This is equivalent to the 3-state “standard” Potts model, which has $q_i \in \{0, 1, 2\}$. In fact, one can use the standard Potts model code to simulate the 3-state vector Potts model by substituting the factor $2 \rightarrow 3/2$ in the energy function Eq. (2.20). The 4-state clock model with interaction strength J is equivalent to two decoupled Ising models with interaction strengths $J/2$ [42]. That is, $Z_{4,J} = Z_{2,J/2} Z_{2,J/2}$, where $Z_{4,J}$ is the partition function of the 4-state clock model, and $Z_{2,J/2}$ is the partition function of the 2-state clock model (i.e. Ising model) with half the normal interaction

strength. Since we have algorithms that can exactly calculate the Ising lattice partition function, this means we can also calculate it for the 4-state clock model. In each of these cases (i.e. $q = 2, 3, 4$), the clock model has a second-order phase transition. However, for $q \geq 5$, the clock model has a pair of BKT transitions. In the limit $q \rightarrow \infty$, the clock model becomes the XY model with a single BKT transition at finite β .

The clock model has been studied using various approaches including analytical [25, 43], Monte Carlo [44–48], and tensor methods [49–51].

The canonical partition function takes the usual form

$$Z = \sum_i e^{-\beta H_i}, \quad (2.50)$$

where the sum is over lattice configurations.

2.4.2 Thermodynamics and Observables

The results in this section were produced using MCMC with a heatbath updating algorithm. The MCMC code was adapted from publicly available Fortran code developed by Bernd Berg [5].

As usual, the internal energy is the expectation value of the energy function Eq. (2.47) measured on a sample of equilibrium lattice configurations

$$E = -\frac{\partial}{\partial \beta} \ln Z = \langle H \rangle = -\left\langle \sum_{\langle i,j \rangle} \cos(\theta_i - \theta_j) \right\rangle. \quad (2.51)$$

At $\beta = 0$, the system is completely disordered, and the angles allowed by Eq. (2.45) are equally likely. Due to the $\mathbb{Z}_q$ symmetry of these angles,

$$\left. \frac{E}{N} \right|_{\beta=0} = 0. \quad (2.52)$$

The minimum energy density, which occurs when $\beta \rightarrow \infty$ and all spins are aligned, is easily calculated to be

$$\left. \frac{E}{N} \right|_{\beta \rightarrow \infty} = -d, \quad (2.53)$$

where d is the dimension of the system. Unless otherwise noted, all results in this dissertation are for the models in $d = 2$ dimensions. The energy in the large- β regime can be approximated using a




Figure 2.12 Here we show several thermodynamic quantities versus β for the $q = 3, 4, 5$ state clock models at different lattice sizes. In the top row, we show the energy density, and in the second row the specific heat. In the third row, we show the magnetization per lattice site, and in the bottom row, we show the magnetic susceptibility. Note that the 2-state clock model is just the Ising model, with results shown in Figure 2.2.

low-temperature expansion. See Appendix A.4. In the top row of Figure 2.12, we show the energy density for the $q = 3, 4, 5$ clock models on different lattice volumes.

As usual, the specific heat is defined as

$$C = -\frac{\beta^2}{N} \frac{dE}{d\beta} = \frac{\beta^2}{N} (\langle E^2 \rangle - \langle E \rangle^2). \quad (2.54)$$

In the second row of Figure 2.12, we show the specific heat for the $q = 3, 4, 5$ clock models on different lattice volumes.

We define the magnetization vector as [40, 46–48]

$$\vec{M} = (M_x, M_y) = \left(\sum_{i=1}^N \cos \theta_i, \sum_{i=1}^N \sin \theta_i \right) = \sum_{i=1}^N \vec{S}_i, \quad (2.55)$$

where N is the number of sites and $\vec{S}_i$ is the spin-vector at site i . Formally, the magnetization is

$$\langle \vec{M} \rangle = \frac{1}{\beta} \frac{\partial}{\partial \vec{h}} \ln Z = \left\langle \sum_{i=1}^N \vec{S}_i \right\rangle. \quad (2.56)$$

However, as noted in the previous models studied, the formal spontaneous magnetization is uninteresting in MCMC simulations when there is no external field or explicit symmetry breaking. So instead, we use the magnitude of the magnetization vector

$$M = |\vec{M}|, \quad (2.57)$$

as a proxy magnetization. Often it is useful to work with the magnetization per spin M/N . Quantities that are often used as order parameters include $\langle M \rangle$ [40, 45, 50, 52] and $\langle M_x \rangle$ or $\langle M_y \rangle$ or even $\langle M^2 \rangle = \langle M_x^2 \rangle + \langle M_y^2 \rangle$ [47]. In the third row of Figure 2.12, we show the magnetization per site for the $q = 3, 4, 5$ clock models on different lattice volumes. Analytically, we know that $\langle M \rangle = 0$ for $\beta = 0$. However, in an MCMC lattice simulation, when computing an average like $\langle |\vec{M}| \rangle$ from a sample of lattice configurations, one is always adding a (small) positive number, and so in practice one sees $\langle |\vec{M}| \rangle|_{\beta=0} = 0$ only in the infinite volume limit.

One can also define a complex magnetization as in [44, 53, 54]

$$M_C = \sum_i e^{i\theta_i} = |M_C| e^{i\psi}. \quad (2.58)$$

The magnitude $|M_C|$ then serves as an order parameter. The vector magnetization $\vec{M} = M_x\hat{x} + M_y\hat{y}$ given by Eq. (2.55) and the complex magnetization $M_C = M_x + iM_y$ given by Eq. (2.58) are really the same 2-component objects. One is written as a real 2D vector, and the other is written as a complex number. In fact, Eq. (2.57) and $|M_C|$ are exactly the same thing.

In the q -state clock model with $q \geq 5$ there are two BKT phase transitions. The order parameter Eq. (2.57) can readily detect⁴ the small- β transition between the disordered and quasi-ordered phases, but it is unable to detect the large- β transition between the quasi-ordered and ordered phases. To detect the large- β transition, a “rotated magnetization” [44, 45, 53–55]

$$m_\psi = \cos(q\psi). \quad (2.59)$$

is sometimes used. Here, $\psi = \tan^{-1}(M_y/M_x)$.

One can also look at the “population” as in [54]

$$S = \frac{q}{q-1} \left[\frac{1}{N} \max \{n_0, n_1, \dots, n_{q-1}\} - \frac{1}{q} \right], \quad (2.60)$$

where N is the total number of lattice sites, and n_k is the number of sites with spin angle θ_k . If the system is disordered, then $\max\{n_0, n_1, \dots, n_{q-1}\} = N/q$ and $S = 0$. On the other hand, if the system is ordered with all N spin variables having the same angle, then $\max_{k=0, q-1}(n_k) = N$ and $S = 1$. This order parameter is more sensitive to the lower temperature transition, i.e. its susceptibility shows a pronounced peak at the lower temperature transition between the ordered and quasi-ordered phases.

For the magnetic susceptibility, we take the usual definition

$$\chi_M = \frac{1}{N} \left(\langle M^2 \rangle - \langle M \rangle^2 \right) \quad (2.61)$$

Additional notes on magnetic susceptibility can be found in Appendix A.2.

2.4.3 Topological Defects: Domain Walls and Vortices

In the Ising model there are domain walls but no vortices because spin variables in the Ising model are restricted to point in one of two opposite directions. In the XY model, there are

⁴Via e.g. Binder cumulants, which will be discussed in a later section




Figure 2.13 A few examples of equilibrium vortices from independent Monte Carlo simulations of the 5-state clock model performed on 32×32 lattices with periodic boundary conditions. Vortex positions are identified by computing the winding number Eq. (2.42) of each plaquette. A winding number of $+1$ corresponds to a vortex, and a winding number of -1 corresponds to an antivortex. In panel (a), an equilibrium configuration is shown for $\beta = 0.6$, which is in the disordered phase. The vortex-antivortex pair density is high and many are unbound. Panel (b) comes from a simulation performed at $\beta = 1.0$, which is close to the critical region of the 5-state clock model. Panel (c) shows an equilibrium configuration at $\beta = 1.3$, which is in the quasi-ordered phase. Here the vortex pair density is very low. In this example we see no vortices.

vortices but no true domain walls since the spin degrees of freedom are continuous. In the clock model, at least for sufficiently large q , we see both domain walls and vortices. Here, the low-temperature physics is dominated by domain wall excitations, whereas the vortices are relevant in the intermediate and high temperature phases [25].

There seems to be some confusion in the literature regarding vortices in the q -state clock model. Some authors [25] claim that vortices appear only for $q \geq 5$, with the implication that they play a very fundamental role in the BKT transition. However, the definition of a vortex clearly allows for vortices also in the $q = 4$ state clock model which does not have a BKT transition. For example, one can construct a vortex in the $q = 4$ state clock model by having the (ordered) spins around a plaquette having the angles $0, \pi/2, \pi, 3\pi/2$.

Further discussion on topological defects can be found in [45, 47, 52, 56]. Discussion of dynamics can be found in [45, 47, 56, 57].

2.4.4 Phase Diagram

The clock model interpolates between the Ising model ($q = 2$), which has a second-order phase transition and the XY model ($q \rightarrow \infty$), which has a BKT transition.

For $q = 2, 3, 4$, the clock model has a second-order phase transition—so-called because the second derivatives of the free energy (e.g. specific heat) are divergent while the first derivatives (e.g. internal energy) are continuous at the critical point. At the critical point β_c , the correlation length, specific heat, and susceptibilities become infinite (in the infinite-volume limit). We can see this in Figure 2.12, which shows several thermodynamic functions versus β for different lattice sizes. The internal energy and magnetization remain continuous, but the specific heat and magnetic susceptibility diverge with volume near the critical point. The critical point separates the model into two phases. The high temperature paramagnetic phase (i.e. $\beta < \beta_c$) is disordered, i.e. the magnetization order parameter M is zero above the critical temperature. As the system is cooled from the high temperature phase it passes through a second-order phase transition at the critical point $\beta = \beta_c$. Here the symmetry is spontaneously broken and the system chooses a magnetization direction. That is, the order parameter M becomes nonzero below the critical temperature. The low temperature ferromagnetic phase (i.e. $\beta > \beta_c$) features long-range order in the form of spontaneous magnetization.

For $q \geq 5$, there are two Berezinskii-Kosterlitz-Thouless (BKT) transitions. The BKT transition is sometimes referred to as an “infinite-order” phase transition. The free energy has an essential singularity at $\beta = \beta_c$, however, it is smooth and infinitely differentiable. The two BKT transitions, which occur at the critical points $\beta_{c1} < \beta_{c2}$, separate the model into three phases. The low temperature (i.e. $\beta > \beta_{c2}$) features long-range order i.e. magnetization. At $\beta = \beta_{c2}$, there is a phase transition believed to be of the BKT type. As $q \rightarrow \infty$, this critical point moves $\beta_{c2} \rightarrow \infty$. Hence the XY model, which corresponds to $q = \infty$, does not have this transition. The intermediate temperature phase (i.e. $\beta_{c1} < \beta < \beta_{c2}$) goes by many names including “quasi-ordered phase”, “quasi-liquid phase”, “topological phase”, “critical phase”, “spin-wave phase”, “condensed BKT phase”, and “vortex-antivortex phase”. In this phase, the correlation length is infinite, and the spin-




Figure 2.14 The phase diagram of the q -state clock model in two dimensions. The black points indicate critical points associated with second order phase transitions, and red points indicate critical temperatures associated with BKT phase transitions.

spin correlation function decays as a power of the distance. There is no true long-range order in this phase, however, the system is quasi-ordered and on a finite lattice, it is magnetized in this phase. This phase is analogous to the low temperature phase of the XY model. The high temperature paramagnetic phase (i.e. $\beta < \beta_{c1}$) is disordered and correlations decay exponentially. This phase is analogous to the high temperature phase of the XY model.

The phase diagram of the q -state clock model is shown in Figure 2.14. The critical points and some critical exponents are tabulated in Table 2.2.

Table 2.2 Here we show (for reference) the critical exponents for some q -state clock models. For $q = 2$, this is the Ising model. For $q = 3$, the clock model is in the same universality class as the 3-state Potts model. For $q = 4$, it has been shown [58] that the clock model is in the Ising universality class. The final column gives the exact infinite-volume critical point for these three cases. The Ising and 3-state Potts models are two well-studied universality classes, and their exact critical exponents can be found in textbooks such as [5]. The critical point of the Ising Model can similarly be found in textbooks such as [5]. The exact critical points for $q = 3, 4$ and the estimated critical points for the $q = 5, 6$ clock models is taken from [51] (Supplemental Material). The estimated critical point for the XY model (i.e. $q = \infty$) is taken from [39].

q	α	β	γ	δ	η	ν	β_c
2	0	1/8	7/4	15	1/4	1	$\frac{1}{2} \ln(1 + \sqrt{2})$
3	1/3	1/9	13/9	14	4/15	5/6	$\frac{2}{3} \ln(1 + \sqrt{3})$
4	0	1/8	7/4	15	1/4	1	$\ln(1 + \sqrt{2})$
5							1.0503(2), 1.1039(2)
6							1.0957(2), 1.4491(8)
$\vdots$	$\vdots$					$\vdots$	$\vdots$
∞							1.11995(6)

CHAPTER 3

MCMC SIMULATIONS AND THEIR ANALYSIS

3.1 Markov Chain Monte Carlo

The canonical ensemble in statistical mechanics is the ensemble that describes the distribution of the possible (micro)states of a mechanical system in thermal equilibrium at a fixed temperature, a fixed volume, and with a fixed number of degrees of freedom. The internal energy E of such a system can vary with the microstate. The probability $P(U)$ for the system to be in a microstate U with energy $E(U)$ is given by the Boltzmann distribution

$$P(U) = \frac{e^{-\beta E(U)}}{Z}, \quad (3.1)$$

where $\beta = 1/(k_B T)$ is the inverse temperature with k_B being Boltzmann's constant and T being the temperature. The partition function of the system has the form

$$Z = \sum_U e^{-\beta E(U)}, \quad (3.2)$$

where the sum is over all possible states of the system. Then the equilibrium expectation value of an observable O is

$$\langle O \rangle = \sum_U O(U) P(U) = \frac{\sum_U O(U) e^{-\beta E(U)}}{\sum_U e^{-\beta E(U)}}. \quad (3.3)$$

More on the canonical ensemble can be found in standard textbooks such as [59].

In principle, one can compute the expectation value $\langle O \rangle$ by enumerating all possible microstates of the system, computing the internal energy $E(U)$ and the observable $O(U)$ for each microstate U , and then computing Eq. (3.3). In practice, this approach is infeasible for even small systems. Consider the Ising model on a 10×10 lattice. There are 100 sites each containing a spin variable which can be in one of two possible states—so the number of possible microstates is 2^{100} (which is larger than 10^{30}). Computing this using a naive approach would take the current fastest supercomputer millions of years¹.

¹With 2^{100} microstates and assuming the calculation of O_i and P_i takes 2×100 floating point operations per 100-site microstate, this is $\sim 10^{32}$ floating point operations. The current fastest supercomputer—Frontier—has an Rmax value of 1102 PetaFLOPS. So total computation time is 10^{32} FLOP / 1102 petaFLOPS ~ 2.9 million years.

Instead of enumerating all states, one could estimate Eq. (3.3) by taking the mean of a randomly selected sample of the microstates. This would be an example of a “Monte Carlo” method. However, a relatively tiny fraction of states actually contributes meaningfully to the sum, so if one were to select the microstates uniformly, one would still need to sample an infeasibly large number of states to get an accurate average. The key is to use “importance sampling” where the microstates are selected according to their underlying probability distribution given in Eq. (3.1). If N microstates are sampled according to this distribution, then an approximation of Eq. (3.3) is given by the average of the sample

$$\langle O \rangle \approx \frac{1}{N} \sum_{i=1}^N O(U_i), \quad (3.4)$$

and the statistical error of the approximation goes as $1/\sqrt{N}$, meaning that only a “small” number of microstates need to be sampled to get a good estimate. The question becomes how to sample (or more accurately how to *generate*) microstates, which in our case will be lattice configurations, according to the probability distribution Eq. (3.1). This is where we come to Markov chain Monte Carlo (MCMC).

The idea is to start with some arbitrary configuration U_0 and construct a chain of configurations, where each configuration can be generated from the previous configuration

$$U_0 \rightarrow U_1 \rightarrow U_2 \rightarrow \cdots, \quad (3.5)$$

such that this Markov chain eventually reaches an equilibrium distribution that matches our desired probability distribution Eq. (3.1). After discarding some number of the initial configurations to ensure equilibration, we can estimate our observables using Eq. (3.4). An obvious complication is that the configurations in such a chain will be correlated with previous configurations. We will discuss in Section 3.3 how to handle these autocorrelations. Another problem that can occur is that of the Markov chain becoming stuck in one sector of configuration space and failing to sample the whole space. This is analogous to a minimizer becoming stuck in a local minimum instead of finding the global minimum. An example of this is discussed in Sec. 4.3.7.

Given a configuration U_t in the Markov chain time series, some updating algorithm (several will be discussed in Section 3.2) produces the next configuration U_{t+1} . Typically, this is done by making some small change, e.g. flipping a single spin in the lattice.

Let $T(U'|U)$ be the conditional transition probability for the updating algorithm to take us from some configuration U to a configuration U' . Since it is a probability, $0 \leq T(U'|U) \leq 1$, and we have the normalization condition $\sum_{U'} T(U'|U) = 1$. We require that every possible configuration in our distribution is reachable by our Markov chain in a finite number of steps. This is the requirement of *ergodicity*. We need our Markov process to eventually reach the equilibrium distribution $P(U)$ and then stay there, i.e. we require that there be no sources or sinks of probability. A sufficient condition, known as the *detailed balance condition*, is

$$T(U'|U)P(U) = T(U|U')P(U'), \quad (3.6)$$

for every pair of configurations U and U' . For a detailed introduction to Markov chain Monte Carlo, see for example [5, 60].

3.2 Updating Algorithms

In this section, we discuss several algorithms that can be used to generate the next lattice configuration U_{t+1} from the current configuration U_t in a Markov chain. Our examples will be classical spin systems, but the same algorithms can be used for quantum field theories in the path integral representation on the lattice.

3.2.1 Metropolis

The Metropolis updating algorithm [61] is one of the simplest MCMC updating algorithms. In most modern applications it has been replaced with more efficient methods, however, it remains a popular choice due to its simplicity and generality. It is a simple-to-implement updating algorithm that works whenever one is able to compute the energy $E(U)$ of the state U .

The detailed balance equation is

$$T(U'|U)e^{-\beta E(U)} = T(U|U')e^{-\beta E(U')}. \quad (3.7)$$

We can rearrange this as

$$\frac{T(U'|U)}{T(U|U')} = e^{-\beta\Delta E}, \quad (3.8)$$

where $\Delta E = E(U') - E(U)$ is the internal energy difference between a candidate configuration U' and the old configuration U .

The Metropolis updating algorithm, in its simplest form, proceeds as follows:

1. Given the current configuration U_t , generate a candidate configuration U' by some random process that selects uniformly on the group. For example, let U' be U_t , but with a single spin flipped to a random value from a uniform distribution
2. Accept this candidate as the new configuration (i.e. let $U_{t+1} = U'$) with probability

$$T(U'|U_t) = \min\left(1, e^{-\beta\Delta E}\right). \quad (3.9)$$

Otherwise, keep the old configuration (i.e. let $U_{t+1} = U_t$)

3. Repeat these steps

It is easy to show that detailed balance is satisfied.

$$\begin{aligned} T(U'|U)e^{-\beta E(U)} &= \min\left(1, e^{-\beta\Delta E}\right) e^{-\beta E(U)} \\ &= \min\left(e^{-\beta E(U)}, \frac{e^{-\beta E(U')}}{e^{-\beta E(U)}} e^{-\beta E(U)}\right) \\ &= \min\left(\frac{e^{-\beta E(U)}}{e^{-\beta E(U')}}, 1\right) e^{-\beta E(U')} \\ &= T(U|U')e^{-\beta E(U')}. \end{aligned} \quad (3.10)$$

All of the spin systems that we have considered in this thesis have been “local” in the sense that if the spin variable at a single site is changed, then the energy change ΔE can be computed by looking only at the *stencil* composed of that spin and its nearest neighbors. As a concrete example, let us consider the Ising model with energy function Eq. (2.2). We will assume $J = 1$ and $h = 0$, so the energy $E(U)$ of a configuration is simply $-\sum_{\langle i,j \rangle} S_i S_j$ where the sum is over nearest neighbor pairs, and the spin variables S_i at a lattice site i can take the value $+1$ or -1 . The number

of terms in the sum $E(U) = -\sum_{\langle i,j \rangle} S_i S_j$ is Nd where N is the number of lattice sites and d is the number of dimensions. Any specific site appears in $2d$ of these terms, which can be grouped together and called a “stencil”. For example, in two dimensions, any given site variable S_k appears in 4 terms which we can denote: $S_k S_r + S_k S_l + S_k S_u + S_k S_d$ where, for example, S_r refers to the nearest neighbor to the right. Respectively, S_l , S_u , and S_d are the nearest neighbors to the left and to the up and down directions. Thus, when a single spin S_k at the site k is changed to S'_k in the two-dimensional Ising model with no external field, then the energy change is

$$\Delta E = E(U') - E(U) = (S_k - S'_k)(S_r + S_l + S_u + S_d). \quad (3.11)$$

All other terms cancel out. For the other classical spin models we’ve considered, ΔE can be computed similarly. The only change is the energy function $E(U)$. Keep in mind, that the number of terms in the stencil depends on the number of dimensions one is operating in.

A Metropolis “sweep” occurs when all the sites in the lattice have (on average) had a chance of being updated once. The order in which lattice sites are visited and updated is a matter of choice—a common choice is to sequentially update all of the sites. Pseudocode for a sequential Metropolis sweep for the Ising model is as follows:

```

for SITE in LATTICE do

    # Get the current value of the spin variable
    OLDSPIN = get_site_variable(SITE)

    # Get a random candidate spin variable
    NEWSPIN = random_spin_value()

    # Compute the change in energy if NEWSPIN were accepted
    DELTAE = 0
    for DIM in NDIMENSIONS do # Computing stencil
        DELTAE += get_site_variable(SITE + 1 unit in forward dir. in this DIM)
        DELTAE += get_site_variable(SITE + 1 unit in backward dir. in this DIM)
    DELTAE *= (OLDSPIN - NEWSPIN)

    # Accept or reject the candidate
    R = random_float[0,1)
    if R <= exp(-BETA*DELTAE) then # Accepting the candidate
        set_site_variable(SITE) = NEWSPIN

```

```
# If rejected, the site variable remains OLDSPIN
```

For other classical spin models one would only need to change the ΔE calculation.

To sample an equilibrium distribution, an MCMC updating algorithm samples states around the minimum of the energy function. In the Metropolis updating scheme, a candidate for the next configuration in the Markov chain is generated by randomly flipping a spin in the current configuration. If the candidate has a lower internal energy, then it is accepted unconditionally. Otherwise, it is accepted probabilistically. This approach makes only small changes at a time—a single spin flip per updating step—resulting in a rather slow advancement through the space of possible configurations.

Depending on the spin system and the temperature regime, the Metropolis acceptance rate can become very low. Whenever a candidate is rejected, the current configuration is repeated in the Markov chain. Hence, a low acceptance rate implies a long autocorrelation time, and this means the MCMC simulation must be run for a longer time in order to get a statistically-independent sample of configurations. In the next section, we look at an updating algorithm that always has a 100% acceptance rate.

3.2.2 Heatbath

The heatbath algorithm effectively combines steps 1 and 2 of the Metropolis algorithm to more directly produce configurations distributed according to the desired probability distribution. Instead of randomly flipping a spin and seeing if the energy is reduced, the heatbath algorithm looks at all possible values for that spin and flips it probabilistically to the value that locally (considering it and its nearest neighbors) minimizes the energy. The result is a much faster advancement through state space, that is, one gets a Markov chain with a much smaller autocorrelation time, which means one can generate the same effective statistics with a much shorter simulation.

The heatbath updating algorithm can be used when it's possible to efficiently invert the cumulative distribution function (CDF). In some cases, inverting the CDF may be unacceptably slow or it may be *a priori* unknown, and then one has to rely on other updating algorithms.

For an equilibrium MCMC simulation, the probability to be in the configuration U is given

by the probability density function (PDF) $P(U) = C' e^{-\beta E(U)}$ where β is the inverse temperature, $E(U)$ is the internal energy of configuration U , and C' is determined by the normalization condition $\sum_U P(U) = 1$. The normalization constant C' is the reciprocal of the partition function, but in practice, this quantity is not known during an MCMC simulation.

During an updating sweep, one site at a time is updated, so we can consider things at the level of an individual site rather than an entire configuration. Suppose the spin system being considered has site variables which can be in q discrete states labeled by k . For example, for the Ising model, $q = 2$ with site variable $S_k = \pm 1$. For the q -state Potts models, the site variable takes the values $S_k \in \{0, 1, 2, \dots, q-1\}$. For the q -state clock models in two dimensions, one can parameterize the vector site variables by a single angle $\theta_k = 2\pi k/q$ with $k \in \{0, 1, 2, \dots, q-1\}$. For an equilibrium configuration, the probability for a particular spin variable to be in the state k is

$$P(k) = C e^{-\beta E_k}, \quad (3.12)$$

where the constant C is determined by the normalization condition

$$\sum_{k=0}^{q-1} P(k) = 1 \implies C = \frac{1}{\sum_{k=0}^{q-1} e^{-\beta E_k}}. \quad (3.13)$$

The energy E_k is the interaction energy of the particular spin variable with its nearest neighbors. For example, for the Ising (see Eq. (2.2)), Potts (see Eq. (2.20)), and clock models (see Eq. (2.47)), we have

$$E_k^{Ising} = -J \sum_{\mu=1}^{ndims} \sum_{j=\pm} S_k S_{\mu_j} \quad (3.14)$$

$$E_k^{Potts} = -2 \sum_{\mu=1}^{ndims} \sum_{j=\pm} \delta_{S_k, S_{\mu_j}} \quad (3.15)$$

$$E_k^{Clock} = - \sum_{\mu=1}^{ndims} \sum_{j=\pm} \vec{S}_k \cdot \vec{S}_{\mu_j} = - \sum_{\mu=1}^{ndims} \sum_{j=\pm} \cos(\theta_k - \theta_{\mu_j}), \quad (3.16)$$

where the first sum is over the number of dimensions, the second sum is over the forward and backward directions, and S_{μ_j} is the nearest neighbor site variable in the μ_j direction. The other spin variables contribute of course to the total energy of the lattice, but when we are considering a

single spin, the contributions from the rest of the lattice can be absorbed into the overall constant C . The CDF associated with Eq. (3.12) is

$$z = F(k) = \sum_{j=0}^k P(j). \quad (3.17)$$

The normalization condition implies that

$$F(q - 1) = 1. \quad (3.18)$$

The heatbath update selects k distributed according to $P(k)$ by converting a uniformly distributed random number $0 \leq z < 1$ into

$$k = F^{-1}(z). \quad (3.19)$$

This is the value that our particular site variable should be set to.

To perform a sequential heatbath sweep, loop over all sites in the lattice. For each site:

1. Loop over the q possible values of the spin variable, and for each one, calculate what would be the local contribution E_k to the energy given the fixed values of the neighboring spins
2. Calculate the CDF Eq. (3.17). It can be normalized using the condition Eq. (3.18)
3. Choose a random number z uniformly in $[0, 1)$
4. Set the value S_k of the spin variable using the smallest k for which $F(k) > z$

With this procedure, one is setting the spin value to S_k with probability $P(k)$ by effectively inverting the CDF $F(k)$.

Here is pseudocode for performing a sequential heatbath sweep:

```
# Cumulative distribution 'function'
CUMULDIST[-1] = 0

# Loop over all sites in the lattice
for SITE in LATTICE do

    # For this site, calculate the energy contribution
    # for each possible value of the site variable
```

```

for K = 0 to Q-1:
    EK = 0

    # Computing stencil
    for DIM in NDIMENSIONS do
        EK += ...MODEL DEPENDENT CALCULATION...

    # Calculate the cumulative distribution function
    CUMULDIST[K] = CUMULDIST[K-1] + exp(-BETA*EK)

# Set the value of the spin variable according to the probability distribution
R = random_float[0,1)
for K = 0 to Q-2:
    P = CUMULDIST[K]/CUMULDIST[Q-1] # Normalized cum. dist.
    if R < P:
        set_site_variable(SITE) = set site variable to value K

# Otherwise
if(not yet set)
    set_site_variable(SITE) = set site variable to value Q-1

```

In the heatbath algorithm, for each site update, one has to compute the CDF as a function of the neighboring sites. For a lattice with N sites, this procedure must be done N times per “sweep”. The new site variable is then chosen by effectively inverting the CDF. For discrete models such as the Ising, Potts, and clock models—whose site variables take only q discrete values—one can efficiently compute and store the CDF. For a continuous model, such as the XY model—whose site variables take continuous values $0 \leq \theta < 2\pi$ —it is not possible to compute the exact CDF. It can still be approximated to any desired precision by discretizing it, however, this approach is not efficient since it would have to be repeated N times per heatbath sweep.

3.2.3 Biased Metropolis Heatbath Algorithm (BMHA)

The Biased Metropolis Heatbath Algorithm (BMHA) [62–65] is the solution to the problem of applying the heatbath algorithm to models with continuous degrees of freedom. Instead of calculating a CDF specific to each site update, one makes a discrete approximation of the general CDF and stores it as a lookup table. To keep memory access efficient, this must be a rather poor approximation of the CDF, however, the discretization error can be completely removed by adding a Metropolis accept/reject step. In a very real sense, the BMHA interpolates between the Metropolis

and heatbath algorithms. Improving the CDF approximation in BMHA improves the acceptance rate but also increases the table size which slows the lookup procedures. There is a balance to be achieved. In practice, one can achieve efficient BMHA algorithms with acceptance rates $\sim 90\%$ even in regimes where Metropolis has very low acceptance.

At its simplest, we want to sample random variables $x \in [x_1, x_2]$ from some PDF $P(x)$ with associated CDF $F(x)$. We do this via a Markov chain where each x_{new} is related to the previously sampled random variable x_{old} . The Metropolis algorithm does this as follows:

1. Uniformly generate $x_{new} \in [x_1, x_2]$

2. Accept with probability

$$P_{metr} = \min \left(1, \frac{P(x_{new})}{P(x_{old})} \right). \quad (3.20)$$

In contrast, the BMHA proceeds as follows (we use the notation of [65]):

1. Discretize x into $nvar$ bins: $x_1 = x^0 < x^1 < \dots < x^{nvar} = x_2$, where the bin widths are $\Delta x^j = x^j - x^{j-1}$ with $j = 1, \dots, nvar$. The Δx^j are chosen nonuniformly such that $x^j = F^{-1}(j/nvar)$

2. Propose x_{new} by uniformly picking a bin j_{new} and proposing x_{new} uniformly within that bin

3. Locate the bin j_{old} to which x_{old} belongs

4. Accept x_{new} with probability

$$P_{bmha} = \min \left(1, \frac{P(x_{new})}{P(x_{old})} \frac{\Delta x^{j_{new}}}{\Delta x^{j_{old}}} \right). \quad (3.21)$$

In the limit $nvar \rightarrow \infty$, $P_{bmha} \rightarrow 1$, i.e. we recover the 100% acceptance rate of the heatbath algorithm. Crucially, this modification of the Metropolis algorithm influences only the acceptance rate—the same detailed balance conditions are satisfied as by the Metropolis algorithm. Note that in an MCMC simulation Step 1. is done only once at the start of the simulation and saved as a lookup table.

In a slightly more complicated example, suppose we want to sample random variables $x \in [x_1, x_2]$ from a PDF $P(x; \alpha)$ that includes a parameter $\alpha \in [\alpha_1, \alpha_2]$, which depends on the environment of the site being updated. In that case, one can uniformly discretize the parameter $\alpha_1 \leq \alpha \leq \alpha_2$ with $npar$ bins. Within each of these bins, the variable $x_1 \leq x \leq x_2$ is discretized into $nvar$ bins as in Step 1. of the BMHA. Each parameter bin will contain a different discretization of the variable x , since the bin widths for the variable depend on the inverse CDF which itself depends on the parameter α through the PDF $P(x; \alpha)$.

To make this more concrete, we consider the $O(2)$ or XY model as a specific example. This is a “local” Hamiltonian, so we can factor any particular site $\vec{S}_i$ out of the interaction term. For example, for the $O(2)$ model on a lattice with N sites and no external field,

$$H = - \sum_{\langle i, j \rangle} \vec{S}_i \cdot \vec{S}_j = - \sum_{i=1}^N \vec{S}_i \cdot \sum_{\mu=1}^{ndim} \vec{S}_{i+\hat{\mu}} = - \frac{1}{2} \sum_{i=1}^N \vec{S}_i \cdot \vec{S}_{i,+}, \quad (3.22)$$

where

$$\vec{S}_{i,+} = \sum_{\mu=1}^{ndim} \left(\vec{S}_{i-\hat{\mu}} + \vec{S}_{i+\hat{\mu}} \right), \quad (3.23)$$

is the “stencil” (i.e. the sum of the nearest neighbors) at site i . The site variable $\vec{S}_i$ is a unit vector, however, the stencil is not since in d dimensions, it is the sum of $2d$ unit vectors. For the Metropolis accept/reject step, we compute the ratio of probabilities

$$\frac{P(x_{new})}{P(x_{old})} = \frac{e^{-\beta E_{new}}}{e^{-\beta E_{old}}} = e^{-\beta(E_{new} - E_{old})} = e^{-\beta \Delta E}, \quad (3.24)$$

where $\Delta E \equiv E_{new} - E_{old}$, and so we only care about the energy *difference* when a single site variable $\vec{S}_i$ is changed to $\vec{S}'_i$

$$\Delta E = - \sum_{\mu} \left(\vec{S}'_i \cdot \vec{S}_{i+\hat{\mu}} + \vec{S}'_i \cdot \vec{S}_{i-\hat{\mu}} \right) + \sum_{\mu} \left(\vec{S}_i \cdot \vec{S}_{i+\hat{\mu}} + \vec{S}_i \cdot \vec{S}_{i-\hat{\mu}} \right) = - \left(\vec{S}'_i - \vec{S}_i \right) \cdot \vec{S}_{i,+}. \quad (3.25)$$

So the ratio of probabilities can be written as $e^{\beta \vec{S}'_i \cdot \vec{S}_{i,+}} / e^{\beta \vec{S}_i \cdot \vec{S}_{i,+}}$, which implies the probability distribution to be sampled is $P(\vec{S}) = e^{\beta \vec{S} \cdot \vec{S}_+}$. We can parameterize the site vector with an angle θ ,

so the PDF² can be written as

$$P(\theta; \alpha) = e^{\alpha \cos(\theta - \theta_+)}, \quad \alpha \equiv \beta |\vec{S}_+|. \quad (3.26)$$

where θ is the random variable and θ_+ is the angle made by the stencil vector. Note that α and θ_+ are parameters which depend on the environment/stencil of the site, and β is fixed for a given simulation. In practice, we do not need to treat θ_+ as a parameter since we can define $\tilde{\theta} = \theta - \theta_+$, sample the simpler PDF $P(\tilde{\theta}; \alpha) = C e^{\alpha \cos(\tilde{\theta})}$, and then calculate $\theta = \tilde{\theta} + \theta_+$. Since θ and θ_+ both span $[0, 2\pi)$, we want $-2\pi \leq \tilde{\theta} < 2\pi$. Note that the stencil $\vec{S}_+$ is a sum of unit vectors, so in d -dimensions, $0 \leq |\vec{S}_+| \leq 2d$. This implies that the parameter α lies in $0 \leq \alpha \leq 2\beta d$. The CDF associated with Eq. (3.26) is

$$z = F(\tilde{\theta}; \alpha) = \int_{-2\pi}^{\tilde{\theta}} P(\theta'; \alpha) d\theta' \quad (3.27)$$

We can simulate the $O(2)$ model with “1 variable / 1 parameter” BMHA. The two-dimensional lookup table is constructed by discretizing the parameter $0 \leq \alpha \leq 2\beta d$ into $npar$ bins with uniform bin widths $\Delta\alpha^i = 2\pi/npar$. Within each of those bins, the variable $-2\pi \leq \tilde{\theta} \leq 2\pi$ is discretized into $nvar$ bins with bin widths $\Delta x^j = x^j - x^{j-1}$ with $j = 1, \dots, nvar$ and the bin widths Δx^j chosen nonuniformly such that $x^j = F^{-1}(z = j/nvar)$. An updating sweep of the lattice would proceed as follows: Loop sequentially over all lattice sites. For each site,

1. Get the angle θ_{old} associated with the current site variable
2. Compute the stencil magnitude $|\vec{S}_+|$ and angle θ_+ and $\alpha = \beta |\vec{S}_+|$
3. The shifted angle is $\tilde{\theta}_{old} \equiv \theta_{old} - \theta_+$
4. Get the bin i to which α belongs
5. Within that α bin, locate the bin j_{old} to which $\tilde{\theta}_{old}$ belongs

²This is an unnormalized PDF, which is fine when computing the ratio of probabilities. However, when using the PDF to construct the CDF, one must normalize it via numerical integration.

6. Still within that α bin, propose $\tilde{\theta}_{new}$ by uniformly picking a $\tilde{\theta}$ bin j_{new} and proposing $\tilde{\theta}_{new}$ uniformly within that bin

7. Accept $\tilde{\theta}_{new}$ with probability

$$P_{bmha} = \min \left(1, \frac{e^{\alpha \cos(\tilde{\theta}_{new})} \Delta \tilde{\theta}^{j_{new}}}{e^{\alpha \cos(\tilde{\theta}_{old})} \Delta \tilde{\theta}^{j_{old}}} \right). \quad (3.28)$$

Otherwise, set $\tilde{\theta}_{new} = \tilde{\theta}_{old}$

8. Unshift the angle $\theta_{new} = \tilde{\theta}_{new} + \theta_+$ with $\theta_{new} \in [0, 2\pi)$

9. Update the site variable so that it points in the direction θ_{new}

The BMHA can be generalized to a PDF $P(x, y, z, \dots; \alpha, \beta, \gamma, \dots)$ with N_{var} random variables $x, y, z, \dots$ and N_{par} parameters $\alpha, \beta, \gamma, \dots$. Increasing the number of parameters does not significantly complicate things since one just uniformly discretizes the space of the parameters and within each resulting “cell”, discretizes the PDF. So a “1 variable / N_{par} parameter” PDF would result in a $(N_{par} + 1)$ -dimensional lookup table. Increasing the number of variables is more complicated. One approach would be to uniformly discretize each dimension of the parameter space so that one has an N_{par} -dimensional grid. Then within each of these cells, discretize the N_{var} -dimensional variable space into “volumes” of equal probability. For more details, see Section 3 of [65].

To perform BMHA updating for the general $O(n)$ model with total energy function $H = -\sum_{\langle i,j \rangle} \vec{S}_i \cdot \vec{S}_j$, where $\vec{S}_i$ and $\vec{S}_j$ are n -dimensional unit vectors, the PDF to be sampled is $P(\{\theta_i\}; \{\alpha_i\}) = e^{\beta \vec{S} \cdot \vec{S}_+}$. Again, $\vec{S}$ is the random variable and $\vec{S}_+$ is the stencil vector fixed by the environment of $\vec{S}$. When expanded in the trigonometric functions of generalized spherical coordinates, the dot product $\vec{S} \cdot \vec{S}_+$ contains n terms containing a total of $n - 1$ angles (these will be the random variables) $\theta_1, \dots, \theta_{n-1}$ and n parameters $\alpha_1, \dots, \alpha_n$. However, one can use the trigonometric identity $\cos(a - b) = \cos a \cos b + \sin a \sin b$ to eliminate one of the parameters in return for shifting one of the angles. Thus for the $O(n)$ model, we need “ $n - 1$ variable / $n - 1$ parameter” BMHA. For example, for the $O(3)$ model, one can write the

PDF as $P(\theta_1, \theta_2; \alpha_1, \alpha_2) = \exp(\alpha_1 \cos \theta_1 + \alpha_2 \sin \theta_1 \cos(\theta_2 - \theta_{+,2}))$, where $\alpha_1 = \beta |\vec{S}_+| \cos \theta_{+,1}$, $\alpha_2 = \beta |\vec{S}_+| \sin \theta_{+,1}$, and $\theta_{+,i}$ is the i th component of the stencil vector $\vec{S}_+$.

3.3 Autocorrelations

In a Markov chain, where configurations are separated by a (computer) time step, consecutive configurations tend to be correlated with each other. This “autocorrelation” of lattice configurations results in the autocorrelation of any observables calculated on those configurations. In other words, observables calculated in MCMC are generally not independent, and this complicates their statistical analysis. A good discussion of autocorrelation is given in [5] and a briefer discussion is given in [66].

For an observable X

$$C_X(X_i, X_{i+t}) = \langle X_i X_{i+t} \rangle - \langle X_i \rangle \langle X_{i+t} \rangle, \quad (3.29)$$

is the correlation function between the observable X measured at Markov times i and $i + t$. In equilibrium, this autocorrelation function depends only on the time separation t , and one writes $C_X(t)$ instead of $C_X(X_i, X_{i+t})$. Typically, the normalized correlation function $c_X(t)$ has exponential fall-off with t

$$c_X(t) \equiv \frac{C_X(t)}{C_X(0)} \sim e^{-t/\tau_{X,exp}}, \quad (3.30)$$

where $\tau_{X,exp}$ is the *exponential* autocorrelation time for the observable X . In general, the autocorrelation time depends on the observable X , the lattice volume, and the updating algorithm.

An estimator for the *integrated* autocorrelation time for an observable X is given by

$$\tilde{\tau}_{X,int} = 1 + 2 \sum_{t=1}^T c_X(t). \quad (3.31)$$

The integrated autocorrelation time $\tau_{X,int}$, which is a measure of the self-correlation between consecutive configurations in a Markov chain, is estimated by finding a window in t for which $\tilde{\tau}_{X,int}$ is nearly independent of t . In other words, one looks for a plateau in a plot of $\tilde{\tau}_{X,int}$ versus t , and the value of $\tilde{\tau}_{X,int}$ at the plateau is taken to be the estimate $\tau_{X,int}$. See Figure 3.1. In practice, it may be hard to find a plateau. One often has to do a controlled search for the plateau by increasing T and the statistics (number of sweeps) used in the simulation. See Fig. (3.2). The integrated




Figure 3.1 Plots of the integrated autocorrelation curve of the energy for the two-dimensional Ising model at four different lattice sizes. Each simulation is performed at the critical point β_c using a heatbath algorithm. The curves are computed using Eq. (3.31), and the plateau value gives the estimate of the integrated autocorrelation time τ_{int} . Notice that τ_{int} grows with lattice volume, and to obtain the plateau value, one must go to much larger values of t as the lattice size is increased.

autocorrelation time is observable dependent and there can be orders of magnitude difference between the integrated autocorrelation times of different observables X . In general, it also depends on the lattice size and the updating algorithm.

In plots of $\tilde{\tau}_{X,int}$ versus t as in Figure 3.2, the error bars at nearby values of t are correlated and roughly the same size. To estimate $\tau_{X,int}$, it usually suffices to take the maximum value and its error bar. One should always check such plots by eye to make sure that one is getting the maximum. For example, in the left panel of Figure 3.2, one is clearly not obtaining the maximum since the curve is increasing beyond the extent of the plot.

In practice, analyzing the autocorrelations in a time series can be a subtle business. For example, in systems with explicitly broken symmetry as we will see in the next chapter, there can be very “long wavelength” autocorrelations coming from magnetization flips (when the entire lattice flips magnetization) as well as ordinary autocorrelations coming from individual spin flips. If the Markov chain is too short, an autocorrelation analysis may completely miss the long wavelength autocorrelations. The safest approach is to start with very small lattices and very large statistics, and slowly work up to larger lattices.

The effect of autocorrelations is to decrease the number of statistically-independent measure-




Figure 3.2 Here we show the integrated autocorrelation curve of the energy for the Ising model on a 16×16 lattice near the critical point calculated from a time series of 2^{16} heatbath sweeps with one measurement per sweep. The integrated autocorrelations are estimated using Eq. (3.31), and we are trying to find the maximum/plateau value. In the left panel, the time extent $T = 32$ is clearly insufficient to resolve the plateau. In the middle panel, we increased T and located a peak which gives an estimate of the integrated autocorrelation time. However, in this case we have large error bars and don't have a nice plateau, which indicates that our data was generated with insufficient statistics to accurately resolve the integrated autocorrelation time. In the final panel, we increased the statistics of the simulation by a factor of 256 to 2^{24} heatbath sweeps, and we increased T to 512. Now we see a nice plateau with small error bars. In many practical examples, it would be difficult to get sufficient statistics to resolve the plateau this well and one might be forced to estimate the integrated autocorrelation time from plots more like the middle one.

ments in one's sample. If autocorrelation is not properly accounted for, the naively calculated statistical error bars will be improperly small. For a sample $\{X_1, X_2, \dots, X_N\}$ of N autocorrelated measurements, the number of effectively independent measurements is

$$N_{indep} = \frac{N}{\tau_{X,int}}. \quad (3.32)$$

For uncorrelated measurements, the variance of the mean is $\sigma_{\bar{X}}^2 = \sigma_X^2/N$. The variance of the mean, corrected for autocorrelations, is

$$\sigma_{\bar{X},corr}^2 = \tau_{X,int} \sigma_X^2 = \tau_{X,int} \frac{\sigma_X^2}{N}. \quad (3.33)$$

Thus, the estimate that one would quote for the sample of correlated measurements $\{X_1, X_2, \dots, X_N\}$ is

$$\hat{X} \pm \sigma_{\bar{X},corr}, \quad (3.34)$$

where $\hat{X}$ is an estimator of the mean.




Figure 3.3 The integrated autocorrelation time (of the energy here) blows up at a phase transition as shown here for the XY model using a Biased Metropolis Heatbath Algorithm on lattices of various size. This is the cause of the *critical slowing down* that plagues MCMC simulations near critical points. The high autocorrelation time means many configurations must be discarded between measurements in order to get an independent sample, and so the overall computational time required gets very large at the critical point.

According to [66], for equilibration, one should discard at least $20\tau_{X,int}$ configurations. To end up with a sample of independent measurements, one should discard a similar quantity of configurations between consecutive measurements. For good statistics, increase this to $1000\tau_{X,int}$, and for statistics at the 1% level, increase it to $10,000\tau_{X,int}$. In practice, however, for computationally intensive simulations as done in lattice QCD, one has to be satisfied with $\sim 1\tau_{X,int}$ configurations discarded between measurements.

For a system with a phase transition, $\tau_{X,int}$ grows large at the critical point. Furthermore, this growth is volume dependent so $\tau_{X,int}$ blows up doubly as the critical point is approached and the lattice volume is increased. See Figure 3.3, to see how $\tau_{X,int}$ varies with β and lattice size as one approaches the critical point in the XY model. At the critical point, the integrated autocorrelation time increases as [5, 66]

$$\tau_{X,int} \propto L^z, \quad (3.35)$$

where L is the lattice extent in each direction, and z is a dynamical critical exponent. This increase as a power of the lattice volume as the critical point is neared is the cause of the critical slowing down that plagues MCMC studies. Keep in mind that some number of configurations ($> \tau_{X,int}$)




Figure 3.4 The integrated autocorrelation time of the energy is calculated using the heatbath algorithm for the two-dimensional Ising model at the critical point on four different lattice sizes. See Figure 3.1. The results are fitted to Eq. (3.35) to get an estimate of the dynamical critical exponent z .

must be discarded between saved configurations to mitigate the autocorrelations. This means as the critical point is approached, the MCMC simulation must be run for longer and longer to obtain equivalent statistics. For the two-dimensional Ising model with a heatbath algorithm, we find

$$\tau_{E,int} \approx 0.1733L^{1.87}, \quad (3.36)$$

with dynamical critical exponent $z = 1.873 \pm 0.015$. See Figure 3.4.

It must be noted that computing the integrated autocorrelation time is expensive. Typically, it requires orders of magnitude more statistics to precisely estimate $\tau_{X,int}$, for some observable X , than it takes to get a seemingly precise estimate of X itself. Nevertheless, if one wants a reliable estimate of some observable X , one must also have some understanding of its integrated autocorrelation time $\tau_{X,int}$.

One obvious way to mitigate autocorrelations is to throw away most of the data. For example, if $\tau_{int} = 10$ for some time series, then saving only every 10th data point should in principle remove the autocorrelations. Unfortunately, since the sample size is now 10 times smaller, the error bar of the sample mean is now $\sqrt{10}$ times larger. Another approach is to do blocking and preaveraging of the data. Chunking the time series into blocks larger than τ_{int} and taking the average of each block produces a new (much shorter) time series which in principle is free of autocorrelations. See




Figure 3.5 Here we illustrate how blocking and preaveraging can be used to mitigate autocorrelations. We construct a stochastic time series containing autocorrelation. This is done by choosing 100K uniformly random numbers $\{r\}$ in $[-0.5, 0.5)$. The n th element a_n of the autocorrelated time series (plotted in top left) is $a_n = r + a_{n-1}$ where r is a random number, however, if n is a multiple of 100, then $a_n = r$. The center left panel shows the autocorrelation function $c(t)$ of the series, and the bottom left panel shows that the integrated autocorrelation time is ~ 65 . Next, we do blocking and preaveraging with block size 100. The resulting shorter time series is shown in the top right. The associated correlation function (center right) and τ_{int} (bottom right), shows that the preaveraged time series is not autocorrelated.

Figure 3.5. Although this new sample is smaller, its variance is also smaller (since each number is already an average), and so the error bar does not increase by as much as if one simply threw away most of the data. The better approach, discussed in the next section, is to jackknife with blocking.

3.4 Jackknife Error Analysis

In this section we follow the notation and approach of [5]. Suppose we have a sample of N estimates of a quantity x

$$x_1, x_2, x_3, \dots, x_N. \quad (3.37)$$

This sample may or may not be an ordered time series. Assuming the x_i are uncorrelated, a correct (i.e. unbiased) estimator of the expectation value of x is the mean

$$\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i, \quad (3.38)$$

and a correct estimator of the variance of this sample mean is

$$s_{\bar{x}}^2 = \frac{1}{N(N-1)} \sum_{i=1}^N (x_i - \bar{x})^2. \quad (3.39)$$

Then our best estimate of x is $\bar{x}$ and our uncertainty is the standard deviation of the mean $s_{\bar{x}}$, and so we would quote the result

$$\bar{x} \pm s_{\bar{x}}. \quad (3.40)$$

If we are interested in a function $f(x)$, as we often are, we find that a correct estimator is $\bar{f} = f(\bar{x})$, i.e. f calculated at the mean $\bar{x}$. To calculate the uncertainty associated with $\bar{f}$, the conventional approach is to use “error propagation” formulae to deduce the uncertainty from the uncertainty of $\bar{x}$. However, this approach becomes difficult and even unreliable when $f(x)$ is a nonlinear function of x , and so we need a different approach.

The jackknife approach to error analysis is easy to implement and robust. See textbooks like [5] for a complete introduction. We define jackknife estimators of $f(x)$ as $f_i^J \equiv f(x_i^J)$ where the jackknife variable is

$$x_i^J = \frac{1}{N-1} \sum_{k \neq i} x_k, \quad (3.41)$$

that is, x_i^J is the mean of the sample when the i th data point is excluded. Then the jackknife mean is

$$\bar{f}^J = \frac{1}{N} \sum_{i=1}^N f_i^J, \quad (3.42)$$

Note that the superscript J is not an index nor a power. Rather, it is just a label to indicate that the quantity is from jackknife. An estimator of the variance of $\bar{f}^J$ is

$$s_{\bar{f}^J}^2 = \frac{N-1}{N} \sum_{i=1}^N (f_i^J - \bar{f}^J)^2. \quad (3.43)$$

Then the estimate that one would quote is

$$\bar{f}^J \pm s_{\bar{f}^J}. \quad (3.44)$$

An example, let us consider a jackknife computation of the magnetic susceptibility

$$\chi = \frac{1}{V} \left(\langle M^2 \rangle - \langle M \rangle^2 \right), \quad (3.45)$$

where we use V here to indicate the number of lattice sites since we are using N in this section to denote the length of the time series. We can think of $\chi(M, M^2)$ as a two-variable function of M and M^2 . Assume that a MCMC simulation has given us a time series $M_1, M_2, \dots, M_N$ of the magnetization measured on N uncorrelated lattice configurations. We create a second time series by squaring each of these to get $M_1^2, M_2^2, \dots, M_N^2$. We define the jackknife variables

$$M_i^J = \frac{1}{N-1} \sum_{k \neq i} M_k, \quad (M^2)_i^J = \frac{1}{N-1} \sum_{k \neq i} M_k^2 \quad (3.46)$$

Then our estimator for the susceptibility is the jackknife mean

$$\bar{\chi}^J = \frac{1}{N} \sum_{i=1}^N \chi_i^J, \quad \chi_i^J \equiv \frac{1}{V} \left((M^2)_i^J - (M_i^J)^2 \right), \quad (3.47)$$

and our estimator for the variance of $\bar{\chi}^J$ is

$$s_{\bar{\chi}^J}^2 = \frac{N-1}{N} \sum_{i=1}^N \left(\chi_i^J - \bar{\chi}^J \right)^2. \quad (3.48)$$

So far we've assumed that the sample data are uncorrelated. However, we know that time series data from an MCMC simulation will be autocorrelated. The simplest way to handle this is to remove the autocorrelations by blocking and pre-averaging the data, and then applying the jackknife machinery to the (de-correlated) pre-averaged data. Suppose we have a time series of N data points $x_1, x_2, x_3, \dots, x_N$ with an integrated autocorrelation time of $\tau_{x,int}$. We chunk the time series into N_b blocks, labeled $j = 1, 2, \dots, N_b$, each containing N/N_b of the original data points. To mitigate the autocorrelations, the block size should be $\gg \tau_{x,int}$ (if $\tau_{x,int}$ is unknown, we can vary the block size, study the error versus block size, and extrapolate to infinite block size in the end). We form a new time series $y_1, y_2, \dots, y_{N_b}$, where y_j is the average computed from block j .

This new time series forms a sample of nearly statistically independent data points and becomes the input for our jackknife procedure.

The more robust and general approach is to do jackknife with blocking. Instead of blocking and then doing jackknife on a shorter pre-averaged time series, one integrates blocking into the jackknife procedure. This approach can be used to handle autocorrelations (by choosing a block size $\gg \tau_{x,int}$, but also to handle the cases where the length N of the original time series is very large and the function being computed is too costly to be repeated N times on the full time series. The difference from the jackknife procedure detailed in Eq. (3.42)–(3.44), is simply that one now excludes entire blocks in Eq. (3.41) rather than single data points. Once again, suppose our original time series is $x_1, x_2, x_3, \dots, x_N$. We choose some number N_b of blocks to chunk the time series into. To maintain “Gaussian-ness”, one should use $N_b > 30$. To mitigate autocorrelations, one should choose N_b such that the block size is much larger than the integrated autocorrelation time, i.e. $N/N_b \gg \tau_{int}$. I assume here that N is a multiple of N_b . We label the blocks as $\mathcal{B}_i$ with $i = 1, 2, \dots, N_b$. This time, we do not pre-average. The main difference is that we define the i th jackknife variable now as

$$x_i^J = \frac{1}{N - N/N_b} \sum_{k \text{ for } x_k \notin \mathcal{B}_i}^{N - N/N_b} x_k. \quad (3.49)$$

where the sum is over all data points not in block $\mathcal{B}_i$, i.e. x_i^J is simply the mean of the time series when block $\mathcal{B}_i$ is excluded. The jackknife estimators of $f(x)$ are now $f_i^J \equiv f(x_i^J)$, the jackknife mean is

$$\bar{f}^J = \frac{1}{N_b} \sum_{i=1}^{N_b} f_i^J, \quad (3.50)$$

and an estimator of the variance of $\bar{f}^J$ is

$$s_{\bar{f}^J}^2 = \frac{N_b - 1}{N_b} \sum_{i=1}^{N_b} (f_i^J - \bar{f}^J)^2. \quad (3.51)$$

Superficially, this may seem equivalent to pre-averaging followed by the jackknife procedure Eq. (3.42)–(3.44), and indeed it is when $\bar{f}$ is a simple function of $\bar{x}$. More commonly, however, we have to deal with some multivariable nonlinear mapping where $\bar{f}$ depends not only on $\bar{x}, \bar{y}, \dots$, but also on the original data points $\{x_i\}, \{y_i\}, \dots$. We will see an example of this in Section 3.6.

3.5 Reweighting

In the canonical ensemble, we commonly describe the statistical properties of a system in thermodynamic equilibrium by the partition function Eq. (3.2). In this form, the partition function is a function of the inverse temperature β and is a sum over all possible lattice configurations U . An MCMC simulation, performed at a particular $\beta = \beta_0$, is used to generate a chain of configurations $U_1, U_2, \dots$, with frequency proportional to the Boltzmann weight $e^{-\beta_0 E(U)}$. After accounting for equilibration and autocorrelations, this chain provides a representative sample of the interesting configurations and is used to estimate various thermodynamic averages at β_0 . What about estimating the thermodynamic averages at a different value β ? Naively, this requires us to repeat the MCMC simulation at β even if β is very close to β_0 .

Alternatively, we can write the partition function as a sum over energies

$$Z(\beta) = \sum_E \Omega(E) e^{-\beta E}, \quad (3.52)$$

where the density of states $\Omega(E)$ gives the number of configurations with energy E . The probability of observing the system with energy E is

$$P_\beta(E) = \frac{\Omega(E)}{Z(\beta)} e^{-\beta E}. \quad (3.53)$$

An MCMC simulation at a fixed $\beta = \beta_0$ produces for us a Markov chain of n configurations $U_1, U_2, \dots, U_n$, with some distribution of energies E . A histogram $\mathcal{H}_{\beta_0}(E)$ of these energies gives an estimate of the equilibrium probability distribution

$$\frac{\mathcal{H}_{\beta_0}(E)}{n} = \frac{\tilde{\Omega}(E)}{Z(\beta_0)} e^{-\beta_0 E}. \quad (3.54)$$

where $\tilde{\Omega}(E)$ is an estimate of the density of states. We can solve Eq. (3.54) for $\tilde{\Omega}(E)$ and use it to replace the two instances of $\Omega(E)$ in Eq. (3.53) to get

$$P_\beta(E) = \frac{\Omega(E) e^{-\beta E}}{\sum_E \Omega(E) e^{-\beta E}} = \frac{\left[\frac{\mathcal{H}_{\beta_0}(E)}{n} Z(\beta_0) e^{\beta_0 E} \right] e^{-\beta E}}{\sum_E \left[\frac{\mathcal{H}_{\beta_0}(E)}{n} Z(\beta_0) e^{\beta_0 E} \right] e^{-\beta E}} = \frac{\mathcal{H}_{\beta_0}(E) e^{-(\beta - \beta_0)E}}{\sum_E \mathcal{H}_{\beta_0}(E) e^{-(\beta - \beta_0)E}}. \quad (3.55)$$

Thus, the histogram from a simulation performed at β_0 , can be used to estimate the probability distribution at β . This is the basis of the “single histogram” reweighting method. See textbooks




Figure 3.6 Here we show an illustration of how useful multihistogram reweighting can be to precisely locate the peak positions and heights for finite size scaling. The five black points are from canonical MCMC simulations, and the gray points are computed via multihistogram reweighting of the canonical time series. In this example, we are looking at the specific heat from the Extended- $O(2)$ model (discussed in a later chapter) with $q = 3$ and $h_q = 1$ simulated on a 64×64 lattice.

such as [5, 67]. Given a canonical simulation performed at β_0 , we can estimate the value of an observable O at a new β via the relation

$$\langle O \rangle_\beta = \sum_E O(E) P_\beta(E) = \frac{\sum_E O(E) \mathcal{H}_{\beta_0}(E) e^{-(\beta - \beta_0)E}}{\sum_E \mathcal{H}_{\beta_0}(E) e^{-(\beta - \beta_0)E}}. \quad (3.56)$$

To get an accurate estimate, there must be sufficient overlap between the histograms associated with β_0 and β . In practice, the result is only accurate for small $\Delta\beta \equiv \beta - \beta_0$. To maintain accuracy, $\Delta\beta$ generally must be taken smaller as a critical point is approached or as the lattice size is increased. To accurately cover a wider range in β , one can use “multihistogram” reweighting.

In multihistogram reweighting [68, 69], multiple histograms (from multiple canonical simulations) are stitched together to approximate the density of states. This improves the accuracy of the reweighted observables, and typically allows one to reweight to a much larger range of β values.

Reweighting is very useful in finite-size scaling [70] to obtain precise peak positions without having to do a large number of canonical simulations at different β . An example is shown in Figure 3.6 for illustration.

3.6 Curve Fitting

3.6.1 Uncorrelated Fitting

Suppose we have N uncorrelated data points $y_i(x_i)$ with $i = 1, 2, \dots, N$ and their standard deviations σ_i . We want to model the data by fitting it to some function $y(x_i; \vec{a})$ with free parameters $\vec{a}$. The goal is to estimate the best fit parameters $\vec{a}$ and their error bars. For linear fits, this problem is easily solved with the linear least squares approach. However, for nonlinear fits, such as power law fits which occur in finite size scaling, the fitting procedure is slightly more complicated. Here, we only briefly sketch the general procedure since one can find the details in references like [5, 71].

The foundation of most fitting methods is the chi-squared function

$$\chi^2(\vec{a}) = \sum_{i=1}^N \left(\frac{y_i - y(x_i; \vec{a})}{\sigma_i} \right)^2. \quad (3.57)$$

Here, y_i are the data points with standard deviations³ σ_i , and $y(x_i; \vec{a})$ is the function we are fitting. The components of $\vec{a} = (a_1, \dots, a_M)$ are the free parameters to be estimated. The best estimates of the free parameters are those which minimize Eq. (3.57). Thus, the task of a fitting procedure is reduced to the problem of minimizing Eq. (3.57).

For linear or polynomial fitting, the minimization of Eq. (3.57) is straightforward. For nonlinear fitting, it is more difficult, but there are standard algorithms such as the Levenberg-Marquardt algorithm proposed by Levenberg [72] and improved by Marquardt [73]. The Levenberg-Marquardt algorithm uses a weighted average of the Newton-Raphson and the steepest descent methods for minimizing Eq. (3.57). Both of these are iterative methods. Initially, the weight is biased toward the steepest descent method which converges towards a (possibly local) minimum even when the initial guess $\vec{a}_0$ for the free parameters is poor. The steepest descent method computes the next iteration point, $\vec{a}_{n+1}$, from

$$\vec{a}_{n+1} = \vec{a}_n - \gamma \nabla \chi^2(\vec{a}_n), \quad (3.58)$$

where γ is a small constant. The disadvantage of the steepest descent method is slow convergence, so once convergence is detected, the Levenberg-Marquardt algorithm shifts the weight to the

³Including the error bars of the data points is important so that the best fit properly gives more weight to points with smaller error bars.

Newton-Raphson method, which converges rapidly but needs a good initial approximation. In the Newton-Raphson method, the next iteration point is computed as

$$\vec{a}_{n+1} = \vec{a}_n - H^{-1} \nabla \chi^2(\vec{a}_n), \quad (3.59)$$

where H is an $M \times M$ Hessian matrix with components

$$H_{jk} = 2 \sum_{i=1}^N \frac{1}{\sigma_i^2} \left[\frac{\partial y(x_i; \vec{a})}{\partial a_j} \frac{\partial y(x_i; \vec{a})}{\partial a_k} - [y_i - y(x_i; \vec{a})] \frac{\partial^2 y(x_i; \vec{a})}{\partial a_j \partial a_k} \right] \Big|_{\vec{a}=\vec{a}_n}. \quad (3.60)$$

After finding the parameters $\vec{a}$ which minimize Eq. (3.57), one computes the inverse of half of the Hessian

$$C = \left[\frac{1}{2} H \right]^{-1}. \quad (3.61)$$

This is the covariance matrix of the standard errors in the fitted parameters $\vec{a}$.

In practice, there are a variety of software packages available which implement fitting procedures such as the one described above. In Python, a common choice is to use `scipy.optimize.curve_fit`, which defaults to the Levenberg-Marquardt algorithm for unconstrained fitting. Such a fitting package will return the best fit parameters $\vec{a}$ along with a covariance matrix. For uncorrelated data, the covariance matrix is diagonal, with the diagonal elements being the estimated variances of the fit parameters.

The purpose of the curve fitting procedure is typically to estimate some figure of interest A and to quantify the uncertainty σ_A of that estimate. If the figure of interest is one of the fit parameters (e.g. suppose we are fitting to the form Bx^A , and we are interested in the exponent A), then we can simply use a standard fitting package, which gives us the best estimates of A and B as well as the covariance matrix. The square roots of the diagonal elements of the covariance matrix give us the uncertainties σ_A and σ_B of the fit parameters. On the other hand, if our figure of interest is some nonlinear function of the fit parameters, then the fitting package will not calculate the uncertainty σ_A for us. In that case, one can use a form of “bootstrap” to estimate σ_A . For example, suppose we are fitting to a model $y(x_i; \vec{a})$ with fit parameters $\vec{a}$, and the figure of interest is the maximum value of the best fit model, i.e. $A = \max(y(x_i; \vec{a}))$. The fitting package returns the best fit parameters $\vec{a}$

along with the uncertainties $\sigma_{a_1}, \sigma_{a_2}, \dots$, but we want A and σ_A . Calculating A is straightforward as one can use the usual calculus approach to obtain the maximum of a given function. To estimate σ_A using a bootstrap approach, we treat each original data point $y_i(x_i)$ as coming from a Gaussian distribution with standard deviation given by the error bar σ_i associated with that point. For each data point $y_i(x_i)$, we select a Gaussian random from that distribution, fit the curve to that set of N random numbers, and calculate the maximum. This is repeated a large number of times so that one ends up with many estimates of the maximum A . The result that one quotes is the mean of this sample, and the error bar is the standard deviation of the sample.

3.6.2 Jackknife Fitting

When fitting correlated data, one can in principle use the approach detailed above, but now the chi-squared function is the more general

$$\chi^2(\vec{a}) = \sum_{i,j} (y_i - y(x_i; \vec{a})) \hat{C}_{ij}^{-1} (y_j - y(x_j; \vec{a})) . \quad (3.62)$$

where the elements $\hat{C}_{ij}^{-1}$ of the inverse covariance matrix are themselves estimated from the data. In practice, this matrix often has very small eigenvalues—resulting in an unstable $\chi^2(\vec{a})$. As a result, this approach is often unworkable in practice. Besides, this approach wouldn't help us to get the error bar associated with a more complicated property (e.g. the maximum) of the best fit. The solution is to use the jackknife method.

The jackknife approach to fitting is a very general method that works for both uncorrelated data and correlated data (provided one uses appropriate block sizes). Rather than attempt a formal description of this procedure, we detail here an example. In finite-size scaling, one needs to extract peak locations and heights from various thermodynamic quantities, e.g. the specific heat $C = \beta^2(\langle E^2 \rangle - \langle E \rangle^2)/V$, where V is the number of lattice sites, and E is the internal energy. Consider Figure 3.6. We want to extract the peak (location and height) along with proper error bars. Our procedure is as follows:

1. Choose the set of data points $C(\beta_1), C(\beta_2), \dots, C(\beta_N)$ to include in the fit. In practice, most or all of these data points are computed via a reweighting procedure from a few canonical

simulations, and so these points are correlated with each other. For simplicity, we assume here that there is no reweighting procedure. Instead, we will assume that the $C(\beta_i)$ data point will be coming from an energy time series of length M generated by a canonical MCMC simulation performed at β_i . So we end up with N time series of the energy E each of length M . Let $E_{i,j}$ be the j th element of the i th time series ($i = 1, \dots, N, j = 1, \dots, M$)

2. We choose the number of jackknife blocks N_b to use when blocking the time series. One should use $N_b > 30$, but also ensure that the number of elements M/N_b in a block is much larger than the integrated autocorrelation time τ_{int} . Let $\mathcal{B}_{i,k}$ represent the k th block in the i th time series
3. Then the k th jackknife estimator of the specific heat from the i th time series ($C_{ik}^J \equiv C_k^J(\beta_i)$) is obtained by computing the specific heat on the entire time series E_i excluding the k th block. That is,

$$C_{ik}^J = \frac{\beta_i^2}{V} \left((E_i^2)_k^J - \left((E_i)_k^J \right)^2 \right) \quad (3.63)$$

where

$$(E_i)_k^J = \frac{1}{M - M/N_b} \sum_{k \text{ for } E_{i,k} \notin \mathcal{B}_{i,k}}^{M-M/N_b} E_{i,k}, \quad (E_i^2)_k^J = \frac{1}{M - M/N_b} \sum_{k \text{ for } E_{i,k}^2 \notin \mathcal{B}_{i,k}}^{M-M/N_b} E_{i,k}^2 \quad (3.64)$$

So for each of the N time series, we end up with N_b estimators of the specific heat

4. Our fitting procedure requires that our data points have proper error bars. At this stage, we can get those error bars by calculating the jackknife variance of the specific heat estimators

$$s_{\overline{C}_i^J}^2 = \frac{N_b - 1}{N_b} \sum_{k=1}^{N_b} \left(C_{ik}^J - \overline{C}_i^J \right)^2, \quad \overline{C}_i^J = \frac{1}{N_b} \sum_{k=1}^{N_b} C_{ik}^J. \quad (3.65)$$

5. For $k = 1$, we plot the N specific heat estimators C_{i1}^J with the error bars given by $s_{\overline{C}_i^J}$. We perform our fitting procedure (e.g. fitting a cubic polynomial using a least squares algorithm) on this set of points and extract the maximum which we label M_1^J . We repeat this for $k = 2, \dots, N_b$, to end up with N_b jackknife estimators M_k^J , ($k = 1, \dots, N_b$) of the maximum

6. We then calculate the jackknife mean and variance of the maximum as

$$\overline{M}^J = \frac{1}{N_b} \sum_{k=1}^{N_b} M_k^J, \quad s_{\overline{M}^J}^2 = \frac{N_b - 1}{N_b} \sum_{k=1}^{N_b} \left(M_k^J - \overline{M}^J \right)^2. \quad (3.66)$$

The final result we quote for the maximum is then

$$\overline{M}^J \pm s_{\overline{M}^J}. \quad (3.67)$$

For a recent paper on alternative approach to correlated fitting, see [74] and references therein.

3.6.3 Goodness-of-fit Measures

After obtaining the best fit for a given model, one needs a way to quantify how good the fit is since a bad best fit suggests that one is using a poor model. This brings us to “goodness-of-fit” measures. A common goodness-of-fit measure is the reduced chi-squared statistic or “chi-squared per degrees of freedom” computed as

$$\chi_\nu^2 = \frac{\chi^2}{\nu}, \quad (3.68)$$

where χ^2 is Eq. (3.57) computed at the best fit parameters, and $\nu = N - M$ is the number of degrees of freedom calculated as the number of points being fitted minus the number of fitting parameters. Typically, a good fit will result in $\chi_\nu^2 \sim 1$, with $\chi_\nu^2 > 1$ implying a bad fit and $\chi_\nu^2 < 1$ suggesting an overfit.

An alternative goodness-of-fit measure described in [5, 71] is

$$Q = 1 - P\left(\frac{\nu}{2}, \frac{\chi^2}{2}\right), \quad (3.69)$$

where $P(a, x) = (1/\Gamma(a)) \int_0^x e^{-t} t^{a-1} dt$ is the lower incomplete gamma function. Again, χ^2 is Eq. (3.57) computed at the best fit parameters, and ν is the number of degrees of freedom. The number Q is the probability that a chi-square as poor as χ^2 will occur by chance. In other words, it gives the likelihood that the difference between the fit and the data is due to chance. According to [71], if $Q > 0.1$ then the fit is believable. If $Q < 0.001$, then the model should be called into question. If $0.001 < Q < 0.1$ then the fit may be acceptable if the errors are nonnormal or have been moderately underestimated.

3.7 Finite Size Scaling

To obtain critical exponents for second-order transitions, we consider the following leading finite-size-scaling ansätze [75–77]

$$\left. \frac{dU_M}{d\beta} \right|_{max} = U_0 + U_1 L^{1/\nu} \quad (3.70)$$

$$C|_{max} = C_0 + C_1 L^{\alpha/\nu} \quad (3.71)$$

$$\langle M \rangle|_{infl} = M_0 + M_1 L^{-\beta/\nu} \quad (3.72)$$

$$\chi_M|_{max} = \chi_0 + \chi_1 L^{\gamma/\nu} \quad (3.73)$$

$$F(\vec{q})|_{max} = F_0 + F_1 L^{2-\eta}. \quad (3.74)$$

where $U_{|\vec{M}|}$ is the Binder cumulant (to be defined and discussed in the next chapter) with respect to the proxy magnetization, C is the specific heat, M is the magnitude of the proxy magnetization, χ_M is the magnetic susceptibility of the proxy magnetization, and $F(\vec{q})$ is the structure factor (to be defined and discussed in the next chapter). Note, when $\alpha = 0$, as it is for transitions of the Ising class, the ansatz Eq. (3.71) must be modified to

$$C|_{max} = C_0 + C_1 L^{\alpha/\nu} \ln L. \quad (3.75)$$

The maxima can be extracted by performing an MCMC scan in β to locate the approximate position of the peaks followed by a higher resolution scan centered on the peaks. The integrated autocorrelation time of the energy should be estimated from each canonical time series. The canonical runs can then be “stitched” together using multihistogram reweighting to interpolate to β values between the canonical MCMC runs. The maxima can be extracted from jackknife bins by fitting a cubic polynomial to the peak. The jackknife method should be used in order to get appropriate error bars for the final results. This is repeated at different lattice sizes and can be fitted to the above finite-size-scaling ansätze in Python using the `curve_fit` function in the `scipy.optimize` package.

For the exponent $1/\nu$, one can fit the derivative of the Binder cumulant to the form given by Eq. (3.70). Instead of taking a finite difference derivative, one can follow [77] and use the analytical

form

$$\frac{dU_M}{d\beta} = \frac{2\langle M^4 \rangle [\langle M^2 \rangle \langle E \rangle - \langle M^2 E \rangle]}{3\langle M^2 \rangle^3} - \frac{\langle M^4 \rangle \langle E \rangle - \langle M^4 E \rangle}{3\langle M^2 \rangle^2}. \quad (3.76)$$

where E is the internal energy and M is the proxy magnetization.

For the exponent $-\beta/\nu$, one can fit Eq. (3.72), where $\langle M \rangle|_{infl}$ is the value of the proxy magnetization at its inflection point. To determine the location of the inflection point, one can use the maximum of the derivative, which can be computed without finite difference as

$$\frac{d\langle M \rangle}{d\beta} = \langle E \rangle \langle M \rangle - \langle EM \rangle, \quad (3.77)$$

where E is the internal energy. As with other maxima, a cubic polynomial can be fitted to the peak to extract the maximum $d\langle M \rangle/d\beta|_{max}$ and its location β_{pc} . We then interpolate the proxy magnetization to the point β_{pc} using a cubic polynomial to obtain $\langle M \rangle|_{infl} = \langle M \rangle(\beta_{pc})$.

The pseudocritical points β_{pc} (i.e. the peak locations) for second-order transitions should approach the infinite volume critical point β_c as

$$\beta_{pc} = \beta_c + \frac{a}{L^{1/\nu}}. \quad (3.78)$$

To distinguish between second-order and BKT transitions, one can check whether the specific heat diverges with volume. For a BKT transition, the specific heat does not diverge with volume, and so Eq. (3.71) is not suitable for BKT transitions. However, this check may not be sufficient. If the critical exponent α is negative, then the specific heat plateaus also for second-order transitions. See for example, [40]. For more on distinguishing BKT from first- and second-order transitions, see [40, 41].

There are well-known scaling and hyperscaling relations, which relate the different critical exponents. Their derivations can be found in textbooks such as Ref. [59]. In particular, we will consider the relations

$$\frac{\gamma}{\nu} = 2 - \eta \quad (3.79)$$

$$\frac{\alpha}{\nu} = \frac{2}{\nu} - d \quad (3.80)$$

$$d - \frac{\gamma}{\nu} = 2\frac{\beta}{\nu}, \quad (3.81)$$

where $d = 2$ in our case.

In a second-order transition, the critical exponent ν is defined by the scaling of the correlation length $\xi \propto \tau^{-\nu}$, where $\tau \equiv (T - T_c)/T_c = \beta_c/\beta - 1$. However, for a BKT transition, ξ diverges faster than any power of τ , and so we cannot define ν in the conventional way. See for example, [15], where it is shown that for the XY model, $\xi \propto e^{-b\tau^{-1/2}}$ or [47], where it is stated that ν should be ∞ for $q = 5, 6$.

One can define the exponent $\tilde{\nu}$ for a BKT transition [54] (we use tilde to distinguish from the ordinary ν) by the scaling

$$\xi \propto e^{b\tau^{-\tilde{\nu}}}, \quad (3.82)$$

and use this to label the universality class of the system. For the XY model, $\tilde{\nu} = 1/2$. Then the pseudocritical point for a BKT transition should approach the infinite volume critical point as

$$\beta_{pc} = \beta_c + \frac{a}{(\ln L + b)^{1/\tilde{\nu}}}. \quad (3.83)$$

This is different from second-order scaling. However, as noted in [54], fits to this form may be unstable, and so this scaling relation may not be a good way to extract the critical exponent $\tilde{\nu}$. To investigate the $q = 5$ state clock model, they ended up using this equation with $\tilde{\nu}$ fixed at $1/2$

$$\beta_{pc} = \beta_c + \frac{a}{(\ln L + b)^2}. \quad (3.84)$$

This form was used to find β_c , but even then it seemed unsuitable as it was only reliable for very large volumes.

To extract the β , γ , and η critical exponents from BKT transitions, one can use the scaling ansätze Eq. (3.72), (3.73), and (3.74), however, in principle, one should replace $\nu \rightarrow \tilde{\nu}$. One has to be careful with the two BKT transitions in the clock models when $q \geq 5$. One may not be able to use the proxy magnetization in all cases with Eq. (3.72) or (3.73). For example, for the second transition in the $q = 5$ clock model, [54] uses a “rotated magnetization” like Eq. (2.59) instead of the proxy magnetization.

Note, for BKT transitions in the $q \geq 5$ clock models, the number of peaks in the specific heat corresponds to the number of phase transitions, however, the location of the peaks do not correspond

to the transition temperatures [51]. Thus, the specific heat alone cannot be used to locate the critical temperatures. For more details on finite-size scaling for BKT transitions, see [78].

3.8 MCMC in Practice

3.8.1 Outline of an MCMC Simulation

An MCMC simulation typically follows an outline like:

```
# Initialize everything
initialize()

# Equilibrate the lattice
for iequi=1 to nequi:
    updating_sweep()

# Measurement sweeps
for imeas=1 to nmeas:

    # Discards
    for idisc=1 to ndisc:
        updating_sweep()

    save_or_measure()

# Finish up
finalize()
```

In the initialization stage, the model parameters (e.g. β , lattice dimensions, boundary conditions, etc.) and run parameters (e.g. `nequi`, `nmeas`, `ndisc`, `start_type`, etc) must be read in, and the lattice arrays must be initialized. At this stage, one chooses whether to do a hot start (i.e. lattice with random values), a cold start (i.e. a lattice with uniform values), or to continue from a previously saved lattice.

It takes some number of equilibrating sweeps for the Markov chain to reach the equilibrium distribution. There are several ways to estimate the number of equilibrating sweeps that are needed including:

1. Measure some observable and see how many sweeps it takes to settle down to some value
2. Even better, do a hot start and a cold start and see how many sweeps it takes for the




Figure 3.7 In these illustrations, the 2-state Potts model (aka the Ising model) is simulated on a 100x100 lattice using both cold (ordered) and hot (disordered) starts. On the left, the simulation is performed at $\beta = 0.4$, and we see that the lattice equilibrates quickly. On the right, the simulation is performed in the ordered phase at $\beta = 0.8$, and we see that at least one of the start types is having trouble reaching equilibration. Since we know the exact energy density for the two-dimensional Ising model on a finite lattice, we can tell that the ordered start is at the correct value. The random start has significantly more trouble reaching equilibration. In general, it seems that for a simulation performed in the ordered phase, the random start may take significantly longer to reach equilibration than the ordered start [5].

measurements from these two runs to converge. See Figure 3.7

3. Do multiple runs with different RNG seeds and see how many sweeps it takes for the measurements to converge

Measurements taken during the equilibration stage are generally discarded—they do not form part of the sample of equilibrium measurements.

After equilibration, one can begin taking measurements. Here one has the choice of *in situ* or *ex situ* measurements. In situ measurements are taken while the MCMC simulation is running. For example, after an updating sweep, the energy, magnetization, and other measurements are taken before proceeding to the next updating sweep. Ex situ measurements are taken outside of the MCMC simulation. This requires the lattice configuration itself be written to file after each measurement sweep. The choice of measurement type often comes down to the computational cost of the updating sweep, the cost of each measurement, and I/O and file storage considerations. For classical spin




Figure 3.8 A time series study of magnetization for the 2-state Potts model (i.e. Ising model) on a 20x20 lattice. At the top left, we perform 165K heatbath sweeps at $\beta = 0.20$. We see that the lattice is not magnetized because 50% of the spins are in the $q = 0$ state and 50% are in the $q = 1$ state. At the top right, we repeat the simulation but for $\beta = 0.44$, which is the critical value of β in the infinite volume limit. Here we see a partially magnetized lattice (the spins tend to be aligned 80/20 instead of 50/50). We see that the lattice magnetization flips often between $q = 0$ (red) dominating 80% of the lattice and $q = 1$ (blue) dominating 80% of the lattice. At the bottom left, we repeat with $\beta = 0.48$. Now the lattice is magnetized more strongly with spins aligned 95/5. We still see magnetization flips, but not as frequently. Finally, at the bottom right, we repeat the simulation with $\beta = 0.80$. Now the lattice is completely magnetized with $\sim 100\%$ of the spins staying in the $q = 0$ state.

systems in two dimensions, one typically does in situ measurements. The updating sweep and the measurements are fairly cheap, so a large number of them can be performed quickly. Writing the lattice to file after each sweep might slow it down significantly, and the storage requirements would be much larger. On the other hand, in lattice QCD for example, the updating sweeps and measurements are much more costly, and so it makes sense to save the lattice configuration to file after an updating sweep, and perform the measurements separately.

Some researchers prefer to immediately histogram in situ measurements. The benefit is that it can reduce the memory and storage requirements, and most things of interest—e.g. means and standard deviations—can be calculated from the histograms. Furthermore, the shape of the histograms can signal the presence of phase transitions as was seen in Figure 2.7. Nevertheless,

some useful information is lost, and it may be preferable to store the entire time series observables rather than histogramming them. To understand how the Markov chain is behaving it is sometimes useful to examine the time series observable (i.e. the observable plotted versus Monte Carlo sweep number). An example of this is shown in Figure 3.8 for the 2-state Potts model. Such an examination can help inform the researcher on the statistical properties of the MCMC simulation. For example, looking at the time series is a quick way to check that the simulation has reached equilibrium. One can always histogram the time series later if needed.

When there is a large autocorrelation time, the effective statistics is significantly reduced, and it may not be worthwhile to measure and save the observable on every sweep. In such cases, one often discards some number of sweeps for each measured sweep. In practice, the autocorrelation time may not be known. If an estimate is known, the discards per measurement sweep can be set to mitigate the autocorrelations, but one should still perform an autocorrelation analysis of the saved measurements. If there remains any residual autocorrelation, it can be mitigated by blocking techniques as discussed in Sections 3.3 and 3.4.

3.8.2 Planning and Organization

The goal of MCMC lattice studies is varied and ranges from studying and developing new techniques and tools to characterizing some new spin or gauge model. Part of the process of understanding a new lattice model is to map out the phase diagram. This requires locating the critical values of the parameters and then doing a finite-size scaling analysis after performing simulations at those critical points on increasingly large lattices. Characterizing the phase diagram of even a relatively simple spin system can take hundreds of thousands of CPU hours and generate terabytes of data to be stored and analyzed. For example, the work described in Chapter 4 required the following resources:

- 30K+ individual MCMC simulations were required just to perform a basic scan of the parameter space at a single lattice size
- Up to 800 nodes on a computing cluster were running simultaneously (trivial parallelization)

- Large autocorrelations required Markov chains of length billions in some cases
- Several terabytes of hard disk space were used for the saved time series observables (many more were discarded)
- 500K+ CPU hours were used on the computing cluster

This is only the resources used from the MCMC side, i.e. it does not include the computing resources needed for the complimentary approach which involved tensor methods. Such a study therefore requires some planning and setup. Given the computing time involved, such work typically takes place on a computing cluster where many nodes can be used in a trivial kind of parallelization to simultaneously run many MCMC simulations at different parameter values. For runs on large lattices, which take a long time, real parallelization may be employed to spread the work of a single simulation over multiple cores and/or nodes. This of course assumes a lattice codebase capable of this kind of parallelization.

Typically, the main simulation code is written in a fast compiled language like C++. During or when the simulation completes, the time series measurements are written to a file on disk. Later, they can be analyzed by reading the time series into another script which computes e.g. the integrated autocorrelation time and the jackknife mean and error bar of the observable. These analysis scripts may be a combination of shell scripts, C++ scripts, and Python scripts. During the simulation, one typically records some fundamental observables which can be used to construct other observables. For example, for a two-dimensional spin system, one might record the time series magnetization components M_x and M_y . During the analysis stage, additional observables such as the magnetization and the Binder cumulant can be constructed from these components and analyzed.

File organization becomes important when one has to deal with gigabytes or terabytes of data. The author found it useful to use the following top-level directories:

- `/CODEBASE/`: Located in the user's home directory on the computing cluster, this directory contains the lattice codebase itself and all analysis scripts

- `/RUNS/`: Located on the cluster’s scratch space, this is the directory where the actual simulations are run. Intermediate files and error outputs are left here to be auto-deleted by the cluster after some number of days. If a simulation fails, the error files should be examined before deletion
- `/STORAGE/`: This is the long-term storage space where the simulation writes the main time series observables. This is where the vast majority of storage space is needed—space which may exceed the home directory allotment
- `/OUTPUTS/`: Located in the user’s home directory, this directory contains summary information obtained from the analysis scripts. That is, after analyzing a time series (stored in `/STORAGE/`), the results would be written to `/OUTPUTS/`

It may be useful to use a common directory structure within `/RUNS/`, `/STORAGE/`, and `/OUTPUTS/`. For example, for the Ising model with two parameters—inverse temperature β and external field h —the following structure could be used: `[modelname]_[ndim]_[length]_[id]/[beta]_[h]/`. For example, a run on a 16×16 lattice at $\beta = 1.3$ and $h = 0$ would take place in the “run_path” directory `/ising_nd2_n16_1/1p3000_0p0000/`. The id tag serves to distinguish runs which are otherwise identical, e.g. they differ only in the random seed used. This directory structure makes it easy to analyze and gather results versus β . For example, to get energy versus β for $h = 0$, one would simply cd into `/STORAGE/ising_nd2_n16_1/`, loop over all directories `/*_0p0000/` therein, and extract the energy from each time series file. The run itself takes place at `/RUNS/run_path` with the time series observables written to `/STORAGE/run_path`. During analysis, summary results are written to `/OUTPUTS/run_path`.

CHAPTER 4

THE EXTENDED-O(2) MODEL

4.1 Introduction

In recent years, the idea of using quantum computers or quantum simulation experiments to approach the real-time evolution or the finite-density behavior of lattice models of interest for high-energy physics has gained considerable interest [79–88]. As the current noisy intermediate scale quantum (NISQ) devices that are available to implement this research program have a very limited number of quantum computing units, such as qubits, trapped ions or Rydberg atoms, it is essential to optimize the discretization procedure. Starting from the standard Lagrangian formulation of lattice field theory models with continuous field variables, one can either discretize the field variables [89–91] used in the path integral, expand the Boltzmann weights using character expansions [81, 92, 93], or use the quantum link method [94, 95].

Models with continuous Abelian symmetries are of great physical interest. Besides the electromagnetic interactions of charged particles in 3+1 dimensions, this also includes models where a mass gap is dynamically generated [96, 97] or a Berezinskii-Kosterlitz-Thouless (BKT) transition [15, 36, 37] occurs. For models with a $U(1)$ symmetry, the character expansion mentioned above is simply the Fourier series. It has been shown [3, 98] that the truncation of these series preserves the original symmetry. On the other hand, the $\mathbb{Z}_q$ clock approximation of the integration over the circle only preserves the $\mathbb{Z}_q$ discrete subgroup. A recent proposal applies the $\mathbb{Z}_q$ clock approximation to the simulation of the Abelian gauge theory in 2+1 dimensions, where transformations between the electric representation and the magnetic representation can significantly reduce the required computational resources [99]. In order to decide how good the $\mathbb{Z}_q$ approximation is in a variety of situations, it is useful to build a continuous family of models interpolating among the various possibilities.

Here we focus on the case of the O(2) nonlinear sigma model in 1+1 dimensions. This model was key to understanding the BKT transition [15, 36, 37, 100] and the corresponding $\mathbb{Z}_q$ clock model has been studied extensively [25, 43–57, 78, 101–104]. We propose to interpolate among

these models by starting with the standard $O(2)$ Hamiltonian and introducing a symmetry-breaking term,

$$\Delta H(h_q, q) = -h_q \sum_x \cos(q\varphi_x). \quad (4.1)$$

When q is an integer, if we take the limit $h_q \rightarrow \infty$, we recover the $\mathbb{Z}_q$ clock model. For the rest of the discussion, it is important to realize that the $O(2)$ -symmetric Hamiltonian is 2π -periodic for all the φ_x variables. In contrast, ΔH has a $2\pi/q$ periodicity. When q is an integer, if we apply the shift q times we obtain the periodicity of the $O(2)$ action. In order to interpolate among the clock models, we will consider noninteger values of q while keeping a fixed φ interval of length 2π . The model and the effect of the symmetry breaking are discussed later both in the standard Lagrangian and tensor formulations.

The idea of having a doubly continuous set of models is interesting from a theoretical point of view but also from a quantum simulation point of view. If we attempt to quantum simulate these models using Rydberg atoms as in Refs. [105–107], it is possible to tune the ratio R_b/a of the radius for the Rydberg blockade and the lattice spacing, as well as local chemical potentials continuously. This allowed interpolations among $\mathbb{Z}_q$ phases for different integer values of q [107]. Sequences of clock models also appear in models for nuclear matter when the number of colors is varied [108].

It is often a difficult task to detect BKT transitions in the quantum Hamiltonian approach, as it is hard to find a good indicator of BKT transitions that has a clear discontinuity, peak or dip. The equivalence of the path integral formulation and statistical mechanics can be used to access universal features and detect phase transitions in statistical mechanics based on Monte Carlo (MC) simulations and tensor renormalization group (TRG) calculations. The Markov chain MC (MCMC) simulations efficiently explore the typical set of the physical configurations. MC calculations use the universal jump in the helicity modulus [109] as an indicator for BKT transitions. But ambiguities in the definition of the helicity modulus in the $\mathbb{Z}_5$ clock model can result in controversial conclusions [44, 110]. The TRG [111, 112] calculations provide a coarse-grained theory where the size of the lattice spacing doubles at each step. If the truncations performed are under control, one can go to the thermodynamic limit quickly. Calculation of the magnetic susceptibility in the

presence of a weak external field is a universal method to detect critical points that can be easily implemented in the TRG. However, it does not show a peak to indicate the large- β BKT transition in the five-state clock model [50, 113]. The study of the $h_q \rightarrow \infty$ limit with fractional q not only provides us a clear picture of what phases the symmetry-breaking term will drive the XY model to, paving the way to discussions for the full phase diagram at finite h_q , but also brings us a new tool to detect BKT transitions in $\mathbb{Z}_n$ models. In contrast, the calculation of the specific heat at increasing volume allows us to discriminate between a second-order phase transition—where it diverges logarithmically with the volume in the Ising case—and a BKT transition or a crossover.

To define the models that we consider in this chapter we start with the two-dimensional classical $O(2)$ nonlinear sigma model, or XY model, where the spin degrees of freedom $\vec{S}$ are unit vectors whose possible directions are confined to a plane. They reside on the sites of a two-dimensional lattice of N sites. The Hamiltonian is

$$H_{O(2)} = - \sum_{x=1}^N \left(\sum_{\mu=1}^2 \vec{S}_x \cdot \vec{S}_{x+\hat{\mu}} + \vec{h} \cdot \vec{S}_x \right), \quad (4.2)$$

where the sum on x is over the sites of the two-dimensional lattice, and on $\mu = 1, 2$ over the directions. The field, $\vec{h}$, is a uniform constant external magnetic field. It is convenient to parameterize the spins $\vec{S}$ with a single angle $\varphi \in [\varphi_0, \varphi_0 + 2\pi)$. The Hamiltonian then takes the form

$$H_{O(2)} = - \sum_{x,\mu} \cos(\varphi_{x+\hat{\mu}} - \varphi_x) - h \sum_x \cos(\varphi_x - \varphi_h), \quad (4.3)$$

where $h = |\vec{h}|$ and φ_h is the direction of the external field that in the absence of other symmetry-breaking terms can be set to zero for convenience.

Next, let us extend the model by introducing a term that can favor certain values of the angle:

$$H_{\text{ext-}O(2)} = - \sum_{x,\mu} \cos(\varphi_{x+\hat{\mu}} - \varphi_x) - h_q \sum_x \cos(q\varphi_x) - h \sum_x \cos(\varphi_x - \varphi_h). \quad (4.4)$$

We call the model with the Hamiltonian (4.4) the “extended- $O(2)$ ” model.

The canonical partition function is

$$Z = \int dU \, e^{-\beta H(U)}, \quad (4.5)$$

where the integration is over lattice configurations U . After parameterizing the spins $\vec{S}$ by an angle φ , the partition function is

$$Z = \int_{\varphi_0}^{\varphi_0+2\pi} \prod_x \frac{d\varphi_x}{2\pi} e^{-\beta H}. \quad (4.6)$$

Specific cases of the Extended- $O(2)$ model have been studied fairly extensively. For $q = 2, 3$, an early normalization group (RG) analysis [114] showed that the model at finite h_q should be in the Ising and 3-state Potts universality class respectively. However, in [115] they show that the critical exponents vary with h_q , and conclude that for $q = 3$, there is a BKT phase for sufficiently small h_q . The $q = 4$ with finite h_q case, sometimes referred to as the XYh_4 model, has been studied extensively due to its relevance to physical systems [116, 117]. The RG analysis [114] showed there to be a second-order phase transition with nonuniversal critical exponents which vary with h_q . The BKT phase of the $O(2)$ was recovered only for $h_q = 0$. In [118], numerical MC studies of this case showed that the critical exponent ν does vary with h_q . This was supported further by [40, 119], and more discussion can be found in [116, 120]. However, in [41], it was found that for sufficiently small h_q , the second-order phase transition for $q = 4$ is replaced by BKT transitions. This is supported by [115]. However, recent work [117, 121] suggests that the BKT phase at finite h_q is essentially a finite-size effect and might not persist to infinite volumes. For $q \geq 5$, the RG analysis [114], showed there should be a BKT phase between a low-temperature ferromagnetic phase and a high-temperature paramagnetic phase. For $q = 6$, this is supported by numerical MC studies on triangular lattices [40]. The $q = 5$ case was shown [122] to have two transitions for any finite h_q .

In the case $h_q = 0$, the Extended- $O(2)$ model reduces to the ordinary $O(2)$ model, which has been studied extensively [18, 21, 24, 39, 123–126]. In the limit $h_q \rightarrow \infty$, the Extended- $O(2)$ model reduces to the Extended q -state clock model. With $h_q \rightarrow \infty$ and integer q , it reduces to the ordinary q -state clock model, which has been studied extensively [25, 43–57, 78, 101–104].

Previous work (cited above) was restricted to integer values of q , and typically focused only on one or two values. Furthermore, many focused only on large or small values of h_q . In this thesis, we perform a comprehensive study of the Extended- $O(2)$ model over a large parameter space with

$\beta, h_q \in \mathbb{R} \geq 0$ and with $q \in \mathbb{R}$. Much of the material in this chapter has been published [127–129].

4.2 The $h_q \rightarrow \infty$ Phase Diagram

4.2.1 Introduction

For integer q the limit $h_q \rightarrow \infty$ forces the spin angles to take the values $\varphi_x^{(k)} = 2\pi k/q$ with $k \in \mathbb{Z}$. Thus, while $h_q = 0$ corresponds to the O(2) model, $h_q \rightarrow \infty$ corresponds to the q -state clock model. The action defined in Eq. (4.4) is also valid for noninteger q and we therefore consider Eq. (4.4) as our definition of the extension of the q -state clock model to noninteger q in the $h_q \rightarrow \infty$ limit. In this case the angle $\varphi_x^{(k)}$ takes the values

$$\varphi_0 \leq \varphi_x^{(k)} = \frac{2\pi k}{q} < \varphi_0 + 2\pi, \quad (4.7)$$

with $k \in \mathbb{Z}$ and some choice of domain $[\varphi_0, \varphi_0 + 2\pi)$. By varying φ_0 , we can obtain different sets of angles that are equivalent to either $k = 0, 1, \dots, \lfloor q \rfloor$ (case 1) or $k = 0, 1, \dots, \lfloor q \rfloor - 1$ (case 2), since in the $h_q \rightarrow \infty$ limit—in the absence of an external field—the action only depends on the relative angle between nearest-neighbor sites $\Delta\varphi_{x,\mu}^{(k)} = \varphi_{x+\hat{\mu}}^{(k)} - \varphi_x^{(k)}$ (see Appendix B.1). Case 2 just has one fewer angle than case 1. As shown in Fig. 4.1, the angular distance between two adjacent values of φ_x on a circle takes two values: $2\pi(q - \lfloor q \rfloor)/q < 2\pi/q$ for case 1, and $2\pi/q < 2\pi(1 + q - \lfloor q \rfloor)/q$ for case 2. These values including 0 have the largest Boltzmann weights in the partition function. The “leftover angle” is

$$\tilde{\phi} \equiv 2\pi \left(1 - \frac{\lfloor q \rfloor}{q} \right), \quad (4.8)$$

in case 1 and $2\pi/q + \tilde{\phi}$ in case 2. With the choice $\varphi_0 = 0$, we have case 1, while choosing $\varphi_0 = -\pi$ is equivalent to case 2 (1) for odd (even) $\lfloor q \rfloor$. At noninteger q the $\mathbb{Z}_q$ symmetry is explicitly broken since the action is not invariant under the operation $k \rightarrow k + 1 \pmod{\lfloor q \rfloor}$. But there is still a $\mathbb{Z}_2$ symmetry because the action is invariant under the operation $k \rightarrow \lfloor q \rfloor - k$.

We also consider the limit $h_q \rightarrow \infty$ directly by simply restricting the values of the originally continuous angle φ to the values given in Eq. (4.7):

$$H_{\text{ext-}q} = -\beta \sum_{x,\mu} \cos(\varphi_{x+\hat{\mu}}^{(k)} - \varphi_x^{(k)}) - h \sum_x \cos(\varphi_x^{(k)} - \varphi_h). \quad (4.9)$$




Figure 4.1 Arrows indicate the allowed spin orientations for the extended q -state clock model with the choice $\varphi_0 = 0$ (left) and $\varphi_0 = -\pi$ (right). In this example, $q = 5.5$. Blue and red denote the leftover angles in the two cases. See Eq. (4.8).

We call the model (4.9) the “extended q -state” clock model for all values of q and the “fractional- q -state” clock model for fractional values of q . For integer q the extended q -state clock model reduces to the ordinary q -state clock model. Numerical results presented in later sections are from the extended q -state clock model.

For the extended q -state clock model, the degrees of freedom are discrete, so we write the partition function now as

$$Z = \sum_{\varphi^{(k)}} \prod_x \frac{d\varphi_x}{2\pi} e^{-\beta H}. \quad (4.10)$$

With the models defined we turn to observables. The main observables that we compute to study the critical behavior are the internal energy, magnetization, and their corresponding susceptibilities. These quantities are defined in the same way for both the continuous and discrete angle cases. The internal energy is defined as

$$E = \langle H_{\text{ext-}q} \rangle = \left\langle - \sum_{x,\mu} \cos(\varphi_{x+\hat{\mu}} - \varphi_x) \right\rangle = -\frac{\partial}{\partial \beta} \ln Z, \quad (4.11)$$

where $\langle \cdots \rangle$ denotes the ensemble average. The specific heat is

$$C = \frac{-\beta^2}{N} \frac{\partial E}{\partial \beta} = \frac{\beta^2}{N} (\langle E^2 \rangle - \langle E \rangle^2). \quad (4.12)$$

In addition we consider the magnetization

$$\vec{M} = \frac{\partial}{\partial \vec{h}} \ln Z = \left\langle \sum_x \vec{S}_x \right\rangle. \quad (4.13)$$

The magnetic susceptibility defined in a manifestly O(2)-invariant way is

$$\chi_{\vec{M}} = \frac{1}{N} \frac{\partial \vec{M}}{\partial \vec{h}} = \frac{1}{N} \left(\langle \vec{M} \cdot \vec{M} \rangle - \langle \vec{M} \rangle \cdot \langle \vec{M} \rangle \right). \quad (4.14)$$

We note that in Monte Carlo simulations at zero external field in a finite system in the absence of explicit symmetry breaking terms the definition of the spontaneous magnetization (4.13) gives $\langle \vec{M} \rangle = 0$. In such situations one often resorts to using a proxy observable [45–47, 50, 52, 78]

$$|\vec{M}| = \left\langle \left| \sum_x \vec{S}_x \right| \right\rangle \quad (4.15)$$

in place of $\vec{M}$. The corresponding susceptibility is

$$\chi_{|\vec{M}|} = \frac{1}{N} \left(\langle |\vec{M}|^2 \rangle - \langle |\vec{M}| \rangle^2 \right). \quad (4.16)$$

While one expects that $|\vec{M}|$ indicates the same critical behavior, in general, $|\vec{M}|$ is numerically different from $\vec{M}$ except deep in the ordered phase. Nevertheless, we expect both definitions of the magnetic susceptibility—Eqs. (4.14) and (4.16)—possess the same critical behavior, and can be relied upon to extract universal features. In the next section we detail the methods used to study the observables defined above.

4.2.2 Methods

4.2.2.1 Monte Carlo

The allowed spin orientations in the extended q -state clock model, given by Eq. (4.7), are discrete, and the model can be studied using a heatbath algorithm. The heatbath algorithm is a MCMC algorithm that drives the lattice toward equilibrium configurations by choosing the new spin at each update according to the probability distribution defined by its neighboring spins. We adapted FORTRAN code developed by Bernd Berg for the standard Potts model [5].

Initial exploration of the extended q -state clock model was performed via MC on a 4×4 lattice with zero external magnetic field.

We used a heatbath algorithm to study the extended q -state clock model on a 4×4 lattice. The general structure of the Monte Carlo algorithm is as follows:

```

# Equilibration:
for iequi = 1 to nequi do
    heatbath sweep
end

# Measurements:
for irpt = 1 to nrpt do
    for imeas = 1 to nmeas do
        for idisc = 1 to ndisc do
            heatbath sweep
        end
        measure observables
    end
    save measurements to file
end

```

Parameters used in the primary data production for the extended q -state clock model are given in Table 4.1. The total number of heatbath sweeps performed (not including equilibration) is $\text{nrpt} \times \text{nmeas} \times \text{ndisc}$. The number of measurements taken is $\text{nrpt} \times \text{nmeas}$. Ensemble averages and error bars for the energy and magnetization were calculated after binning with nrpt bins each of size nmeas . For specific heat and magnetic susceptibility, the measurements were binned and jackknifed.

We studied the extended q -state clock model on a 4×4 lattice with zero external magnetic field. For each β , we initialized to a random lattice (hot start) then we performed 2^{15} equilibrating sweeps followed by 2^{22} measurement sweeps. Each measurement sweep was followed by 2^8 discarded sweeps. We calculated the energy density and specific heat as defined in Eqs. (4.11) and (4.12). We calculated the proxy magnetization and susceptibility as defined in Eqs. (4.15) and (4.16). Results for $1 < q < 6$ are shown in Figs. 4.2–4.6.

In the large- β (low temperature) regime, the heatbath algorithm has difficulty appropriately sampling the configuration space. At large- β , the lattice freezes along a particular magnetization direction. When $q \in \mathbb{Z}$ all magnetization directions are equivalent, however, when $q \notin \mathbb{Z}$, the discrete rotational symmetry is broken and the direction of magnetization matters. The configuration space is split into two sectors with different thermodynamic properties. In one sector, the magnetization is in the direction 0 or $-\tilde{\phi}$, defined in Eq. (4.8), where relatively large fluctuations

Table 4.1 The Monte Carlo parameters used in the primary data production for the extended q -state clock model. See Figs. 4.2–4.6. Here, q refers to the q -state clock model, h is the external magnetic field, $nequi$ is the number of equilibrating heatbath sweeps used, $nrpt$ is the repetitions of measurement sweeps, $nmeas$ is the number of measurement sweeps, and $ndisc$ is the number of sweeps discarded between each measurement sweep.

q	$[1.1, 6.0]$ with $\Delta q = 0.1$
β	$[0.0, 2.0]$ with $\Delta\beta = 0.01$
h	0.0
Lattice	4×4
Start type	Random (hot)
$nequi$	2^{15}
$nrpt$	2^6
$nmeas$	2^{16}
$ndisc$	2^8




Figure 4.2 Monte Carlo results for the extended q -state clock model on a 4×4 lattice for $1.1 \leq q \leq 2.0$. The top panel shows energy density and specific heat, and the bottom panel shows proxy magnetization and magnetic susceptibility. Statistical error bars are omitted since they are smaller than the line thickness.




Figure 4.3 Monte Carlo results for the extended q -state clock model on a 4×4 lattice for $2.1 \leq q \leq 3.0$. The top panel shows energy density and specific heat, and the bottom panel shows proxy magnetization and magnetic susceptibility. Statistical error bars are omitted since they are smaller than the line thickness. Dashed lines indicate regions where we have data but we do not have the uncertainty fully under control.




Figure 4.4 Monte Carlo results for the extended q -state clock model on a 4×4 lattice for $3.1 \leq q \leq 4.0$. The top panel shows energy density and specific heat, and the bottom panel shows proxy magnetization and magnetic susceptibility. Statistical error bars are omitted since they are smaller than the line thickness. Dashed lines indicate regions where we have data but we do not have the uncertainty fully under control.




Figure 4.5 Monte Carlo results for the extended q -state clock model on a 4×4 lattice for $4.1 \leq q \leq 5.0$. The top panel shows energy density and specific heat, and the bottom panel shows proxy magnetization and magnetic susceptibility. Statistical error bars are omitted since they are smaller than the line thickness. Dashed lines indicate regions where we have data but we do not have the uncertainty fully under control.




Figure 4.6 Monte Carlo results for the extended q -state clock model on a 4×4 lattice for $5.1 \leq q \leq 6.0$. The top panel shows energy density and specific heat, and the bottom panel shows proxy magnetization and magnetic susceptibility. Statistical error bars are omitted since they are smaller than the line thickness.




Figure 4.7 The measured energy density for $q = 4.5$ at different number of heatbath sweeps for different random number generator (RNG) seeds. The vertical axis is the energy density. The horizontal axis gives the number of heatbath sweeps used (not including equilibration sweeps). The different colors indicate runs with different seeds for the RNG. All runs were performed at $\beta = 2.5$ with hot-started 4×4 lattices. This shows that the configuration space consists of two distinct sectors, and the heatbath algorithm has difficulty appropriately sampling both sectors unless the number of heatbath sweeps is taken very large.

may still be possible due to the spins being able to flip over the relatively small “leftover angle”. In the other sector, the magnetization is in one of the other directions where fluctuations are less likely. To appropriately sample the configuration space one has to use very large statistics or run multiple heatbath streams at the same parameters but with different seeds for the random number generator. In Fig. 4.7, we show an example of this phenomenon for $q = 4.5$ and $\beta = 2.5$ on a 4×4 lattice. In this example, we need $\gtrsim 2^{32}$ heatbath sweeps to adequately sample both sectors. This MC slowdown makes it difficult to study larger lattices, and is a strong motivation for using an alternative method such as TRG.

The Monte Carlo slowdown is illustrated by an explosion of the integrated autocorrelation time in the intermediate- β regime. For an observable O , an estimator of the integrated autocorrelation time is given by

$$\tilde{\tau}_{O,int} = 1 + 2 \sum_{t=1}^T \frac{C(t)}{C(0)} \quad (4.17)$$

where $C(t) = \langle O_i O_{i+t} \rangle - \langle O_i \rangle \langle O_{i+t} \rangle$ is the correlation function between the observable O measured at Markov times i and $i + t$. The integrated autocorrelation time $\tau_{O,int}$ is estimated by finding a window in t for which $\tilde{\tau}_{O,int}$ is nearly independent of t . The integrated autocorrelation time for




Figure 4.8 The integrated autocorrelation time of the energy density for several q on a 4×4 lattice using a heatbath algorithm. At large β , the integrated autocorrelation time τ_{int} grows abruptly when $q \notin \mathbb{Z}$. Note the log scale on the vertical axis. Connecting lines are included to guide the eyes. The rapid increase of the integrated autocorrelation time for noninteger q occurs because the MCMC approach is having difficulty sampling the full configuration space. Details are given in the main text.

$4 < q \leq 5$ is shown in Fig. 4.8. The values of T needed to extract these points are given in Fig. 4.9. The rapid increase of the integrated autocorrelation times for noninteger q we believe occurs because the MCMC approach is having difficulty sampling the full configuration space as described above. This effect is discussed in much more detail in Section 4.3.7.

To mitigate the effect of autocorrelation in our results, we discarded 2^8 heatbath sweeps between each saved measurement. The saved measurements were then binned (i.e. preaveraged) with bin size 2^{16} before calculating the means and variances.

4.2.2.2 Tensor Renormalization Group (TRG)

For this model, the heatbath approach suffers from a slowdown that makes it difficult to study the large- β regime already on very small lattices. An alternative approach, the tensor renormalization group (TRG), which does not suffer from this slowdown, was used to study the model on much larger lattices and in the thermodynamic limit. This allows us to perform finite-size scaling and characterize the phase transitions in the extended q -state clock models. The TRG results are validated by comparison with exact and Monte Carlo results on small lattices (see Appendix B.2).




Figure 4.9 The values of T defined in Eq. (4.17) needed to extract the integrated autocorrelation times shown in Fig. 4.8. The vertical axis gives the values of q and the horizontal axis gives the values of β . Blank/white regions in this heatmap indicate cases where $T > 2^{16}$ is needed to get a reliable estimate of the integrated autocorrelation time. These were not attempted due to the computational cost.

The TRG methods are not exact. Because the bond dimension, D_{bond} , of the coarse-grained tensor increases exponentially with the renormalization group (RG) steps, a truncation for D_{bond} must be applied to avoid uncontrolled growth of memory needs on classical computers. For noncritical phases, the fixed-point tensor of the RG flow has small D_{bond} , but a larger bond dimension is needed to have the correct RG flow near the critical point. It has been reported that TRG methods using $D_{\text{bond}} = 40$ can locate the phase transition point with an error of order 10^{-4} for the Ising model [113], and of order 10^{-3} for the O(2) model [24, 113] and the clock models [50, 113].

To perform TRG calculations, we need to express the partition function as a contraction of a tensor network. We rewrite the weight of each link by a singular value decomposition (SVD):

$$e^{\beta \cos\left[\frac{2\pi}{q}(k_{x+\hat{\mu}} - k_x)\right]} = \sum_{n_{x,\mu}=0}^{\lfloor q \rfloor} U_{k_{x+\hat{\mu}} n_{x,\mu}} G_{n_{x,\mu}} V_{k_x n_{x,\mu}}. \quad (4.18)$$

Then we sum over the original k indices and the partition function can be expressed in the dual space from the expansion in terms of n indices

$$Z = \sum_{\{n\}} \prod_x T_{lrd}, \quad (4.19)$$

where $l = n_{x-\hat{s},s}$, $r = n_{x,s}$, $d = n_{x-\hat{\tau},\tau}$, $u = n_{x,\tau}$ for each site x , and the local rank-four tensor is defined as

$$T_{lrdu} = \sqrt{G_l G_r G_d G_u} C_{lrdu}, \quad (4.20)$$

$$C_{lrdu} = \sum_{k_x=0}^{\lfloor q \rfloor} e^{h \cos\left(\frac{2\pi}{q} k_x - \psi_h\right)} U_{k_x l} V_{k_x r} U_{k_x d} V_{k_x u}. \quad (4.21)$$

Notice that for integer q , the matrix $U = V^{-1}$ can be chosen as $U_{kn} = \exp(i2\pi kn/q)$, then if $h = 0$, the tensor C_{lrdu} becomes a δ -function that gives a $\mathbb{Z}_q$ selection rule for values of n :

$$\text{mod } (n_{x-\hat{s},s} + n_{x-\hat{\tau},\tau} - n_{x,s} - n_{x,\tau}, q) = 0. \quad (4.22)$$

The tensor reformulation of the expectation value of a local observable can be obtained in the same way. For example, the first component of the magnetization is equal to the expectation value of $\cos(\varphi)$ at an arbitrary site x_0 , $m_1 = \langle \cos(\varphi_{x_0}^{(k)}) \rangle$, which can be expressed as

$$m_1 = \frac{\sum_{\{n\}} T_{x_0,lrdu}^i \prod_{x \neq x_0} T_{lrdu}}{\sum_{\{n\}} \prod_x T_{lrdu}}, \quad (4.23)$$

where $T_{x_0,lrdu}^i = \sqrt{G_l G_r G_d G_u} C_{x_0,lrdu}^i$ is an impure tensor residing at site x_0 , and

$$C_{x_0,lrdu}^i = \sum_{k_{x_0}=0}^{\lfloor q \rfloor} \cos\left(\frac{2\pi}{q} k_{x_0}\right) e^{h \cos\left(\frac{2\pi}{q} k_{x_0} - \psi_h\right)} U_{k_{x_0} l} V_{k_{x_0} r} U_{k_{x_0} d} V_{k_{x_0} u}. \quad (4.24)$$

To compute the internal energy, we need to calculate the expectation values of link interactions $\epsilon_\mu = \langle \cos(\varphi_{x_0+\hat{\mu}}^{(k)} - \varphi_{x_0}^{(k)}) \rangle$. Taking $\mu = s$ as an example, we perform another SVD for the target link

$$\cos\left[\frac{2\pi}{q} (k_{x_0+\hat{s}} - k_{x_0})\right] e^{\beta \cos\left[\frac{2\pi}{q} (k_{x_0+\hat{s}} - k_{x_0})\right]} = \sum_{n_{x_0,s}=0}^{\lfloor q \rfloor} U_{k_{x_0+\hat{s}} n_{x_0,s}}^i G_{n_{x_0,s}}^i V_{k_{x_0} n_{x_0,s}}^i, \quad (4.25)$$

and introduce two impure tensors $\tilde{T}_{x_0+\hat{s},lrdu}^i, \tilde{T}_{x_0,lrdu}^i$ residing at nearest neighbor sites $x_0 + \hat{s}, x_0$, by replacing $U_{k_{x_0+\hat{s}},l}, G_l$ with $U_{k_{x_0+\hat{s}},l}^i, G_l^i$ and replacing $V_{k_{x_0},r}, G_r$ with $V_{k_{x_0},r}^i, G_r^i$ in Eq. (4.20), respectively. Thus the tensor reformulation of the expectation value of the link interaction is written as

$$\epsilon_\mu = \frac{\sum_{\{n\}} \tilde{T}_{x_0+\hat{\mu},lrdu}^i \tilde{T}_{x_0,lrdu}^i \prod_{x \neq x_0+\hat{\mu}, x_0} T_{lrdu}}{\sum_{\{n\}} \prod_x T_{lrdu}}. \quad (4.26)$$

In the following, we use TRG and higher-order TRG (HOTRG) to contract tensor networks with impure tensors [130, 131] up to a volume $N = L^2 = 2^{24} \times 2^{24}$, calculate the first component of the magnetization $\vec{m} = (m_1, m_2)$ and internal energy $E = \epsilon_s + \epsilon_\tau$, and take derivatives of $\vec{m}$ and E with respect to $\vec{h}$ and β , respectively, to find the magnetic susceptibility and specific heat.^{1,2} The locations and heights of the peaks of χ_M and C_v are obtained via a spline interpolation on datasets with $\Delta\beta = 10^{-3}$. The tensor contraction in HOTRG is performed with ITENSORS JULIA LIBRARY [132].

4.2.3 Thermodynamics

In the extended q -state clock model, there is a $\mathbb{Z}_q$ symmetry when $q \in \mathbb{Z}$. When $q \notin \mathbb{Z}$, this symmetry is explicitly broken. We choose $\varphi_0 = 0$, so the allowed spin orientations divide the unit circle into $[q]$ arcs of which $[q] - 1$ have measure $2\pi/q$. The remainder has measure $\tilde{\phi}$ given in Eq. (4.8) and illustrated in Fig. 4.1. There remains a $\mathbb{Z}_2$ symmetry and an approximate $\mathbb{Z}_{[q]}$ symmetry.

Monte Carlo results obtained with a heatbath algorithm on a 4×4 lattice with zero external field are shown for $1.1 \leq q \leq 6.0$ in Fig. 4.2–4.6. The four panels show the energy density and the specific heat defined in Eqs. (4.11) and (4.12) as well as the proxy magnetization and susceptibility defined in Eqs. (4.15) and (4.16). For example, for $q = 5$, the energy density is zero at $\beta = 0$ because there is no linear term in the series expansion of the partition function due to the $\mathbb{Z}_5$ symmetry. The nonzero energy density at $\beta = 0$ for $4 < q < 5$ is consistent with the explicit $\mathbb{Z}_5$ symmetry breaking. As $q \rightarrow 4^+$, the stronger symmetry breaking results in a more negative energy density. There is a double-peak structure in the specific heat, where the large- β peak moves toward $\beta = \infty$ as q is decreased. The proxy magnetization also increases for smaller values of q with stronger symmetry breaking, and the peak of the magnetic susceptibility moves toward smaller β values. Note that the double-peak structure of this proxy magnetic susceptibility will appear at larger system sizes $L \geq 12$ [46]. The true magnetic susceptibility with an external field at large volumes will show the

¹One can assume that the TRG and the HOTRG return the same results if the bond dimension is sufficiently large.

²If it is not declared in the main text or in the captions, the bond dimension is set to be sufficiently large so that the outputs converges.




Figure 4.10 The energy density for the fractional- q -state clock model from TRG for $q = 4.1, 4.3, 4.5, 4.7$, and 4.9 . These results were obtained with a volume 1024×1024 . The energy is shifted vertically by 2 (to make it positive) and plotted on a log scale to better illustrate the difference between the curves.

double-peak structure (see below).

In Fig. 4.10, we show the logarithm of the energy density of this model from TRG for $q = 4.1, 4.3, 4.5, 4.7$, and 4.9 with $0 \leq \beta \leq 16$ and $N = 1024 \times 1024$. For large enough β , we have

$$Z = e^{2N\beta} \left[[q] + 2Ne^{-4\beta(1-\cos \tilde{\phi})} + \dots \right], \quad (4.27)$$

so that the energy density converges exponentially with β . The results in Fig. 4.10 confirm this behavior. We also notice that for smaller q , there is a larger range of β where the energy density does not change much. In this range, $1 - \cos \tilde{\phi}$ is close to zero for q close to 4 from above, and the inverse temperature β is large enough that terms containing larger angular distances are negligible, but β is still not large enough to change the values of $\exp[-\beta(1 - \cos \tilde{\phi})]$ significantly from 1 and ignore higher orders. The result for $q = 4.1$ shown in Fig. 4.10 indicates that the specific heat is almost zero for $\beta > 2.5$, which is confirmed in Fig. 4.11.

In Fig. 4.11, we show the specific heat for $q = 4.1, 4.5, 4.9$, and 5.0 at volumes ranging from 4×4 to 128×128 . For generic q , there are two peaks in the specific heat.³ In Fig. 4.11 we see only

³There is only a single peak for $q = 2, 3, 4$ and for fractional q just below these integers.




Figure 4.11 Specific heat of the extended q -state clock model from TRG for $q = 4.1, 4.5, 4.9$, and 5.0 at volumes from 4×4 up to 128×128 . All vertical axes use a shared scale, and all horizontal axes use a shared scale. In general, there is a double-peak structure in the specific heat (for $q = 4.1$, the second peak is at $\beta \sim 75$). Insets show the height of the second peak versus the linear system size $L = \sqrt{N}$ plotted on a log scale, where the solid line is a linear fit.

a single peak for $q = 4.1$ since the second peak is at much larger β . For $q \notin \mathbb{Z}$ and not too close to 5 , the first peak shows little or no dependence on volume. The second peak grows logarithmically with volume, as shown in the insets for $q = 4.5, 4.9$. This is in contrast to the integer case $q = 5$ where there are two BKT transitions and both peaks show little dependence on volume for lattice sizes larger than 32×32 . Because the specific heat is the second-order derivative of free energy, the results in Fig. 4.11 indicate that the first peak is associated with either a crossover or a phase transition with an order larger than 2 , and the second peak is associated with a second-order phase transition. To conclusively characterize the phase transitions, if any, associated with these two peaks in the fractional- q -state clock model, we study the magnetic susceptibility in the next two subsections.

We find that the thermodynamic curves vary smoothly for $n < q \leq n + 1$ where n is an integer. When q is taken slightly larger than n from below, these curves change abruptly since an additional degree of freedom is introduced. The specific heat exhibits a double-peak structure with the second




Figure 4.12 The magnetic susceptibility as a function of β for finite volumes. $D_{\text{bond}} = 40$.

peak at very large β . As q is increased further, this second peak moves toward small β , until at $q = n + 1$, the thermodynamic curves of the integer- $(n + 1)$ -state clock model are recovered.

In the small- β (high temperature) regime, all allowed angles are nearly equally accessible, and the model behaves approximately like a $[q]$ -state clock model. The model is dominated by the approximate $\mathbb{Z}_{[q]}$ symmetry, and there is a peak in the specific heat. In Fig. 4.8, we show that at intermediate β , an explosion of the integrated autocorrelation time of the energy is observed in the MC simulation as the model quickly reduces to a rescaled Ising model. At large β , the configuration space separates into thermodynamically distinct sectors, and the Markov chain has trouble adequately sampling both sectors. This is discussed further in Section 4.2.2.1. At large- β , spin flips across the “leftover angle” $\tilde{\phi}$ are strongly favored relative to spin flips across the other angles. Thus, in the large- β regime, the model behaves as a rescaled Ising model. The existence of an Ising critical point is conclusively established via TRG in Sec. 4.2.5.

We next present our TRG results and discuss the phase transitions in the fractional- q -state clock model in the rest of this section. We first present the results for the magnetic susceptibility without an external field at small volumes in Fig. 4.12. For $q < 5$, there is a small- β peak converging




Figure 4.13 The magnetic susceptibility as a function of β at $N = 2^{24} \times 2^{24}$. For $q = 4.995$, the external field is $h = 4 \times 10^{-5}$. For $q = 4.999, 4.9999$, the external field is $h = 4 \times 10^{-5}$ for the lower curve and $h = 2 \times 10^{-5}$ for the upper curve.

quickly with volume, which means the peak is associated with a crossover. As q is increased, there is a high plateau moving toward small β . The height of the plateau increases with volume as a power law (notice the logarithmic scale in the y-axis). The divergent plateau signals a phase transition. As there is no spontaneous symmetry breaking in small volumes, the response of the system to an external field remains high at low temperatures, so we cannot see the transition from a critical phase to the symmetry-breaking phase where there is a small magnetic susceptibility based on these results for $h = 0$ and small volumes. But notice that for fractional q like $q = 4.9$, there is a higher plateau after the first one for each volume. This is due to an “approximate $\mathbb{Z}_{[q]}$ symmetry breaking” after a $\mathbb{Z}_2$ symmetry breaking. This “approximate $\mathbb{Z}_{[q]}$ symmetry breaking” is not a true phase transition so it moves to $\beta = \infty$ in the thermodynamic limit as shown in the results for $q = 4.9$. There is only one $\mathbb{Z}_5$ symmetry breaking for $q = 5$ so there is a single plateau for each volume. We will use the magnetic susceptibility with a weak external field in the thermodynamic limit to detect the phase transitions.

In Fig. 4.13 we present the magnetic susceptibility in the thermodynamic limit as a function of β for $q = 4.995, 4.999, 4.9999$ with a small external field $h = 4 \times 10^{-5}, 2 \times 10^{-5}$. The height of the large- β peak is large for all three values of q , indicating a phase transition near this peak. The small-

β peak for $q = 4.995$, $h = 4 \times 10^{-5}$ is invisible because it is very small. For $q = 4.999$, $h = 4 \times 10^{-5}$, closer to 5, the small- β peak is higher, and the large- β peak moves toward the large- β BKT transition point $\beta_{q=5,c2}^{\text{BKT}}$ for $q = 5$. When the external field is decreased to $h = 2 \times 10^{-5}$, the large- β peak height is almost doubled, while the small- β peak height does not change, which means there is no phase transition near the small- β peak. For $q = 4.9999$ (closer to 5) with $h = 4 \times 10^{-5}$, one can see that the small- β peak becomes higher than the large- β one, and the large- β peak is fading away, which is consistent with the results in Refs. [50, 113] for the five-state clock model. When the external field is decreased to $h = 2 \times 10^{-5}$, the large- β peak grows a much larger amount of height than the small- β one does and becomes higher than the small- β one. We have confirmed that the small- β peak will eventually converge for small enough h . All these behaviors are evidence that for all fractional $q > 4$, the small- β peak in χ_M does not diverge, and the large- β peak diverges at $h = 0$ indicating a phase transition.

In the following, we discuss the behavior of the two peaks in the thermodynamic limit in detail. The main observations are that for $\varphi_0 = 0$ and $q > 4$, there are two peaks in the specific heat and the magnetic susceptibility, the small- β one is finite and is associated with a crossover, and the large- β one diverges which is characteristic of an Ising critical point. When q is approaching an integer from below, the height of the crossover peak of the magnetic susceptibility diverges as $[q] - q$ goes to zero with a power law:

$$\chi_M^* \sim ([q] - q)^{-y}. \quad (4.28)$$

We can formulate the scaling hypothesis with $\Delta q = [q] - q$ and h ,

$$f(\lambda^p \Delta q, \lambda^r h) = \lambda^d f(\Delta q, h), \quad (4.29)$$

where λ parameterizes a scale transformation, p and r are the scaling dimensions, and $d = 2$ in two-dimensional space. Notice that the reduced temperature should not enter the homogeneous function independently because there is an essential singularity in the correlation length as a function of temperature for BKT transitions. We assume Eq. (4.29) holds for any critical temperature, in particular, the BKT crossover peak position $\beta^{\text{BKT}}(\Delta q, h)$, which is a power law function of Δq and

h , considered in the following calculations. Then the magnetization and the magnetic susceptibility satisfy the following relations:

$$\lambda^r M(\lambda^p \Delta q, \lambda^r h) = \lambda^d M(\Delta q, h), \quad (4.30)$$

$$\lambda^{2r} \chi_M(\lambda^p \Delta q, \lambda^r h) = \lambda^d \chi_M(\Delta q, h), \quad (4.31)$$

from which one can obtain

$$M \sim h^{\frac{1}{\delta}} \sim (\Delta q)^{\frac{1}{\delta'}}, \quad (4.32)$$

$$\chi_M \sim h^{-1+\frac{1}{\delta}} \sim (\Delta q)^{-\frac{\delta-1}{\delta'}}, \quad (4.33)$$

where $1/\delta = (d-r)/r$, $1/\delta' = (d-r)/p$. We then have $\delta' = (\delta-1)/y = 14/y$, where we have used the fact that $\delta = 15$ for BKT transitions. Note that the expansion of the action to the first order in Δq is

$$S_{\text{ext-}q} = S_{\mathbb{Z}_{[q]}} + \beta \frac{\Delta q}{q} \Delta \varphi_{x,\mu} \sin(\Delta \varphi_{x,\mu}), \quad (4.34)$$

where $S_{\mathbb{Z}_{[q]}}$ is the action for the integer- q clock model. The perturbation term that breaks the $\mathbb{Z}_{[q]}$ symmetry has a very different form from the one for the external field $h \cos(\varphi_x)$. We numerically show in the following that δ' is equal to the magnetic critical exponent $\delta = 15$ for the BKT transitions and Ising critical points in two dimensions. Thus if we define a new susceptibility as $\partial M / \partial \Delta q$, which should scale as $(\Delta q)^{-1+1/\delta'}$, the exponent $y' = 1 - 1/\delta' = 14/15$, still the same as y . However, calculating δ' from y' requires higher accuracy since $\delta' = 1/(1-y')$ where one significant digit is subtracted in the denominator. Both the peak position of the crossover and the Ising critical point go to the two BKT transition points at integer q with the same power law, which provides us a new way to locate the phase transitions in models with $\mathbb{Z}_n$ symmetry. However, for $\varphi_0 = -\pi$ and odd $[q]$, both peaks of χ_M are finite for fractional q so there are no critical points. For $q \rightarrow 5^+$, only the small- β peak can be used to extract the BKT transition of $q = 5$ since the large- β peak fades away.

4.2.4 Small- β peak: Crossover

We have shown in Fig. 4.13 that for $q = 4.999$, the height of the small- β peak converges to about 400 for small enough external field. The dependence of the peak height on the external field




Figure 4.14 The small- β peak of magnetic susceptibility as a function of β for $q = 4.3$. The peak height converges when the external field is taken to zero.

is larger for values of q closer to an integer from below. In Fig. 4.14, we show another example for $q = 4.3$. One can see that the h dependence of the peak height is much smaller. The peak height at $h = 10^{-2}$ differs from the value in the $h \rightarrow 0$ limit by only about 0.04. As the external field is decreased, the peak height converges to a constant $\chi_M \approx 1.2$, implying that there is no phase transition around this peak. This is true for all fractional q . Thus, for fractional q , the first peak in the specific heat is associated with a crossover rather than a true phase transition. As q approaches 5 from below, we expect the small- β peak height to diverge because there is a BKT transition for $q = 5$, and we expect the location of the peak to go to the small- β BKT transition point.

To check this, we calculate the converged peak height χ_M^* of the magnetic susceptibility for values of $5 - q \leq 10^{-4}$ with $D_{\text{bond}} = 40$ and $D_{\text{bond}} = 50$, where we use $dh = 10^{-10}$ in the numerical differentiation. We plot χ_M^* versus $5 - q$ in Fig. 4.15, where we see that the dependence of the peak height on D_{bond} is very small for $D_{\text{bond}} \geq 40$. Applying a linear fit to $\ln(\chi_M^*)$ versus $\ln(5 - q)$ shows that the peak height diverges as $q \rightarrow 5^-$ with a power law

$$\chi_M^* \sim \begin{cases} (5 - q)^{-0.948(6)} & \text{if } D_{\text{bond}} = 40 \\ (5 - q)^{-0.9389(6)} & \text{if } D_{\text{bond}} = 50, \end{cases} \quad (4.35)$$

from which we obtain $\delta' = 14.77(9)$ for $D_{\text{bond}} = 40$ and $\delta' = 14.911(10)$ for $D_{\text{bond}} = 50$. The




Figure 4.15 The log-log plot of the maximal value of the magnetic susceptibility χ_M for the small- β peak as a function of $5 - q$. The peak height diverges with a power law when $q \rightarrow 5^-$.




Figure 4.16 Same as Fig. 4.15, but for $q \rightarrow 6^-$.

value of δ' is close to the magnetic critical exponent $\delta = 15$ for BKT transitions and Ising critical points in two dimensions, but now $5 - q$ (rather than the external field h) is playing the role of the symmetry-breaking parameter. Since both h and $5 - q$ break the $\mathbb{Z}_5$ symmetry to a $\mathbb{Z}_2$ symmetry, the agreement on the critical exponents is reasonable. The value of δ' is also checked for $q \rightarrow 6^-$ in Fig. 4.16, where a larger D_{bond} is applied in HOTRG. In this case, the linear fit gives

$$\chi_M^* \sim \begin{cases} (6 - q)^{-0.937(2)} & \text{if } D_{\text{bond}} = 50 \\ (6 - q)^{-0.9329(16)} & \text{if } D_{\text{bond}} = 60, \end{cases} \quad (4.36)$$

from which we obtain $\delta' = 14.94(3)$ with $D_{\text{bond}} = 50$ and $\delta' = 15.007(26)$ with $D_{\text{bond}} = 60$. Within uncertainty, the results of $q \rightarrow 6^-$, $D_{\text{bond}} = 60$ agree perfectly with the value of the magnetic critical exponent $\delta = 15$.

The location of the converged peak β_p as a function of $5 - q$ is depicted in Fig. 4.17. The discrepancy of the peak positions between $D_{\text{bond}} = 40$ and $D_{\text{bond}} = 50$ is invisible at $5 - q = 10^{-4}$ and increases as $5 - q$ is decreased, which is reasonable because a larger bond dimension is needed for systems closer to a critical point. The overall discrepancy is small. We extrapolate the peak position to $q = 5$ with a power law and obtain the small- β BKT transition point $\beta_{q=5,c1}^{\text{BKT}} = 1.0494(6)$ with $D_{\text{bond}} = 40$ and $\beta_{q=5,c1}^{\text{BKT}} = 1.0506(4)$ with $D_{\text{bond}} = 50$ for the five-state clock model. The results are consistent with 1.0503(2) in Ref. [51]. The same procedure is performed for $q \rightarrow 6^-$ in Fig. 4.18. The peak positions β_p for $D_{\text{bond}} = 50$ and those for $D_{\text{bond}} = 60$ have little discrepancy for $6 - q \geq 7 \times 10^{-7}$. For $6 - q < 7 \times 10^{-7}$, β_p for $D_{\text{bond}} = 60$ becomes larger than that for $D_{\text{bond}} = 50$. The power law fit gives us $\beta_{q=6,c1}^{\text{BKT}} = 1.0983(4)$ with $D_{\text{bond}} = 50$ and $\beta_{q=6,c1}^{\text{BKT}} = 1.1019(5)$ for $D_{\text{bond}} = 60$. The results are consistent with 1.101(4) in Ref. [133].

We have shown that $[q] - q$ plays the same role as an external magnetic field for the small- β BKT transition in integer- q -state clock models. The magnetic susceptibility is always finite for fractional $q < [q]$, and diverges as $q \rightarrow [q]^-$ with a critical exponent $y = 14/15$. This provides us an alternative way to extract the locations of BKT transitions in clock models. However, the situation is very different for the large- β peak of the magnetic susceptibility.

4.2.5 Large- β peak: Ising criticality

To understand the large- β peak in the specific heat, we again study the magnetic susceptibility χ_M . For a fixed q , the critical point, if any, is given by the location of the peak of χ_M in the limit $h \rightarrow 0$ where the height χ_M^* of the peak is infinite. A power law extrapolation to $h = 0$ is performed on peak positions of χ_M for small values of h . In Fig. 4.19, we present the peak height χ_M^* of the susceptibility as a function of the external field h . The linear fit of $\ln(\chi_M^*)$ versus $\ln(h)$ gives $\chi_M^* \sim h^{-0.93318(16)}$, from which we obtain the magnetic critical exponent $\delta = 14.97(4)$. This value is consistent with the value $\delta = 15$ of the BKT transitions and Ising critical points. We have shown




Figure 4.17 The power law extrapolation of the small- β peak position to $q = 5$ from below. Here the extrapolation gives $\beta_c = 1.0506(4)$.




Figure 4.18 Same as Fig. 4.17, but for $q \rightarrow 6^-$.

before that there is no phase transition around the small- β peak of χ_M . A BKT transition should be accompanied with a continuous critical region, so the divergent large- β peak of χ_M must be an Ising critical point. In the Ising universality class,

$$\begin{aligned}
 \xi &\sim |\beta - \beta_c|^{-\nu} \sim L, \\
 M &\sim (\beta - \beta_c)^\beta \sim L^{-1/8} \sim h^{1/15}, \\
 \chi_M &\sim |\beta - \beta_c|^{-\gamma} \sim L^{7/4},
 \end{aligned} \tag{4.37}$$




Figure 4.19 The maximal magnetic susceptibility in Fig. 4.22 as a function of the external field for $q = 4.9$. A linear fit of the log-log plot gives the exponent $\delta = 14.97(4)$. This is consistent with the value $\delta = 15$ of the Ising universality class. $D_{\text{bond}} = 40$.

where $\nu = 1, \beta = 1/8, \gamma = 7/4$ are the universal critical exponents. There is a universal function relating $\chi_M/L^{1.75}$ and $L(\beta - \beta_c)$ with fixed $hL^{15/8}$. In Fig. 4.21, we plot $\chi_M/L^{1.75}$ versus $L(\beta - \beta_c)$ for various lattice sizes around the large- β peak of χ_M , where the value of β_c is obtained from the Ising approximation in Eq. (4.40) described below. One can see that all the data collapse onto a single curve, which gives strong evidence that this is a critical point of the Ising universality class.

In Fig. 4.20, we show the logarithm of χ_M as a function of $\log_2(L)$ for $\beta = 1.1$. For $q = 5$, $\beta = 1.1$ is between two BKT points so is a critical point, and χ_M keeps increasing with a positive power of L as expected. When $q < 5$, even for a very small $5 - q = 10^{-6}$, χ_M always saturates at a large enough volume. These results again prove that there are no BKT transitions in fractional q , no matter how close to $[q]$ the q is. Because the maximal value of χ_M should diverge as a negative power of $5 - q$ when $q \rightarrow 5^-$, the increment of the height of the plateaus between $5 - q = 10^{-n}$ and $5 - q = 10^{-n-1}$ should be a constant for any integer n , which is also confirmed in Fig. 4.20.

We next obtain the Ising critical point by extrapolating the peak position of χ_M to $h = 0$. An example for $q = 4.9$ is shown in Fig. 4.22. The power law extrapolation gives $\beta_c = 1.44614(2)$. We can repeat the same procedure for other values of q . But notice that we need a larger bond dimension in TRG when q is very close to an integer from below, because more degrees of freedom become important, and the critical point is close to a BKT transition point. The phase transition in




Figure 4.20 The dependence of magnetic susceptibility on the size of the system L for $\beta = 1.1$ with $D_{\text{bond}} = 40$. The line on the circles is a linear fit for the first ten points: $1.275(3) \log_2(L) - 0.74(2)$.




Figure 4.21 Data collapse of the rescaled magnetic susceptibility versus the rescaled inverse temperature for $q = 4.3$. The reduced external field is $hL^{15/8} = 40$, and $\beta_c \approx 9.3216$, which is the approximate critical point of the model with $q = 4.3$, is obtained from Eq. (4.40).




Figure 4.22 The power law extrapolation of the peak position of χ_M to zero external field for $q = 4.9$. Here the extrapolation gives $\beta_c = 1.44614(2)$. $D_{\text{bond}} = 40$.

the Ising universality class is a transition from a disordered phase to a symmetry-breaking phase. The structure of the fixed-point tensor in TRG can easily characterize this phase transition. As proposed in Ref. [134], the symmetry-breaking indicator

$$X = \frac{(\sum_{ru} T_{rruu})^2}{\sum_{lrdu} T_{lrdu} T_{rldd}} \quad (4.38)$$

should be 1 in the disordered phase and 2 in the $\mathbb{Z}_2$ symmetry breaking phase. Thus the discontinuity in X for the fixed point tensor as a function of β can be used to locate the phase transition. An example for $q = 4.9$ is shown in Fig. 4.23, where the value of X changes from 1 to 2 at $1.4461 < \beta < 1.4462$, consistent with the result from the extrapolation of the peak position of χ_M in Fig. 4.22. The advantage of this method is that we only need to contract a single tensor network for each value of β and scan a β range once to locate the phase transition point. This saves us a lot of computational effort when we extrapolate the Ising critical point to $q = \lceil q \rceil$ from below.

Now we use the value of X to locate the Ising critical point and find the large- β BKT transitions for $q = 5, 6$. Notice that for $q = 5$, although both a small external field and a small deviation of q from an integer break the $\mathbb{Z}_{\lceil q \rceil}$ symmetry to a $\mathbb{Z}_2$ symmetry, the magnetic susceptibility with a weak external field does not have a peak around the large- β BKT transition, so it fails to predict the location of the phase transition [50, 113], but here we always have an Ising critical point for fractional q . In Fig. 4.24(a), we calculate the Ising critical point for $q \rightarrow 5^-$ with $D_{\text{bond}} = 40$ and




Figure 4.23 The β dependence of X , defined in Eq. (4.38), from the fixed point tensor for $q = 4.9$. The discontinuity is located between 1.4461 and 1.4462, consistent with the result from the extrapolation of the peak position of χ_M to zero external field in Fig. 4.22. The tensor bond dimension is 40.

extrapolate the result to $q = 5$ with a power law. The value of $5 - q$ is between 6×10^{-4} and 10^{-3} , where the dependence of β_c on D_{bond} is small. The extrapolated BKT transition point for $q = 5$ is $\beta_{q=5,c2}^{\text{BKT}} = 1.1027(14)$, consistent with the result 1.1039(2) obtained in Ref. [51]. As a comparison, we also present the extrapolation of the crossover peak with the same D_{bond} in the same figure. The exponents of the two power law scalings are the same within uncertainties, and the values of the exponents are consistent with 0.2677(84) obtained in Ref. [50] for magnetic susceptibility with an external field $h \leq 10^{-3}$.

The results for $q \rightarrow 6^-$ are shown in Fig. 4.24(b), where we use a larger bond dimension $D_{\text{bond}} = 60$ and data with $6 - q \leq 10^{-4}$. The extrapolated BKT transition point is $\beta_{q=6,c2}^{\text{BKT}} = 1.435(3)$, consistent with 1.441(6) in Ref. [133]. Comparing the Ising critical points and the crossover peak positions, we find that the power law scaling parts are exactly the same except for a minus sign within uncertainties, which means they approach the two BKT transition points in the same manner. We believe this behavior can be seen for all $q \rightarrow n^-$ for integer $n \geq 5$. To determine the exponent accurately, we need to use a larger D_{bond} and data closer to an integer, which is beyond the scope of this work. However, this exponent should be the same as the power law scaling of β_p of χ_M in a weak external field for all clock models with integer $q \geq 5$, where there is always an




Figure 4.24 Power law extrapolation of the small- β peak position (triangles) of χ_M with zero external field to $q = 5$ [lower curve in (a)] and to $q = 6$ [lower curve in (b)]. The same procedure is performed for the Ising critical points (circles) at higher β [upper curves in (a) and (b)]. The horizontal dashed lines in (a) are the locations of two BKT transitions for $q = 5$ obtained from Ref. [51]. The horizontal dash-dotted lines in (b) are the locations of two BKT transitions for $q = 6$ obtained from Ref. [133]. $D_{\text{bond}} = 40$ in (a) and $D_{\text{bond}} = 60$ in (b).

emergent $O(2)$ symmetry [25, 135]. In the limit of $O(2)$ model, this exponent is found to be around 0.162 [24, 39, 136].

At large β , the fractional- q -state clock model is a rescaled Ising ($q = 2$) model because the link interactions for the two angles 0 and $\tilde{\phi}$ dominate the weights in the partition function. There are two peaks in the specific heat, and if these peaks are sufficiently separated, then the second peak is that of an Ising model where β is rescaled as $\beta \rightarrow \alpha\beta$, with a rescaling factor

$$\alpha = \frac{1 - \cos \tilde{\phi}}{2}, \quad (4.39)$$

where the small angular distance $\tilde{\phi}$ in the model depends on q and is defined in Eq. (4.8). Thus the critical point β_c of the fractional- q -state clock model can be approximated by the critical point β_{rIsing} of the two-dimensional rescaled Ising model,

$$\beta_c \simeq \beta_{\text{rIsing}} \equiv \frac{\ln(1 + \sqrt{2})}{1 - \cos \tilde{\phi}}. \quad (4.40)$$

In Table 4.2 we list some of these critical points for different values of q . These points give approximately the location of the large- β peak in the specific heat for the extended q clock model.

The critical point β_{rIsing} of the rescaled Ising model (see Table 4.2) is a good approximation for the true critical point β_c for values of q that are not too close to an integer from below. For example,

Table 4.2 The critical points of the rescaled Ising model for different values of q . They are calculated by Eq. (4.40) and give the approximate location of the large- β peak in the specific heat of the fractional q -state clock model in the infinite-volume limit.

q	β_{rIsing}	q	β_{rIsing}
1.1	5.5521	3.6	1.7627
1.2	1.7627	3.7	1.4054
1.3	1.0022	3.8	1.1681
1.4	0.7209	3.9	1.0022
1.5	0.5876	4.0	∞
1.6	0.5163	4.1	75.2052
1.7	0.4764	4.2	19.8386
1.8	0.4544	4.3	9.3216
1.9	0.4437	4.4	5.5521
2.0	∞	4.5	3.7673
2.1	19.8386	4.6	2.7764
2.2	5.5521	4.7	2.1665
2.3	2.7764	4.8	1.7627
2.4	1.7627	4.9	1.4808
2.5	1.2755	5.0	∞
2.6	1.0022	5.1	116.284
2.7	0.8329	5.2	30.3313
2.8	0.7209	5.3	14.0839
2.9	0.6433	5.4	8.2861
3.0	∞	5.5	5.5521
3.1	43.0567	5.6	4.0399
3.2	11.5787	5.7	3.1120
3.3	5.5521	5.8	2.4995
3.4	3.3770	5.9	2.0728
3.5	2.3409	6.0	∞




Figure 4.25 The specific heat for $q = 4.5$ for several different lattice sizes. The dotted curves are from TRG and the solid curves are from the rescaled Ising model. For 16×16 and larger lattices, the rescaled Ising model accurately captures the second peak in the specific heat.

in Fig. 4.25 we compare the specific heat from TRG with the specific heat of the rescaled Ising model for $q = 4.5$ and several different lattice sizes. For lattices 16×16 and larger, the specific heat for the rescaled Ising model accurately captures the large- β peak of the fractional- q -state clock model. As q approaches an integer from below, the approximation begins to fail. In Fig. 4.26, we compare the rescaled Ising critical points with the true critical points as $q \rightarrow 5$ from below. Most of the true critical points are obtained from X defined in Eq. (4.38), and five points are from χ_M and agree perfectly with those from X . For $4.0 < q \lesssim 4.7$, the difference between the true critical point and the Ising approximation is less than 0.01. As $q \rightarrow 5^-$, the difference becomes larger and is around 0.175 for $q = 5$.

4.2.6 Integration interval

For the extended q -state clock model, the allowed angles are restricted to the integration interval $\varphi \in [\varphi_0, \varphi_0 + 2\pi)$. All results presented so far used $\varphi_0 = 0$, which is in the so-called case 1. We also considered the possibility $\varphi_0 = -\pi$. As shown in Fig. 4.27, for $q = 4.5$, the number of allowed spin orientations and their relative sizes are the same. In fact, the models with $\varphi_0 = 0$ and $\varphi_0 = -\pi$ are equivalent and in case 1 for all even $\lfloor q \rfloor$. For odd $\lfloor q \rfloor$, the choice of $\varphi_0 = 0$ or $\varphi_0 = -\pi$ results




Figure 4.26 The difference between the critical point of the rescaled Ising model β_{rIsing} and the true critical point β_c for fractional- q -state clock models as a function of q . The circles are from the discontinuity of values of X , β_{cX} , the stars are from the extrapolations of the peak position of χ_M to zero external field, $\beta_{c\chi}$. The inset shows the values of the true critical point as a function of q and the solid line is Eq. (4.40).




Figure 4.27 The allowed spin orientations for the extended q -state clock model when $\varphi_0 = 0$ (left column) and $\varphi_0 = -\pi$ (right column) for two different values of q . When $q = 4.5$ (top row), the number of allowed spin orientations and their relative sizes are the same, and the models are equivalent. When $q = 5.5$ (bottom row), there are six allowed spin orientations when $\varphi_0 = 0$, but only five allowed orientations when $\varphi_0 = -\pi$, so the two models are different. The dashed gray lines indicate the $\mathbb{Z}_2$ symmetry axes.




Figure 4.28 The magnetic susceptibility at $N = 2^{24} \times 2^{24}$, $h = 0$, and $D_{\text{bond}} = 40$, as a function of β for $\varphi_0 = -\pi$. The susceptibility does not diverge with the volume, which implies there is no phase transition here. This is different from the model with $\varphi_0 = 0$, which has a divergent peak in the susceptibility corresponding to a second-order phase transition.




Figure 4.29 Same as Fig. 4.28, but for $q = 5.001, 5.00001$. The large- β peak fades away and moves across the BKT phase transition point at $q = 5$.

in a different number of allowed spin orientations and different thermodynamic behaviors. For example, when $5 < q < 6$, there are six allowed orientations when $\varphi_0 = 0$, but only five allowed orientations when $\varphi_0 = -\pi$ (see Fig. 4.27 for $q = 5.5$). One can show that the model with $\varphi_0 = -\pi$ is in case 2 for all odd $\lfloor q \rfloor$. For integer values of q , the extended q -state clock model reduces to the ordinary q -state clock model for both $\varphi_0 = 0$ and $\varphi_0 = -\pi$.

We consider $5 < q < 6$ and $\varphi_0 = -\pi$. In Fig. 4.28, we show the magnetic susceptibility at $h = 0$




Figure 4.30 Same as Fig. 4.15, but for $\varphi_0 = -\pi, q \rightarrow 5^+$. $D_{\text{bond}} = 50$.

as a function of β . First of all, one can see that the magnetic susceptibility is finite for all values of q presented here, which means there are no phase transitions for any $5 < q < 6$ and $\varphi_0 = -\pi$. The maximal values of χ_M are larger for q closer to 5. For $q = 5.1$, we see a double-peak structure in χ_M , with the large- β peak higher than the small- β peak. As q increases toward 6, the first crossover peak fades away and disappears around $q = 5.3$. The magnetic susceptibility eventually converges to a single-peak structure with the peak position around $\beta = 1.5$ and the peak height around 4. In order to see the behavior of χ_M for $q \rightarrow 5^+$, we plot χ_M versus β for $q = 5.001$ and $q = 5.00001$ in Fig. 4.29. We see that the maximal value of χ_M is much larger than those in Fig. 4.28, because we are approaching a $\mathbb{Z}_5$ clock model from above. Unlike $q = 5.1$, the small- β peak becomes higher than the large- β peak for $q = 5.001$, and the large- β peak fades away as we move closer to $q = 5$. The small- β peak moves towards the small- β BKT transition for $q = 5$ so it can be used to extrapolate the value of $\beta_{q=5,c1}^{\text{BKT}}$, while the large- β peak moves across the large- β BKT transition for $q = 5$ from right to left so it fails to predict the value of $\beta_{q=5,c2}^{\text{BKT}}$. This means that the magnetic susceptibility cannot capture the crossover going to the large- β BKT transition point for $q \rightarrow 5^+$. But the crossover should go to the BKT transition point as the $\mathbb{Z}_{[q]}$ symmetry is recovered. The cross derivative of the free energy should be able to capture this [113], but we will just focus on the magnetic susceptibility.

We then check the divergent behavior of the small- β peak of χ_M as $q \rightarrow 5^+$ using $D_{\text{bond}} = 50$.




Figure 4.31 Same as Fig. 4.17, but for $\varphi_0 = -\pi$, $q \rightarrow 5^+$. $D_{\text{bond}} = 50$.

Figure 4.30 shows the linear fit for $\ln(\chi_M^*)$ versus $\ln(q - 5)$. We see that the peak height of χ_M diverges as a power law $\chi_M^* \sim 1/(q - 5)^{0.9275(15)}$, which again gives a value of the critical exponent $\delta' = 15.09(2)$, close to the expected value 15. We extrapolate the peak position to $q = 5$ from above in Fig. 4.31. One can see that the power law scaling is the same as case 1 where $q \rightarrow 5^-$, $\varphi_0 = 0$, and gives $\beta_{q=5, c1}^{\text{BKT}} = 1.0514(5)$, consistent with the result in Fig. 4.17.

4.2.7 Phase diagram

The clock model with integer q has been studied extensively [25, 43–57, 78, 101–104]. For $q = 2, 3, 4$, there is a disordered phase and a $\mathbb{Z}_q$ symmetry-breaking phase separated by a second-order phase transition. For $q \geq 5$, there is a disordered phase at small- β and a $\mathbb{Z}_q$ symmetry-breaking phase at large- β with a critical phase for intermediate β between them. The boundaries of the critical phase are two BKT transition points of infinite order [25]. In our extended q -state clock model, one must make a choice of the integration interval $\varphi \in [\varphi_0, \varphi_0 + 2\pi)$. For the choice $\varphi_0 = 0$ and fractional q , both the specific heat and the magnetic susceptibility have a double-peak structure. We have shown that the small- β peak is associated with a crossover, and the large- β peak is associated with a phase transition of the Ising universality class. For the choice $\varphi_0 = -\pi$ and fractional q , the phase structure is a little more complicated. For even $\lfloor q \rfloor$, we get the same behavior as with $\varphi_0 = 0$, but for odd $\lfloor q \rfloor$, we get a trivial case with no critical point.




Figure 4.32 The phase diagram of the extended q -state clock model [i.e. the $h_q = \infty$ plane of the extended-O(2) model] for $\varphi_0 = 0$ (left) and $\varphi_0 = -\pi$ (right). For $q = 2, 3, 4$, there is a second-order phase transition with a $\mathbb{Z}_q$ ordered phase at large β . For finite integer $q \geq 5$, there is a critical phase between a pair of BKT transitions and a $\mathbb{Z}_q$ ordered phase at large β . For fractional $q > 2$ with $\varphi_0 = 0$, there is a crossover line, a second-order transition line, and a $\mathbb{Z}_2$ ordered region between every consecutive pair of integers. For $\varphi_0 = -\pi$, the same is true for every other consecutive pair of integers.

The phase diagrams for both $\varphi_0 = 0$ and $\varphi_0 = -\pi$ are shown in Fig. 4.32. For $\varphi_0 = 0$, as $q \rightarrow [q]^-$, the $\mathbb{Z}_{[q]}$ symmetry is recovered. For $q > 4$ and $\varphi_0 = 0$, both the small- β crossover line and the large- β Ising critical point are smoothly connected to the small- β BKT transition point and the large- β BKT transition point for integer q from below, respectively. Notice that for $q < 4$ and $\varphi_0 = 0$, only the large- β Ising critical point is smoothly connected to the second-order phase transition for integer q from below, while the crossover peak fades away for q close enough (around 3.9 for $3 < q < 4$) to an integer from below. When $q \rightarrow [q]^+$ and $\varphi_0 = 0$, the Ising critical point goes to infinity, while the crossover line goes to a smaller value than the phase transition point for the $[q]$ -state clock model because there is 1 more degree of freedom than the $\mathbb{Z}_{[q]}$ clock model. For even $[q]$ and $\varphi_0 = -\pi$, the phase diagram is the same as that for even $[q]$ and $\varphi_0 = 0$. For odd $[q]$ and $\varphi_0 = -\pi$, the $\mathbb{Z}_{[q]}$ symmetry is recovered when $q \rightarrow [q]^+$. For $5 < q < 6$, both crossover lines are smoothly connected to the two BKT transition points for integer q from above. When q is increased toward an even $[q]$ and $\varphi_0 = -\pi$, the small- β crossover line fades away and the large- β crossover line goes to a larger value than the large- β BKT transition for integer q because there is 1 fewer degree of freedom than the $[q]$ -state clock model. For $3 < q < 4$, there is only one crossover




Figure 4.33 For the extended-O(2) model, the phase diagram is three-dimensional. In the $h_q = 0$ plane, it is the XY model for all values of q . The XY model has a disordered phase at small β , a single BKT transition at $\beta_c = 1.11995(6)$ [39], and a critical phase at large β . In the $h_q = \infty$ plane, it is the extended q -state clock model, which has the phase diagram shown in Fig. 4.32. In this example we have $\varphi_0 = 0$. Establishing the phase diagram at finite- h_q will be addressed in future work.

line that is smoothly connected to the second-order phase transition point for $q = 3$ and goes to around 0.77 when $q \rightarrow 4^-$.

4.2.8 Summary and Outlook

Interpolations among $\mathbb{Z}_n$ clock models have been realized experimentally using a simple Rydberg simulator, where $\mathbb{Z}_n$ ($n \geq 2$) symmetries emerge by tuning continuous parameters, the detuning and Rabi frequency of the laser coupling, and the interaction strength between Rydberg atoms [107]. This paves the way to quantum simulation of lattice field theory with discretized field variables. We are interested in a theory that can interpolate among the O(2) model and $\mathbb{Z}_n$ clock models. We define an extended-O(2) model by adding a symmetry breaking term $h_q \cos(q\varphi_x)$ to the action of the two-dimensional O(2) model. For integer q , $\mathbb{Z}_q$ clock models emerge for large enough h_q . For fractional q , we believe there exists a much more interesting phase structure. The first step to graph the full phase diagram in the (h_q, q, β) cube is to consider the limit $h_q \rightarrow \infty$. In this section, we studied the fractional- q -state clock model as the $h_q \rightarrow \infty$ limit of the extended-O(2) model with

angular variables in the domain $[\varphi_0, \varphi_0 + 2\pi)$. In this limit, the angular variable takes discrete values $\varphi_{x,k} = 2\pi k/q$ with k integer. By varying φ_0 , the set of values of integer k take either case 1 $(0, 1, \dots, \lfloor q \rfloor)$ or case 2 $(0, 1, \dots, \lfloor q \rfloor - 1)$. In case 1, $\mathbb{Z}_{\lceil q \rceil}$ symmetry is recovered as $q \rightarrow \lceil q \rceil^-$, while $\mathbb{Z}_{\lfloor q \rfloor}$ symmetry is recovered as $q \rightarrow \lfloor q \rfloor^+$ in case 2.

For the integer- q -state clock model, there is a single second-order phase transition when $q = 2, 3, 4$. When $q \geq 5$, there are two BKT transitions with a critical phase between them. We studied the fractional- q -state clock model using Monte Carlo (MC) and tensor renormalization group (TRG) methods. We establish the phase diagram of the model for both $\varphi_0 = 0$ and $\varphi_0 = -\pi$. When $\varphi_0 = 0$, we are in case 1, and analysis of the finite-size scaling shows a crossover and a phase transition of the Ising universality class. When $\varphi_0 = -\pi$, we are in case 1 for even $\lfloor q \rfloor$ and in case 2 for odd $\lfloor q \rfloor$. There are no critical points for case 2.

In case 1, we found that there are two peaks in both the specific heat and the magnetic susceptibility. The height of the small- β peak is always finite for fractional q . The large- β peak diverges and characterizes an Ising critical point. When $q \rightarrow \lceil q \rceil^-$ and $q < 4$, the large- β Ising critical point is smoothly connected to the second-order phase transition point for $\mathbb{Z}_{\lceil q \rceil}$ clock models, while the small- β peak fades away. When $q \rightarrow \lceil q \rceil^-$ and $q > 4$, the large- β Ising critical point and the position of the small- β peak are smoothly connected, with the same power law scaling $\sim (\lceil q \rceil - q)^b$, to the large and small BKT points respectively for $\mathbb{Z}_{\lceil q \rceil}$ clock models. We also found that the height of the small- β peak of the magnetic susceptibility diverges as a power law $1/(\lceil q \rceil - q)^{14/15}$, from which we obtain a critical exponent $\delta' = 15$ in the ansatz of the scaling of the magnetization $M \sim (\lceil q \rceil - q)^{1/\delta'}$. This critical exponent is equal to δ associated with the magnetization with an external field $M \sim h^{1/\delta}$. In case 2, there are no critical points. When $q \rightarrow \lfloor q \rfloor^+$ and $q > 5$, the small- β peak also goes to the small BKT point with the same power law scaling and the same δ' exponent as case 1, while the large- β peak fades away and cannot be used to extrapolate the large BKT point of $\mathbb{Z}_{\lfloor q \rfloor}$ clock models.

To use the magnetic susceptibility to locate a critical point, a weak external field must be applied for the magnetic susceptibility to be finite, and extrapolate the peak position to $h = 0$. This method

works in most cases, but the peak fades away near the large- β BKT point of integer- q -state clock models. Our procedure provides an alternative approach to locate the BKT transitions of clock models, by breaking the $\mathbb{Z}_{[q]}$ symmetry to a $\mathbb{Z}_2$ symmetry in the q direction instead of h direction. This procedure creates an Ising critical point that can be used to extrapolate the large- β BKT point for clock models.

Our results clarify what phases the symmetry-breaking term $h_q \cos(q\varphi_x)$ will drive the system to. These phases should have boundaries in the finite- h_q direction. For small enough h_q , the extended- $O(2)$ model should go back to the same universality class as the ordinary XY model, which has been studied extensively [18, 21, 24, 39, 123–126]. For the ordinary XY model, there is a single BKT transition from a disordered phase to a quasi-long-range-ordered critical phase at $\beta_c = 1.11995(6)$ [39]. Figure 4.33 shows the work that remains to be done to figure out the phase diagram in the (β, h_q, q) space interpolating between the known phase diagram at $h_q = 0$ and the phase diagram at $h_q = \infty$ discussed here. There should be a rich phase diagram in the finite- h_q region, which will be investigated in the next section.

4.3 The Extended- $O(2)$ Model at Finite- h_q

In the previous section, we studied the $h_q \rightarrow \infty$ limit of the Extended- $O(2)$ model. In that limit, we could neglect the symmetry-breaking term in the Hamiltonian and directly restrict the spins to the clock angles. Now we return to the full Hamiltonian

$$H_{\text{ext-}O(2)} = - \sum_{x,\mu} \cos(\varphi_{x+\hat{\mu}} - \varphi_x) - h_q \sum_x \cos(q\varphi_x) - h \sum_x \cos(\varphi_x - \varphi_h). \quad (4.41)$$

When $h_q = 0$, this is the classic XY model which is known to have a BKT transition. When $h_q > 0$, the second term breaks periodicity and we must choose $\varphi \in [\varphi_0, \varphi_0 + 2\pi)$ for some choice φ_0 . From here on, we use $\varphi_0 = 0$. When $h_q \rightarrow \infty$, the continuous angle φ_x at a site x is forced into the discrete “clock” angles

$$\varphi_0 \leq \varphi_{x,k} = \frac{2\pi k}{q} < \varphi_0 + 2\pi, \quad k = 0, 1, \dots, [q], \quad (4.42)$$

as illustrated in Figure 4.34. For $h_q \rightarrow \infty$, these angles are effectively the only *allowed* angles in the model. This limit of the model was studied extensively in Sec. 4.2. For finite but large h_q , these




Figure 4.34 Here we illustrate how the spin angle distribution changes with h_q for the example $q = 4.5$. From left to right, we have $h_q = 0, 1, 4, 64$. For large values of h_q , the spins strongly prefer the “clock” angles defined by $\varphi = 2\pi k/q$ for $k = 0, 1, \dots, \lfloor q \rfloor$. Note the “leftover” angle, which in this case is $\tilde{\phi} = \pi/4.5$. These results are obtained by setting $\beta = 0$ in the model where β is treated as a coupling instead of as an inverse temperature.

angles are *preferred* angles.

For integer q and $h_q > 0$, the $O(2)$ symmetry is broken down to a $\mathbb{Z}_q$ symmetry. For noninteger q and $h_q > 0$, the $O(2)$ symmetry is broken down to a $\mathbb{Z}_2$ symmetry. This is because for noninteger q , there is a “leftover” angle

$$\tilde{\phi} = 2\pi \left(1 - \frac{\lfloor q \rfloor}{q} \right). \quad (4.43)$$

See Fig. 4.34 for an illustration of this leftover angle for the case $q = 4.5$. To “see” the residual $\mathbb{Z}_2$ symmetry, imagine drawing a line that bisects the leftover angle. The angles above this line are symmetric with the ones below the line. Because of the leftover angle’s relation to the residual $\mathbb{Z}_2$ symmetry it plays an important role in this model.

We consider a variety of observables. We define the internal energy as

$$E = -\frac{\partial}{\partial \beta} \ln Z = \langle H_{\text{ext-}O(2)} \rangle, \quad (4.44)$$

where $\langle \dots \rangle$ denotes the ensemble average. We define the specific heat

$$C_V = \frac{-\beta^2}{V} \frac{\partial E}{\partial \beta} = \frac{\beta^2}{V} (\langle E^2 \rangle - \langle E \rangle^2). \quad (4.45)$$

We consider also the “interaction energy” as the expectation value of only the interaction term in the Hamiltonian

$$\tilde{E} = \left\langle - \sum_{x,\mu} \cos(\varphi_{x+\hat{\mu}} - \varphi_x) \right\rangle. \quad (4.46)$$

The corresponding specific heat

$$\tilde{C}_V = \frac{\beta^2}{V} (\langle \tilde{E}^2 \rangle - \langle \tilde{E} \rangle^2), \quad (4.47)$$

is expected to have the same critical behavior as the full specific heat Eq. (4.45).

With zero external field, $h = 0$, we measure a proxy magnetization

$$\langle |\vec{M}| \rangle = \left\langle \left| \sum_x \vec{S}_x \right| \right\rangle. \quad (4.48)$$

where $\vec{S}_x$ is the two-component unit vector sitting at the lattice site x . The corresponding magnetic susceptibility is

$$\chi_{|\vec{M}|} = \frac{1}{V} (\langle |\vec{M}|^2 \rangle - \langle |\vec{M}| \rangle^2). \quad (4.49)$$

For the proxy order parameter defined by Eq. 4.48, we have the corresponding Binder cumulant

$$U_{|\vec{M}|} = 1 - \frac{\langle |\vec{M}|^4 \rangle}{3 \langle |\vec{M}|^2 \rangle^2}. \quad (4.50)$$

In the limit $\beta \rightarrow \infty$, everything is frozen and so $\langle |\vec{M}|^4 \rangle = \langle |\vec{M}|^2 \rangle^2$, and this Binder cumulant goes to the trivial value $2/3$. In the limit $\beta \rightarrow 0$, the magnetization $\vec{M}$ is a 2-dimensional Gaussian distribution centered at zero. The magnitude $|\vec{M}|$ of such a distribution is itself a Rayleigh distribution, which has a fourth moment $\langle |\vec{M}|^4 \rangle = 2 \langle |\vec{M}|^2 \rangle^2$. Hence, in the limit $\beta \rightarrow 0$, this Binder cumulant goes to the trivial value $1/3$. The Binder cumulant varies rapidly with β and varies with lattice size. However, it has been shown that at a critical value of β , the Binder cumulant takes a universal value independent of lattice size [137], and as such, it is a useful quantity for identifying phase transitions and locating their critical points.

We also consider the Binder cumulant

$$U_\phi = 1 - \frac{\langle M_\phi^4 \rangle}{3 \langle M_\phi^2 \rangle^2}, \quad (4.51)$$

of the “rotated magnetization” $M_\phi = \langle \cos(q\phi) \rangle$, where $\phi \equiv \arctan(M_y/M_x)$, and where M_x and M_y are the x and y -components of the magnetization $\vec{M}$.

We consider also the structure factor

$$F(\vec{p}) = \frac{1}{L^2} \sum_i \sum_j e^{i(\vec{x}_i - \vec{x}_j) \cdot \vec{p}} \langle \vec{S}_i \cdot \vec{S}_j \rangle, \quad (4.52)$$

which is the Fourier transform of the spin-spin correlator $\langle \vec{S}_i \cdot \vec{S}_j \rangle$. For an $L \times L$ lattice, we define $\vec{p} = 2\pi\vec{n}/L$, with $\vec{n} \in \{(0, 0), (0, 1), (1, 0), \dots, (L/2, L/2)\}$. The position vector $\vec{x}_i$ is the 2-component Cartesian vector corresponding to the i th lattice site. For periodic boundary conditions, we can write Eq. 4.52 as

$$F(\vec{p}) = \frac{1}{L^2} \left| \sum_i e^{-i\vec{p} \cdot \vec{x}_i} \vec{S}_i \right|^2. \quad (4.53)$$

Now there is only a single sum over the lattice—making this an efficient observable from which to obtain the critical exponent η associated with the correlation function.

4.3.1 An Initial Survey

For the Extended- $O(2)$ model, there is a three-parameter (β, q, h_q) phase diagram. For $h_q \rightarrow \infty$, the Extended- $O(2)$ model reduces to the Extended- q -state Clock Model—a phase diagram that we established in the previous section and in [127]. For integer q , it is the well-studied ordinary clock model. For noninteger q , there is a crossover at small β and a second-order phase transition of the Ising universality class at large β . On the other hand, when $h_q \rightarrow 0$, the Extended- $O(2)$ model reduces to the XY model which has a BKT transition at $\beta_c \approx 1.12$. This phase transition occurs for all values of q since q is irrelevant when $h_q = 0$. Thus, we have the picture given by Figure 4.33. Now, we want to replace the question mark in that figure. What is the phase diagram between $h_q = 0$ and $h_q = \infty$?

We started with an initial survey of the parameter space on small 4×4 lattices. Using Monte Carlo, we did a scan with $\beta \in [0, 2]$, $\Delta\beta = 0.02$, $q \in [1.1, 6.0]$, $\Delta q = 0.1$, and $h_q = 0.1, 1, 2, 4, 16, 64$. This required a total of 30,300 Monte Carlo simulations, so despite the small lattice size, this survey required a nontrivial number of CPU hours. In Figure 4.35 we plot heatmaps of the specific heat. The specific heat often diverges with volume at phase transitions and shows a peak at fixed volumes. Therefore, a heatmap of the specific heat can serve as a proxy for the phase diagram of the model, and this one is very suggestive. Each panel in this figure is at a different value of the symmetry-breaking parameter h_q . In the bottom-right panel, $h_q = 64$. Compare this with the phase diagram of the Extended- q -state clock model shown in the left panel of Figure 4.32. The picture is very similar, which leads us to conclude that $h_q = 64$ is already effectively $h_q = \infty$.




Figure 4.35 Here we show a heatmap of the specific heat obtained via Monte Carlo on a 4×4 lattice. Each panel shows the result at a different value of the symmetry-breaking parameter h_q . For general noninteger q , we see two peaks in the specific heat—one at small- β near $\beta \in [0.5, 1]$ and one at large- β . Note that each pixel corresponds to a distinct MCMC simulation.

Notice that there are discontinuities as one moves left or right in this picture across integer values of q . This should not be surprising. When q is dialed across an integer, the number of *preferred* angles in the model increases or decreases by one. For example, for $q = \dots, 4.8, 4.9, 5.0$, there are five preferred angles in the model (these are the *allowed* angles when $h_q = \infty$), but when $q = 5.0 + \epsilon$, there are six preferred angles in the model. Thus, it is not surprising that the phase diagram shows discontinuities at large h_q . As h_q is dialed to smaller values (e.g. $h_q = 2$ or 1), the discontinuities smooth out and the proxy phase diagram becomes of a set of smooth curves. At small h_q (e.g. $h_q = 0.1$), most of the curves have faded away and what remains is a horizontal line that connects to the BKT transition of the XY model at $h_q = 0$.




Figure 4.36 Here we show a heatmap of the proxy magnetization obtained via Monte Carlo on a 4×4 lattice. Each panel shows the result at a different value of the symmetry-breaking parameter h_q . Note that each pixel corresponds to a distinct MCMC simulation.

For a generic noninteger q , we see two peaks (for sufficiently large lattices) in the specific heat—one at small- β near $\beta \in [0.5, 1]$ and one at large- β with its β -position depending on q . The main question for this section of this thesis is whether these peaks correspond to phase transitions, and if so, what kind of transitions are they?

In Figure 4.36, we show similar heatmaps of the proxy magnetization. Since this quantity illustrates the ordered and disordered phases, it gives us further information about the phase diagram. Finally, in Figure 4.37, we show heatmaps of the susceptibility of the proxy magnetization. The magnetic susceptibility generally diverges at a phase transition, however, to see all of the expected peaks in this quantity, one has to go to larger lattices.




Figure 4.37 Here we show a heatmap of the susceptibility of the proxy magnetization obtained via Monte Carlo on a 4×4 lattice. Each panel shows the result at a different value of the symmetry-breaking parameter h_q . Note that each pixel corresponds to a distinct MCMC simulation.

4.3.2 Autocorrelations and the Implementation of BMHA

In the $h_q \rightarrow \infty$ limit, discussed in previous sections, the degrees of freedom of the Extended- $O(2)$ model can be treated as discrete. From the MCMC side, this allowed us to use an efficient heatbath algorithm. Furthermore, it allowed us to use a TRG method to study the system at large volumes. The model is more difficult to study at finite h_q . The degrees of freedom are now continuous, and MCMC heatbath is not an option. We are left with the Metropolis algorithm, which suffers from low acceptance rates and leads to large autocorrelation times as shown in Fig. 4.38. The TRG algorithm we had worked only for the $h_q \rightarrow \infty$ limit. We needed to make some algorithm advancements.




Figure 4.38 The integrated autocorrelation time of the internal energy (left) and the proxy magnetization (right) for several lattice sizes L . The three different colors correspond to three different parameter sets. The solid lines are from the BMHA algorithm, and the dashed lines are from the Metropolis algorithm. We see that the integrated autocorrelation time is one to two orders of magnitude smaller for the BMHA algorithm. All results shown are from simulations performed at $\beta = 1$. In practice, we find that the integrated autocorrelation time increases with larger β . Note, connecting lines were added only to help guide the eyes.

From the MCMC side, we can implement a BMHA algorithm which is designed to approach heatbath acceptance rates. The result is integrated autocorrelation times which are orders of magnitude better than the Metropolis algorithm as shown in Fig. 4.38. The BMHA implementation of the $O(2)$ model described in Sec. 3.2.3 uses a “1 variable / 1 parameter” BMHA algorithm. For the Extended- $O(2)$ model, the PDF can be written in the form

$$P(\varphi; \alpha, \varphi_+) = e^{\alpha \cos(\varphi - \varphi_+) + \beta h_q \cos(q\varphi)}, \quad \alpha \equiv \beta \left| \vec{S}_+ \right| \quad (4.54)$$

where φ is the BMHA variable and α and φ_+ are BMHA parameters. The model parameters β , h_q , and q are fixed for a given simulation. Recall that $\vec{S}_+$ is the “stencil”, i.e. the sum of the nearest neighbors at a given site, and φ_+ is the angle of the stencil vector. The presence of the second term in the exponential means that, unlike in the $O(2)$ case, we cannot eliminate the angle φ_+ by a shift of the variable φ . Instead, we now have to use a “1 variable / 2 parameter” BMHA algorithm. This increases the dimensionality of the BMHA lookup table, but it does not significantly complicate things since each parameter can be uniformly discretized.

4.3.3 β as a Coupling or as Inverse Temperature

We define the Extended- $O(2)$ model to be the model with Hamiltonian

$$H_{\text{ext-}O(2)} = - \sum_{x,\mu} \cos(\varphi_{x+\hat{\mu}} - \varphi_x) - h_q \sum_x \cos(q\varphi_x) - h \sum_x \cos(\varphi_x - \varphi_h), \quad (4.55)$$

and canonical partition function

$$Z = \int_{\varphi_0}^{\varphi_0+2\pi} \prod_x \frac{d\varphi_x}{2\pi} e^{-\beta H}. \quad (4.56)$$

With this definition, β , which appears only in the Boltzmann factor and not in the Hamiltonian, plays the role of the thermodynamic inverse temperature.

An alternative definition of the model is to write the Hamiltonian as

$$\tilde{H}_{\text{ext-}O(2)} = -\beta \sum_{x,\mu} \cos(\varphi_{x+\hat{\mu}} - \varphi_x) - \tilde{h}_q \sum_x \cos(q\varphi_x) - h \sum_x \cos(\varphi_x - \varphi_h), \quad (4.57)$$

and then write the partition function as

$$\tilde{Z} = \int_{\varphi_0}^{\varphi_0+2\pi} \prod_x \frac{d\varphi_x}{2\pi} e^{-\tilde{H}}. \quad (4.58)$$

Notice that β is now attached to the first term of the Hamiltonian rather than in the Boltzmann factor of the partition function. In this field-theoretic definition, β , $\tilde{h}_q$, and h are coupling parameters. This was the definition that was used in [127].

Not surprisingly, there are slight differences between these two models as shown in Fig. 4.39. In this figure, the left panel shows the interaction energy of the Extended- $O(2)$ model with $q = 4.9$ on a 4×4 lattice. The right panel shows the corresponding specific heat. The results in blue are from the model with β as the inverse temperature and $h_q = 5$, and the results in red are from the model where β is treated as a coupling and with $\tilde{h}_q = 5$.

Although there are slight differences between the two models, the general critical behavior should be the same since results from one model can be mapped to the other with $\tilde{h}_q = \beta h_q$. Note that the two models coincide when $\beta = 1$ and $h_q = \tilde{h}_q$. For a given q , the same kind of phase transition(s) should occur in both models although the critical temperatures may differ. In fact,




Figure 4.39 (Left) The interaction energy of the Extended-O(2) model with $q = 4.9$ and $h_q = \tilde{h}_q = 5$ on a 4×4 lattice. (Right) The corresponding specific heat. The results in blue are from the model with β as the inverse temperature, and the results in red are from the model where β is treated as a coupling. The results coincide for $\beta = 1$ but differ slightly elsewhere. The connecting lines are only to guide the eyes.

performing the same kind of survey as in Figure 4.35 but with β treated as a coupling, shows the same kind of pictures but stretched a bit in the β dimension.

From a field theory perspective, it is more natural to treat β as a coupling parameter attached to the interaction term of the Hamiltonian. On the other hand, from a thermodynamics perspective, it is more natural to have β factored out of the Hamiltonian. In the latter approach, one can use the standard definition Eq. (4.44) of the internal energy as the derivative of the free energy with respect to β , and $1/\beta$ is the true thermodynamic temperature. In the end, our need to use temperature reweighting when doing the finite-size-scaling analysis forced us to use the choice Eq. (4.55) such that β is the true inverse temperature. One could use temperature reweighting also with the alternative definition, however, it would not be as straightforward since one would have to factor β out of the action. But this would leave a nontrivial interplay between β and $\tilde{h}_q$ in the sense that any change in β would effectively also change $\tilde{h}_q$. Similar considerations would apply to the external field coupling h , however, for the present Monte Carlo studies, we have $h = 0$. Unless otherwise noted, the results shown in this thesis are from the model with Hamiltonian Eq. (4.55)

where β serves as the inverse temperature.

4.3.4 The Clock Model Limit

In the clock model, we have the energy function

$$H = - \sum_{x,\mu} \cos(\varphi_{x+\hat{\mu}} - \varphi_x), \quad (4.59)$$

with the angles $\varphi_x^{(k)}$ selected discretely as

$$\varphi_0 \leq \varphi_x^{(k)} = \frac{2\pi k}{q} < \varphi_0 + 2\pi, \quad (4.60)$$

with $k \in \mathbb{Z}$. For example, if $\varphi_0 = 0$, then $k = 0, 1, 2, \dots, \lfloor q \rfloor$ with $q \in \mathbb{R}$. In general, φ_0 must be specified because, for example, the number of k 's in a 2π domain varies depending on φ_0 .

In the zero-field Extended- $O(2)$ model, we have the energy function

$$H = - \sum_{x,\mu} \cos(\varphi_{x+\hat{\mu}} - \varphi_x) - h_q \sum_x \cos(q\varphi_x). \quad (4.61)$$

with the angles φ_x selected continuously in

$$\varphi_0 \leq \varphi \in \mathbb{R} < \varphi_0 + 2\pi. \quad (4.62)$$

Now, the allowed angles are the continuous values $\varphi \in [\varphi_0, \varphi_0 + 2\pi)$, but when h_q is taken large, the Boltzmann factor $e^{-\beta H} \sim e^{\beta h_q \sum \cos(q\varphi)}$ forces the angles to take the values $\varphi_x = 2\pi k/q$ with $k \in \mathbb{Z}$. Hence, we claim the Extended- $O(2)$ model reduces to the clock model in the limit $h_q \rightarrow \infty$. This is true for integer values of q , however, for noninteger q , this claim fails as we can see in Fig. 4.40.

For non-integer q , the 2π -rotational symmetry is broken. Thus, we have to more carefully define the domain of the angles φ_x . For example, one could choose $\varphi_0 = -\pi$ and then $\varphi_x \in [-\pi, \pi)$ or one could choose $\varphi_0 = 0$ and then $\varphi_x \in [0, 2\pi)$. These different choices lead to very different phase diagrams as we showed for the clock model in Fig. 4.32. We prefer the choice $\varphi_x \in [0, 2\pi)$, simply because the phase diagram gives a more consistent periodic picture as q is increased⁴. However, with the choice $\varphi_x \in [0, 2\pi)$, there is a hard cutoff at the angle $\varphi = 0$. When $h_q = \infty$, this is not a

⁴See the left panel of Fig. 4.32 as opposed to the right panel in the same figure.




Figure 4.40 Here we compare the q -state clock model, which has been extended to noninteger values of q and has energy function Eq. (4.59) (black curves), with the Extended-O(2) model which has energy function Eq. (4.61) (colored curves) on a 4×4 lattice. For the Extended-O(2) model, we consider increasing values of finite h_q in order to understand the limit $h_q \rightarrow \infty$. Naively, we expect the Extended-O(2) model to limit to the clock model when h_q is taken large. The Extended-O(2) model results are from lower statistics runs. The top row compares the “interaction” energy, the second row compares the corresponding specific heat, the third row compares the proxy magnetization, and the bottom row compares the susceptibility of the proxy magnetization. The different columns correspond to different values of q with $q = 4.1, 4.7, 5.0$ from left to right. In the naive approach ($\varphi \in [0, 2\pi)$) illustrated here, the $h_q \rightarrow \infty$ limit of the Extended-O(2) model is the clock model only when q is integer. For noninteger values of q , the $h_q \rightarrow \infty$ limit of the Extended-O(2) model differs from the clock model.

problem, but at finite h_q , this cutoff skews the distribution of angles severely enough that the model at finite h_q does not smoothly connect with the model at infinite h_q when $h_q \rightarrow \infty$. To fix this, we need to slightly shift the angle domain such that $\varphi_x \in [-\varepsilon, 2\pi - \varepsilon)$. To match the clock model (i.e. $h_q = \infty$ case), one needs an ε that varies with q and satisfies the condition $0 < \varepsilon < 2\pi(1 - \lfloor q \rfloor / q)$. In general, we choose

$$\varepsilon = \pi \left(1 - \frac{\lfloor q \rfloor}{q} \right). \quad (4.63)$$

That is, the partition function Eq. (4.6) is modified to be

$$Z = \int_{-\varepsilon}^{2\pi-\varepsilon} \prod_x \frac{d\varphi_x}{2\pi} e^{-\beta H}. \quad (4.64)$$

After taking the $h_q \rightarrow \infty$ limit, the $\varepsilon \rightarrow 0$ limit can be taken to connect with the clock models. To understand why such a shift is needed, see Appendix B.3. In Fig. 4.41, we see the effect of implementing such a shift in the angles. Now the Extended- $O(2)$ model connects to the clock model in the $h_q \rightarrow \infty$ limit for both integer and noninteger q .

There remains one more subtlety. When q is noninteger, one may find that the Extended $O(2)$ model with $h_q \rightarrow \infty$ and the clock model do not agree when $\beta \rightarrow 0$. This can be seen in Figure 4.41. For non-integer q , the $\mathbb{Z}_q$ symmetry is explicitly broken, and for example, the energy density of the clock model does not go to zero in the limit $\beta \rightarrow 0$ as can be seen in Figs. 4.2–4.6. In contrast, the energy density in the Extended $O(2)$ model does go to zero in the limit $\beta \rightarrow 0$ for all values of h_q . The discrepancy can be understood by considering the Boltzmann factor of the Extended $O(2)$ model, which has the form

$$e^{-\beta H} = e^{\beta \sum \cos(\varphi_{x+\hat{\mu}} - \varphi_x) + \beta h_q \sum \cos(q\varphi_x)} \quad (4.65)$$

The clock limit occurs when $h_q \rightarrow \infty$ and the second term becomes very large—forcing the angles into the discrete values $\varphi_x = 2\pi k/q$. When $\beta \rightarrow 0$, the two limits compete against each other. To match the clock model in the $h_q \rightarrow \infty$ limit for all values of β , one has to adjust h_q such that βh_q stays large as one approaches $\beta = 0$. Incidentally, when β is treated as a coupling attached only to the interaction term instead of as a proper inverse temperature, then this discrepancy at small β does




Figure 4.41 Here we compare the q -state clock model (black curves) with the $h_q \rightarrow \infty$ limit of the Extended-O(2) model (colored curves) on a 4×4 lattice. The Extended-O(2) model results are from lower statistics runs. The top row compares the “interaction” energy, the second row compares the corresponding specific heat, the third row compares the proxy magnetization, and the bottom row compares the susceptibility of the proxy magnetization. The different columns correspond to different values of q with $q = 4.1, 4.7, 5.0$ from left to right. Here, $\varphi \in [-\varepsilon, 2\pi - \varepsilon]$ with ε given by Eq. (4.63), and now the Extended-O(2) model becomes the clock model in the $h_q \rightarrow \infty$ limit. Notice that the matching breaks down at very small β in this illustration when q is noninteger. As explained in the text, the proper matching procedure would be to adjust h_q such that βh_q stays large as one approaches $\beta = 0$.




Figure 4.42 (Left) The phase diagram of a quantum simulator with continuously tunable parameters composed of Rydberg atoms [106, 107]. *Reproduced with permission from Springer Nature.* Note the similarity with a heatmap (right) of the specific heat from the Extended- $O(2)$ model with finite h_q obtained from TRG on a $L = 1024$ lattice. Note that each pixel corresponds to a distinct TRG calculation.

not occur because in the Boltzmann factor, $e^{-\tilde{H}} = e^{\beta \sum \cos(\varphi_{x+\hat{\mu}} - \varphi_x) + h_q \sum \cos(q\varphi_x)}$, β is not multiplying the symmetry-breaking term.

4.3.5 The Search for Exotic Phases

As shown in Fig. 4.42, the heatmap of the specific heat for the Extended- $O(2)$ model at finite h_q is similar to the phase diagram of a quantum simulator composed of Rydberg atoms [106, 107]. The two models both contain crystalline $\mathbb{Z}_q$ ordered phases for $q = 2, 3, 4$, and both have a small set of continuously tunable parameters. For the Extended- $O(2)$ model, these parameters are β , q , and h_q . For the Rydberg simulator, these parameters include the interaction range R_B/a and the detuning Δ/Ω . However, the physics of these two models is quite different. The Rydberg model, for example, contains long-range interactions whereas the Extended- $O(2)$ model contains only nearest-neighbor interactions. As such, the similarity between the Rydberg phase diagram and our proxy phase diagram may be purely coincidental. However, we considered it worth investigating.

The heatmaps of the specific heat at large h_q show discontinuities similar to the clock model

phase diagram shown in Fig. 4.32. However, at small h_q , these discontinuities seem to disappear, and the heatmaps show smooth lobes. This motivated us to explore these regions in parameter space to look for exotic phases such as the floating phases in Rydberg chains [138] or the commensurate-incommensurate phase transition in the Pokrovsky-Talapov model [139].

We start by looking at several thermodynamic functions in the vicinity of the “hotspot” near $q \sim 3$, $\beta \sim 0.8$ and $h_q \sim 1$ (see the bright yellow point in the right panel of Fig. 4.42). This point is chosen because it appears as a “triple point” in the proxy phase diagram formed by the heatmap of the specific heat. That is, $q = 3$ is where three curves converge to a point. Similar “triple point” features are seen near $q = 4$, $q = 5$, and $q = 6$. If exotic phases exist near any of these points, we expect them to exist near all of these points. We could just as well have chosen to study one of those other points. From left to right, the columns in Fig. 4.43 correspond to $q = 2.8, 2.9, 3.0, 3.1, 3.2$. From top to bottom, the rows in this figure illustrate the energy density, specific heat, proxy magnetization, and susceptibility of the proxy magnetization. For $q = 3$, we see the classic signals of a second order phase transition. The specific heat and susceptibility diverge with volume near $\beta = 0.8$. For $q = 2.8$ and 2.9 , there are two features of interest, which manifest as two peaks in the magnetic susceptibility. For $q = 2.9$, the peaks are near $\beta = 0.7$ and $\beta = 0.85$. For $q = 2.8$, the peaks are near $\beta = 0.6$ and $\beta = 1$. For both $q = 2.8$ and 2.9 , the specific heat and magnetic susceptibility appear to plateau with increasing volume near the small- β peak. This implies that there is only a crossover here and not a true transition. For the large- β peak, both the specific heat and magnetic susceptibility seem to diverge with volume. This suggests a phase transition here. For $q = 3.1$ and 3.2 , we see only a single peak in the specific heat and susceptibility, and both quantities appear to plateau here, which suggests that this is a crossover rather than a phase transition. For $q = 3.1$ and 3.2 , there seem to be phase transitions at much larger β —beyond the range of these plots and far from this “hotspot”.

The heatmaps of the specific heat, magnetization, and magnetic susceptibility show smooth curves for intermediate values of h_q (e.g. for $h_q \sim 1$). This suggests the possibility that the phase diagram will have smooth phase transition curves even as one crosses integer values of q . This




Figure 4.43 From left to right, the columns correspond to $q = 2.8, 2.9, 3.0, 3.1, 3.2$. From top to bottom, the rows illustrate the energy density, specific heat, proxy magnetization, and susceptibility of the proxy magnetization. The different colors correspond to the different lattice sizes $L = 4, 8, 16, 32$. Connecting lines are drawn to help guide the eyes. For $q = 2.8$ and 2.9 we see evidence of a crossover at smaller β and a phase transition at larger β . For $q = 3.0$, there is only a phase transition, and for $q = 3.1$ and 3.2 we see only a crossover for the illustrated β range.




Figure 4.44 The specific heat (left) and the entanglement entropy (right) from TRG for various q near $q = 3$ with $\beta \in [0, 2]$ and $h_q = 1$. These results were obtained from lattices of size $L = 1024$. Note that each pixel corresponds to a distinct TRG calculation. The specific heat suggests a smooth curve in the phase diagram, however, the entanglement entropy shows a clear discontinuity across $q = 3$.

would be very different from the model in the $h_q \rightarrow \infty$ limit where there are clear discontinuities at integer values of q in the phase diagram. Furthermore, comparison of these heatmaps at intermediate h_q with the phase diagrams of Rydberg atom chains suggests the possibility of exotic phases. However, a closer look at how the thermodynamic functions behave with lattice size near $q = 3$ (see Fig. 4.43), suggests a more mundane picture. In particular, it suggests that for $q = 3$ and $q \lesssim 3$, near $h_q \sim 1$, there is a second-order phase transition near $\beta = 0.8$, however for $q \gtrsim 3$, there is only a crossover in this region. This is similar to the model at $h_q \rightarrow \infty$ —suggesting that the phase diagram at $h_q \rightarrow \infty$ persists to finite h_q .

This conclusion is supported by some large-volume results obtained from TRG. See Fig. 4.44. Here we have the specific heat (left panel) and the entanglement entropy (right panel) for various q near $q = 3$ with $\beta \in [0, 2]$ and $h_q = 1$. These results were obtained from lattices of size $L = 1024$. The specific heat suggests a smooth curve in the phase diagram, however, the entanglement entropy shows a clear discontinuity across $q = 3$. For $q = 3$, the entanglement entropy for large β is $\approx \ln 3$, which suggests a $\mathbb{Z}_3$ ordered phase. Inside the “lobe,” the entanglement entropy is $\approx \ln 2$, which suggests $\mathbb{Z}_2$ order.




Figure 4.45 Here we show the distribution of spin angles φ via histograms. These results were obtained on 4×4 lattices with $\beta = h_q = 1$ and $q = 2.8$ (left), $q = 3.0$ (center), and $q = 3.2$ (right).

Another investigative technique we can use is to study the distribution of the spin angles φ . In Fig. 4.45, we show histograms of the spin angles for several different values of q . From left to right, the histograms correspond to $q = 2.8, 3.0, 3.2$. All three examples are from 4×4 lattices and with $\beta = h_q = 1$. Compare with Fig. 4.34, which uses polar histograms. The peaks in the distribution correspond to the preferred “clock angles” $2\pi k/q$ with $k = 0, 1, \dots, \lfloor q \rfloor$. For a given set of parameters (q, β, h_q) , a Monte Carlo simulation was run using a BMHA algorithm with 2^{20} equilibrating sweeps followed by 2^{13} measurement sweeps each of which was separated by 2^8 discarded sweeps. For each measurement sweep, the lattice configuration was saved to file. For a 4×4 system, each configuration contains 16 angles—one for each site. These $2^{13} \times 16$ angles were histogrammed to obtain a pictures like those in Fig. 4.45. For $q = 3$, there are three equal and equidistant peaks as expected for a system with $\mathbb{Z}_3$ symmetry. When q is noninteger, that symmetry is broken and the peaks are no longer equal or evenly spaced. Now, the system seems to be dominated by the two preferred angles which are closest to each other: $\varphi = 0$ (or 2π) and $\varphi = 2\pi \lfloor q \rfloor / q = -\tilde{\phi}$. Interestingly, the effect is stronger for q just above an integer (e.g. $q = 3.2$) than for q just below an integer (e.g. $q = 2.8$).

To illustrate angle histograms over a much larger parameter space, we use a heatmap, which converts the histogram height into color as shown in Fig. 4.46. These results are again with $h_q = 1$ and $L = 4$, but now with $\beta \in [0, 1.2]$ and for $q = 2.5, 2.8, 2.9, 3.0, 3.1, 3.2$, and they further demonstrate that the system is strongly affected by the “leftover” angle when q is noninteger. However, it also shows that this effect becomes weaker as q approaches the next integer from below.




Figure 4.46 Heatmaps of the angle histograms (see Fig. 4.45) for a 4×4 lattice with $h_q = 1$, $\beta \in [0, 1.2]$, and $q = 2.5, 2.8, 2.9, 3.0, 3.1, 3.2$. Note that each pixel corresponds to a distinct MCMC simulation.

For example, the effect is large for $q = 3.1$ but small for $q = 2.9$. An alternative representation is shown in Fig. 4.47, where we plot a heatmap for constant $\beta = 1$ and varying q . We looked also at histogram heatmaps for smaller $h_q = 0.1$ and for even and odd lattice sites separately, but nothing fundamentally different was seen.

Since the angle histogram for e.g. $q = 2.9$ is very similar to that of $q = 3.0$ (likewise 3.9 is similar to 4.0 and so on), this suggests the possibility that there is a $\mathbb{Z}_3$ ordered phase also for $q = 2.9$. In other words, the lines of $\mathbb{Z}_q$ order in the phase diagram may have some width at finite h_q . However, we were unable to find any evidence of such “thick” $\mathbb{Z}_q$ order lines.

At this point, we have not found evidence for exotic phases or really anything different from the $h_q \rightarrow \infty$ case. However, given the similarity between the phase diagram of a Rydberg atom system and our heatmap of the specific heat (see Fig. 4.42), we would be remiss if we did not check for the floating phases that are found in Rydberg chains. Evidence for these phases includes nontrivial structures (curves) in heatmaps of the structure factor $F(p)$ when plotted with momenta




Figure 4.47 Heatmaps of the angle histograms for constant $\beta = 1$ and varying $q \in [1.1, 6.0]$. Results are from 4×4 lattices with $h_q = 1$. The red labels on the right edge correspond to the q -values shown in Fig. 4.45. Note that each pixel corresponds to a distinct MCMC simulation.




Figure 4.48 Heatmaps of the structure factor $F(p)$ with momenta p on the vertical axis and with q on the horizontal axis. The different panels correspond to different β with $\beta = 0.7, 0.9, 1.1$ from left to right. These results are from 32×32 lattices with $h_q = 1$ and are from the model where β is treated as a coupling. Note that each pixel corresponds to a distinct MCMC simulation.




Figure 4.49 Heatmaps of the *staggered* structure factor $F(p)$ with momenta p on the vertical axis and with q on the horizontal axis. The different panels correspond to different β with $\beta = 0.7, 0.9, 1.1$ from left to right. These results are from 32×32 lattices with $h_q = 1$ and are from the model where β is treated as a coupling. Note that each pixel corresponds to a distinct MCMC simulation.

on one axis and the tunable parameter on the other axis [138]. In our case, we studied the structure factor $F(p)$ on $L = 32$ lattices with $p = 2\pi n/L$ and $n = 0, 1, \dots, L/2$. To keep things simple, we considered only on-axis modes such as $\vec{p} = (0, 0), (0, 1), (0, 2), \dots, (0, L/2)$. We used $h_q = 1$ and looked at several β -values. See Fig. 4.48. No evidence of floating phases is seen. We consider also the “staggered” structure factor, which can be computed using the same definition of the structure factor after rotating all spin vectors on odd-parity sites by the angle π . The results given in Fig. 4.49 show no evidence of exotic phases.

4.3.6 Binder Cumulants

The Binder cumulant

$$U_M = 1 - \frac{\langle M^4 \rangle}{3\langle M^2 \rangle^2}. \quad (4.66)$$

is the fourth order cumulant of the magnetization. Here, $M = |\vec{M}|$ is the magnitude of the magnetization vector. In the limit $\beta \rightarrow \infty$, everything is frozen and so $\langle |\vec{M}|^4 \rangle = \langle |\vec{M}|^2 \rangle^2$, and this Binder cumulant goes to the trivial value $2/3$. In the limit $\beta \rightarrow 0$, the magnetization $\vec{M}$ is a 2-dimensional Gaussian distribution centered at zero. The magnitude $|\vec{M}|$ of such a distribution is itself a Rayleigh distribution, which has a fourth moment $\langle |\vec{M}|^4 \rangle = 2\langle |\vec{M}|^2 \rangle^2$. Hence, in the limit $\beta \rightarrow 0$, this Binder cumulant goes to the trivial value $1/3$. The Binder cumulant varies rapidly with β and varies with lattice size. However, it has been shown that at a critical value of β , the

Binder cumulant takes a universal value independent of lattice size [137], and as such, it is a useful quantity for identifying phase transitions and locating their critical points.

There are several things that one can do with the Binder cumulant. The first is to identify the critical points associated with phase transitions. Since the Binder cumulant takes a universal value at a critical point regardless of lattice size, one can perform simulations on different lattice sizes and identify the critical temperature by seeing where the Binder cumulant curves cross. See for example, [140, 141].

The Binder cumulant can also be used to identify the type of phase transition. For second order phase transitions, the temperature dependency of the Binder cumulants at different lattice sizes have a distinct intersection point. In other words, the Binder cumulant curves intersect at a single point. See for example, [140]. In contrast, at a BKT transition the Binder cumulants intersect at all critical temperatures, that is, the Binder cumulant curves lie on top of each other in the critical phase. Binder cumulants can be used to distinguish first- and second-order phase transitions by considering also the cumulant of the energy [142]. See also [143, 144].

Finally, the Binder cumulant can be used to extract the critical exponent ν via finite-size scaling. See for example, [46, 140, 145]. One can also show “data collapse” of Binder cumulants to provide support for a particular value of ν [54, 146–152]. Other finite-size-scaling of Binder cumulants is discussed in [153–156].

The Binder cumulant has been used to study the Ising model [157] and the XY model [158, 159]. It has also been used to study the clock models [46, 52, 54, 78]. For the q -state clock model with $q \geq 5$, there are two BKT transitions. The small- β transition separates the disordered and quasi-long-range-ordered phases, and the large- β transition separates the quasi-long-range-ordered and long-range-ordered phases. The Binder cumulants as defined in Eq. (4.66) for different lattice sizes merge at the small- β transition and run together for larger values of β . Thus, this Binder cumulant is useful for obtaining the small- β critical point, but not for obtaining the large- β critical point. To




Figure 4.50 Here, we show the specific heat (top row), magnetic susceptibility (second row), Binder cumulant (third row), and the Binder cumulant of the rotated magnetization (bottom row). From left to right, the columns correspond to $q = 2, 3, 4, 5, 6$. In all cases, $h_q = 1$. The different color shades correspond to different lattice sizes ranging from $L = 4$ (light) to $L = 128$ (dark). Connecting lines are added to help guide the eyes. The vertical dashed gray lines show where the Binder cumulants intersect.

extract the latter, one can define the Binder cumulant [44, 45, 54, 55, 160]

$$U_\phi = 1 - \frac{\langle M_\phi^4 \rangle}{3\langle M_\phi^2 \rangle^2}, \quad (4.67)$$

of the “rotated magnetization” $M_\phi = \langle \cos(q\phi) \rangle$, where $\phi \equiv \arctan(M_y/M_x)$, and where M_x and M_y are the x and y -components of the magnetization $\vec{M}$. This Binder cumulant allows one to extract the large- β transition in the clock models but not the small- β transition.

In Fig. 4.50, we show some Binder cumulants for the Extended- $O(2)$ model with $h_q = 1$ and integer q . The columns correspond to different values of q with $q = 2, 3, 4, 5, 6$ from left to right. The first two rows show the specific heat and magnetic susceptibility. For $q = 2, 3, 4$ the specific heat and susceptibility both seem to diverge with volume—suggesting a second-order phase transition. For $q = 5, 6$ we see two peaks in the specific heat. Both peaks converge to a finite height even as the volume is increased. This suggests the presence of two BKT transitions as in the clock model (i.e. $h_q = \infty$) case. This is supported by the behavior of the susceptibility which seems to diverge with volume at all β -values between the two peaks. The small- β transition is marked by a




Figure 4.51 As in Fig. 4.50 but with $h_q = 0.1$.

prominent peak in the susceptibility while the critical region up to the large- β transition is marked by a “shoulder” in the susceptibility. The third row of the figure shows the Binder cumulant given by Eq. (4.66). The bottom row of the figure shows the Binder cumulant of the rotated magnetization given by Eq. (4.67). Notice that for $q = 5, 6$, the Binder cumulants defined by Eq. (4.66) cross at the small- β transition and those defined by Eq. (4.67) cross at the large- β transition. The crossing or intersection points are highlighted by the vertical dashed gray lines. For $q = 2, 3, 4$ where there seems to be a single transition, both kinds of Binder cumulants cross at the same β . In Fig. 4.51, we show the same set of plots but for a smaller symmetry-breaking parameter $h_q = 0.1$.

For noninteger values of q , the Binder cumulant behaves in an unexpected manner. In Fig. 4.52, we show some Binder cumulants for noninteger q with $h_q = 1$ near the “hotspot”. Note that this data is from the model where β is treated as a coupling, and so the transition points may be slightly shifted from our usual model where β is treated as an inverse temperature. The qualitative picture should be the same regardless of whether β is treated as a coupling or as an inverse temperature. From left to right, the columns correspond to $q = 2.8, 2.9, 3.0, 3.1, 3.2$. The top row corresponds to the ordinary Binder cumulant defined in Eq. (4.66). The different shades correspond to different lattice sizes from $L = 4$ (light) to $L = 64$ (dark). Notice that for $q = 3$, the Binder cumulants cross near the expected transition point. Compare with Fig. 4.43 where other thermodynamic




Figure 4.52 Here we look at some Binder cumulants for the model with noninteger q and $h_q = 1$ near the “hotspot”. From left to right, the columns correspond to $q = 2.8, 2.9, 3.0, 3.1, 3.2$. The different shades correspond to different lattice sizes from $L = 4$ (light) to $L = 64$ (dark). The top row corresponds to the ordinary Binder cumulant defined in Eq. (4.66). For noninteger q , the Binder cumulants pinch together instead of cross at a second-order phase transition. In the bottom row, we show the Binder cumulant defined in Eq. (4.69) that has been modified to emphasize the Ising-ness of the model. This Binder cumulant shows again the classic crossing phenomenon at the phase transition. This data is from the model where β is treated as a coupling, and so the transition points may be slightly shifted from our usual model where β is treated as an inverse temperature.

quantities are plotted for the same values of q and h_q . For $q = 3.1$ and 3.2 the Binder cumulants do not cross—further evidence that there is only a crossover in this small- β regime. The expected second-order phase transition for $q = 3.1$ and 3.2 is at much larger β . For $q = 2.8$ and 2.9 , the Binder cumulants seem to intersect near the expected second-order transition point, however, they “pinch” together instead of crossing each other. We have been unable to find other examples of such Binder cumulant “pinching” in the literature.

To understand the Binder cumulant “pinch” as seen in Fig. 4.52 for $q = 2.8$ and 2.9 , we looked at the time series proxy magnetization $|\vec{M}|$ and the angle φ that $\vec{M}$ makes with the x -axis for a simulation run in the pinched region. We found that the magnetization angle tended to point in the direction $-\tilde{\phi} < \varphi < 0$ where $\tilde{\phi}$ is the “leftover” angle defined in Eq. (4.43). For example, for

$q = 2.9$, the *preferred* angles of the model are $\varphi = 0, 2\pi/2.9, 4\pi/2.9$, and the leftover angle is $\tilde{\phi} = 2\pi - 4\pi/2.9$. In this example, the system is spending its time magnetized mostly in the $\varphi = 0$ direction or in the nearby $\varphi = 4\pi/2.9 \equiv -\tilde{\phi}$ direction. This suggests that one should be able to turn it into an Ising model by projecting the magnetization $\vec{M}$ onto a coordinate system that is rotated by $\tilde{\phi}/2$. That is, we define a new magnetization

$$\tilde{M} = M_x \sin\left(\frac{\tilde{\phi}}{2}\right) + M_y \cos\left(\frac{\tilde{\phi}}{2}\right), \quad (4.68)$$

where M_x and M_y are the components of the ordinary magnetization vector $\vec{M}$. Then we use a Binder cumulant

$$U_{\tilde{M}} = 1 - \frac{\langle \tilde{M}^4 \rangle}{3\langle \tilde{M}^2 \rangle^2}. \quad (4.69)$$

In the bottom row of panels in Fig. 4.52, we show examples of this quantity for $q = 2.8, 2.9, 3.0, 3.1, 3.2$ for the model with $h_q = 1$. For $q = 2.8$ and 2.9 , the pinching phenomenon is now replaced with the classic crossing behavior. For $q = 3.0$, the expected crossing behavior is also seen at the second-order transition point. For β larger than the critical point, the Binder cumulant Eq. (4.69), seems to occasionally get stuck at $U_{\tilde{M}} = 0$ yielding unreliable results. For $q = 3.1$ and 3.2 , the expected Ising transition is at β much larger than the values shown in Fig. 4.52, thus no crossing is seen. Instead, we see $U_{\tilde{M}} \approx 0$ for sufficiently large lattices.

4.3.7 Broken Symmetry and Ultra-long Autocorrelations: A Cautionary Tale

When q is noninteger, the $\mathbb{Z}_q$ symmetry is broken. This results in a practical and pernicious difficulty for MCMC approaches to studying this model. In practice, we found that a naive MCMC approach leads to misleading conclusions about the scaling behavior and therefore to an incorrect determination of critical exponents. This was only discovered after we found apparent scaling law violations for some critical exponents.

To understand what happens when q is noninteger, it will help to consider the integer case first, for example $q = 5$. Furthermore, for simplicity we will consider the clock limit (i.e. $h_q \rightarrow \infty$) where the so-called “clock angles” are the only allowed angles for the spin variables. When $q = 5$, the five allowed angles form a $\mathbb{Z}_5$ symmetry group. See Fig. 4.53. In the low temperature limit, the $\mathbb{Z}_5$




Figure 4.53 The clock angles for $q = 5.0$ (gray) and $q = 4.9$ (black). For $q = 5$, there are five equidistant angles which divide the unit circle into five equal sectors. For $q = 4.9$, there are again five angles, but the $\mathbb{Z}_5$ symmetry is broken. Now there are four equal sectors and one slightly smaller “leftover” sector.

symmetry is spontaneously broken, and the system chooses one of the five possible magnetization directions. See Fig. 4.54. For sufficiently low temperature (i.e. large β), the lattice is mostly magnetized in one direction. In a MCMC simulation, an individual spin variable may easily flip to a neighboring direction, however, the probability for the whole lattice to flip magnetization direction is very small. The MCMC chain is effectively *stuck* or *frozen* in a particular magnetization direction. Fortunately, this is not a problem because the five magnetization directions are thermodynamically equivalent because, for example:

- In a fully magnetized lattice, the probability of any given spin flipping to either neighboring direction is the same regardless of the overall magnetization direction
- The much smaller probability of the entire lattice flipping to either neighboring direction is the same regardless of the overall magnetization direction

If the five directions were not equivalent, then the ensemble average of a thermodynamic observable from a Markov chain stuck in one magnetization direction would be different from the ensemble average from one stuck in a different direction. In that case, to get an accurate estimate of the true




Figure 4.54 At very low temperatures, the spin system tends to be fully or nearly fully magnetized. For the $q = 5.0$ or $q = 4.9$ state clock models, there are five possible directions for this magnetization—corresponding to the five “clock angles” (See Fig. 4.53). When $q = 5$, there is $\mathbb{Z}_5$ symmetry, and these directions are equivalent. When $q = 4.9$, this is not true.

value of the observable, one would need a Markov chain long enough to sample all magnetization directions many times. In the limit where temperature goes to zero, this would require an infinitely long Markov chain. Note, we refer here to a fully magnetized lattice only because that is easier to visualize and describe. However, this symmetry effect occurs also for nonmagnetized lattices. This may all be obvious and seemingly trivial, but it’s worth pointing out because this equivalence of the magnetization directions does not hold when the $\mathbb{Z}_q$ symmetry is broken, which is what happens when q is noninteger.

Consider what happens with $q = 4.9$. Now the clock angles $\varphi = 0, 2\pi/4.9, 4\pi/4.9, 6\pi/4.9, 8\pi/4.9$ are no longer equivalent. The clock angles divide the unit circle into four equal sectors and one slightly smaller leftover sector. See Fig. 4.53. As usual, at low temperature the lattice will tend to be magnetized in one of the five angles. However, these angles are no longer equivalent. Suppose *Case 1* is when the lattice is magnetized in the $\varphi = 0$ or $\varphi = 8\pi/4.9$ and *Case 2* is when it is magnetized in any of the other three directions. In Case 1, the lattice is magnetized in one of the directions adjacent to the slightly smaller “leftover” sector. As such, it is easier for a given spin to flip to a neighboring direction in Case 1 than it is in Case 2. Similarly, it is easier for the entire lattice to flip magnetization directions in Case 1 than it is in Case 2. Since Case 1 and Case 2 result in different spin-flipping probabilities, they define two thermodynamically distinct sectors. At large β , a $q = 4.9$ lattice will tend to freeze in one of the two Case 1 directions or in one of the three Case 2 directions. Two different MCMC simulations at the same parameters could give completely different ensemble averages if the Markov chain is too short and each simulation gets




Figure 4.55 Here we show some thermodynamic results for the model with $q = 4.9$, $h_q = 4$ on a 16×16 lattice. Clockwise from the top-left panel we show the energy density, magnetic susceptibility, structure factor, and specific heat. Data points are connected with lines to help guide the eyes. Near $\beta = 1.5$ we see that all four thermodynamic quantities become unstable in a manner that suggests the Markov chain became stuck in one of two thermodynamically distinct sectors. Clearly, the naive error bars (here the statistical variance only) are not enough to capture the true uncertainty.

stuck in a different sector. The problem gets more severe with increasing h_q , increasing β , or increasing lattice size. The end result is that for noninteger q in the low temperature limit or the large lattice size limit, with the MCMC approach one needs infinite statistics to get an accurate estimate of thermodynamic observables. This freezing effect was briefly discussed in Sec. 4.2.2.1 in the context of the $h_q \rightarrow \infty$ limit of the model, but the effect occurs also at finite and even at small h_q .

This effect can be seen in thermodynamic curves as in Fig. 4.55. In this figure, we show the energy density, magnetic susceptibility, structure factor, and specific heat for the model with $q = 4.9$ and $h_q = 4$ on a 16×16 lattice. Near $\beta = 1.5$ we see that all four thermodynamic quantities become unstable in a manner that suggests the Markov chain became stuck in one of two thermodynamically distinct sectors. Clearly, the naive error bars (here the statistical variance only) are not enough to capture the true uncertainty. Even a standard autocorrelation analysis does not capture this effect since the Markov chain may be stuck in one sector for the entire simulation. The

only way out seems to be to increase the statistics until the Markov chain repeatedly samples both thermodynamic sectors. However, the statistics needed makes the MCMC approach infeasible for even moderate lattices of size $L > 32$.

One way to clearly show the impact of the freezing phenomenon on observables is to run parallel simulations which differ only in the random number generator (RNG) seed. By chance, some of the runs may end up stuck in one sector while other runs may end up stuck in the other sector. See Fig. 4.56. Here we show observables estimated on Markov chains of different lengths. The data is from the model with $q = 4.9$, $h_q = 1$, and $\beta = 1.66$. The left column comes from an 8×8 lattice, and the right column comes from a 16×16 lattice. The rows, from top to bottom, are energy density, specific heat, proxy magnetization, susceptibility of the proxy magnetization, and structure factor. The horizontal axis is the log base-2 of the number of updating (BMHA) sweeps performed. That is, ‘12’ on the horizontal axis corresponds to 2^{12} updating sweeps. The different color points (with connecting lines to guide the eyes) correspond to different seeds for the random number generator. The error bars on the data points come only from the statistical variance of the data. These results show that the system tends to freeze in one of two thermodynamically distinct sectors, and a reliable estimate of the true value of an observable can only be obtained with very large statistics such that both sectors are sampled. For example, for the 8×8 results shown in the left column, we see that one needs a Markov chain of 2^{18} to 2^{20} BMHA sweeps to properly sample both sectors. Even there the error bars need to be much larger to account for the ultra-long autocorrelations caused by the freezing phenomenon. The problem worsens dramatically with increasing lattice size as we see for the 16×16 results in the right column where one needs statistics of 2^{27} to get reliable averages. For lattices larger than 32×32 , it becomes completely infeasible to reach the necessary statistics. Keep in mind that this phenomenon is a result of the broken $\mathbb{Z}_q$ symmetry and only occurs when q is not integer.

A practical question is, can one still do finite size scaling with Monte Carlo to extract critical exponents despite this freezing phenomenon? For noninteger q , the model exhibits two potential phase transitions. These manifest as two peaks in the specific heat, for example. Since this freezing




Figure 4.56 Here we show observables estimated on Markov chains of different lengths. The data is from the model with $q = 4.9$, $h_q = 1$, and $\beta = 1.66$. The left column comes from an 8×8 lattice, and the right column comes from a 16×16 lattice. The rows, from top to bottom, are energy density, specific heat, proxy magnetization, susceptibility of the proxy magnetization, and structure factor. The horizontal axis is the log base-2 of the number of updating (BMHA) sweeps performed. That is, ‘12’ on the horizontal axis corresponds to 2^{12} updating sweeps. The different color points (with connecting lines to guide the eyes) correspond to different seeds for the random number generator. These results indicate that the system tends to freeze in one of two thermodynamically distinct sectors, and a reliable estimate of the true value of an observable can only be obtained with very large statistics such that both sectors are sampled. This problem worsens dramatically with increasing lattice size. It is a result of the broken $\mathbb{Z}_q$ symmetry and only occurs when q is not integer.




Figure 4.57 Finite size scaling of the magnetic susceptibility (left) and the structure factor (right) for the model with $q = 4.9$ and $h_q = 1$ at the large- β peak where we expect a phase transition. According to the scaling law Eq. (3.79), these two quantities must scale with the same exponent. That scaling law is clearly violated here since $\gamma/\nu = 1.566(12)$ and $2 - \eta = 1.729(10)$. We believe this apparent violation is due to the freezing effect which occurs with MCMC at large lattice sizes. The result is that with MCMC alone, we cannot perform reliable finite size scaling for the Extended- $O(2)$ model with noninteger q .

phenomenon seems to occur just after the second peak (see Fig. 4.55), one might optimistically think there is sufficient reliable data to perform finite size scaling. That is what we attempted to do. See Fig. 4.57. In practice, the finite-size scaling seemed to work very well with nice peaks and power law scaling. It wasn't until we had enough data to check the scaling laws such as Eq. (3.79) that we found them to be severely violated for noninteger q . Such a scaling law violation implies that there is no phase transition here or there is something wrong with the MCMC results, and we believe it is the latter. To reliably study the model at noninteger q , we are forced to use an alternative approach such as tensor renormalization group (TRG) methods which do not suffer from this freezing effect.

4.3.8 Tensor Renormalization Group (TRG)

Using the character expansion

$$e^{\beta \cos(\theta_1 - \theta_2)} = \sum_{m=-\infty}^{\infty} I_m(\beta) e^{im(\theta_1 - \theta_2)}, \quad (4.70)$$

where I is the modified Bessel function of the first kind, the partition function Z can be expressed as

$$Z = \sum_{\{x,t\}} \prod_n A_{x_n t_n x_{n-1} t_{n-2}} \quad (4.71)$$

where

$$A_{ijkl} = \sqrt{I_i(\beta) I_j(\beta) I_k(\beta) I_l(\beta)} \int_0^{2\pi} \frac{d\theta}{2\pi} e^{i\theta(i+j-k-l)} e^{\beta h \cos \theta}. \quad (4.72)$$

If we do not have the symmetry-breaking term $h_q \cos \theta$, there is the “selection rule”, so that we can analytically evaluate the elements of the tensor. In the presence of the symmetry-breaking term, one may evaluate the integral using the Gaussian quadrature rule. There are some possible choices for polynomial functions to be used. For the noncompact scalar theories, Hermite polynomials are used. In our case the integral for each angle has a finite range, so Legendre polynomials are suitable. Then the elements of the tensor are defined by

$$A_{ijkl} \approx \sqrt{I_i(\beta) I_j(\beta) I_k(\beta) I_l(\beta)} \sum_{\nu=1}^K \frac{d\theta}{2\pi} e^{i\theta(i+j-k-l)} e^{\beta h \cos \theta}. \quad (4.73)$$

Note that the elements of A can be complex.

There is another (rather simpler) way to discretize the angles. By simply replacing the integral by summations, we obtain

$$Z \approx \left(\prod_n \sum_{\alpha_n=1}^K \right) \prod_n \frac{w_{\alpha_n}}{2} \prod_{\nu=1}^2 h(\pi x_{\alpha_n} + \pi, \pi x_{\alpha_{n+\hat{\nu}}} + \pi). \quad (4.74)$$

The local Boltzmann factor h can be regarded as a $K \times K$ matrix now and can be decomposed numerically by singular values:

$$h_{ij} = \sum_{k=1}^K U_{ik} \sigma_k V_{kj}^\dagger, \quad (4.75)$$

where $\{\sigma\}$ is the singular values and U and V are unitary matrices. Using them, we can approximately make a tensor network representation of Z :

$$Z \approx \left(\prod_n \sum_{\alpha_n=1}^K \right) \prod_n T_{x_n t_n x_{n-1} t_{n-2}}, \quad (4.76)$$

where the tensor is defined by


$$T_{ijkl} = \sum_{\alpha=1}^K \frac{w_{\alpha}}{2} \sqrt{\sigma_i \sigma_j \sigma_k \sigma_l} U_{\alpha i} U_{\alpha j} V_{k\alpha}^{\dagger} V_{l\alpha}^{\dagger}. \quad (4.77)$$

To balance between accuracy and efficiency, one may take a large K and initially truncate the bond dimension of the tensor. For the present work, K is set to 1024, and the accuracy is examined through a comparison with MC for some parameters. Note that in this construction the elements of T are real. In our work we employ the simpler discretization method.


Whereas Monte Carlo is a well-established approach with easily quantifiable uncertainties, it is an approach that sometimes struggles with critical slowing down. In the Extended-O(2) model studied in this paper, the Monte Carlo approach has no problem when q is integer. However, when q is noninteger, the Markov chain suffers from large autocorrelation times which worsen with increasing lattice size. The result is that for noninteger q , it becomes infeasible to produce de-correlated samples for large volumes e.g. for $L \gg 32$. In contrast, the TRG approach does not suffer from autocorrelation or critical slowing down. However, because of the truncations used, there are systematic uncertainties that are not always well-quantified. To be confident that our TRG methods are reliable, we compare them against MC in the regimes which are accessible to Monte Carlo.

In Fig. 4.58, we consider several different values of q and h_q at two different volumes. For noninteger values of q , we compare small volumes $L = 16, 32$. For integer values of q , we are able to compare larger volumes like $L = 128, 256$, since autocorrelations in Monte Carlo are much less severe for integer q . In the top panel of each figure, we compare the specific heat from Monte Carlo and TRG. With TRG, we compute the free energy and apply smoothing splines before computing the specific heat as the second derivative of the free energy. These figures show very good agreement between Monte Carlo and TRG for the specific heat.


With Monte Carlo, a zero external field ($h = 0$) is used, and so the magnetization $\vec{M}$ averaged over many equilibrium configurations would average to zero for all β . Instead, we measure a proxy magnetization $|\vec{M}|$ defined in Eq. (4.15). The corresponding susceptibility is defined in Eq. (4.16). With TRG, we use a different approach. With TRG, external magnetic fields are applied, and the




(a) $q = 3.9, h_q = 0.1$




(b) $q = 3.9, h_q = 1$




(c) $q = 4.1, h_q = 0.1$



(d) $q = 4.1, h_q = 1$



(e) $q = 5, h_q = 0.1$



(f) $q = 5, h_q = 1$

Figure 4.58 In each figure, we compare the specific heat from Monte Carlo (error bars and no connecting lines) with the specific heat from TRG (solid line). For the top two figures, we show the case $q = 3.9$ with $h_q = 0.1$ (left) and $h_q = 1$ (right). The β range was chosen to show both peaks in the specific heat, although at these small volumes, only one peak is visible for $h_q = 0.1$. In the middle two figures, we show the case $q = 4.1$ with $h_q = 0.1$ (left) and $h_q = 1$ (right). The β range was chosen to show small- β peak in the specific heat. For the case $h_q = 1$, the specific heats of the two different volumes lie directly on top of each other. In the bottom two figures, we show the case $q = 5$ with $h_q = 0.1$ (left) and $h_q = 1$ (right). At integer q , we are able to reliably go to larger volumes with Monte Carlo, so we compare $L = 128$ and $L = 256$ here. The β range was chosen to show the small- β peak in the specific heat.




Figure 4.59 We compute the susceptibility for $q = 3.9$ and $h_q = 1$ in the vicinity of the second-order transition in two different ways. In the top panel, we see how the susceptibility peaks from Monte Carlo approach the critical point, and in the bottom panel we look at the susceptibility peaks from TRG. From the Monte Carlo peak positions, we estimate $\beta_c = 1.19511(32)$ (vertical blue line) by fitting the peak locations to the finite-size-scaling form Eq. (3.78). From the TRG peak positions, we estimate $\beta_c = 1.196$ (vertical dashed line) via the same finite-size scaling form. As described in the main text, because of large autocorrelations, the Monte Carlo results at noninteger q are not strictly reliable for $L > 32$.

susceptibility is computed by taking the numerical derivative of the free energy with respect to the external field. Because of the different approaches used by Monte Carlo and TRG, it is not meaningful to directly compare the magnetization or the magnetic susceptibility. However, we expect the critical behavior to be the same in the thermodynamic limit. We show an example of this in Fig. 4.59 for the case $q = 3.9$ and $h_q = 1$. In the top panel, we show how the susceptibility (calculated from Monte Carlo) peaks approach the critical point with increasing volume. The critical point $\beta_c = 1.19511(32)$ (vertical blue line) was estimated by fitting the Monte Carlo peak locations to the finite-size scaling form Eq. (3.78). In the bottom panel of Fig. 4.59, we do the same but with the susceptibility calculated using TRG. Here, we estimate $\beta_c = 1.196$ (vertical dashed line) using the same finite-size scaling form. Note, that the Monte Carlo results are not strictly reliable for $L > 32$ in this figure since for $q = 3.9$ there are very large and unmitigated autocorrelations. Nevertheless, it seems the finite-size-scaling shift of the peaks are in general

agreement with the finite-size-scaling shift of the TRG peaks.

4.3.9 Mapping the Phase Diagram

4.3.9.1 Integer q

For integer values of q , the Extended- $O(2)$ model retains a rotational symmetry. When $h_q = 0$, this is the full $O(2)$ symmetry. When h_q is finite, this symmetry reduces to a $\mathbb{Z}_q$ symmetry. The model is easier to study at integer q because the autocorrelations remain manageable even at large β .

To get a good survey of the model, we study $q = 2, 3, \dots, 6$ at $h_q = 0.1, 1.0$. At infinite h_q , the Extended- $O(2)$ model reduces to the ordinary clock model which for $q = 2, 3, 4$ has a single second-order transition, and for $q \geq 5$, has a pair of infinite-order BKT transitions. See Table 2.2. We start with the assumption that the same kind of transitions persist to finite h_q . We show then that this assumption fails to hold at small h_q .

We performed MCMC simulations at $L = 16, 32, 64, 96, 128, 160, 192, 224, 256$. A BMHA updating algorithm was used with an acceptance rate around 90%. Each run started with 2^{15} equilibration sweeps followed by 2^{16} measurement sweeps. Each measurement sweep was separated by 2^6 (or more) discarded sweeps to help mitigate autocorrelation. Later during the analysis stage, the residual integrated autocorrelation time was estimated and the effect removed by binning the observables. For some of the larger volumes, more runs were performed to increase the statistics.

Using the FSS fit forms Eq. (3.70)–(3.75), we extracted the exponents $1/\nu$, α/ν , $-\beta/\nu$, γ/ν , and $2 - \eta$ for integer q and $h_q = 0.1, 1.0$. These are tabulated in Table B.1. The critical exponents ν , α , β , γ , and η are then extracted and listed in Table 4.3. For example, the exponent α is obtained by taking the ratio of measured exponents α/ν and $1/\nu$. This is done at the level of jackknife bins in order to get reliable error bars on the final critical exponents. For $q = 5, 6$, we do not include results from Eq. (3.70) since the quantity seems unstable and the fits tend to be bad. As a result, the bare exponents for $q = 5, 6$, which depend on $1/\nu$ are not listed in Table 4.3.

For $q = 2, 3, 4$, we fit the thermodynamic quantities to Eq. (3.78) to estimate the critical exponent

Table 4.3 The first two columns refer to the model parameters, and the remaining columns list the critical exponents obtained for the model. They are obtained from the ratio exponents (e.g. Table B.1, which are in turn obtained from finite-size-scaling. Some exponents could not be reliably extracted with our data. For $q = 5, 6$, the results refer to the transition at small inverse temperature.

q	h_q	ν	α	β	γ	η
2	0.1	1.044(53)	0.021(18)	0.319(94)	1.859(96)	0.287(13)
2	1.0	1.022(71)	-0.010(16)	N/A	1.78(12)	0.251(14)
3	0.1	0.658(24)	0.354(35)	N/A	1.291(45)	0.3651(93)
3	1.0	0.809(26)	0.311(21)	N/A	1.411(36)	0.295(16)
4	0.1	2.49(89)	-0.30(12)	0.76(87)	4.3(1.5)	0.2570(85)
4	1.0	1.20(14)	-0.162(19)	0.78(61)	2.09(24)	0.2532(93)
5	0.1	N/A	N/A	N/A	N/A	0.2716(72)
5	1.0	N/A	N/A	N/A	N/A	0.2653(94)
6	0.1	N/A	N/A	N/A	N/A	0.2880(86)
6	1.0	N/A	N/A	N/A	N/A	0.2654(92)

Table 4.4 Here we tabulate the jackknife average of the estimates of ν and β_c listed in Tables B.2 and B.3. For $q = 5, 6$, the results refer to the transition at small inverse temperature.

q	h_q	ν	β_c
2	0.1	1.059(50)	0.87756(46)
2	1.0	0.985(58)	0.65448(24)
3	0.1	0.878(21)	1.00044(56)
3	1.0	0.812(22)	0.82584(11)
4	0.1	2.04(18)	1.0841(84)
4	1.0	1.361(74)	0.9916(18)
5	0.1	N/A	1.136(11)
5	1.0	N/A	1.128(12)
6	0.1	N/A	1.1251(83)
6	1.0	N/A	1.1143(92)




Figure 4.60 Here we look at data for some volumes ($L = 16, \dots, 256$) for the model with $q = 2$ and $h_q = 0.1$. The top panel shows the magnetic susceptibility, the bottom panel shows the specific heat, and the inset within the top panel shows the Binder cumulant. The data is from reweighted Monte Carlo with error bars included. The vertical dashed line is at $\beta_c \approx 0.8773$ where the Binder cumulant curves cross. This figure illustrates the behavior typical of a second-order phase transition—the Binder cumulants cross at the transition point, the magnetic susceptibility diverges, and the specific heat diverges logarithmically.

ν and critical points β_c . The results are tabulated in Table B.2. Similarly, we use Eq. (3.84), to estimate the BKT critical points for $q = 5, 6$, and list them in Table B.3. The estimates of ν and β_c are averaged (at the level of jackknife bins) to obtain the final estimates listed in Table 4.4. Compare the estimates of the critical exponent ν in Table 4.3, which were obtained from Eq. (3.70) with the estimates in Table 4.4, which were obtained from Eq. (3.78). The results are consistent except for the model with $q = 3$ and $h_q = 0.1$. The cause of this discrepancy is not clear.

For $q = 2$ with $h_q = \infty$ (i.e. the Ising model), there is a second-order phase transition at $\beta_c = \ln(1 + \sqrt{2})/2 \approx 0.4407$ with critical exponents $\nu = 1$, $\alpha = 0$, $\gamma = 7/4$, $\beta = 1/8$, and $\eta = 1/4$ (see Table 2.2). With $h_q = 1.0$, the thermodynamic quantities show Ising-like divergences near a critical point, which we estimate to be $\beta_c = 0.65448(24)$ in the infinite-volume limit. From finite-size-scaling, we find $1/\nu = 0.978(68)$, $\alpha/\nu = -0.010(16)$, $\gamma/\nu = 1.745(19)$, and $2 - \eta = 1.749(14)$. The bare exponents are $\nu = 1.022(71)$, $\alpha = -0.010(16)$, $\gamma = 1.78(12)$, and

$\eta = 0.251(14)$. These are consistent with the critical exponents of the Ising universality class. We were unable to extract β/ν from our data. When we decrease h_q to 0.1, the picture changes a little. The thermodynamic quantities show similar Ising-like divergences, however, the model is no longer in the Ising universality class. The specific diverges logarithmically (see Fig. 4.60). The critical point is now at $\beta_c = 0.87756(46)$. From finite-size-scaling, we find $1/\nu = 0.958(49)$, $\alpha/\nu = 0.020(17)$, $\gamma/\nu = 1.780(16)$, $\beta/\nu = 0.306(81)$, and $2 - \eta = 1.713(13)$. The bare exponents $\nu = 1.044(53)$, $\alpha = 0.021(18)$, $\beta = 0.319(94)$, $\gamma = 1.859(96)$, and $\eta = 0.287(13)$, are not consistent with the Ising universality class, however, the scaling relations (in particular Eq. (3.79) and (3.81)) are violated, so this case is somewhat inconclusive. As $h_q \rightarrow 0$, the critical point moves to larger β and seems to connect with the BKT transition near $\beta_c = 1.12$ of the XY model when $h_q = 0$. So for $q = 2$ with large h_q , there is a second-order phase transition of the Ising universality class. However, at small h_q , the second-order transition is no longer of the Ising class—contradicting an early renormalization group analysis [114].

For $q = 3$ with $h_q = \infty$ (i.e. the 3-state Potts model), there is a second-order phase transition at $\beta_c = 2 \ln(1 + \sqrt{3})/3 \approx 0.6700$, with critical exponents $\nu = 5/6$, $\alpha = 1/3$, $\gamma = 13/9$, $\beta = 1/9$, and $\eta = 4/15$ (see Table 2.2). At $h_q = 1.0$, the thermodynamic functions show second-order divergences near a critical point, which we estimate to be $\beta_c = 0.82584(11)$ in the infinite-volume limit. From finite-size-scaling, we find $1/\nu = 1.236(40)$, $\alpha/\nu = 0.385(21)$, $\gamma/\nu = 1.743(16)$, and $2 - \eta = 1.705(16)$ which are consistent with the exponents at $h_q = \infty$. The bare exponents are $\nu = 0.809(26)$, $\alpha = 0.311(21)$, $\gamma = 1.411(36)$, and $\eta = 0.295(16)$. When we reduce h_q to 0.1, the critical point moves to $\beta_c = 1.00044(56)$. Assuming a second-order transition, we find from finite-size-scaling $1/\nu = 1.521(57)$, $\alpha/\nu = 0.538(48)$, $\gamma/\nu = 1.962(16)$, and $2 - \eta = 1.6349(93)$, which are very different from the 3-state Potts universality class. The bare exponents are $\nu = 0.658(24)$, $\alpha = 0.354(35)$, $\gamma = 1.291(45)$, and $\eta = 0.3651(93)$. The scaling relations are violated in this case. As $h_q \rightarrow 0$, the critical point moves to larger β and seems to connect with the BKT transition of the XY model when $h_q = 0$.

For $q = 4$ with $h_q = \infty$, the model is known to be in the Ising universality class with critical point




Figure 4.61 Here we look at data for some volumes ($L = 16, \dots, 256$) for the model with $q = 4$ and $h_q = 0.1$. The top panel shows the magnetic susceptibility, the bottom panel shows the specific heat, and the inset within the top panel shows the Binder cumulant. The data is from reweighted Monte Carlo with error bars included. The vertical dashed line is at $\beta_c \approx 1.09$ where the Binder cumulant curves merge. This figure illustrates the behavior typical of a BKT phase transition—the Binder cumulants merge at the transition point, the magnetic susceptibility diverges, and the specific heat plateaus.

at $\beta_c = \ln(1+\sqrt{2}) \approx 0.8814$ and critical exponents given in Table 2.2. At $h_q = 1$, the thermodynamic functions seem to diverge near a critical point which we estimate to be $\beta_c = 0.9916(18)$ in the infinite-volume limit. From finite-size-scaling, we find $1/\nu = 0.834(96)$, $\gamma/\nu = 1.742(14)$, $-\beta/\nu = -0.65(51)$, and $2-\eta = 1.7468(93)$. When extracting the exponent α/ν , attempts to fit to the form Eq. (3.71) failed. However, a fit to the form Eq. (3.75) yielded $\alpha/\nu = -0.1349(90)$. The bare exponents are $\nu = 1.20(14)$, $\alpha = -0.162(19)$, $\beta = 0.78(61)$, $\gamma = 2.09(24)$, and $\eta = 0.2532(93)$. With the large error bars, these exponents are still consistent with the Ising universality class. When we reduce h_q to 0.1, the critical point moves to $\beta_c = 1.0841(84)$. Assuming a second-order transition, we find from finite-size-scaling $1/\nu = 0.40(14)$, $\alpha/\nu = -0.120(19)$, $\gamma/\nu = 1.728(16)$, $-\beta/\nu = -0.31(31)$, and $2-\eta = 1.7430(85)$. The bare exponents are $\nu = 2.49(89)$, $\alpha = -0.30(12)$, $\beta = 0.76(87)$, $\gamma = 4.3(1.5)$, and $\eta = 0.2570(85)$. These are not consistent with the Ising universality class. Since α is negative and ν is finite, the specific heat decreases with increasing




Figure 4.62 Here we look at data for some volumes ($L = 16, \dots, 256$) for the model with $q = 5$ and $h_q = 1.0$. The top panel shows the magnetic susceptibility, and the bottom panel shows the specific heat. The inset within the top panel shows the Binder cumulant of the proxy magnetization. The vertical dashed line is at $\beta_{c1} \approx 1.08$, where the Binder cumulant curves cross. The inset within the bottom panel shows the Binder cumulant of the rotated magnetization defined in Eq. (4.51). The vertical dotted line is at $\beta_{c2} \approx 1.22$, where the Binder cumulant curves cross. The data is from reweighted Monte Carlo with error bars included. This figure illustrates the behavior typical of the two BKT transitions that occur in the ordinary clock model for $q \geq 5$. The magnetic susceptibility diverges at both transitions, and the specific heat plateaus at both transitions. For $L = 256$, we show only results near the first transition point. Already here the error bars are rather large.

volume. This can be seen in Fig. 4.61. Other thermodynamic quantities diverge, and the Binder cumulants cross suggesting a BKT transition instead of a second-order transition. The scaling relation Eq. (3.80) does not hold in this case, but the other two relations do hold. As $h_q \rightarrow 0$, the critical point moves to slightly larger β and seems to connect with the BKT transition of the XY model when $h_q = 0$.

Finally, we consider $q = 5, 6$. At $h_q = \infty$, these models both show a pair of BKT transitions. At finite h_q there also seems to be two BKT transitions for $q = 5, 6$. See Fig. 4.62. To extract critical exponents, we look at maxima in the thermodynamic quantities. This works well for the low- β transition, but not for the large- β transition since many of the thermodynamic quantities do not show a well-defined peak near the second transition. Nevertheless, the transition point can be located

by the Binder cumulant crossings. In a BKT transition, the specific heat does not diverge with volume, and one cannot use Eq. (3.71) to extract a critical exponent α . Furthermore, the correlation length diverges faster than any power, and so Eq. (3.70) cannot be used to extract ν . In fact, in the conventional sense, ν must be infinite for a BKT transition. Fits to the forms Eq. (3.72)–(3.74) are performed and exponents recorded in Table B.1. The results seem generally consistent with BKT transitions. For $q = 5, 6$, estimates of the critical points for the small- β transition obtained by fitting to Eq. (3.84) are listed in Table B.3. The average values are listed in Table 4.4. However, it is not clear that Eq. (3.84) is a reliable way to estimate the critical point for BKT transitions, and the method of Binder cumulant crossings may be preferred. For $q = 5, 6$, the small- β transition connects to the BKT transition of the XY model when $h_q = 0$. The large- β transition moves to larger β as $h_q \rightarrow 0$. Recently, it was estimated that for $q = 5$, the large- β transition limits to $\beta_{c2} \approx 2.27$ as $h_q \rightarrow 0$ [122].

We summarize now the results for integer q . Recall that for $h_q = \infty$, this model becomes the ordinary q -state clock model, which is known to have a second-order transition for $q = 2, 3, 4$, and a pair of BKT transitions for $q \geq 5$. In this work, we’ve studied what happens at finite h_q . For $q \geq 5$, there seems to be a pair of BKT transitions for all values of $h_q > 0$. For $q = 2, 3, 4$, there appears to be a second-order transition for large and intermediate values of h_q . The critical exponents, and hence the universality class, seems to vary with h_q , and this transition appears to turn into a BKT transition in the small h_q limit. There are very few results in the literature even for integer q (we found no literature on noninteger q) that give specific numbers such as critical exponents or critical points. Of the few results that we found, our findings seem to agree with, for example, [40, 41, 117] but disagree with [115]. The reason for this is not clear to us.

4.3.9.2 Noninteger q

At noninteger q , the explicitly broken $\mathbb{Z}_q$ symmetry causes difficulties for the MCMC approach. In fact, we are unable to get reliable results on lattices larger than 32×32 due to the freezing effect described in Sec. 4.3.7. This makes finite-size scaling infeasible with the MCMC approach. Instead, we use TRG to investigate the noninteger q regime. We validate the TRG approach by




Figure 4.63 Heatmaps from TRG of the specific heat (top row) and entanglement entropy (bottom row) for the Extended-O(2) model with $h_q = 1$ (left column) and $h_q = 0.1$ (right column). Each pixel is a distinct TRG calculation performed with $L = 1024$ and bond dimension 40. There are a few missing points (white) in the top-left panel.

comparison with MCMC on smaller lattices. See Sec. 4.3.8.

With TRG, we started with a low precision scan of the parameter space for $q \in [1, 6]$, $\beta \in [0, 2]$, $h_q = 0.1, 1.0$, and $L = 4, 8, \dots, 1024$. For example, we show in Fig. 4.63 heatmaps of the specific heat (top row) and entanglement entropy (bottom row) for $h_q = 1$ (left column) and $h_q = 0.1$ (right column). Each pixel is a data point obtained from TRG performed with $L = 1024$ and bond dimension 40. The color in each heatmap ranges from dark to light and this corresponds to a value ranging from 0 to 2.5. The cutoff choice of 2.5 (which truncates some of the specific heat values) was made to increase the contrast in the heatmaps. For $h_q = 1$, the heatmap of the specific heat shows smooth lobes suggesting smooth lines in the phase diagram even as one crosses integer values of q . When h_q is reduced to 0.1, the lines at large β fade away leaving a thick horizontal line that connects to the BKT transition at $h_q = 0$. However, heatmaps of the entanglement entropy show that there remain definite discontinuities as one crosses integer values of q . This suggests a phase diagram at finite- h_q that is similar to the phase diagram at infinite h_q .




Figure 4.64 Here we look at the behavior of the magnetic susceptibility near the small- β peak for noninteger q . In this example, $q = 4.1$. We obtain the susceptibility from TRG with bond dimension 32 after applying small external magnetic fields $h = 10^{-5}$, $\delta h = 10^{-5}$. In the left panel, $h_q = 0.01$ and in the right panel, $h_q = 0.001$. For sufficiently large volumes, the magnetic susceptibility peaks plateau—implying a crossover. For a true phase transition, the peaks would diverge with increasing volume. As the symmetry-breaking parameter h_q is decreased toward zero, one must go to larger volumes to show that the peaks are not diverging.

At noninteger values of q , the specific heat generally shows two peaks. In the $h_q = \infty$ limit, we found that the first peak is only a crossover, but the second peak is associated with a second-order phase transition of the Ising universality class. Here, we investigate the finite h_q regime.

First we consider the small- β peak, which we know is a crossover at $h_q = \infty$. On the other hand, this peak connects to the BKT transition of the XY model when $h_q = 0$. The question is; what happens at finite h_q ? In Fig. 4.64, we find that the magnetic susceptibility plateaus with volume—indicating a crossover—since a phase transition here would result in divergence of the susceptibility with volume. However, as we take $h_q \rightarrow 0$, it becomes more and more difficult to distinguish between crossover and BKT transition as it takes larger and larger lattices for the susceptibility to plateau. We conclude that most likely the small- β peak corresponds to a crossover even as $h_q \rightarrow 0$ and only becomes a BKT transition exactly at $h_q = 0$. However, we suspect that to prove this for values of h_q very close but not equal to zero one would require very large lattices.

Next, we consider the large- β peak, which corresponds to a second-order transition of the Ising universality class when $h_q = \infty$. As q approaches an integer from below, the second-order transition (i.e. the second peak) occurs at relatively small values of β . As q is decreased toward the next




Figure 4.65 In the top panel, data collapse of the magnetic susceptibility from TRG for the large- β peak for the model with $q = 3.9$ and $h_q = 1$. The estimate $\beta_c = 1.196$ was obtained by fitting the peak positions to the finite-size scaling form Eq. (3.78). To extract the magnetic susceptibility, an external field $h = 40/L^{15/8}$ was imposed ($dh = 10^{-5}$ for the numerical differentiation). In the bottom panel, data collapse of the specific heat from TRG for the same model. These results are consistent with there being a second-order phase transition here of the Ising universality class.

integer, the critical point moves to large values of β . In Figure 4.65, we show TRG results for $q = 3.9$ and $h_q = 1$ near the large- β peak. In the top panel, we show the data collapse of the magnetic susceptibility from TRG. The estimate $\beta_c = 1.196$ was obtained by fitting the peak positions to the finite-size scaling form Eq. (3.78). To extract the magnetic susceptibility, an external field $h = 40/L^{15/8}$ was imposed ($dh = 10^{-5}$ for the numerical differentiation). In the bottom panel, we show the data collapse of the specific heat from TRG for the same model. The specific heat was computed from the second derivative of the free energy after applying smoothing splines to the free energy. The data collapse curves are consistent with a second-order phase transition of the Ising universality class.

4.3.9.3 The Phase Diagram at Finite h_q

In Fig. 4.66 we present four two-dimensional conjectures for the full three-parameter phase diagram. Each slice is at a different value of the symmetry-breaking parameter h_q . In the top-left panel, the phase diagram at $h_q = 0$. This is the well-studied XY model at all values of q . There is a




Figure 4.66 Here we present four two-dimensional conjectures for the full three-parameter phase diagram. Each slice is at a different value of the symmetry-breaking parameter h_q . In the top-left panel, the phase diagram at $h_q = 0$. This is the well-studied XY model at all values of q . There is a single BKT transition near $\beta = 1.12$ and a BKT critical phase at larger β . The top-right panel gives the conjectured phase diagram at small $h_q > 0$. For noninteger q , there is generally a crossover at small β and an Ising phase transition at larger β with an Ising ordered phase at large β . For integer q , our data is consistent with BKT transitions for all integer q given sufficiently small (but positive) h_q . For $q = 2, 3, 4$ our conjecture that these are BKT transitions is rather tenuous as it is possible that the apparent BKT transition is really the result of studying the model on too small lattices. For integer $q \geq 5$, there are two BKT transitions with the second transition appearing at larger β . This transition does not seem to connect to the BKT line at $h_q = 0$ in the limit $h_q \rightarrow 0$ [122]. In the bottom-left panel, h_q is increased to intermediate values. The most significant change is that for $q = 2, 3, 4$ the transitions are now second-order, however, they are not in the same universality class as the clock model (i.e. $h_q = \infty$) transitions. Finally, in the bottom-right panel, we show the phase diagram for the $h_q \rightarrow \infty$ case. For integer q , this is the ordinary q -state clock model, and for noninteger q , this is what we call the “Extended q -state clock model”.

single BKT transition near $\beta = 1.12$ and a BKT critical phase at larger β . The top-right panel gives the conjectured phase diagram at small $h_q > 0$. For noninteger q , there is generally a crossover at small β and an Ising phase transition at larger β with an Ising ordered phase at large β . For integer q , our data is consistent with BKT transitions for all integer q given sufficiently small (but positive) h_q . For $q = 2, 3, 4$ our conjecture that these are BKT transitions is rather tenuous as it is possible that the apparent BKT transition is really the result of studying the model on too small lattices. For integer $q \geq 5$, there are two BKT transitions with the second transition appearing at larger β . This transition does not seem to connect to the BKT line at $h_q = 0$ in the limit $h_q \rightarrow 0$ [122]. In the bottom-left panel, h_q is increased to intermediate values. The most significant change is that for $q = 2, 3, 4$ the transitions are now second-order, however, they are not in the same universality class as the clock model (i.e. $h_q = \infty$) transitions. Finally, in the bottom-right panel, we show the phase diagram for the $h_q \rightarrow \infty$ case. For integer q , this is the ordinary q -state clock model, and for noninteger q , this is what we call the “Extended q -state clock model”.

4.4 Future Questions

There are many unanswered questions about the models studied in this thesis. Some of the ones that the author finds most interesting include the following list of questions which are asked in no particular order:

- What is really happening in the $h_q \rightarrow 0$ limit for $q = 2, 3, 4$? Is it a second-order phase transition for all $h_q > 0$ or does it turn into a BKT transition at finite h_q ? This should be a straightforward question to answer, but it will likely require very large lattices.
- In the $O(2)$ model and in the q -state clock models, vortices are often claimed to play a significant role in the BKT phase transitions. What role do vortices play in the Extended- $O(2)$ model? This model at noninteger q might help us to clarify the role of vortices in phase transitions in two-dimensional systems.
- The Extended- $O(2)$ model offers multiple ways in which one can continuously vary a parameter to go from a BKT transition to a second-order transition and vice versa. For example,

by varying h_q (for many values of q), one can go from the XY model with a BKT transition at $h_q = 0$ to a model with a second-order transition at $h_q > 0$. Alternatively, one can go from a second-order transition to a BKT transition for example by dialing from $q = 4.999999$ to $q = 5$. There are many other examples in the Extended- $O(2)$ model in which one can vary h_q or q to switch between a model with a BKT transition and one with a second-order transition. This makes the model a nice sandbox within which to study these two kinds of phase transitions and their relationship.

- It is interesting to note that the BKT critical point found in the $O(2)$ model—and here at the limit of the extended phase diagram—can also be reached through a completely different interpolation. By considering the $O(3)$ nonlinear sigma model with an additional symmetry breaking term which breaks the $O(3)$ symmetry down to an $O(2)$ symmetry, one can interpolate between $\mathbb{Z}_2$ to $O(3)$, and from $O(3)$ to $O(2)$ by tuning the sign, and magnitude, of the additional symmetry-breaking term [161]. Further additional symmetry breaking terms could be interesting. Positive-definite worm algorithms have been constructed for the $O(3)$ nonlinear sigma model, and could be used to simulate the model efficiently [162, 163].
- One of the original motivations for studying these models was related to making $\mathbb{Z}_q$ approximations of $U(1)$ gauge theories. It would be interesting to revisit that motivation especially to see how q -state clock models with noninteger values of q might be useful when making discrete approximations of $U(1)$.
- In this thesis, we have focused on the Extended q -state clock model and the Extended- $O(2)$ model in two dimensions. It might be interesting to study these models in other dimensions as well.
- In this thesis, we interpolate between the $O(2)$ model and its discrete $\mathbb{Z}_q$ approximations by adding a term to the $O(2)$ model which contains a symmetry-breaking parameter h_q and a discretization parameter q . We go further and consider q -state clock models with noninteger q . It might be interesting to consider a similar approach to interpolate between $O(3)$ and

various discrete approximations of $O(3)$ and further to consider the possible generalizations of such models to noninteger values of the discretization parameter.

CHAPTER 5

QUANTUM FIELD THEORIES ON THE LATTICE

5.1 Introduction

The previous chapters of this thesis focused on the study of spin systems using primarily MCMC methods. Although these general methods were developed for and work well for classical systems, one can also apply them to the study of quantum field theories. In fact, the field of lattice QCD uses many of the same numerical methods that were described in this thesis. As such, in the concluding chapter of this thesis we give a brief sketch of QCD and how it is studied on the lattice.

The author's study of the Extended- $O(2)$ model, which was detailed extensively in the previous chapter, has several connections to the quantum simulation of quantum field theories. In fact, the original motivation for studying the model stemmed from recent attempts to quantum simulate lattice models with continuous Abelian symmetries using discrete approximations. Thus, as a natural step beyond conventional lattice field theory we conclude with a sketch of the next generation approach—quantum simulation. But first, we introduce QCD in the continuum.

5.2 QCD in the Continuum

Quantum chromodynamics (QCD), a key component of the Standard Model of particle physics, is the theory of the strong interaction which occurs between color-charged particles like quarks and gluons.

The QCD Lagrangian is easy to write down, and it looks like the Lagrangian of simpler theories such as quantum electrodynamics. However, it is a non-Abelian theory and the gauge field interacts with itself, making the theory a lot more difficult to deal with. The QCD Lagrangian density is

$$\mathcal{L} = \sum_f \bar{\psi}_i^f \left(i\gamma_\mu D_{ij}^\mu - m_f \delta_{ij} \right) \psi_j^f - \frac{1}{4} F_{\mu\nu}^a F_a^{\mu\nu}, \quad (5.1)$$

where $D_{ij}^\mu = \partial_\mu \delta_{ij} + ig(T^a)_{ij} A_a^\mu$ is the gauge covariant derivative, and $F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a - gf_{abc} A_\mu^b A_\nu^c$ is the gluon field strength tensor. The quark fields $\psi_i^f(x)$ come in six flavors $f = \{\text{up, down, charm, strange, top, bottom}\}$ which vary in mass m_f . The gluon fields $A_a^\mu(x)$ are coupled to the quark fields with coupling strength g via the $SU(3)$ generators T^a and to each other via




Figure 5.1 In QCD, the coupling constant $\alpha_s = g^2/4\pi$ varies with the energy scale—a phenomenon referred to as the “running of the coupling”. As the energy and momentum transfer Q^2 increases, the coupling decreases logarithmically—a phenomenon known as asymptotic freedom. On the other hand, the coupling increases with decreasing energy, resulting in confinement.

the $SU(3)$ structure constants f_{abc} . Repeated indices indicate summation with $a, b, c = 1, \dots, 8$ with $i, j = 1, 2, 3$ (color indices) and with $\mu, \nu = 0, 1, 2, 3$ (Dirac indices). The QCD action is the Lagrangian density integrated over all spacetime

$$S = \int d^4x \mathcal{L}. \quad (5.2)$$

At high energies and small length scales, the coupling between color-charged particles becomes weak. This is the phenomenon of *asymptotic freedom* [164, 165]. See Figure 5.1. In this regime, the coupling is small, and perturbation theory is applicable. At low energies and larger length scales, the coupling between particles becomes strong, resulting in the phenomenon of *confinement*. To study processes in this regime, one must use nonperturbative methods such as lattice QCD [166].

The Green’s functions, i.e. the time-ordered vacuum expectation values

$$G(x_1, \dots, y_1, \dots, z_1, \dots) = \left\langle 0 \left| \mathcal{T} \left\{ \hat{A}(x_1) \cdots \hat{\psi}(y_1) \cdots \hat{\bar{\psi}}(z_1) \cdots \right\} \right| 0 \right\rangle, \quad (5.3)$$

encode everything in the theory. The 2-point Green’s function, for example, is the propagator. The Green’s functions can be written as path integrals

$$G(x_1, \dots, y_1, \dots, z_1, \dots) = \frac{\int \mathcal{D}A \mathcal{D}\bar{\psi} \mathcal{D}\psi A(x_1) \cdots \psi(y_1) \cdots \bar{\psi}(z_1) \cdots e^{iS[A, \bar{\psi}, \psi]}}{\int \mathcal{D}A \mathcal{D}\bar{\psi} \mathcal{D}\psi e^{iS[A, \bar{\psi}, \psi]}} \quad (5.4)$$

where the integration goes over all field configurations.

The QCD generating functional takes the form¹

$$Z[J, J_\psi, J_{\bar{\psi}}] = \frac{\int \mathcal{D}A \mathcal{D}\bar{\psi} \mathcal{D}\psi e^{i \int d^4x \left(\mathcal{L}[A, \bar{\psi}, \psi] + J_\mu A^\mu + J_\psi \psi + \bar{\psi} J_{\bar{\psi}} \right)}}{\int \mathcal{D}A \mathcal{D}\bar{\psi} \mathcal{D}\psi e^{iS[A, \bar{\psi}, \psi]}}. \quad (5.5)$$

Notice the similarity between the QCD generating functional and the partition function of a classical field theory. The generating functional is so-called because one can generate the Green's functions by taking functional derivatives as in

$$G(x_1, \dots, y_1, \dots, z_1, \dots) = \frac{\delta \dots}{\dots i \delta J(x_1) \dots i \delta J_\psi(y_1) \dots i \delta J_{\bar{\psi}}(z_1)} Z[J, J_\psi, J_{\bar{\psi}}] \Big|_{J=J_\psi=J_{\bar{\psi}}=0}. \quad (5.6)$$

5.2.1 Perturbative QCD

Here, we give only the barest sketch of the established field of perturbative QCD, mainly to contrast it later with the ab initio approach of lattice QCD.

The non-interacting Lagrangian density $\mathcal{L}_0$ (e.g. $\mathcal{L}$ with $g = 0$) is quadratic in each of the fields. This allows one to “complete the square” and evaluate the field integrals as Gaussian integrals, to obtain²

$$Z_0[J, J_\psi, J_{\bar{\psi}}] = \mathcal{N} e^{i \int d^4x \left(\frac{1}{2} J^\mu(x) D_{\mu\nu}(x-y) J^\nu(y) + J_{\bar{\psi}}(x) S_F(x-y) J_\psi(y) \right)}, \quad (5.7)$$

where $D_{\mu\nu}(x-y)$ is the free gauge field propagator, $S_F(x-y)$ is the free fermion propagator, and $\mathcal{N}$ is an undetermined constant factor. Then the generating functional of the full theory can be obtained perturbatively by expanding in g

$$Z[J, J_\psi, J_{\bar{\psi}}] = \mathcal{N} \exp \left\{ i \int d^4z \mathcal{L}_{interacting} \left(-i \frac{\delta}{\delta J}, i \frac{\delta}{\delta J_\psi}, -i \frac{\delta}{\delta J_{\bar{\psi}}} \right) \right\} Z_0[J, J_\psi, J_{\bar{\psi}}] \quad (5.8)$$

See, for example, [167]. Notice that the infinite-dimensional integrals over the field configurations (with measures $\mathcal{D}A \mathcal{D}\bar{\psi} \mathcal{D}\psi$) are gone now, so there is a possibility of actually evaluating some of these integrals.

¹In practice, there may also be “ghost” sources, and the Lagrangian may contain ghost fields and a gauge fixing term.

²This is simply a schematic picture as we are neglecting things like gauge fixing and Fadeev-Popov terms.

To connect to collider experiments, one may be interested in computing the elements of the scattering matrix, as these can be related to collider observables such as cross-sections and decay widths. For scattering, one assumes an “incoming” state $|\{p_i\}\rangle$ prepared in the distant past with a definite number of particles with definite momenta $\{p_i\}$. The particles interact in a localized volume and a localized time and scatter from each other. In the far future, the “outgoing” state $|\{k_j\}\rangle$ particles are assumed to no longer be interacting. Then the elements of the scattering matrix are denoted $S_{ji} = \langle\{k_j\}|\{p_i\}\rangle$, and these are related to the momentum space Green’s functions. One approach of “perturbative QCD” is to expand such momentum space path integrals in powers of the coupling constant g . It turns out that there is an extremely useful diagrammatic representation—called “Feynman diagrams”—of the resulting integrals.

Unfortunately, many of the integrals in the perturbative expansion are divergent. Some diverge in the IR (energy $\rightarrow 0$) limit and others diverge in the UV (momentum $\rightarrow \infty$) limit. To resolve the divergences, the integrals must be regularized in some manner. There are several regularization approaches including Pauli-Villars which regulates the UV divergences, adding small masses which regulates the IR divergences, and dimensional regularization which cures both types of divergences. In dimensional regularization, one exploits the fact that the integrals converge in higher dimensions to solve the integrals in arbitrary d dimensions and then take the limit $d \rightarrow 4$ in the end. This approach is often preferred because it preserves both Lorentz invariance and gauge invariance.

Dimensional regularization allows one to evaluate the integrals, which are divergent in $d = 4$ dimensions, by going to arbitrary dimensions d . However, the results remain divergent when one takes the limit $d \rightarrow 4$. To connect the theory with the physical world, one can directly or indirectly redefine the fields or parameters of the theory via physical observables. In effect, one can absorb the divergent terms into a redefinition of the fields or parameters in a process called “renormalization”. At the lowest order in perturbation theory, the parameters of the theory (e.g. quark masses) are physically meaningful. However, at higher orders where the divergences occur, these “bare” parameters are no longer the physical parameters. Rather, it is the renormalized parameters that are physically meaningful.

Perturbative QCD at small α_s has proven very useful, however, at some point, the coupling is no longer a small parameter and perturbation theory can no longer be used. Even in the regimes where it does work, convergence is slow—even at large momentum—and this makes perturbative QCD calculations very difficult.

5.3 Lattice Regularization

Here, we follow closely Ref. [66]. Recall that the vacuum expectation value of an observable O can be written as a path integral

$$\langle O \rangle = \frac{\int \mathcal{D}A \mathcal{D}\bar{\psi} \mathcal{D}\psi O[A, \bar{\psi}, \psi] e^{iS[A, \bar{\psi}, \psi]}}{\int \mathcal{D}A \mathcal{D}\bar{\psi} \mathcal{D}\psi e^{iS[A, \bar{\psi}, \psi]}}. \quad (5.9)$$

See for example, Eq. (5.3) and (5.4). The question is how to evaluate such integrals. These are infinite-dimensional integrals over the field configurations. Even worse, the integrands are strongly oscillating (note the imaginary unit in the exponentials), and so they behave very poorly under numerical integration. In the approach of perturbative QCD, such path integrals can be expanded in powers of the coupling $\alpha_s = g^2/4\pi$. However, this only works in the very high energy perturbative regime and even there, the convergence is slow.

Recall that in the canonical ensemble of classical statistical mechanics, the equilibrium expectation value of an observable O is

$$\langle O \rangle = \frac{\int \mathcal{D}U O(U) e^{-\beta E(U)}}{\int \mathcal{D}U e^{-\beta E(U)}}, \quad (5.10)$$

where the integral is over all possible microstates U of the system, $E(U)$ is the energy of the microstate, $O(U)$ is the value of the observable on that microstate, and β is the inverse temperature. See Eq. (3.3). In lattice QCD, we exploit the similarity between Eq. (5.9) and (5.10) to reinterpret the QCD expectation value as an expectation value in classical statistical physics. That is, we plan to evaluate objects like Eq. (5.9) with some kind of Monte Carlo integration. However, first we need to put QCD on a lattice.

Defining QCD on a finite lattice has several important consequences:

1. The infinite-dimensional integrals become finite-dimensional integrals

2. The lattice serves as a regularizer—removing the divergences from the theory. The finite lattice spacing removes the UV divergences by limiting the momenta, and the finite lattice volume removes the IR divergences
3. To connect to the physical world, one must carefully take the continuum (lattice spacing $\rightarrow 0$) and thermodynamic (volume $\rightarrow \infty$) limits of any observable calculated on the lattice
4. Much like the continuum theory, the lattice theory renormalizes the parameters. To connect lattice results with the physical world, one must find a way to relate a lattice observable to a real world quantity—a process called “setting the scale”

For later convenience, we rescale the gluon fields as

$$A_\mu(x) \rightarrow \frac{1}{g} A_\mu(x). \quad (5.11)$$

Then the QCD action becomes

$$S = \int d^4x \sum_f \bar{\psi}^f (i\gamma_\mu(\partial^\mu + iA^\mu) - m_f) \psi^f - \int d^4x \frac{1}{4g^2} F_{\mu\nu}^a F_a^{\mu\nu}, \quad (5.12)$$

with $F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a - f_{abc} A_\mu^b A_\nu^c$. Now the coupling g appears only as a factor in front of the gauge term. From here on, we neglect the color indices i, j that appear in Eq. (5.1).

One key difference between Eq. (5.9) and (5.10) is the complex exponential in the first. To reinterpret the QCD expectation value as a classical expectation value, we need to perform a Wick rotation to imaginary time ($t \rightarrow it$). This yields the Euclidean QCD action

$$S_E = \underbrace{\int d^4x \sum_f \bar{\psi}_f (\gamma_\mu(\partial_\mu + iA_\mu) + m_f) \psi_f}_{S_F} + \underbrace{\int d^4x \frac{1}{4g^2} F_{\mu\nu}^a F_{\mu\nu}^a}_{S_G}, \quad (5.13)$$

where S_F is the fermion part of the action and S_G is the gauge part of the action. Note that the lattice works as a regulator whether it's Euclidean or Minkowski. The reason to go to a Euclidean lattice is so that $e^{iS} \rightarrow e^{-S_E}$ becomes real and can be interpreted as a probability weight. Perhaps confusingly, the gamma matrices γ_μ in this expression are now the Euclidean




Figure 5.2 The variable $U_{\mu}(n)$ is associated with the link pointing in the direction $\hat{\mu}$ from the lattice site n to the site $n + \hat{\mu}$.

gamma matrices which obey the anti-commutation relations $\{\gamma_u, \gamma_v\} = 2\delta_{uv}\mathbb{I}$ instead of the usual Minkowski gamma matrices which obey the anti-commutation relations $\{\gamma_u, \gamma_v\} = 2g_{uv}\mathbb{I}$ where g_{uv} is the Minkowski metric tensor $g_{uv} = \text{diag}(1, -1, -1, -1)$.

We want to discretize Eq. (5.13) for a lattice. We construct a 4-dimensional Euclidean lattice with sites labeled by $n = (n_1, n_2, n_3, n_4)$ and separated by lattice spacing a . We define the quark fields on the lattice sites

$$\psi(x) \longrightarrow \psi(n). \quad (5.14)$$

We can use a symmetric finite-difference approximation of the derivative

$$\partial_{\mu}\psi(n) \longrightarrow \frac{\psi(n + \hat{\mu}) - \psi(n - \hat{\mu})}{2a}. \quad (5.15)$$

Thus for a single free fermion (i.e. no gauge fields), we obtain the discretized fermion action

$$S_F^0 = a^4 \sum_n \bar{\psi}(n) \left(\sum_{\mu=1}^4 \gamma_{\mu} \frac{\psi(n + \hat{\mu}) - \psi(n - \hat{\mu})}{2a} + m\psi(n) \right). \quad (5.16)$$

In the continuum, gauge invariance of the free fermion action requires the introduction of a direction-dependent gauge field

$$A_{\mu}(x) = \sum_{i=1}^8 A_{\mu}^{(i)}(x) T_i \in su(3), \quad (5.17)$$

where the T_i are generators of the Lie group $SU(3)$, and the $A_{\mu}(x)$ are elements of the Lie algebra $su(3)$. On the lattice, one plays a similar game to enforce gauge invariance, but now it turns out that one needs a field of direction-dependent objects which are elements $U_{\mu}(n)$ of the $SU(3)$ Lie group instead of the algebra. They are related to the original fields $A_{\mu}(x)$ as

$$U_{\mu}(n) = e^{iaA_{\mu}(n)} \in SU(3). \quad (5.18)$$




Figure 5.3 In lattice QCD, the quarks fields are defined on the lattice sites and the gauge fields are defined on the links between the sites. Gauge invariant objects can be constructed from closed loops and from paths going from one quark field to another. The ordered product of link variables along an elementary square is called a *plaquette*. The image on the right is originally from JICFus.

By convention, we associate these direction-dependent objects with the “links” between the lattice sites as illustrated in Figure 5.2.

This leads us to the so-called “naive” fermion action for a single flavor of fermion

$$S_F = a^4 \sum_n \bar{\psi}(n) \left[\sum_{\mu=1}^4 \gamma_\mu \frac{U_\mu(n)\psi(n + \hat{\mu}) - U_{-\mu}(n)\psi(n - \hat{\mu})}{2a} + m\psi(n) \right]. \quad (5.19)$$

Although this lattice fermion action has the right continuum limit, it generates unwanted additional fermions. This “doubling problem” can be resolved by adding an additional term to obtain the “Wilson” fermion action. This is beyond the scope of this thesis, but detailed discussions can be found in any textbook on LQCD.

For the lattice gauge action, our first priority is to construct something that reduces to the continuum gauge action in the continuum limit $a \rightarrow 0$. To do this, we need gauge-invariant objects. It turns out that taking the ordered product of link variables along a closed loop and then tracing over the final matrix results in a gauge-invariant object. The ordered product of link variables around the simplest closed loop, called the *plaquette* (see Figure 5.3), is

$$\square_{\mu\nu}(n) = U_\mu(n)U_\nu(n + \hat{\mu})U_\mu^\dagger(n + \hat{\mu} + \hat{\nu})U_\nu^\dagger(n). \quad (5.20)$$

The trace of plaquettes is one of the simplest gauge-invariant objects on the lattice, and one can use them to construct the Wilson gauge action

$$S_G = \frac{2}{g^2} \sum_n \sum_{\mu < \nu} \text{Re Tr} \left(\mathbb{1} - \square_{\mu\nu}(n) \right). \quad (5.21)$$

It can be shown that in the continuum limit, this recovers the continuum gauge action, e.g. Ref. [66].

With Eq. (5.19) and (5.21) we now have a full albeit rudimentary lattice definition of QCD. Next we shall see what we can do with this.

5.4 MCMC

As before, the main objects are correlators i.e. Green's functions as in Eq. (5.9). But now, instead of doing a perturbative expansion, Feynman diagrams, and loop integrals, we reinterpret an equilibrium QFT as a classical statistical physics problem. That is, we estimate the *Euclidean* path integrals (obtained by Wick rotation $iS \rightarrow -S_E$)

$$\langle O \rangle = \frac{\int \mathcal{D}U \mathcal{D}\bar{\psi} \mathcal{D}\psi O[U, \bar{\psi}, \psi] e^{-S_E[U, \bar{\psi}, \psi]}}{\int \mathcal{D}U \mathcal{D}\bar{\psi} \mathcal{D}\psi e^{-S_E[U, \bar{\psi}, \psi]}}. \quad (5.22)$$

via Markov chain Monte Carlo (MCMC). At this stage, all quantities are defined in terms of a discrete spacetime lattice.

5.4.1 MCMC with Pure Gauge

We start with the simpler example of “pure gauge”, in which the theory consists only of gluons—no quarks. The vacuum expectation value of observables is

$$\langle O \rangle = \frac{\int \mathcal{D}U O[U] e^{-S_G[U]}}{\int \mathcal{D}U e^{-S_G[U]}}. \quad (5.23)$$

Compare this with the classical path integral Eq. (5.10). We can interpret Eq. (5.23) as the classical expectation value of $O[U]$ where the U are importance-sampled with Boltzmann factor $e^{-S_G[U]}$. That is, the action $S_G[U]$ plays the role of the classical energy quantity $\beta E(U)$. This reinterpretation is only possible because the Wick rotation converted the complex exponential into a real exponential which can be interpreted as a probability density. Using the approach detailed in Section 3.1 for classical spin systems, we can estimate Eq. (5.23) with a Monte Carlo simulation

$$\langle O \rangle \approx \frac{1}{N} \sum_{U_n} O[U_n], \quad (5.24)$$

provided a sample of N gauge configurations U_n , distributed with probability $\propto e^{-S_G[U]}$. The statistical error of the mean is

$$\propto \frac{1}{\sqrt{N}}. \quad (5.25)$$

As we did for classical spin systems in Section 3.1, we sample lattice configurations U by using a Markov chain. Now instead of configurations having e.g. a 2-component spin at each lattice site, we have a complex 3×3 $SU(3)$ matrix at each lattice *link*.

To get $\{U_n\}$ distributed with probability $\propto e^{-S_G[U]}$, we want a Markov process (aka updating algorithm) that starts with an arbitrary configuration U_0 and stochastically generates more

$$U_0 \xrightarrow{\text{update}} U_1 \xrightarrow{\text{update}} U_2 \xrightarrow{\text{update}} \dots \quad (5.26)$$

To sample the desired distribution, it is sufficient to show that for every pair of configurations U and U' , the Markov process satisfies the *detailed balance condition* [60, pp 289]

$$e^{-S_G[U]} P(U \rightarrow U') = e^{-S_G[U']} P(U' \rightarrow U), \quad (5.27)$$

where $P(U \rightarrow U')$ is the probability that the process generates U' given U . The process must also be *ergodic* i.e. it must be able to access the whole space of configurations.

The prototypical Markov process is the Metropolis algorithm:

1. Given current state U , generate a new candidate U' via some random process that selects uniformly on the group
2. Accept U' as the new configuration with probability

$$P_A(U \rightarrow U') = \min \left(1, e^{S_G[U'] - S_G[U]} \right). \quad (5.28)$$

3. Repeat

Alternatives include the heatbath algorithm [168–171] in which Steps 1 and 2 are combined in an optimized way. The implementation depends on the details of the gauge group. Another approach is that of hybrid Monte Carlo (HMC) [172], which uses molecular dynamics evolution, and ensures detailed balance by a Metropolis accept/reject step at the end. The HMC algorithm is the standard gauge configuration updating algorithm for fermionic lattice QFTs.

5.4.2 MCMC with Fermions

To include fermions in our lattice MCMC, we return to the full (Euclidean) path integral

$$\langle O \rangle = \frac{\int \mathcal{D}U \mathcal{D}\bar{\psi} \mathcal{D}\psi O[U, \bar{\psi}, \psi] e^{-S_E[U, \bar{\psi}, \psi]}}{\int \mathcal{D}U \mathcal{D}\bar{\psi} \mathcal{D}\psi e^{-S_E[U, \bar{\psi}, \psi]}}. \quad (5.29)$$

The Euclidean action is composed of the gauge action and fermionic action, and it can be written in the form

$$S_E[U, \bar{\psi}, \psi] = S_G[U] + S_F[U, \bar{\psi}, \psi] = S_G[U] + \bar{\psi} M[U] \psi \quad (5.30)$$

where the “Dirac operator” $M[U]$ depends on the gauge fields but not the quark fields $\bar{\psi}$ and ψ . The fermionic variables ψ and $\bar{\psi}$ are not ordinary numbers or matrices like we are used to dealing with in classical MCMC. Rather, they are non-commuting Grassman variables. These must be integrated out in some way before we can proceed with MCMC.

We start with the *Mathews-Salam formula* [173, 174]. Given a Grassman algebra with $2N$ generators, $\theta_i, \bar{\theta}_i$ for $i = 1, 2, \dots, N$ and some complex $N \times N$ matrix Q , it can be shown that

$$\int \mathcal{D}\Theta \mathcal{D}\bar{\Theta} e^{\bar{\Theta} Q \Theta} \equiv \left(\prod_i \int d\theta_i d\bar{\theta}_i \right) e^{\sum_{i,j} \bar{\theta}_i Q_{ij} \theta_j} = \det[Q]. \quad (5.31)$$

We can use this to rewrite the denominator of Eq. (5.29) as

$$Z = \int \mathcal{D}U \mathcal{D}\bar{\psi} \mathcal{D}\psi e^{-S_E[U, \bar{\psi}, \psi]} = \int \mathcal{D}U e^{-S_G[U]} \det M[U]. \quad (5.32)$$

Unless the observable O is independent of the fermion fields $\bar{\psi}$ and ψ , the same result cannot be used to simplify the numerator of Eq. (5.29). However, if O is bilinear in the fermion fields, then one can use a generalization of Eq. (5.31). The end result (see LQCD textbooks such as [66] for the details) is that Eq. (5.29) can be written as

$$\langle O \rangle = \frac{1}{Z} \int \mathcal{D}U O[U, M^{-1}[U]] e^{-S_G[U]} \det M[U], \quad (5.33)$$

where Z is given by Eq. (5.32).

The Grassman variables are now gone from Eq. (5.33), so provided that $\det M[U]$ is real and nonnegative, then we can once again interpret $\langle O \rangle$ as a classical path integral and use MCMC to

estimate the observable O . Now, we sample the configurations U from the probability distribution

$$P(U) = \frac{1}{Z} e^{-S_G[U]} \det M[U]. \quad (5.34)$$

Then as before,

$$\langle O \rangle = \int \mathcal{D}U P(U) O[U, M^{-1}[U]] \approx \frac{1}{N} \sum_{U_n} O[U_n]. \quad (5.35)$$

We cannot generally ensure that $\det M[U]$ is real and nonnegative. However, we can for the case of two mass-degenerate quark flavors. Fortunately, QCD with two mass-degenerate dynamical light quarks is a pretty good approximation of full QCD, so we will use that example here. For two degenerate flavors, Eq. (5.34) becomes

$$P(U) = \frac{1}{Z} e^{-S_G[U]} \det M_u[U] \det M_d[U] = \frac{1}{Z} e^{-S_G[U]} \det [MM^\dagger], \quad (5.36)$$

where M_u refers to the up quark and M_d refers to the down quark. Here, we combined the two quark determinants since for degenerate flavors, they are the same. Thus, we are assured that $\det[MM^\dagger]$ is real and nonnegative and that $P(U)$ can be interpreted as a probability distribution function.

Note that $M[U]$ is a huge matrix of size proportional to $V \times V$ where V is the volume of the lattice. So for even a moderately sized lattice with $V = 32^3 \times 64$, $M[U]$ has millions of rows and millions of columns. This makes the direct calculation of $\det M$ infeasible. Also, note that $M[U]$ is a functional of the lattice configuration U , and so it has to be recomputed for each lattice configuration. It turns out that one can express the fermion determinant as a Gaussian integral of a bosonic “pseudofermion” field (see LQCD textbooks e.g. [66] for details)

$$\det [MM^\dagger] \propto \int \mathcal{D}\phi^\dagger \mathcal{D}\phi e^{-\phi^\dagger (MM^\dagger)^{-1} \phi}, \quad (5.37)$$

where $\phi^\dagger$ and ϕ are bosonic fields. We have effectively traded the determinant $\det[MM^\dagger]$ for the inverse $(MM^\dagger)^{-1}$. The inverse $(MM^\dagger)^{-1}$ is a very large matrix of size proportional to $V \times V$ and highly nonlocal, however, now there is the possibility of using iterative methods to approximate V -length vectors of the form $(MM^\dagger)^{-1} \phi$. With this, our path integral becomes

$$\langle O \rangle = \int \mathcal{D}U \mathcal{D}\phi^\dagger \mathcal{D}\phi P(U) O[U, M^{-1}] \quad \text{with} \quad P(U) = \frac{1}{Z} e^{-S_G[U] - \phi^\dagger (MM^\dagger)^{-1} \phi}. \quad (5.38)$$

In principle, we could use the Metropolis algorithm as before. However, Metropolis is a local algorithm and applying it to a problem like this would be abysmally inefficient. The solution is to use a global updating algorithm like Hybrid Monte Carlo (HMC) or Rational Hybrid Monte Carlo (RHMC) [175, 176].

5.4.3 The Hybrid Monte Carlo (HMC) Algorithm

For simplicity, we consider a real scalar field Q . The vacuum expectation value of an observable O is given by

$$\langle O \rangle = \frac{\int \mathcal{D}Q O[Q] e^{-S[Q]}}{\int \mathcal{D}Q e^{-S[Q]}}. \quad (5.39)$$

We introduce a field P of momenta conjugate to the field variables Q . Since the observable O does not depend on P , we can just as well compute $\langle O \rangle$ as

$$\langle O \rangle = \frac{\int \mathcal{D}Q \mathcal{D}P O[Q] e^{-\frac{1}{2}P^2 - S[Q]}}{\int \mathcal{D}Q \mathcal{D}P e^{-\frac{1}{2}P^2 - S[Q]}}. \quad (5.40)$$

Note that Eq. (5.40) is equivalent to Eq. (5.39) since the Gaussian integrals over P cancel out in the numerator and denominator. However, Eq. (5.40) can be interpreted as the vacuum expectation value of an observable $O[Q]$ in a classical and nonrelativistic microcanonical system with Hamiltonian

$$H[Q, P] = \frac{1}{2}P^2 + S[Q]. \quad (5.41)$$

For such a system, the classical equations of motion (with respect to a fictitious “computer” time) are

$$\dot{P} = -\frac{\partial H}{\partial Q} = -\frac{\partial S}{\partial Q}, \quad \dot{Q} = \frac{\partial H}{\partial P} = P. \quad (5.42)$$

Thus, given a configuration Q , we can invent some conjugate field P and then evolve the two forward in fictitious time by integrating the “molecular dynamics equations” Eq. (5.42) to get a new configuration Q' . Doing this repeatedly results in an ensemble of configurations, upon which one can estimate $\langle O \rangle$.

The evolution of Eq. (5.42) is along a surface of constant energy. To obtain configurations at different energy, the conjugate field P is periodically refreshed with random values. In practice, the integration of Eq. (5.42) cannot be done exactly due to discretization errors. To correct for this,

a Metropolis accept-reject step is performed at the end of each evolution “trajectory”—accepting the new configuration with a probability that depends on the change of the total Boltzman factor.

5.4.4 Example: Hadron Spectroscopy

Before we conclude this brief overview of lattice QCD, it will be helpful to consider an example of the kind of ab initio quantities that can be computed in LQCD. We sketch a simple example here—measuring the mass of a pion. Specifically, we will focus on the π^+ meson. We follow closely Ref. [66].

We start by identifying lattice interpolating operators O and $\bar{O}$ for the pions such that the corresponding Hilbert space operators $\hat{O}$ and $\hat{O}^\dagger$ annihilate and create the pions. These operators are functionals of the lattice fields, and have the quantum numbers of the pion. We can use

$$O_{\pi^+}(n) = \bar{d}(n)\gamma_5 u(n), \quad \bar{O}_{\pi^+}(m) = \bar{u}(m)\gamma_5 d(m), \quad (5.43)$$

where u , $\bar{u}$, d , and $\bar{d}$ are the *up* and *down* quark spinors.

Next, we construct 2-point correlators from these interpolating operators. For example, the π^+ correlator

$$\langle O_{\pi^+}(n)\bar{O}_{\pi^+}(m) \rangle = \langle \bar{d}(n)\gamma_5 u(n)\bar{u}(m)\gamma_5 d(m) \rangle. \quad (5.44)$$

The fermionic part of the expectation value can be computed in closed form. After simplifying and making use of Wick’s theorem, it can be shown that (see any LQCD textbook for details)

$$\langle O_{\pi^+}(n)\bar{O}_{\pi^+}(m) \rangle_F = -\text{tr} [\gamma_5 M_u^{-1}(n, m)\gamma_5 M_d^{-1}(m, n)], \quad (5.45)$$

where M^{-1} is the inverse of the Dirac operator, and can be understood as a quark propagator. The object on the right-hand side is something that can be measured on the lattice. However, given that M^{-1} is a highly nonlocal matrix of size proportional to $V \times V$ where V is the volume of the lattice, measuring this quantity is a highly nontrivial endeavor. It is done by inserting “quark sources” and then applying iterative methods such as the conjugate gradient algorithm. The subscript F on the left indicates that this is only the fermionic part of the expectation value. The gauge part remains.

The full expectation value is the path integral

$$\left\langle O_{\pi^+}(n) \overline{O}_{\pi^+}(m) \right\rangle = -\frac{1}{Z} \int \mathcal{D}U e^{-S_G[U]} \det[M_u] \det[M_d] \text{tr} \left[\gamma_5 M_u^{-1}(n, m) \gamma_5 M_d^{-1}(m, n) \right], \quad (5.46)$$

where

$$Z = \int \mathcal{D}U e^{-S_G[U]} \det[M_u] \det[M_d]. \quad (5.47)$$

As before, we assume mass-degenerate quarks, and we express the quark determinants in terms of pseudofermion fields. Then,

$$\left\langle O_{\pi^+}(n) \overline{O}_{\pi^+}(m) \right\rangle = -\frac{1}{Z} \int \mathcal{D}U \mathcal{D}\phi^\dagger \mathcal{D}\phi e^{-S_G[U] - \phi^\dagger (MM^\dagger)^{-1} \phi} \text{tr} \left[\gamma_5 M^{-1}(n, m) \gamma_5 M^{-1}(m, n) \right], \quad (5.48)$$

where

$$Z = \int \mathcal{D}U \mathcal{D}\phi^\dagger \mathcal{D}\phi e^{-S_G[U] - \phi^\dagger (MM^\dagger)^{-1} \phi}. \quad (5.49)$$

On the other hand, after appropriate momentum projection, one can write such correlators as a spectral decomposition

$$\left\langle O_{\pi^+}(\vec{0}, n_t) \overline{O}_{\pi^+}(\vec{0}, 0) \right\rangle = \sum_k \langle 0 | \hat{O} | k \rangle \langle k | \hat{O}^\dagger | 0 \rangle e^{-n_t E_k} = A_0 e^{-n_t E_0} + A_1 e^{-n_t E_1} + \dots. \quad (5.50)$$

Since the lowest energy E_0 corresponds to the rest mass of the particle, this gives us a way to connect the pion correlator to the mass of the pion.

So to compute the π^+ pion mass using LQCD, one would proceed as follows:

1. Generate, using e.g. the HMC algorithm, an ensemble of N de-correlated gauge configurations U using the distribution weight

$$P(U) = \frac{1}{Z} e^{-S_G[U] - \phi^\dagger (MM^\dagger)^{-1} \phi}. \quad (5.51)$$

2. On each of these configurations, measure the pion correlator

$$C(n_t) = -\text{tr} \left[\gamma_5 M^{-1}(\vec{0}, n_t, \vec{0}, 0) \gamma_5 M^{-1}(\vec{0}, 0, \vec{0}, n_t) \right] \quad (5.52)$$

for different n_t

3. Extract the mass m from fits to exponential decays

$$C(n_t) \approx A_0 e^{-n_t m} \quad (5.53)$$

4. Average the measurements over the ensemble of configurations (Eq. (5.24)) to obtain the best estimate
5. Repeat for different lattice sizes and extrapolate to zero lattice spacing and infinite lattice volume

5.4.5 Challenges

As we've seen in the preceding section, for lattice QCD with dynamical fermions, the distribution weight for the gauge fields is proportional to determinants of huge matrices. Similarly, the calculation of correlators requires the inversion of these huge matrices. Because of their size, neither the determinants nor the inverses of these matrices can be calculated directly, and so the state-of-the-art in lattice QCD is algorithms which allow us to estimate these quantities indirectly. Even the most basic calculations are difficult and necessitate the use of the latest technologies (GPUs, TPUs, etc.) and high performance algorithms such as multigrid methods.

Conventional lattice QCD works because by Wick-rotating the path integrals to imaginary time, we are able to reinterpret those path integrals as classical partition functions of a system in a heat bath at some particular temperature and thereby employ our extensive classical MCMC methods. However, this means we are always studying a system in equilibrium. The study of real-time phenomena such as evolution and decay dynamics is completely out of reach with this approach. To study real-time lattice QCD, we will need completely new approaches such as quantum simulation.

A similar problem occurs with nonzero chemical potential. In the vacuum, the chemical potential (i.e. the net number of baryons or quarks) is zero. However, this is not true in matter, and especially not in extreme situations such as in heavy-ion collisions or in ultra-dense matter such as neutron stars. In conventional lattice QCD, calculations are done at zero (or very nearly zero)

chemical potential and for good reason. It turns out that introducing a nonzero chemical potential results in a complex fermion determinant, which means we can no longer interpret the path integral as a classical partition function.

5.5 Quantum Simulation of Lattice Quantum Field Theories

5.5.1 Introduction

In a quantum many-body system, the Hilbert space grows exponentially with the system size. The simulation of such systems in the strongly-coupled regime is an extremely hard problem for classical computers. Even when we simplify to equilibrium systems in Euclidean time, as we do in lattice QCD, our conventional methods suffer from sign problems. Real-time dynamics are laughably out of reach. But what if we could construct a simple quantum system that emulates a quantum field theory in some way? By using quantum objects instead of classical bits, the exponential growth of the Hilbert space with the system size would match that of the QFT. Furthermore, nothing would prevent us from evolving the simple system forward in time to gain insight into the dynamics of the QFT.

The idea of quantum simulation seems to have originated with Richard Feynman [177]. He gave a talk in 1982 in which he noted that phenomena in quantum field theory, assuming a theory in which spacetime is discrete, can be imitated by phenomena in solid state physics. In a way, a lattice of atoms can imitate quantum field theory. He postulated, therefore, that one could build quantum simulators in the lab which could tell us things about quantum field theory:

“I[. . .] believe it’s true that with a suitable class of quantum machines you could imitate any quantum system, including the physical world.”

In this section, we consider the quantum simulation of quantum field theories. That is, we consider quantum systems that can be made to emulate a quantum field theory in some manner such that a measurement of the quantum system gives us information about the quantum field theory. There are really two distinct classes of quantum simulators. An *analog* quantum simulator is a special-purpose device—built and tuned specifically to emulate the Hamiltonian of a specific

quantum field theory. In an analog simulation, time is not discretized—the simulator evolves continuously in real time. This may be accomplished by setting up and tuning the device, and then letting it evolve naturally on its own. An example would be cold atoms hopping around on an optical lattice. In a *digital* quantum simulator, the time evolution proceeds stroboscopically typically via the successive application of quantum logic gates. The simulator could be a special-purpose device or it could be a universal quantum computer. The drawback of digital simulations is a need for error correction which means that useful results may be decades in the future.

In lattice QCD, spacetime is discretized and one typically works in the Lagrangian formalism. In both digital and analog quantum simulations, only space is discretized at the beginning, and one works in the Hamiltonian formalism. In a digital simulation, time is eventually discretized and time evolution proceeds stroboscopically. In an analog simulation, time never gets discretized. Rather, the system is prepared in some initial state and then allowed to evolve in real time. General reviews of both analog and digital quantum simulation include [79, 80, 178, 179], and there is even a recent textbook on the subject [180].

5.5.2 Analog Simulation

In lieu of the preferred analytical solution, the modern scientist often resorts to computational simulation. But even that can be so difficult that one must occasionally resort to physical models. In general, the simulation of natural phenomena, such as fluid dynamics, is computationally difficult. So difficult that scale models and wind tunnels are still used to study aerodynamic forces on vehicles. The validity of this kind of analog simulation rests on there being some kind of isomorphism between the problem of interest (e.g. vehicle aerodynamics) and the physical model (e.g. scale model in a wind tunnel). In the case of a scale model, the similarity is obvious. But in analog simulation, the similarity is not always obvious. Consider for example that electric circuits have been used to model neural systems [181] and that the flow of water in a set of tanks and pipes has been used to model the economy of the United Kingdom [182]. The general idea of analog simulation is to map a complex system of interest to a simpler physical system which can be built and studied in the lab. For example, the analog simulation of black holes using Bose-Einstein

condensates has been proposed [183, 184]. In our case, we want to map complicated systems like lattice quantum field theories to simpler quantum-mechanical systems that we can build and study in the lab.

In an analog quantum simulation, a Hamiltonian of interest (e.g. of a gauge theory) is mapped to a simpler and controllable quantum system (e.g. an optical lattice). After setup and initialization, the simple system is evolved forward in time. At the end, some quantity is measured on the simple system and interpreted in terms of the system of interest. Analog simulation of gauge theories is promising because it would bypass the sign problems inherent in conventional lattice QCD and allow us to study the real-time dynamics of the theory [86].

Analog quantum simulators are purpose-built devices—they are not universal quantum computers. They are allowed to evolve in real time as opposed to digital simulators which are stroboscopically evolved in computer time by the application of quantum gates. By mapping the degrees of freedom onto continuous interactions, an analog simulator requires fewer qubits than a digital algorithm. This makes the analog simulator potentially more efficient than digital simulators in the current noisy intermediate-scale quantum (NISQ) era [178]. Given their purpose-built nature, analog simulators are limited to simpler interactions, however, they are also more easily scalable to large system size [185]. Finally, with analog simulators, we may be able to perform qualitative experiments to determine, e.g. if a given model has a phase transition, even before we are able to perform precision measurements.

Schematically, the general approach for an analog simulation is as follows:

1. Map the Hamiltonian of the system of interest H_{sys} directly onto the Hamiltonian H_{sim} of a simpler and controllable simulator. This assumes that you can find a simulator and a mapping for which this is possible. This may turn out to be too difficult, so instead of studying H_{sys} , you may need to study a simpler effective many-body model of the system of interest. Often the mapping will require some cleverness [186]
2. Prepare the state to be measured. In the simplest case, this simply means to let the simulator

settle into equilibrium

3. Measure the state. In the simplest case, measuring the simulator in its prepared state will yield information about the system of interest
4. Repeat the previous two steps many times to get a statistical average of the observables of interest

The most common technologies proposed for the analog quantum simulation of quantum field theories include optical lattices, trapped ion chains, and superconducting circuits. In an optical lattice, overlapping laser beams are used to create interference patterns, which can be used to trap atoms via the optical dipole force in the regions where destructive interference occurs. Atoms can hop from site to site and atoms that meet at the same site interact with each other. Various lattice geometries can be realized by superimposing laser fields at different angles. The optical lattice can be tuned by adjusting intensity, frequency, and phase of the lasers. The atoms used in an optical lattice would typically be neutral alkaline atoms since they have only a single valence electron—simplifying things. An individual atom in the lattice can be addressed and manipulated, for example, flipping its spin via a focused laser beam. In the end, the lattice configuration can be measured, for example, by fluorescence. For excellent review articles that include information on analog simulations with optical lattices, see [80, 86, 87, 185–190]. In a trapped ion chain, cold ions are contained in electromagnetic traps. A qubit can be encoded in two internal states of the ion and manipulated by laser pulses and measured via fluorescence. For excellent review articles that include information on using trapped ions to perform analog quantum simulations of LGTs, see [80, 84, 186, 187, 189, 191]. In superconducting circuits, qubits are constructed from circuit elements and Josephson junctions. Qubits can be coupled via resonators and manipulated using microwave radiation. Superconducting circuits are very tunable and seem promising due to their apparent scalability—they can be fabricated using the existing technology of integrated circuits. For excellent review articles that include information on using superconducting circuits to perform analog quantum simulations of LGTs, see [80, 186, 187, 189, 192].

Analog quantum simulation has been around for a while in the condensed matter field. In [193], it was shown that an ultracold dilute gas of bosonic atoms in an optical lattice could be described by a Bose-Hubbard model. This model describes “the hopping of bosonic atoms between the lowest vibrational states of the optical lattice sites”. The Bose-Hubbard Hamiltonian is

$$H = -J \sum_{\langle i,j \rangle} b_i^\dagger b_j + \sum_i \epsilon_i \hat{n}_i + \frac{1}{2} U \sum_i \hat{n}_i (\hat{n}_i - 1), \quad (5.54)$$

where b_i and $b_i^\dagger$ are the annihilation and creation operators, and the number operator $\hat{n}_i = b_i^\dagger b_i$ counts the number of bosonic atoms at lattice site i . The parameter U gives the strength of the repulsion of two atoms on a lattice site, J is the hopping matrix element between adjacent sites, and ϵ_i describes the energy offset of each lattice site. In [194], this was experimentally implemented by loading Rb-87 atoms into a three-dimensional optical lattice. The optical lattice potential had the form

$$V(x, y, z) = V_0 [\sin^2(kx) + \sin^2(ky) + \sin^2(kz)]. \quad (5.55)$$

In each individual well (or trap), i.e. at each lattice site, the confining potential for an atom can be approximated by a harmonic potential. In their setup, there are ~ 65 lattice sites in each dimension with an average atom number of up to 2.5 atoms per lattice site near the center of the 3D lattice. The researchers were able to observe and study the quantum phase transition from superfluid to Mott insulator. Simulating the Bose-Hubbard model on an optical lattice is straightforward since the Bose-Hubbard Hamiltonian can be directly related to the Hamiltonian of interacting atoms on a lattice. However, for other models, it is generally not that straightforward, and work must be done to relate a Hamiltonian of interest to a Hamiltonian that is simple enough to simulate on an optical lattice but not so simple that it becomes a useless analogue.

The quantum simulation of gauge theories on analog simulators is significantly more challenging than for condensed matter systems such as the Bose-Hubbard model. This is largely because the gauge theories of high energy physics realize certain symmetries that are usually not naturally realized in an analog simulator. In order to simulate a gauge theory, there are two particularly

problematic requirements; simulator must be able to include both fermionic and bosonic degrees of freedom, and it must have local gauge invariance [86].

Because the gauge theories of interest to us contain both fermionic and gauge fields, optical lattices seem attractive because they can hold both fermionic and bosonic atoms. However, there remains a complication. In Wilson’s formulation of lattice gauge theory, the gauge fields are defined on the links of the lattice. These continuous gauge degrees of freedom give rise to an infinite-dimensional Hilbert space at each link, and it is not obvious how to represent such objects on an analog simulator. One solution is to truncate the link Hilbert space in some way such that the model with truncated gauge fields still approximates the full theory. For example, in a theory with $U(1)$ gauge fields, the fields may be approximated with $\mathbb{Z}_q$ [195]. To optimize such a $\mathbb{Z}_q$ approximation, it is useful to have a continuous family of models that interpolate among the various possibilities. Doing this was one of the original motivations for our study of the Extended- $O(2)$ model [127–129], which resulted in this thesis. Similarly, $SU(N)$ gauge fields may be “digitized” in some way to represent them with discrete subgroups. This has been studied for both $SU(2)$ [196] and $SU(3)$ [197–199].

Several strategies have been proposed to include local gauge invariance in an analog simulation. One possibility is to add a penalty term to the Hamiltonian that suppresses symmetry-violating transitions. This is straightforward for Abelian models, but not for non-Abelian models. Another option is to exploit microscopic intrinsic symmetries of the atomic matter and combine with spatially arranged potentials to engineer gauge invariance up to high energy. See reviews such as [187, 189].

5.5.3 Digital Simulation

For a quantum simulation to be *digital*, the only requirement is that the time evolution proceeds stroboscopically. In this sense, digital quantum simulation can be done on special-purpose machines such as optical lattices of Rydberg atoms or on trapped ion chains. Typically, however, when referring to “digital” simulation, a narrower definition is being used in that the simulation is being performed on a universal gate-based quantum computer. This is a more specific kind of quantum device in which the time evolution is discretized and implemented by the application of quantum

logic gates, for example [200–207]. This is the approach that we focus on here. For review articles on digital quantum simulation, see [80, 87, 186, 187, 208].

In general, the digital simulation of a quantum field theory involves three steps:

1. *Initial state preparation*: The first step is to initialize the set of qubits to the state $|\psi(0)\rangle$. This may be challenging, and finding an efficient algorithm may not be possible. If $|\psi(0)\rangle$ is the ground state of the theory, one approach called “adiabatic state preparation” is to start the system in the ground state of the simpler non-interacting theory and then to adiabatically turn on the interactions. The author and others recently studied a variety of state preparation methods for the Schwinger model [209, 210].

2. *Unitary evolution*: The initial state is time-evolved to the final state

$$|\psi(t)\rangle = U(t) |\psi(0)\rangle, \quad U(t) = e^{-i\hbar H t},$$

by applying the unitary operator U to the initial state. The challenge here is to write U as a sequence of quantum gate operations. This may be achieved by writing H in terms of Pauli matrices, which we know how to implement via quantum gates. Then by discretizing time t in a process known as *Trotterization*, one may approximate U as a product of quantum gate operations

$$U(t) \approx \prod_n e^{-i\hbar H_n \delta t}.$$

More generally, any unitary operation can be quantum simulated—the challenge is to do it efficiently.

3. *Measurement*: Once the system is in the desired final state, it is measured. One possibility is to use *quantum state tomography*. However, this approach requires a lot of resources, and more direct alternatives are preferred.

In practice, the three steps above give a single data point since the act of measuring a quantum state will also destroy it. Therefore, the three steps are repeated many times to obtain a statistical mean and error bar.

The main challenge in implementing digital quantum simulations is the same as the main challenge facing universal quantum computation—error correction and its need for an order of magnitude more qubits. In the NISQ era, there may be little that we can accomplish unless we can significantly reduce the need for error correction in our digital simulations. Another challenge is how to represent fermionic degrees of freedom on a gate-based quantum computer. One approach is the Jordan-Wigner transformation [211] in which fermionic modes are mapped to strings of Pauli matrices [212,213]. However, the Jordan-Wigner transformation is a non-local mapping and is very inefficient in more than $1 + 1$ dimensions. An alternative is the Bravyi-Kitaev mapping [214,215], which reduces the non-locality that is induced. As in analog simulation, there is also the challenge of discretizing and truncating the infinite-dimensional Hilbert space of the gauge links. Finally, gauge invariance, which imposes constraints relating the matter and gauge degrees of freedom at each lattice site, poses a complication. If these constraints are not satisfied, then one is simulating a theory that contains many undesired non-physical states. One way to solve two challenges at once is to use these gauge constraints to integrate out the fermionic or gauge degrees of freedom. This approach may be used to eliminate the fermionic matter in any dimension or it may be used to eliminate the gauge fields if one is in a single spatial dimension [216,217].

CHAPTER 6

SUMMARY

Lattice QCD, which is a nonperturbative ab initio approach, is our most powerful tool for extracting predictions from QCD. However, there are significant limitations. In particular, lattice QCD is limited to equilibrium physics at zero baryon density. Because of so-called “sign problems”, lattice QCD is not a viable means for studying QCD dynamics or QCD at finite density. We need a new approach. Quantum simulation can, in principle, simulate quantum field theories at finite density and real time with no sign problem. However, the realization of a quantum simulator for QCD will require significant advancements from both the theory and the hardware sides. We need to start with toy models—spin models, Abelian-Higgs in 1+1 dimensions, the Schwinger model, SU(2) with fermions, and so on before eventually reaching full QCD.

Theoretical work in the analog simulation of quantum field theories has shown that the Abelian-Higgs model (i.e. scalar QED) in 1+1 dimensions can be mapped to a Rydberg ladder [1–4]. In this proposed approach, the Abelian-Higgs Hamiltonian is mapped to an angular momentum Hamiltonian, which in turn is truncated to a finite number (e.g. 5) of angular momentum states, which in principle can be simulated on a Rydberg ladder.

In the limit of infinite self-coupling and zero gauge coupling, the Abelian-Higgs model reduces to the classical O(2) spin model. In this dissertation, we study a class of O(2)-like models. We find interesting behavior already on very small lattices, so we think this model would be a good candidate for analog simulation before attempting the simulation of full quantum field theories like the Abelian-Higgs model.

We studied the effects of adding a symmetry breaking term $-h_q \sum_x \cos(q\varphi_x)$ to the Hamiltonian of the ordinary classical O(2) model. This adds two continuously tunable parameters which allows us to explore the effect of symmetry-breaking in O(2)-like models. There is some previous literature on this model, however, it is generally limited to specific integer values of q and specific values of h_q . In this work, we have studied a much larger parameter space, and unlike previous studies, we considered also noninteger values of q .

When q is integer and $h_q > 0$, the $O(2)$ symmetry is reduced to a $\mathbb{Z}_q$ symmetry. In the $h_q \rightarrow \infty$ limit, the model reduces to the ordinary q -state clock model, which for $q = 2, 3, 4$ has a second-order phase transition and for $q \geq 5$ has two BKT transitions. In the limit $q \rightarrow \infty$, the ordinary $O(2)$ model is recovered. For $q \geq 5$, the pair of BKT transitions seems to persist to finite $h_q > 0$. When $h_q \rightarrow 0$, the small- β transition connects to the BKT transition of the $O(2)$ model, whereas the large- β transition goes to a different finite value of β . For $q = 2, 3, 4$, the phase diagram changes more significantly as the symmetry-breaking parameter h_q is changed. At intermediate h_q , there is still a second-order phase transition in each case, however, the critical exponents vary with h_q , and so the universality class is not the same as in the $h_q = \infty$ case. For sufficiently small $h_q > 0$, there is evidence that the phase transitions become BKT instead of second-order. If true, this is the only example, that we are aware of, in which one can vary a smooth parameter to go from a second-order to a BKT transition without breaking a symmetry or introducing a new symmetry to the model. However, the evidence is not conclusive and future studies on larger lattices should be done to confirm this.

For noninteger q , early results suggested smooth “lobes” in the phase diagram reminiscent of the phase diagram seen in Rydberg atom chains (see Fig. 4.42). However, this initial picture did not hold up under scrutiny. In the $h_q \rightarrow \infty$ limit, we find that for noninteger q , there is a crossover at small β and an Ising phase transition at larger β . In the phase diagram, there is a smooth Ising phase transition line that terminates at the next larger integer, such that the phase diagram looks nearly periodic from integer to integer, but with an abrupt discontinuity at each integer. For noninteger q , this picture seems to persist to all finite $h_q > 0$. This model vividly illustrates the strong influence that symmetries exert in these classes of models and even a barely broken symmetry results in completely different behavior. Thus, in an analog simulation, it may be difficult to keep the system in the desired universality class unless the symmetries are somehow protected.

In summary, we found that adding symmetry breaking term $-h_q \sum_x \cos(q\varphi_x)$ to the ordinary classical $O(2)$ model results in a rich phase diagram, illustrated in Fig. 4.66, containing second-order phase transitions of various universality classes and BKT transitions. This extended- $O(2)$

model can serve as a playground in which to explore these different kinds of phase transitions and their relationships. The model interpolates between different $\mathbb{Z}_q$ models via the continuously tunable parameter q , and so it may be useful for optimizing the various $\mathbb{Z}_q$ approximations of $U(1)$ in quantum simulations of Abelian gauge theories.

BIBLIOGRAPHY

- [1] A. Bazavov, Y. Meurice, S.-W. Tsai, J. Unmuth-Yockey, and J. Zhang, “Gauge-invariant implementation of the Abelian-Higgs model on optical lattices,” *Phys. Rev. D* **92**, 076003 (2015).
- [2] J. Zhang, J. Unmuth-Yockey, J. Zeiher, A. Bazavov, S.-W. Tsai, and Y. Meurice, “Quantum Simulation of the Universal Features of the Polyakov Loop,” *Phys. Rev. Lett.* **121**, 223201 (2018).
- [3] Y. Meurice, “Examples of symmetry-preserving truncations in tensor field theory,” *Phys. Rev. D* **100**, 1, 014506 (2019), arXiv:1903.01918.
- [4] Y. Meurice, “Theoretical methods to design and test quantum simulators for the compact Abelian Higgs model,” *Phys. Rev. D* **104**, 094513 (2021).
- [5] B. A. Berg, *Markov Chain Monte Carlo Simulations and Their Statistical Analysis: With Web-Based Fortran Code*, World Scientific (2004), ISBN 978-9812389350.
- [6] E. W. Montroll, R. B. Potts, and J. C. Ward, “Correlations and Spontaneous Magnetization of the Two-Dimensional Ising Model,” *Journal of Mathematical Physics* **4**, 2, 308–322 (2004).
- [7] C. N. Yang, “The Spontaneous Magnetization of a Two-Dimensional Ising Model,” *Phys. Rev.* **85**, 808–816 (1952).
- [8] R. B. Griffiths, “Peierls Proof of Spontaneous Magnetization in a Two-Dimensional Ising Ferromagnet,” *Phys. Rev.* **136**, A437–A439 (1964).
- [9] K. Huang, *Statistical mechanics*, Edition 2, John Wiley & Sons (1987), ISBN 978-0471815181.
- [10] R. Peierls, “On Ising’s model of ferromagnetism,” *Mathematical Proceedings of the Cambridge Philosophical Society* **32**, 3, 477–481 (1936).
- [11] P. Fendley, *Modern Statistical Mechanics (book in progress)* URL: <https://users.ox.ac.uk/~phys1116/book.html>.
- [12] R. J. Baxter, “Potts model at the critical temperature,” *Journal of Physics C: Solid State Physics* **6**, 23, L445 (1973).
- [13] Blatt, M and Wiseman, S. and Domany, E., “Clustering data through an analogy to the Potts model,” *Advances in Neural Information Processing Systems* 8 (NIPS 1995) (1995).
- [14] S. J. K. Jensen and O. G. Mouritsen, “Is the Phase Transition of the Three-State Potts Model Continuous in Three Dimensions?” *Phys. Rev. Lett.* **43**, 1736–1739 (1979).
- [15] J. M. Kosterlitz, “The critical properties of the two-dimensional xy model,” *Journal of Physics C: Solid State Physics* **7**, 6, 1046 (1974).
- [16] R. G. Edwards, J. Goodman, and A. D. Sokal, “Multi-grid Monte Carlo (II). Two-dimensional XY model,” *Nuclear Physics B* **354**, 2, 289–327 (1991).

- [17] J. F. Fernández, M. F. Ferreira, and J. Stankiewicz, “Critical behavior of the two-dimensional XY model: A Monte Carlo simulation,” *Phys. Rev. B* **34**, 292–300 (1986).
- [18] R. Gupta and C. F. Baillie, “Critical behavior of the two-dimensional XY model,” *Phys. Rev. B* **45**, 2883–2898 (1992).
- [19] R. Gupta, J. DeLapp, G. G. Batrouni, G. C. Fox, C. F. Baillie, and J. Apostolakis, “Phase Transition in the 2D XY Model,” *Phys. Rev. Lett.* **61**, 1996–1999 (1988).
- [20] M. Hasenbusch, “The two-dimensional XY model at the transition temperature: a high-precision Monte Carlo study,” *Journal of Physics A: Mathematical and General* **38**, 26, 5869 (2005).
- [21] Y. Komura and Y. Okabe, “Large-Scale Monte Carlo Simulation of Two-Dimensional Classical XY Model Using Multiple GPUs,” *Journal of the Physical Society of Japan* **81**, 11, 113001 (2012).
- [22] G. Palma, T. Meyer, and R. Labbé, “Finite size scaling in the two-dimensional XY model and generalized universality,” *Phys. Rev. E* **66**, 026108 (2002).
- [23] H. Weber and P. Minnhagen, “Monte Carlo determination of the critical temperature for the two-dimensional XY model,” *Phys. Rev. B* **37**, 5986–5989 (1988).
- [24] J. F. Yu, Z. Y. Xie, Y. Meurice, Y. Liu, A. Denblyker, H. Zou, M. P. Qin, J. Chen, and T. Xiang, “Tensor renormalization group study of classical XY model on the square lattice,” *Phys. Rev. E* **89**, 013308 (2014).
- [25] G. Ortiz, E. Cobanera, and Z. Nussinov, “Dualities and the phase diagram of the p-clock model,” *Nuclear Physics B* **854**, 3, 780–814 (2012).
- [26] P. Archambault, S. T. Bramwell, J.-Y. Fortin, P. C. W. Holdsworth, S. Peysson, and J.-F. Pinton, “Universal magnetic fluctuations in the two-dimensional XY model,” *Journal of Applied Physics* **83**, 11, 7234–7236 (1998).
- [27] P. Archambault, S. T. Bramwell, and P. C. W. Holdsworth, “Magnetic fluctuations in a finite two-dimensional model,” *Journal of Physics A: Mathematical and General* **30**, 24, 8363 (1997).
- [28] S. T. Bramwell and P. C. W. Holdsworth, “Magnetization and universal sub-critical behaviour in two-dimensional XY magnets,” *Journal of Physics: Condensed Matter* **5**, 4, L53 (1993).
- [29] S. T. Bramwell and P. C. W. Holdsworth, “Magnetization: A characteristic of the Kosterlitz-Thouless-Berezinskii transition,” *Phys. Rev. B* **49**, 8811–8814 (1994).
- [30] L. Berthier, P. C. W. Holdsworth, and M. Sellitto, “Nonequilibrium critical dynamics of the two-dimensional XY model,” *Journal of Physics A: Mathematical and General* **34**, 9, 1805 (2001).
- [31] C. Kawabata and K. Binder, “Evidence for vortex formation in Monte Carlo studies of the two-dimensional XY-model,” *Solid State Communications* **22**, 11, 705–710 (1977).

- [32] R. Loft and T. A. DeGrand, “Numerical simulation of dynamics in the XY model,” *Phys. Rev. B* **35**, 8528–8541 (1987).
- [33] R. Kenna, “The XY Model and the Berezinskii-Kosterlitz-Thouless Phase Transition,” (2005), arXiv:cond-mat/0512356.
- [34] P. Olsson, “Effective vortex interaction in the two-dimensional XY model,” *Phys. Rev. B* **46**, 14598–14616 (1992).
- [35] J. Tobochnik and G. V. Chester, “Monte Carlo study of the planar spin model,” *Phys. Rev. B* **20**, 3761–3769 (1979).
- [36] V. L. Berezinskii, “Destruction of Long-range Order in One-Dimensional and Two-Dimensional Systems Having a Continuous Symmetry Group I. Classical Systems,” *Sov. Phys. JETP* **32**, 3, 493 (1971) URL: <https://inspirehep.net/files/f55503250f690969aedfda4ceaf9b4f9>.
- [37] J. M. Kosterlitz and D. J. Thouless, “Ordering, metastability and phase transitions in two-dimensional systems,” *Journal of Physics C: Solid State Physics* **6**, 7, 1181–1203 (1973).
- [38] M. B. Einhorn, R. Savit, and E. Rabinovici, “A physical picture for the phase transitions in ZN symmetric models,” *Nuclear Physics B* **170**, 1, 16–31 (1980).
- [39] R. G. Jha, “Critical analysis of two-dimensional classical XY model,” *Journal of Statistical Mechanics: Theory and Experiment* **2020**, 8, 083203 (2020).
- [40] E. Rastelli, S. Regina, and A. Tassi, “Monte Carlo simulation of a planar rotator model with symmetry-breaking fields,” *Phys. Rev. B* **69**, 174407 (2004).
- [41] E. Rastelli, S. Regina, and A. Tassi, “Monte Carlo simulation for square planar model with a small fourfold symmetry-breaking field,” *Phys. Rev. B* **70**, 174447 (2004).
- [42] D. D. Betts, “The Exact Solution of Some Lattice Statistics Models with Four States Per Site,” *Canadian Journal of Physics* **42**, 8, 1564–1572 (1964).
- [43] S. Elitzur, R. B. Pearson, and J. Shigemitsu, “Phase structure of discrete Abelian spin and gauge systems,” *Phys. Rev. D* **19**, 3698–3714 (1979).
- [44] S. K. Baek and P. Minnhagen, “Non-Kosterlitz-Thouless transitions for the q -state clock models,” *Phys. Rev. E* **82**, 031102 (2010).
- [45] S. Chatterjee, S. Puri, and R. Paul, “Ordering kinetics in the q -state clock model: Scaling properties and growth laws,” *Phys. Rev. E* **98**, 032109 (2018).
- [46] M. S. S. Challa and D. P. Landau, “Phase transitions in the six-state vector Potts model in two dimensions,” *Journal of Applied Physics* **55**, 6, 2429–2431 (1984).
- [47] J. Tobochnik, “Properties of the q -state clock model for $q = 4, 5$, and 6 ,” *Phys. Rev. B* **26**, 6201–6207 (1982).

- [48] Y. Tomita and Y. Okabe, “Probability-changing cluster algorithm for two-dimensional XY and clock models,” *Phys. Rev. B* **65**, 184405 (2002).
- [49] J. Chen, H.-J. Liao, H.-D. Xie, X.-J. Han, R.-Z. Huang, S. Cheng, Z.-C. Wei, Z.-Y. Xie, and T. Xiang, “Phase Transition of the q -State Clock Model: Duality and Tensor Renormalization*,” *Chinese Physics Letters* **34**, 5, 050503 (2017).
- [50] Y. Chen, Z.-Y. Xie, and J.-F. Yu, “Phase transitions of the five-state clock model on the square lattice,” *Chinese Physics B* **27**, 8, 080503 (2018).
- [51] Z.-Q. Li, L.-P. Yang, Z. Y. Xie, H.-H. Tu, H.-J. Liao, and T. Xiang, “Critical properties of the two-dimensional q -state clock model,” *Phys. Rev. E* **101**, 060105 (2020).
- [52] C. M. Lapilli, P. Pfeifer, and C. Wexler, “Universality Away from Critical Points in Two-Dimensional Phase Transitions,” *Phys. Rev. Lett.* **96**, 140603 (2006).
- [53] O. Borisenko, V. Chelnokov, G. Cortese, R. Fiore, M. Gravina, and A. Papa, “Phase transitions in two-dimensional $Z(N)$ vector models for $N > 4$,” *Phys. Rev. E* **85**, 021114 (2012).
- [54] O. Borisenko, G. Cortese, R. Fiore, M. Gravina, and A. Papa, “Numerical study of the phase transitions in the two-dimensional $Z(5)$ vector model,” *Phys. Rev. E* **83**, 041120 (2011).
- [55] S. K. Baek, P. Minnhagen, and B. J. Kim, “True and quasi-long-range order in the generalized q -state clock model,” *Phys. Rev. E* **80**, 060101 (2009).
- [56] F. Corberi, E. Lippiello, and M. Zannetti, “Scaling and universality in the aging kinetics of the two-dimensional clock model,” *Phys. Rev. E* **74**, 041106 (2006).
- [57] Y. Leroyer and K. Roudi, “Dynamical phase diagram of the two-dimensional p -state clock model,” *Journal of Physics A: Mathematical and General* **24**, 8, 1931 (1991).
- [58] M. Suzuki, “Solution of Potts Model for Phase Transition,” *Progress of Theoretical Physics* **37**, 4, 770–772 (1967).
- [59] R. Pathria and P. Beale, *Statistical Mechanics*, Edition 3, Elsevier Science (2011), ISBN 9780081026939.
- [60] H. J. Rothe, *Lattice gauge theories: an introduction*, Edition 3, World Scientific Publishing Company (2012).
- [61] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, “Equation of State Calculations by Fast Computing Machines,” *The Journal of Chemical Physics* **21**, 6, 1087–1092 (1953).
- [62] A. Bazavov, “Biased Metropolis-Heatbath Algorithms for Lattice Gauge Theory,” *PoS* **110** (2005).
- [63] A. Bazavov and B. A. Berg, “Heat bath efficiency with a Metropolis-type updating,” *Phys. Rev. D* **71**, 114506 (2005).

- [64] A. Bazavov, B. A. Berg, and U. M. Heller, “Biased Metropolis-heat-bath algorithm for fundamental-adjoint $SU(2)$ lattice gauge theory,” *Phys. Rev. D* **72**, 117501 (2005).
- [65] A. Bazavov, B. A. Berg, and H.-X. Zhou, “Application of Biased Metropolis Algorithms: From protons to proteins,” *Mathematics and Computers in Simulation* **80**, 6, 1056–1067 (2010), fifth IMACS Seminar on Monte Carlo Methods Applications of Computer Algebra 2007 (ACA 2007) special session on Nonstandard Applications of Computer Algebra Computational Biomechanics and Biology, a collection of papers presented at the 1st IMACS International Conference on the Computational Biomechanics and Biology ICCBB 2007.
- [66] C. Gattringer and C. B. Lang, *Quantum Chromodynamics on the Lattice: An Introductory Presentation*, Springer (2010), ISBN 978-3642018497.
- [67] D. Landau and K. Binder, *A guide to Monte Carlo simulations in statistical physics*, Cambridge university press (2021).
- [68] A. M. Ferrenberg, D. P. Landau, and R. H. Swendsen, “Statistical errors in histogram reweighting,” *Phys. Rev. E* **51**, 5092–5100 (1995).
- [69] A. M. Ferrenberg and R. H. Swendsen, “Optimized Monte Carlo data analysis,” *Computers in Physics* **3**, 5, 101–104 (1989).
- [70] A. M. Ferrenberg and R. H. Swendsen, “New Monte Carlo technique for studying phase transitions,” *Phys. Rev. Lett.* **61**, 2635–2638 (1988).
- [71] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical recipes in C*, Edition 2, Cambridge university press Cambridge (1992), ISBN 0-521-43108-5.
- [72] K. Levenberg, “A method for the solution of certain non-linear problems in least squares,” *Quarterly of applied mathematics* **2**, 2, 164–168 (1944).
- [73] D. W. Marquardt, “An Algorithm for Least-Squares Estimation of Nonlinear Parameters,” *Journal of the Society for Industrial and Applied Mathematics* **11**, 2, 431–441 (1963).
- [74] M. Bruno and R. Sommer, “On fits to correlated and auto-correlated data,” *Computer Physics Communications* **285**, 108643 (2023).
- [75] N. C. Bartelt and T. L. Einstein, “Finite-size effects on the critical structure factor of the two-dimensional Ising model,” *Journal of Physics A: Mathematical and General* **19**, 8, 1429 (1986).
- [76] M. Droz and A. Malaspinas, “Computation of structure factors by finite-size scaling and transfer matrix methods,” *Journal of Physics C: Solid State Physics* **16**, 8, L231 (1983).
- [77] M. Weigel and W. Janke, “Error estimation and reduction with cross correlations,” *Phys. Rev. E* **81**, 066701 (2010).
- [78] M. S. S. Challa and D. P. Landau, “Critical behavior of the six-state clock model in two dimensions,” *Phys. Rev. B* **33**, 437–443 (1986).

- [79] J. Preskill, “Simulating quantum field theory with a quantum computer,” PoS **LAT-TICE2018**, 024 (2018), arXiv:1811.10085.
- [80] M. C. Bañuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero, C. A. Muschik, B. Reznik, E. Rico, L. Tagliacozzo, K. Van Acoleyen, F. Verstraete, U.-J. Wiese, M. Wingate, J. Zakrzewski, and P. Zoller, “Simulating lattice gauge theories within quantum technologies,” The European Physical Journal D **74**, 8, 165 (2020).
- [81] Y. Meurice, R. Sakai, and J. Unmuth-Yockey, “Tensor lattice field theory for renormalization and quantum computing,” Rev. Mod. Phys. **94**, 025005 (2022), arXiv:2010.06539v2.
- [82] A. Ciavarella, N. Klco, and M. J. Savage, “A Trailhead for Quantum Simulation of SU(3) Yang-Mills Lattice Gauge Theory in the Local Multiplet Basis,” Phys. Rev. D **103**, 9, 094501 (2021), arXiv:2101.10227.
- [83] Z. Davoudi, I. Raychowdhury, and A. Shaw, “Search for efficient formulations for Hamiltonian simulation of non-Abelian lattice gauge theories,” Phys. Rev. D **104**, 074505 (2021), arXiv:2009.11802.
- [84] Z. Davoudi, M. Hafezi, C. Monroe, G. Pagano, A. Seif, and A. Shaw, “Towards analog quantum simulations of lattice gauge theories with trapped ions,” Phys. Rev. Res. **2**, 2, 023015 (2020), arXiv:1908.03210.
- [85] J. Bender, E. Zohar, A. Farace, and J. I. Cirac, “Digital quantum simulation of lattice gauge theories in three spatial dimensions,” New Journal of Physics **20**, 9, 093001 (2018), arXiv:1804.02082.
- [86] E. Zohar, J. I. Cirac, and B. Reznik, “Quantum simulations of lattice gauge theories using ultracold atoms in optical lattices,” Reports on Progress in Physics **79**, 1, 014401 (2015).
- [87] U.-J. Wiese, “Ultracold quantum gases and lattice systems: quantum simulation of lattice gauge theories,” Annalen der Physik **525**, 10-11, 777–796 (2013).
- [88] H. Singh and S. Chandrasekharan, “Qubit regularization of the $O(3)$ sigma model,” Phys. Rev. D **100**, 054505 (2019).
- [89] H. Lamm and S. Lawrence, “Simulation of Nonequilibrium Dynamics on a Quantum Computer,” Phys. Rev. Lett. **121**, 17, 170501 (2018), arXiv:1806.06649.
- [90] N. Klco, E. F. Dumitrescu, A. J. McCaskey, T. D. Morris, R. C. Pooser, M. Sanz, E. Solano, P. Lougovski, and M. J. Savage, “Quantum-classical computation of Schwinger model dynamics using quantum computers,” Phys. Rev. A **98**, 032331 (2018).
- [91] D. C. Hackett, K. Howe, C. Hughes, W. Jay, E. T. Neil, and J. N. Simone, “Digitizing Gauge Fields: Lattice Monte Carlo Results for Future Quantum Computers,” Phys. Rev. A **99**, 6, 062341 (2019), arXiv:1811.03629.
- [92] L. Tagliacozzo, A. Celi, and M. Lewenstein, “Tensor Networks for Lattice Gauge Theories with Continuous Groups,” Phys. Rev. X **4**, 041024 (2014).

- [93] Y. Liu, Y. Meurice, M. P. Qin, J. Unmuth-Yockey, T. Xiang, Z. Y. Xie, J. F. Yu, and H. Zou, “Exact Blocking Formulas for Spin and Gauge Models,” *Phys. Rev.* **D88**, 056005 (2013), arXiv:1307.6543.
- [94] S. Chandrasekharan and U. J. Wiese, “Quantum link models: A discrete approach to gauge theories,” *Nuclear Physics B* **492**, 1, 455–471 (1997), arXiv:hep-lat/9609042.
- [95] R. Brower, S. Chandrasekharan, and U. J. Wiese, “QCD as a quantum link model,” *Phys. Rev.* **D60**, 094502 (1999), arXiv:hep-th/9704106.
- [96] J. Schwinger, “Gauge Invariance and Mass. II,” *Phys. Rev.* **128**, 2425–2429 (1962).
- [97] A. M. Polyakov, “Compact Gauge Fields and the Infrared Catastrophe,” *Phys. Lett. B* **59**, 82–84 (1975).
- [98] Y. Meurice, “Discrete aspects of continuous symmetries in the tensorial formulation of Abelian gauge theories,” *Phys. Rev. D* **102**, 1, 014506 (2020), arXiv:2003.10986.
- [99] J. F. Haase, L. Dellantonio, A. Celi, D. Paulson, A. Kan, K. Jansen, and C. A. Muschik, “A resource efficient approach for quantum and classical simulations of gauge theories in particle physics,” *Quantum* **5**, 393 (2021).
- [100] J. B. Kogut, “An introduction to lattice gauge theory and spin systems,” *Rev. Mod. Phys.* **51**, 659–713 (1979).
- [101] P. Ruján, G. O. Williams, H. L. Frisch, and G. Forgács, “Phase diagrams of two-dimensional $Z(q)$ models,” *Phys. Rev. B* **23**, 1362–1370 (1981).
- [102] K. Kaski and J. D. Gunton, “Universal dynamical scaling in the clock model,” *Phys. Rev. B* **28**, 5371–5373 (1983).
- [103] C.-O. Hwang, “Six-state clock model on the square lattice: Fisher zero approach with Wang-Landau sampling,” *Phys. Rev. E* **80**, 042103 (2009).
- [104] T. Surungan, S. Masuda, Y. Komura, and Y. Okabe, “Berezinskii-Kosterlitz-Thouless transition on regular and Villain types of q -state clock models,” *Journal of Physics A: Mathematical and Theoretical* **52**, 27, 275002 (2019).
- [105] J. Zhang, J. Unmuth-Yockey, J. Zeiher, A. Bazavov, S. W. Tsai, and Y. Meurice, “Quantum simulation of the universal features of the Polyakov loop,” *Phys. Rev. Lett.* **121**, 22, 223201 (2018), arXiv:1803.11166.
- [106] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, “Probing many-body dynamics on a 51-atom quantum simulator,” *Nature* **551**, 7682, 579–584 (2017), arXiv:1707.04344.
- [107] A. Keesling *et al.*, “Quantum Kibble–Zurek mechanism and critical dynamics on a programmable Rydberg simulator,” *Nature* **568**, 7751, 207–211 (2019), arXiv:1809.05540.

- [108] R. D. Pisarski, “Remarks on nuclear matter: How an ω_0 condensate can spike the speed of sound, and a model of $Z(3)$ baryons,” *Phys. Rev. D* **103**, 7, L071504 (2021), arXiv: 2101.05813.
- [109] M. E. Fisher, M. N. Barber, and D. Jasnow, “Helicity Modulus, Superfluidity, and Scaling in Isotropic Systems,” *Phys. Rev. A* **8**, 1111–1124 (1973).
- [110] Y. Kumano, K. Hukushima, Y. Tomita, and M. Oshikawa, “Response to a twist in systems with Z_p symmetry: The two-dimensional p -state clock model,” *Phys. Rev. B* **88**, 104427 (2013).
- [111] M. Levin and C. P. Nave, “Tensor renormalization group approach to 2D classical lattice models,” *Phys. Rev. Lett.* **99**, 12, 120601 (2007), arXiv:cond-mat/0611687.
- [112] Z. Y. Xie, J. Chen, M. P. Qin, J. W. Zhu, L. P. Yang, and T. Xiang, “Coarse-graining renormalization by higher-order singular value decomposition,” *Phys. Rev. B* **86**, 4, 045139 (2012), arXiv:1201.1144.
- [113] Y. Chen, K. Ji, Z. Y. Xie, and J. F. Yu, “Cross derivative of the Gibbs free energy: A universal and efficient method for phase transitions in classical spin models,” *Phys. Rev. B* **101**, 165123 (2020).
- [114] J. V. José, L. P. Kadanoff, S. Kirkpatrick, and D. R. Nelson, “Renormalization, vortices, and symmetry-breaking perturbations in the two-dimensional planar model,” *Phys. Rev. B* **16**, 1217–1241 (1977).
- [115] T.-L. H. Nguyen and V. T. Ngo, “Study on the critical properties of thin magnetic films using the clock model,” *Advances in Natural Sciences: Nanoscience and Nanotechnology* **8**, 1, 015013 (2017).
- [116] S. T. Bramwell, P. C. W. Holdsworth, and J. Rothman, “Magnetization in Ultrathin Films: Critical Exponent β for the 2D XY Model with 4-Fold Crystal Fields,” *Modern Physics Letters B* **11**, 4, 139–148 (1997).
- [117] A. Taroni, S. T. Bramwell, and P. C. W. Holdsworth, “Universal window for two-dimensional critical exponents,” *Journal of Physics: Condensed Matter* **20**, 27, 275233 (2008).
- [118] D. Landau, “Non-universal critical behavior in the planar XY-model with fourth order anisotropy,” *Journal of Magnetism and Magnetic Materials* **31-34**, 1115–1116 (1983).
- [119] G. Hu and S. Ying, “Monte Carlo and coarse graining renormalization studies for the XY model with cubic anisotropy,” *Physica A: Statistical Mechanics and its Applications* **140**, 3, 585–596 (1987).
- [120] P. Calabrese and A. Celi, “Critical behavior of the two-dimensional N-component Landau-Ginzburg Hamiltonian with cubic anisotropy,” *Phys. Rev. B* **66**, 184410 (2002).
- [121] A. Chlebicki and P. Jakubczyk, “Criticality of the $O(2)$ model with cubic anisotropies from nonperturbative renormalization,” *Phys. Rev. E* **100**, 052106 (2019).

- [122] N. Butt, X.-Y. Jin, J. C. Osborn, and Z. H. Saleem, “Moving from continuous to discrete symmetry in the 2D XY model,” (2022), arXiv:2205.03548.
- [123] L. Vanderstraeten, B. Vanhecke, A. M. Läuchli, and F. Verstraete, “Approaching the Kosterlitz-Thouless transition for the classical XY model with tensor networks,” *Phys. Rev. E* **100**, 062136 (2019).
- [124] P. Butera and M. Comi, “Quantitative study of the Kosterlitz-Thouless phase transition in an XY model of two-dimensional plane rotators: High-temperature expansions to order β^{20} ,” *Phys. Rev. B* **47**, 11969–11979 (1993).
- [125] P. Olsson, “Monte Carlo analysis of the two-dimensional XY model. II. Comparison with the Kosterlitz renormalization-group equations,” *Phys. Rev. B* **52**, 4526–4535 (1995).
- [126] H. Arisue, “High-temperature expansion of the magnetic susceptibility and higher moments of the correlation function for the two-dimensional XY model,” *Phys. Rev. E* **79**, 011107 (2009).
- [127] L. Hostetler, J. Zhang, R. Sakai, J. Unmuth-Yockey, A. Bazavov, and Y. Meurice, “Clock model interpolation and symmetry breaking in O(2) models,” *Phys. Rev. D* **104**, 054505 (2021).
- [128] L. Hostetler, J. Zhang, R. Sakai, J. Unmuth-Yockey, A. Bazavov, and Y. Meurice, “Clock model interpolation and symmetry breaking in O(2) models,” *PoS(LATTICE2021)*353 (2021), arXiv:2110.05527.
- [129] L. Hostetler, R. Sakai, J. Zhang, J. Unmuth-Yockey, A. Bazavov, and Y. Meurice, “Symmetry Breaking in an Extended-O(2) Model,” *PoS(LATTICE2022)*014 (2022), arXiv:2212.06893.
- [130] S. Morita and N. Kawashima, “Calculation of higher-order moments by higher-order tensor renormalization group,” *Computer Physics Communications* **236**, 65–71 (2019).
- [131] S. Akiyama, Y. Kuramashi, T. Yamashita, and Y. Yoshimura, “Phase transition of four-dimensional Ising model with higher-order tensor renormalization group,” *Phys. Rev. D* **100**, 054510 (2019).
- [132] M. Fishman, S. White, and E. Stoudenmire, “The ITensor Software Library for Tensor Network Calculations,” *SciPost Physics Codebases* (2022), arXiv:2007.14822.
- [133] H. Ueda, K. Okunishi, K. Harada, R. Krčmár, A. Gendiar, S. Yunoki, and T. Nishino, “Finite- m scaling analysis of Berezinskii-Kosterlitz-Thouless phase transitions and entanglement spectrum for the six-state clock model,” *Phys. Rev. E* **101**, 062111 (2020).
- [134] Z.-C. Gu and X.-G. Wen, “Tensor-entanglement-filtering renormalization approach and symmetry-protected topological order,” *Phys. Rev. B* **80**, 155131 (2009).
- [135] G. Sun, T. Vekua, E. Cobanera, and G. Ortiz, “Phase transitions in the $\mathbb{Z}_p$ and U(1) clock models,” *Phys. Rev. B* **100**, 094428 (2019).

- [136] J. Zhang, Y. Meurice, and S.-W. Tsai, “Truncation effects in the charge representation of the $O(2)$ model,” *Phys. Rev. B* **103**, 245137 (2021).
- [137] K. Binder, “Critical Properties from Monte Carlo Coarse Graining and Renormalization,” *Phys. Rev. Lett.* **47**, 693–696 (1981) URL: <https://link.aps.org/doi/10.1103/PhysRevLett.47.693>.
- [138] M. Rader and A. M. Läuchli, “Floating Phases in One-Dimensional Rydberg Ising Chains,” (2019), arXiv:1908.02068.
- [139] A. Lazarides, O. Tieleman, and C. Morais Smith, “Pokrovsky-Talapov model at finite temperature: A renormalization-group analysis,” *Phys. Rev. B* **80**, 245418 (2009) URL: <https://link.aps.org/doi/10.1103/PhysRevB.80.245418>.
- [140] M. K. Ramazanov and A. K. Murtazaev, “Phase transitions and critical properties in the antiferromagnetic Heisenberg model on a layered cubic lattice,” *JETP Letters* **106**, 2, 86–91 (2017).
- [141] F. Karsch and S. Stickan, “The three-dimensional, three-state Potts model in an external field,” *Physics Letters B* **488**, 3, 319–325 (2000).
- [142] M. S. S. Challa, D. P. Landau, and K. Binder, “Finite-size effects at temperature-driven first-order transitions,” *Phys. Rev. B* **34**, 1841–1852 (1986) URL: <https://link.aps.org/doi/10.1103/PhysRevB.34.1841>.
- [143] A. Mailhot, M. L. Plumer, and A. Caillé, “Finite-size scaling of the frustrated Heisenberg model on a hexagonal lattice,” *Phys. Rev. B* **50**, 6854–6859 (1994).
- [144] J. Lee and J. M. Kosterlitz, “Finite-size scaling and Monte Carlo simulations of first-order phase transitions,” *Phys. Rev. B* **43**, 3265–3277 (1991).
- [145] D. A. Matoz-Fernandez, D. H. Linares, and A. J. Ramirez-Pastor, “Determination of the critical exponents for the isotropic-nematic phase transition in a system of long rods on two-dimensional lattices: Universality of the transition,” *Europhysics Letters* **82**, 5, 50007 (2008).
- [146] B. Ahrens and A. K. Hartmann, “Critical behavior of the random-field Ising model at and beyond the upper critical dimension,” *Phys. Rev. B* **83**, 014205 (2011).
- [147] S. Biswas, A. Kundu, and A. K. Chandra, “Dynamical percolation transition in the Ising model studied using a pulsed magnetic field,” *Phys. Rev. E* **83**, 021109 (2011).
- [148] B. Ahrens, J. Xiao, A. K. Hartmann, and H. G. Katzgraber, “Diluted antiferromagnets in a field seem to be in a different universality class than the random-field Ising model,” *Phys. Rev. B* **88**, 174408 (2013).
- [149] Z. Merdan and R. Erdem, “The finite-size scaling study of the specific heat and the Binder parameter for the six-dimensional Ising model,” *Physics Letters A* **330**, 6, 403–407 (2004).

- [150] K. S. D. Beach, L. Wang, and A. W. Sandvik, “Data collapse in the critical region using finite-size scaling with subleading corrections,” (2005), arXiv:cond-mat/0505194.
- [151] M. Dolfi, J. Gukelberger, A. Hehn, J. Imriška, K. Pakrouski, T. F. Rønnow, M. Troyer, I. Zintchenko, F. Chirigati, J. Freire, and D. Shasha, “A model project for reproducible papers: critical temperature for the Ising model on a square lattice,” (2014), arXiv:1401.2000.
- [152] J. Houdayer and A. K. Hartmann, “Low-temperature behavior of two-dimensional Gaussian Ising spin glasses,” *Phys. Rev. B* **70**, 014418 (2004).
- [153] R. A. Dumer and M. Godoy, “Ising model on a 2D additive small-world network,” *The European Physical Journal B* **95**, 9 (2022).
- [154] Mon, K. K., “Finite-size scaling of the 5D Ising model,” *Europhys. Lett.* **34**, 6, 399–404 (1996) URL: <https://doi.org/10.1209/epl/i1996-00470-4>.
- [155] P. O. Mari and I. A. Campbell, “Ising spin glasses: Corrections to finite size scaling, freezing temperatures, and critical exponents,” *Phys. Rev. E* **59**, 2653–2658 (1999).
- [156] G. A. Canova, Y. Levin, and J. J. Arenzon, “Kosterlitz-Thouless and Potts transitions in a generalized XY model,” *Phys. Rev. E* **89**, 012126 (2014).
- [157] W. Selke, “Critical Binder cumulant of two-dimensional Ising models,” *The European Physical Journal B - Condensed Matter and Complex Systems* **51**, 2, 223–228 (2006).
- [158] D. Loison, “Binder’s cumulant for the Kosterlitz-Thouless transition,” *Journal of Physics: Condensed Matter* **11**, 34, L401 (1999).
- [159] M. Hasenbusch, “The Binder cumulant at the Kosterlitz–Thouless transition,” *Journal of Statistical Mechanics: Theory and Experiment* **2008**, 08, P08003 (2008).
- [160] G. Cortese, *Critical properties of 2D $Z(N)$ vector models for $N > 4$* , Phd thesis, Università della Calabria (2011), available at <https://dspace.unical.it/bitstream/handle/10955/5380/thesis.pdf>.
- [161] M. Klomfass, U. M. Heller, and H. Flyvbjerg, “Interpolating between Ising, XY-, and non-linear σ -models,” *Nuclear Physics B* **360**, 2, 264–282 (1991).
- [162] U. Wolff, “Simulating the all-order strong coupling expansion III: O(N) sigma/loop models,” *Nuclear Physics B* **824**, 1, 254 – 272 (2010).
- [163] F. Bruckmann, C. Gattringer, T. Kloiber, and T. Sulejmanpasic, “Two-dimensional O(3) model at nonzero density: From dual lattice simulations to repulsive bosons,” *Phys. Rev. D* **94**, 114503 (2016).
- [164] D. J. Gross and F. Wilczek, “Ultraviolet Behavior of Non-Abelian Gauge Theories,” *Phys. Rev. Lett.* **30**, 1343–1346 (1973).
- [165] H. D. Politzer, “Reliable Perturbative Results for Strong Interactions?” *Phys. Rev. Lett.* **30**, 1346–1349 (1973).

- [166] K. G. Wilson, “Confinement of quarks,” *Phys. Rev. D* **10**, 2445–2459 (1974).
- [167] B. Anastasiou, *Quantum Field Theory II* (2018) URL: <https://people.phys.ethz.ch/~babis/Teaching/QFT1/qft2.pdf>.
- [168] M. Creutz, “Monte Carlo study of quantized SU(2) gauge theory,” *Phys. Rev. D* **21**, 2308–2315 (1980).
- [169] N. Cabibbo and E. Marinari, “A new method for updating SU(N) matrices in computer simulations of gauge theories,” *Physics Letters B* **119**, 4, 387–390 (1982).
- [170] K. Fabricius and O. Haan, “Heat bath method for the twisted Eguchi-Kawai model,” *Physics Letters B* **143**, 4, 459–462 (1984).
- [171] A. Kennedy and B. Pendleton, “Improved heatbath method for Monte Carlo calculations in lattice gauge theories,” *Physics Letters B* **156**, 5, 393–399 (1985).
- [172] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth, “Hybrid Monte Carlo,” *Physics Letters B* **195**, 2, 216–222 (1987).
- [173] T. Matthews and A. Salam, “The Green’s functions of quantized fields,” *Il Nuovo Cimento* (1955-1965) **12**, 563–565 (1954).
- [174] T. Matthews and A. Salam, “Propagators of quantized field,” *Il Nuovo Cimento* (1955-1965) **2**, 120–234 (1955).
- [175] I. Horváth, A. Kennedy, and S. Sint, “A new exact method for dynamical fermion computations with non-local actions,” *Nuclear Physics B - Proceedings Supplements* **73**, 1, 834–836 (1999).
- [176] M. Clark and A. Kennedy, “The RHMC algorithm for 2 flavours of dynamical staggered fermions,” *Nuclear Physics B - Proceedings Supplements* **129-130**, 850–852 (2004), lattice 2003.
- [177] R. P. Feynman, “Simulating Physics with Computers,” *Int J Theor Phys* **21**, 467–488 (1982).
- [178] I. C. Cloët, M. R. Dietrich, J. Arrington, A. Bazavov, M. Bishof, A. Freese, A. V. Gorshkov, A. Grassellino, K. Hafidi, Z. Jacob, M. McGuigan, Y. Meurice, Z.-E. Meziani, P. Mueller, C. Muschik, J. Osborn, M. Otten, P. Petreczky, T. Polakovic, A. Poon, R. Pooser, A. Roggero, M. Saffman, B. VanDevender, J. Zhang, and E. Zohar, “Opportunities for Nuclear Physics & Quantum Information Science,” (2019), arXiv:1903.05453.
- [179] C. W. Bauer, Z. Davoudi, A. B. Balantekin, T. Bhattacharya, M. Carena, W. A. de Jong, P. Draper, A. El-Khadra, N. Gemelke, M. Hanada, D. Kharzeev, H. Lamm, Y.-Y. Li, J. Liu, M. Lukin, Y. Meurice, C. Monroe, B. Nachman, G. Pagano, J. Preskill, E. Rinaldi, A. Roggero, D. I. Santiago, M. J. Savage, I. Siddiqi, G. Siopsis, D. V. Zanten, N. Wiebe, Y. Yamauchi, K. Yeter-Aydeniz, and S. Zorzetti, “Quantum Simulation for High Energy Physics,” (2022), arXiv:2204.03381.

- [180] Y. Meurice, *Quantum Field Theory*, 2053-2563, IOP Publishing (2021), ISBN 978-0-7503-2187-7.
- [181] E. Lewis, “Using electronic circuits to model simple neuroelectric interactions,” *Proceedings of the IEEE* **56**, 6, 931–949 (1968).
- [182] A. Corkhill, “A Superb Explanatory Device: The MONIAC, an Early Hydraulic Analog Computer,” *University of Melbourne Collections* **10**, 24–28 (2012).
- [183] R. Dardashti, K. P. Y. Thébault, and E. Winsberg, “Confirmation via Analogue Simulation: What Dumb Holes Could Tell Us about Gravity,” *The British Journal for the Philosophy of Science* **68**, 1, 55–89 (2017).
- [184] R. Dardashti, S. Hartmann, K. Thébault, and E. Winsberg, “Hawking radiation and analogue experiments: A Bayesian analysis,” *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **67**, 1–11 (2019).
- [185] U.-J. Wiese, “Towards quantum simulating QCD,” *Nuclear Physics A* **931**, 246–256 (2014).
- [186] I. Georgescu, S. Ashhab, and F. Nori, “Quantum simulation,” *Reviews of Modern Physics* **86**, 1, 153–185 (2014), arXiv:1308.6253.
- [187] M. C. Bañuls and K. Cichy, “Review on novel methods for lattice gauge theories,” *Reports on Progress in Physics* **83**, 2, 024401 (2020).
- [188] I. Bloch, J. Dalibard, and S. Nascimbène, “Quantum simulations with ultracold quantum gases,” *Nature Physics* **8**, 267–276 (2012).
- [189] M. Dalmonte and S. Montangero, “Lattice gauge theory simulations in the quantum information era,” *Contemporary Physics* **57**, 3, 388–412 (2016), arXiv:1602.03776.
- [190] C. Gross and I. Bloch, “Quantum simulations with ultracold atoms in optical lattices,” *Science* **357**, 6355, 995–1001 (2017).
- [191] R. Blatt and C. F. Roos, “Quantum simulations with trapped ions,” *Nature Physics* **8**, 4, 277–284 (2012).
- [192] A. A. Houck, H. E. Türeci, and J. Koch, “On-chip quantum simulation with superconducting circuits,” *Nature Physics* **8**, 4, 292–299 (2012).
- [193] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, “Cold Bosonic Atoms in Optical Lattices,” *Physical Review Letters* **81**, 15, 3108–3111 (1998), arXiv:cond-mat/9805329.
- [194] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, “Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms,” *Nature* **415**, 6867, 39–44 (2002).
- [195] S. Notarnicola, E. Ercolessi, P. Facchi, G. Marmo, S. Pascazio, and F. V. Pepe, “Discrete Abelian gauge theories for quantum simulations of QED,” *Journal of Physics A: Mathematical and Theoretical* **48**, 30, 30FT01 (2015), arXiv:1503.04340.

- [196] D. Petcher and D. H. Weingarten, “Monte Carlo calculations and a model of the phase structure for gauge theories on discrete subgroups of $SU(2)$,” *Phys. Rev. D* **22**, 2465–2477 (1980).
- [197] G. Bhanot and C. Rebbi, “Monte Carlo simulations of lattice models with finite subgroups of $SU(3)$ as gauge groups,” *Phys. Rev. D* **24**, 3319–3322 (1981).
- [198] H. Grosse and H. Kühnelt, “Phase structure of lattice gauge theories for non-abelian subgroups of $SU(3)$,” *Physics Letters B* **101**, 1, 77–81 (1981).
- [199] A. Alexandru, P. F. Bedaque, S. Harmalkar, H. Lamm, S. Lawrence, and N. C. Warrington (NuQS Collaboration), “Gluon field digitization for quantum computers,” *Phys. Rev. D* **100**, 114501 (2019).
- [200] E. Gustafson, P. Dreher, Z. Hang, and Y. Meurice, “Benchmarking quantum computers for real-time evolution of a $(1 + 1)$ field theory with error mitigation,” *Quantum Sci. Technol.* **6**, 045020 (2021), arXiv:1910.09478.
- [201] E. Gustafson, Y. Meurice, and J. Unmuth-Yockey, “Quantum simulation of scattering in the quantum Ising model,” *Phys. Rev. D* **99**, 094503 (2019).
- [202] S. P. Jordan, K. S. M. Lee, and J. Preskill, “Quantum Algorithms for Quantum Field Theories,” *Science* **336**, 6085, 1130–1133 (2012).
- [203] S. P. Jordan, K. S. M. Lee, and J. Preskill, “Quantum Computation of Scattering in Scalar Quantum Field Theories,” *Quantum Information and Computation* **14**, 1014–1080 (2014), arXiv:1112.4833.
- [204] S. P. Jordan, H. Krovi, K. S. M. Lee, and J. Preskill, “BQP-completeness of scattering in scalar quantum field theory,” *Quantum* **2**, 44 (2018).
- [205] N. Klco and M. J. Savage, “Digitization of scalar fields for quantum computing,” *Physical Review A* **99**, 5 (2019).
- [206] N. Klco and M. J. Savage, “Systematically localizable operators for quantum simulations of quantum field theories,” *Physical Review A* **102**, 1 (2020).
- [207] B. Chakraborty, M. Honda, T. Izubuchi, Y. Kikuchi, and A. Tomiya, “Classically Emulated Digital Quantum Simulation of the Schwinger Model with Topological Term via Adiabatic State Preparation,” (2022), arXiv:2001.00485.
- [208] N. Klco, A. Roggero, and M. J. Savage, “Standard model physics and the digital quantum revolution: thoughts about the interface,” *Reports on Progress in Physics* **85**, 6, 064301 (2022).
- [209] G. Pederiva, A. Bazavov, B. Henke, L. Hostetler, D. Lee, H.-W. Lin, and A. Shindler, “Quantum State Preparation for the Schwinger Model,” (2021), arXiv:2109.11859.
- [210] G. Pederiva, A. Bazavov, B. Henke, L. Hostetler, D. Lee, H.-W. Lin, and A. Shindler, “Comparison of Quantum State Preparation Algorithms for the Schwinger Model,” In progress.

- [211] P. Jordan and E. Wigner, “Über das Paulische Äquivalenzverbot,” *Zeitschrift für Physik* **47**, 9, 631–651 (1928).
- [212] G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, “Quantum algorithms for fermionic simulations,” *Phys. Rev. A* **64**, 022319 (2001).
- [213] R. Somma, G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, “Simulating physical phenomena by quantum networks,” *Phys. Rev. A* **65**, 042323 (2002).
- [214] S. B. Bravyi and A. Y. Kitaev, “Fermionic Quantum Computation,” *Annals of Physics* **298**, 1, 210–226 (2002).
- [215] J. T. Seeley, M. J. Richard, and P. J. Love, “The Bravyi-Kitaev transformation for quantum computation of electronic structure,” *The Journal of Chemical Physics* **137**, 22, 224109 (2012).
- [216] E. Zohar, “Quantum simulation of lattice gauge theories in more than one space dimension—requirements, challenges and methods,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **380**, 2216 (2021).
- [217] G. Pardo, T. Greenberg, A. Fortinsky, N. Katz, and E. Zohar, “Resource-efficient quantum simulation of lattice gauge theories in arbitrary dimensions: Solving for Gauss’s law and fermion elimination,” *Phys. Rev. Res.* **5**, 023077 (2023).
- [218] B. Kaufman, “Crystal Statistics. II. Partition Function Evaluated by Spinor Analysis,” *Phys. Rev.* **76**, 1232–1243 (1949).
- [219] P. D. Beale, “Exact Distribution of Energies in the Two-Dimensional Ising Model,” *Phys. Rev. Lett.* **76**, 78–81 (1996).
- [220] R. Häggkvist, A. Rosengren, D. Andrén, P. Kundrotas, P. H. Lundow, and K. Markström, “Computation of the Ising partition function for two-dimensional square grids,” *Phys. Rev. E* **69**, 046104 (2004).
- [221] H. Au-Yang and J. H. Perk, “Ising correlations at the critical temperature,” *Physics Letters A* **104**, 3, 131–134 (1984).
- [222] B. M. McCoy, “The Romance of the Ising Model,” in K. Iohara, S. Morier-Genoud, and B. Rémy (Editors) *Symmetries, Integrable Systems and Representations*, 263–295, Springer London, London (2013), ISBN 978-1-4471-4863-0.
- [223] T. T. Wu, “Theory of Toeplitz Determinants and the Spin Correlations of the Two-Dimensional Ising Model. I,” *Phys. Rev.* **149**, 380–401 (1966).

APPENDIX A

ADDITIONAL DETAILS ON CLASSICAL SPIN SYSTEMS

A.1 Exact Lattice Results for the Ising Model in Two Dimensions

The exact solution of the Ising model is known. Not only do we have exact results for the Ising model in the infinite volume, but we also have exact results for finite lattices. This makes it a good test model when developing new Monte Carlo methods for statistical physics and quantum field theory. For example, in [218], the exact partition function for the Ising model with interaction strength J on an $m \times n$ lattice is given

$$Z = \frac{1}{2} \left(2 \sinh(2\beta J) \right)^{mn/2} \left[\prod_{r=1}^n \left(2 \cosh \left(\frac{m}{2} \gamma_{2r-1} \right) + \prod_{r=1}^n \left(2 \sinh \left(\frac{m}{2} \gamma_{2r-1} \right) \right. \right. \right. \\ \left. \left. \left. + \prod_{r=1}^n \left(2 \cosh \left(\frac{m}{2} \gamma_{2r} \right) + \prod_{r=1}^n \left(2 \sinh \left(\frac{m}{2} \gamma_{2r} \right) \right) \right) \right], \quad (\text{A.1})$$

where γ_r is defined by the equation

$$\cosh \gamma_r = \coth(2\beta J) \cosh(2\beta J) - \cos \left(\frac{r\pi}{n} \right). \quad (\text{A.2})$$

This “algorithm” for the exact partition function for the two dimensional Ising model on finite lattices has been further improved [219, 220] allowing us to calculate the exact lattice partition function using e.g. Mathematica. From the partition function, one can then calculate things like the internal energy

$$E = -\frac{\partial}{\partial \beta} \ln Z, \quad (\text{A.3})$$

and the specific heat

$$C = -\frac{\beta^2}{nm} \frac{\partial E}{\partial \beta}. \quad (\text{A.4})$$

A.2 Notes on Magnetic Susceptibility

Consider a classical spin system with Hamiltonian

$$H = -J \sum_{\langle i,j \rangle} \vec{S}_i \cdot \vec{S}_j - \vec{h} \cdot \vec{M}, \quad (\text{A.5})$$

where the magnetization vector is

$$\vec{M} = \sum_{i=1}^N \vec{S}_i. \quad (\text{A.6})$$

The canonical partition function is

$$Z = \sum_k e^{-\beta H_k}, \quad (\text{A.7})$$

where the sum is over all possible lattice configurations. The expectation value of an observable O is then

$$\langle O \rangle = \frac{\sum_k O_k e^{-\beta H_k}}{Z}. \quad (\text{A.8})$$

The Helmholtz free energy is

$$F = -\frac{1}{\beta} \ln Z. \quad (\text{A.9})$$

In formal statistical mechanics, a magnetic field $\vec{h}$ is applied, breaking the $\mathbb{Z}_q$ or $O(N)$ symmetry and giving the magnetization vector a preferred direction. Then the average value of the magnetization in a given direction j is related to the derivative of the Helmholtz free energy F with respect to the component of the applied magnetic field in the same direction

$$\langle M_j \rangle = - \left(\frac{\partial F}{\partial h_j} \right)_T. \quad (\text{A.10})$$

Then the isothermal magnetic susceptibility is

$$\chi_{M,j} = \left(\frac{\partial \langle M_j \rangle}{\partial h_j} \right)_T = - \left(\frac{\partial^2 F}{\partial h_j^2} \right)_T. \quad (\text{A.11})$$

The susceptibility can be determined by studying the effect on magnetization as the applied magnetic field is varied. For example, we can approximate the first derivative of the magnetization as

$$\chi_{M,j} \approx \frac{\langle M_j(h_{j,2}) \rangle - \langle M_j(h_{j,1}) \rangle}{h_{j,2} - h_{j,1}}, \quad (\text{A.12})$$

where $h_{j,2}$ is a weak magnetic field measured along the j direction, $h_{j,1}$ is in the same direction but with a small difference in magnitude, and $\langle M_j(h_{j,i}) \rangle$ is the average magnetization when the applied magnetic field is $h_{j,i}$. This is simply a finite difference approximation of the derivative. In the limit of small h , one should recover the results of *spontaneous* magnetization.

Alternatively, by using Eq. (A.8), we can write

$$\chi_{M,j} = \frac{\partial \langle M_j \rangle}{\partial h_j} = \frac{\partial}{\partial h_j} \frac{\sum_k (M_j)_k e^{-\beta H_k}}{\sum_i e^{-\beta H_k}} \quad (\text{A.13})$$

where the subscript k refers to the k th possible configuration. We expand the right-hand side using the quotient rule of derivatives and then use

$$\frac{\partial}{\partial h_j} \sum_k e^{-\beta H_k} = \frac{\partial}{\partial h_j} \sum_k e^{-\beta J S_k + \beta \vec{h} \cdot \vec{M}_k} = \sum_k \beta (M_j)_k e^{-\beta J S_k + \beta \vec{h} \cdot \vec{M}_k} \quad (\text{A.14})$$

where $S_k = \sum_{\langle i,j \rangle} \vec{S}_i \cdot \vec{S}_j$ measured on the k th configuration. The partial derivative brings down a factor of β and the j -th component of $\vec{M}$. After simplifying and applying Eq. (A.8), one ends up with

$$\chi_{M,j} = \beta \left(\langle M_j^2 \rangle - \langle M_j \rangle^2 \right). \quad (\text{A.15})$$

If we are measuring things along the direction of the applied magnetic field and if the magnetization vector $\vec{M}$ is parallel to the applied magnetic field $\vec{h}$, which in general it should be, then we can replace M_j with $M = |\vec{M}|$ and then

$$\chi_M = \beta \left(\langle M^2 \rangle - \langle M \rangle^2 \right). \quad (\text{A.16})$$

One can also write the susceptibility as

$$\chi_M = \beta \langle (\delta M)^2 \rangle, \quad (\text{A.17})$$

where $\delta M = M - \langle M \rangle$. That is, the susceptibility is related to the expectation value of the “subtracted” magnetization squared.

In this dissertation, we generally use the definition

$$\chi_M = \frac{1}{\beta N} \frac{\partial \langle M \rangle}{\partial h} = \frac{1}{N} \left(\langle M^2 \rangle - \langle M \rangle^2 \right). \quad (\text{A.18})$$

In the lattice Monte Carlo approach, we might never apply and vary an external magnetic field and instead study the spontaneous magnetization directly. Now there is no *a priori* preferred direction. Rather, the discrete symmetry, if any, is spontaneously broken when the system magnetizes.

A.3 Spin-spin Correlation Functions

A.3.1 Ising Model

The correlation of spins (within a single lattice configuration) as a function of distance is a measure of the amount of clustering in the lattice.

In this section, we consider the 2d Ising model with no external magnetic field. To make this discussion easier to follow, we switch our notation a bit. We let $\sigma_{ij} \in \{-1, 1\}$ be the spin variable sitting at the site with Cartesian coordinates (i, j) . This notation makes it easier to identify the nearest neighbors of a site. For example, the four nearest neighbors of $\sigma_{i,j}$ are $\sigma_{i,j-1}$, $\sigma_{i,j+1}$, $\sigma_{i-1,j}$, and $\sigma_{i+1,j}$.

We define the spin-spin correlation function schematically as¹

$$G(r) = \langle \sigma_{0,0} \sigma_{0,r} \rangle,$$

where $\sigma_{0,r}$ is the spin variable r units to the right of $\sigma_{0,0}$. By translation invariance, we can take $\sigma_{0,0}$ to be any spin variable in the lattice, and by isotropy, we can take the other spin variable to be r units in any of the four cardinal directions from $\sigma_{0,0}$. Note, that each spin variable has exactly four spin variables at a distance r —above, below, left, and right. We are not considering diagonal correlations—only vertical and horizontal. To maximize our statistics, we include all possible spin-variables and all four directions, such that we really calculate

$$G(r) = \frac{1}{4N} \sum_i \sum_j \left[\sigma_{i,j} \sigma_{i-r,j} + \sigma_{i,j} \sigma_{i+r,j} + \sigma_{i,j} \sigma_{i,j-r} + \sigma_{i,j} \sigma_{i,j+r} \right],$$

where N is the number of sites in the lattice.

For a fully ordered lattice (i.e. $\beta = \infty$), all spin variables have the same value and $G(r) = 1$. For a fully disordered lattice (i.e. $\beta = 0$), the spin variables have random values and $G(r) = 0$. Between these extremes, $G(r)$ tends to be an exponential decay. This is illustrated in Figure A.1. For $\beta > \beta_c$, the correlation function approaches a nonzero asymptotic value, which turns out to be M^2 —the spontaneous magnetization squared. Given the periodic boundary conditions, the

¹Some authors subtract the magnetization, defining the correlation function as $G(r) = \langle \sigma_{0,0} \sigma_{0,r} \rangle - \langle \sigma_{0,0} \rangle^2$.




Figure A.1 Here we look at the temperature dependence of the correlation function for a 512×512 MCMC simulation of the Ising model. At small β , the lattice is disordered and $G(r) \approx 0$. At large β , the lattice is frozen and $G(r) \approx 1$. At intermediate β , $G(r)$ exhibits exponential decay. Note, these are illustrative results without error bars.




Figure A.2 Here we calculate the correlation function $G(r)$ for the 2d Ising model on lattices of size 128^2 , 256^2 , and 512^2 . Simulations were performed close to the critical value of β . For comparison, we include exact values for $N = 1, 2, 3, 4, 5$ from [221]. The line gives the large asymptotic values from Eq. (A.19). Note, these are illustrative results without error bars.

correlation function is symmetric in the sense that $C(r) = C(L - r)$. Thus, when plotting the correlation function, typically, only the first half is plotted. I.e. $0 \leq r \leq L/2$.

At the critical point β_c , we can calculate $G(r)$ exactly via Toeplitz determinants. See [222] and references therein. However, this method of calculating them is inefficient for large r . The first several values are given in Au-Yang and Perks [221]. At β_c , we know that the asymptotic (i.e. large

r) behavior is [222, 223]

$$G(r) \sim \frac{2^{1/12} e^{3\zeta'(-1)} \cosh^{1/4}(2\beta_c)}{r^{1/4}} \simeq \frac{0.70338}{r^{1/4}}, \quad (\text{A.19})$$

where $\zeta'(-1)$ is the value of the derivative of the Riemann zeta function at -1 . Thus, at least for $\beta = \beta_c$, for both small and large r , we can validate our lattice results. See Figure A.2.

A.3.2 Clock Models

As in [25, 50, 56], we define the correlation function

$$G(r_{ij}) = \langle \vec{S}_i \cdot \vec{S}_j \rangle = \langle \cos(\theta_i - \theta_j) \rangle = \langle \cos \theta_i \cos \theta_j \rangle + \langle \sin \theta_i \sin \theta_j \rangle, \quad (\text{A.20})$$

where r_{ij} is the distance between the sites i and j . In this definition, the brackets indicate ensemble average. That is, it indicates the average over many equilibrium configurations. The correlation function is translationally invariant, and so it does not depend separately on i and j , but rather only on the distance r between the sites i and j . Thus, we sum it over all N sites in the lattice. We consider only ‘on-axis’ correlations, i.e., we consider only correlations in the horizontal and vertical directions and not in any diagonal directions. For a given site in a 2D lattice, there are four sites which are r units away—up r units, down r units, left r units, and right r units. So in practice, we calculate the correlation function as

$$G(r) = \frac{1}{4N} \sum_{i=1}^N \sum_j \langle \cos(\theta_i - \theta_j) \rangle, \quad (\text{A.21})$$

where i goes over all sites in the lattice, and j goes over the four relevant sites (up, down, left, and right) which are r units from site i .

In [43], the correlation function is defined in exponential form as

$$F(r_{ij}) = \langle e^{i(\theta_i - \theta_j)} \rangle. \quad (\text{A.22})$$

In the low temperature phase, there is long-range order, so $G(r_{ij})$ is constant. In the high temperature phase, there is no long-range order, and $G(r)$ decays exponentially with r . In the intermediate temperature phase (which exists for $q \geq 5$), there is quasi-long-range order, and the

correlation function decays as a power of the distance

$$G(r_{ij}) \sim r_{ij}^{-\eta}, \quad (\text{A.23})$$

where η is a temperature-dependent exponent.

A.4 Low Temperature Expansion

We now consider the low temperature (large β) series expansion of the clock model in two dimensions. This is an approximation that is helpful in pinning down the asymptotic behavior of the energy density. The idea is to work backwards from the lowest energy configurations. The internal energy of a given q -state clock model configuration can be computed as

$$E = - \sum_{\langle i,j \rangle} \cos(\theta_i - \theta_j), \quad (\text{A.24})$$

where the sum is over nearest-neighbors pairs, and $\theta_i = 2\pi k/q$ with $k = 0, 1, \dots, q-1$. The partition function is

$$Z = \sum_k e^{-\beta E_k}. \quad (\text{A.25})$$

where the sum is over the possible lattice configurations. Once we have the partition function, the energy density is given by

$$\frac{\langle E \rangle}{N} = - \frac{\partial}{\partial \beta} \frac{\ln Z}{N}. \quad (\text{A.26})$$

Keep in mind that Eq. (A.24) is the energy of a single lattice configuration whereas Eq. (A.26) should be thought of as a function of β .

For the low temperature series expansion, we start by considering the lowest energy states, and we expand the partition function as the power series

$$Z = g_0 e^{-\beta E_0} + g_1 e^{-\beta E_1} + g_2 e^{-\beta E_2} + \dots = g_0 e^{-\beta E_0} \left[1 + \frac{g_1}{g_0} e^{-\beta(E_1-E_0)} + \frac{g_2}{g_0} e^{-\beta(E_2-E_0)} + \dots \right], \quad (\text{A.27})$$

where E_k gives the k th lowest possible energy, and g_k is the number of lattice configurations with that energy. This is only a valid approximation if $e^{-\beta(E_i-E_0)} \ll 1$, for all $i > 0$ such that the first term is a good approximation and the remaining terms are small perturbations which die off quickly as the order is increased. For the q -state clock model, we end up with powers of $e^{-\epsilon\beta}$ where

$\epsilon = 1 - \cos(2\pi/q)$. This quantity is small for $\beta \sim 2$ when q is fairly small. However, when $q \rightarrow \infty$ for constant β , we get $e^{-\epsilon\beta} \rightarrow 1$, and the approximation fails. Thus, this kind of expansion will not work for the XY model, which is the clock model with $q \rightarrow \infty$.

For the clock model, the lowest possible energy occurs when all spins are aligned. The energy of such configurations is $E_0 = -2N \cos 0 = -2N$, where $2N$ is the number of nearest neighbor pairs in a two-dimensional lattice. There are a total of $g_0 = q$, configurations with this energy since there are q different directions in which the aligned spins could be pointing.

The next lowest energy occurs when a single spin in the configuration is turned one unit away from alignment. All other spins remain aligned. The energy of such a configuration is $E_1 = -(2N - 4) \cos 0 - 4 \cos(2\pi/q) = -2N + 4(1 - \cos(2\pi/q))$, where $(2N - 4)$ is the number of nearest neighbors which remain aligned and 4 is the number of pairs which are out of alignment. There are $g_1 = 2Nq$, configurations with this energy since the unaligned spin could be at N different places, it could be pointing one unit out of alignment to either the left or the right, and there are q directions for the overall magnetization.

Determining the next lowest energy is a bit harder. Is it when a single spin is misaligned by two units or is it when a pair of spins are misaligned both by a single unit? If we go through the calculations we find that, for $q > 3$, then the next lowest energy occurs when there are two spins misaligned by one unit but those two spins must be neighboring and pointing in the same misaligned direction. In this case, there are a total of $(2N - 6)$ aligned pairs and 6 misaligned pairs, so the energy is $E_2 = -(2N - 6) \cos 0 - 6 \cos(2\pi/q) = -2N + 6(1 - \cos(2\pi/q))$. The first misaligned spin could be put at N different sites, the second spin could be put at 4 different places relative to the first, and there are q different directions for the overall orientation. Thus, there are $g_2 = 4Nq$, configurations with this energy.

One could go on, but we'll stop here. We now write our partition function as

$$Z = qe^{2N\beta} \left[1 + 2Ne^{-4\epsilon\beta} + 4Ne^{-6\epsilon\beta} + \dots \right], \quad (\text{A.28})$$




Figure A.3 The clock model for several q along with the low temperature series expansion (solid lines).

where $\epsilon \equiv 1 - \cos(2\pi/q)$. Taking the logarithm and then differentiating with respect to β gives

$$\frac{\langle E \rangle}{N} = -2 + \epsilon \frac{8u^4 + 24u^6 + \dots}{1 + 2Nu^4 + 4Nu^6 + \dots}, \quad (\text{A.29})$$

where $u \equiv e^{-\epsilon\beta}$. For q and β such that $u = e^{-\epsilon\beta}$ is small, one can make the approximation

$$\frac{\langle E \rangle}{N} \simeq -2 + \epsilon \left[8u^4 + 24u^6 \right]. \quad (\text{A.30})$$

At $\beta = 2$ for the 2-, 3-, 4-, and 5-state clock models, this seems to be a valid approximation. See Fig. (A.3). See also the appendix in [47].

APPENDIX B

TECHNICAL DETAILS OF SIMULATING THE EXTENDED- $O(2)$ MODEL

B.1 Effect of the Angle Cutoff in the Extended q -state Clock Model

In the extended q -state clock model, the spins are allowed to take the angles $\varphi^{(k)} = 2\pi k/q$, where k is an integer. By restricting the values of the angle to the domain $[\varphi_0, \varphi_0 + 2\pi)$, the values of k must satisfy $\varphi_0 q/2\pi \leq k < \varphi_0 q/2\pi + q$. Let $q = \lfloor q \rfloor + \delta q$, $\varphi_0 q/2\pi = p + \epsilon$, p is an integer and $0 \leq \epsilon, \delta q < 1$, then

$$k \in \begin{cases} (p, p+1, \dots, p + \lfloor q \rfloor) & \text{if } \epsilon = 0, \\ (p+1, p+2, \dots, p + \lfloor q \rfloor) & \text{if } \epsilon > 0, \epsilon + \delta q < 1, \\ (p+1, p+2, \dots, p + \lfloor q \rfloor + 1) & \text{if } \epsilon > 0, \epsilon + \delta q \geq 1. \end{cases} \quad (\text{B.1})$$

Because the action only depends on the angular distance, we actually have two cases: $k = 0, 1, 2, \dots, \lfloor q \rfloor$ (case 1), $k = 0, 1, 2, \dots, \lfloor q \rfloor - 1$ (case 2). This defines a *particular* model that interpolates between the integer q 's of the ordinary q -state clock model. If $\varphi_0 = 0$, we are in case 1. For this model, the interpolation is smooth (in the sense that the thermodynamic curves change smoothly) as an integer q is approached from below. In this case, the allowed angles feature a cutoff at $\lfloor q \rfloor 2\pi/q$ that breaks the periodicity. One could choose the allowed angles differently and go to case 2, where the allowed angles feature a cutoff at $2\pi(\lfloor q \rfloor - 1)/q$ that breaks the periodicity, and the interpolation is smooth as an integer q is approached from above.

One could remove the cutoff and restore a form of periodicity by allowing angles

$$\varphi^{(k)} = \frac{2\pi k}{q}, \quad k = 0, 1, 2, \dots, \infty.$$

However, in this case, for a rational $q = r/s$, where r/s is a reduced fraction, the model is the ordinary r -state clock model since $2\pi k/(r/s) = 2\pi ks/r$ and ks is an integer. Thus, removing the cutoff would destroy the smooth interpolation in $n < q \leq n+1$ for integer n . For any irrational q , this model without cutoff becomes equivalent to the ∞ -state clock model, that is, the XY model.




Figure B.1 The specific heat for the extended q -state clock model with $q = 3.141592654$. Results are from Monte Carlo on 4×4 lattices. The different curves correspond to different allowed domains for the spin angles φ . The specific heat of the XY model is included for reference. As the upper limit of the domain goes to infinity, we get the XY model. Statistical error bars are omitted since they are smaller than the line thickness. Dashed lines indicate regions where we have data but we do not have the uncertainty fully under control.

In Fig. B.1 we explore the effect of increasing the cutoff for the model with $q = 3.141592654 \approx \pi$. For example, for the case $\varphi \in [0, 2\pi)$, the allowed angles are $2\pi k/\pi = 2k$ with $k = 0, 1, 2, 3$, whereas for the case $\varphi \in [0, 4\pi)$, we have $k = 0, 1, \dots, 6$. In all of these cases there remains a $\mathbb{Z}_2$ symmetry, and we expect the model to have an Ising transition at very large β for certain cutoffs. As the upper limit of the domain is moved to infinity (i.e. as the cutoff is removed), this Ising transition moves to infinity and the model becomes the XY model. A minor detail is that $3.141592654 = 3141592654/1000000000$ is in fact a rational number, so one would actually get the 3141592654-state clock model if the cutoff were removed completely. However, it would be indistinguishable from the XY model for practical purposes.

B.2 Validating TRG with MC in the Extended q -state Clock Model

Whereas Monte Carlo methods are well understood in the context of classical spin models, TRG is a relatively new approach. We validate TRG results at small β and small volume using exact and Monte Carlo results. Exact results can be computed for $q = 2$ (Ising model) and $q = 4$ (two coupled Ising models). We use Monte Carlo to validate the TRG results for other (including




Figure B.2 An example comparison of the energy density from exact calculation and that from TRG for $q = 4.0$. The top panel shows the energy density for lattices of size 4×4 , 8×8 , and 16×16 . The three panels on the bottom show the difference between exact and TRG results. TRG shows deviations from exact results near the phase transition. Here $D_{\text{bond}} = 64$.




Figure B.3 Same as Fig. B.2, but for the specific heat. TRG shows deviations from exact results near the peak, however, this is mostly due to discretization error from the derivative.

fractional) values of q .

In Fig. B.2, we show that the energy density from TRG agrees very well with the exact calculation, and that the tiny discrepancy appears only around the critical point. In TRG, the specific heat is calculated by taking a finite difference derivative of the energy. In Fig. B.3, we compare the specific heat from TRG with the exact values for $q = 4.0$. TRG deviates from the exact




Figure B.4 An example comparison of TRG and Monte Carlo measurements of the energy density for $q = 4.9$. The top panel shows the energy density for lattices of size 4×4 , 8×8 , and 16×16 . The three panels on the bottom show the difference between Monte Carlo and TRG results. Here, Monte Carlo is taken to be the baseline. Here $D_{\text{bond}} = 64$.




Figure B.5 Same as Fig. B.4, but for the specific heat. TRG shows deviations from Monte Carlo results near the peak, however, this is mostly due to discretization error from the derivative.

results near the peak of the specific heat, but this deviation is mostly due to the discretization error from the derivative.

Exact solutions are not known for fractional- q , so we validate TRG by comparing with results from Monte Carlo at small β and small volume. For example, Fig. B.4 shows that the discrepancy




Figure B.6 The distribution Eq. (B.2) of angles in the clock model when $\varphi_0 = 0$ and $\beta = 0$.

between the energy density from TRG and that from Monte Carlo is only of order 10^{-4} . In Fig. B.5, we compare the specific heat from TRG with that of Monte Carlo for $q = 4.9$. TRG deviates from the Monte Carlo results near the peak of the specific heat. However, this deviation is again almost entirely due to discretization error from the derivative.

B.3 The Need for the ε -shift in the Extended- $O(2)$ Model

In Section 4.3.4, we noted that to match the clock model in the $h_q \rightarrow \infty$ limit of the Extended- $O(2)$ model, we need to shift the angles by a small ε , such that $\varphi \in [-\varepsilon, 2\pi - \varepsilon)$.

To understand the reason for this, we must look more closely at the spin distributions. For the clock model, it is useful to consider the distribution of spins at $\beta = 0$. At this temperature, one can think of the spins as being selected uniformly from the discrete distribution Eq. (4.60). We can write such a distribution as a “Dirac comb”. For example, for the case $\varphi_0 = 0$, the angles are distributed as

$$P_{q, \varphi_0=0}^{clock}(\varphi) \sim \sum_{k=0}^{\lfloor q \rfloor} \delta\left(\varphi - \frac{2\pi k}{q}\right), \quad (\text{B.2})$$

on a finite domain (in this case in a domain from 0 to $2\pi\lfloor q \rfloor/q$). See Fig. B.6. When $\beta > 0$ one can still think of the angles as being selected from this Dirac comb, but now some peaks are more likely to be sampled than others.

For the Extended- $O(2)$ model, when β is fixed and h_q is large, the second term in Eq. (4.61) dominates the energy. The energy wants to be minimized, and this occurs when $\cos(q\varphi_x) = 1$, which occurs when $q\varphi_x = 2\pi k$ or $\varphi_x = 2\pi k/q$. Thus, when h_q is large, the second term in Eq. (4.61) has the effect of enforcing Eq. (4.60). To make this more quantitative, it is useful to




Figure B.7 A generic example plot of the distribution Eq. (B.3) for $h_q = 1$ and $h_q = 8$ with $\beta = 1$. After choosing appropriate end points and taking the limit $h_q \rightarrow \infty$, this becomes the Dirac comb distribution of the clock model angles.

consider the spin distribution

$$P_{q,\varphi_0}^{extO2}(\varphi) \sim e^{\beta h_q \cos(q\varphi)}. \quad (\text{B.3})$$

This is plotted in Fig. B.7.

We claim that in the limit $h_q \rightarrow \infty$ the Extended-O(2) model defined by Eq. (4.61) and (4.62) is equivalent to the clock model defined by Eq. (4.59) and (4.60). However, there is a subtlety. We've shown numerically that for some values of φ_0 (an example being $\varphi_0 = 0$) the Extended-O(2) model converges to the clock model in the limit $h_q \rightarrow \infty$ but only if the spins are shifted by some ε

$$\varphi_0 - \varepsilon \leq \varphi \in \mathbb{R} < \varphi_0 - \varepsilon + 2\pi. \quad (\text{B.4})$$

Compare Fig. 4.40, which used $\varepsilon = 0$ and Fig. 4.41, which used $\varepsilon = \pi(1 - \lfloor q \rfloor / q)$. We obtain completely different critical behavior depending on the choice of ε .

To understand this discrepancy, we need to consider what happens when one enforces boundaries $[\varphi_0, \varphi_0 + 2\pi)$ on the distribution Eq. (B.3) and then takes the limit $h_q \rightarrow \infty$. When h_q is increased, the peaks in the distribution become sharper, but their position does not change. The discrepancy occurs when the boundary φ_0 or $\varphi_0 + 2\pi$ coincides with the center of a peak¹ $2\pi k/q$. This occurs for example at $\varphi = 0$ when one has the left boundary at $\varphi_0 = 0$ as illustrated in Fig. B.8. No matter how large h_q is made in this example, one never recovers the Dirac comb of the clock model

¹When $q \in \mathbb{Z}$, periodicity is recovered, and so if a peak occurs at one boundary then another peak occurs at the other boundary and the discrepancy does not occur. I.e. the half peak missing from the left boundary is recovered at the other boundary. In short, one does not have to worry about this discrepancy when $q \in \mathbb{Z}$.




Figure B.8 A plot of the distribution Eq. (B.3) for $q = 4.5$ and generic h_q . If we naively choose the angles in $[0, 2\pi)$, we get a model that does not match the clock model in the $h_q \rightarrow \infty$ limit. No matter how large we make h_q (i.e. no matter how sharp we make the peaks), we are always cutting off half of that first peak, and so we don't get the Dirac comb of Eq. (B.2) in the $h_q \rightarrow \infty$ limit.




Figure B.9 Same as the previous picture, but now we shift the angle domain from $[0, 2\pi)$ to $[-\varepsilon, 2\pi - \varepsilon)$, so that the domain includes all of the peaks. Now it is obvious that we recover the Dirac comb of Eq. (B.2) in the $h_q \rightarrow \infty$ limit. In this example, $h_q = 8$ and $\varepsilon = 0.35$. However, one can use much smaller ε and still include most of the peaks if one makes h_q sufficiently large.

because one is always cutting off half of the first peak. However, if one shifts the boundaries to $[-\varepsilon, 2\pi - \varepsilon)$, as shown in Fig. B.9 then one recovers the Dirac comb of the clock model when $h_q \rightarrow \infty$.

To match the clock model in the $h_q \rightarrow \infty$ limit, it should be sufficient to choose any ε such that

$$P_{q, \varphi_0}^{extO2}(\varphi) \xrightarrow{h_q \rightarrow \infty} P_{q, \varphi_0}^{clock}(\varphi), \quad (\text{B.5})$$

where for the clock model, angles are selected from $[\varphi_0, \varphi_0 + 2\pi)$, but for the Extended-O(2) model, they are selected from $[\varphi_0 - \varepsilon, \varphi_0 - \varepsilon + 2\pi)$. For the rest of this section, we assume $\varphi_0 = 0$. In that

case, the condition is satisfied for

$$0 < \varepsilon < 2\pi \left(1 - \frac{\lfloor q \rfloor}{q}\right). \quad (\text{B.6})$$

This is simply the condition that all of the desired clock angles occur strictly between $-\varepsilon$ and $-\varepsilon + 2\pi$. Note, that for $q \in \mathbb{Z}$ there is no discrepancy and one can use $\varepsilon = 0$.

What value of ε should we use? One could just choose a constant value like $\varepsilon = 0.1$. However, any constant value will be unsafe for some values of q . For example, the condition Eq. (B.6) is not satisfied for $4 < q < 4.06469$ when $\varepsilon = 0.1$, so for these values of q , the Extended-O(2) model would fail to reproduce the clock model results in the $h_q \rightarrow \infty$ limit. A better approach might be to choose ε so that the $\lfloor q \rfloor$ peaks of the distribution $P_{q,\varphi_0}^{extO2}(\varphi)$ are centered in the domain $[-\varepsilon, 2\pi - \varepsilon)$.

This implies

$$\varepsilon = \pi \left(1 - \frac{\lfloor q \rfloor}{q}\right). \quad (\text{B.7})$$

For example, for $q = 4.5$, we have 5 allowed angles in the clock model and 5 peaks in $P_{q,\varphi_0}^{extO2}(\varphi)$ located at $2\pi k/4.5$ with $k = 0, 1, 2, 3, 4$. The first peak is at $\varphi = 0$, and the fifth peak is at $\varphi = 2\pi \cdot 4/4.5$. Thus the value of ε that centers the five peaks within the domain is $\varepsilon = 2\pi(1 - 4/4.5)/2 \approx 0.35$.

B.4 FSS Results at Integer q

Table B.1 Critical exponents are tabulated for integer $q = 2, 3, 4, 5, 6$ and finite $h_q = 0.1, 1.0$. The Q -value is a goodness-of-fit measure defined in Eq. (3.69). For $q = 5, 6$, there appear to be two BKT transitions as is the case for $h_q = \infty$. Here, we include data only for the transition at small inverse temperature. Empty cells indicate that it was not possible to obtain a trustworthy fit.

q	h_q	Exponent	Fit Form	L_{min}	L_{max}	Q-value	χ^2/dof	Exponent
2.0	0.1	$1/\nu$	Eq. (3.70)	32	256	0.196	1.483	0.958(49)
2.0	0.1	α/ν	Eq. (3.75)	32	256	0.498	0.878	0.020(17)
2.0	0.1	γ/ν	Eq. (3.73)	32	256	0.386	1.058	1.780(16)
2.0	0.1	$-\beta/\nu$	Eq. (3.72)	32	256	0.617	0.711	-0.306(81)
2.0	0.1	$2 - \eta$	Eq. (3.74)	32	256	0.198	1.476	1.713(13)
2.0	1.0	$1/\nu$	Eq. (3.70)	32	256	0.722	0.571	0.978(68)
2.0	1.0	α/ν	Eq. (3.75)	32	256	0.011	2.991	-0.010(16)

Table B.1 (cont'd)

2.0	1.0	γ/ν	Eq. (3.73)	32	256	0.532	0.828	1.745(19)
2.0	1.0	$-\beta/\nu$	Eq. (3.72)					
2.0	1.0	$2 - \eta$	Eq. (3.74)	32	256	0.07	2.06	1.749(14)
3.0	0.1	$1/\nu$	Eq. (3.70)	32	256	0.858	0.384	1.521(57)
3.0	0.1	α/ν	Eq. (3.71)	32	256	0.029	2.514	0.538(48)
3.0	0.1	γ/ν	Eq. (3.73)	32	256	0.264	1.302	1.962(16)
3.0	0.1	$-\beta/\nu$	Eq. (3.72)					
3.0	0.1	$2 - \eta$	Eq. (3.74)	32	256	0.852	0.392	1.6349(93)
3.0	1.0	$1/\nu$	Eq. (3.70)	32	256	0.822	0.436	1.236(40)
3.0	1.0	α/ν	Eq. (3.71)	32	256	0.402	1.032	0.385(21)
3.0	1.0	γ/ν	Eq. (3.73)	32	256	0.873	0.361	1.743(16)
3.0	1.0	$-\beta/\nu$	Eq. (3.72)					
3.0	1.0	$2 - \eta$	Eq. (3.74)	32	256	0.407	1.021	1.705(16)
4.0	0.1	$1/\nu$	Eq. (3.70)	32	256	0.942	0.241	0.40(14)
4.0	0.1	α/ν	Eq. (3.75)	32	256	0	10.332	-0.120(19)
4.0	0.1	γ/ν	Eq. (3.73)	32	256	0.53	0.831	1.728(16)
4.0	0.1	$-\beta/\nu$	Eq. (3.72)	32	256	0.09	1.922	-0.31(31)
4.0	0.1	$2 - \eta$	Eq. (3.74)	32	256	0.398	1.035	1.7430(85)
4.0	1.0	$1/\nu$	Eq. (3.70)	32	256	0.167	1.576	0.834(96)
4.0	1.0	α/ν	Eq. (3.75)	32	256	0.53	0.83	-0.1349(90)
4.0	1.0	γ/ν	Eq. (3.73)	32	256	0.568	0.777	1.742(14)
4.0	1.0	$-\beta/\nu$	Eq. (3.72)	32	256	0.39	1.049	-0.65(51)
4.0	1.0	$2 - \eta$	Eq. (3.74)	32	256	0.073	2.032	1.7468(93)
5.0	0.1	γ/ν	Eq. (3.73)	32	256	0.039	2.379	1.727(15)
5.0	0.1	$-\beta/\nu$	Eq. (3.72)	32	256	0.067	2.083	-0.35(22)
5.0	0.1	$2 - \eta$	Eq. (3.74)	32	256	0.028	2.541	1.7284(72)
5.0	1.0	γ/ν	Eq. (3.73)	32	256	0.042	2.339	1.722(14)
5.0	1.0	$-\beta/\nu$	Eq. (3.72)	32	256	0.06	2.148	-0.44(24)
5.0	1.0	$2 - \eta$	Eq. (3.74)	32	256	0.354	1.117	1.7347(94)
6.0	0.1	γ/ν	Eq. (3.73)	32	256	0.678	0.629	1.710(11)
6.0	0.1	$-\beta/\nu$	Eq. (3.72)	32	256	0.009	3.08	-0.15(16)
6.0	0.1	$2 - \eta$	Eq. (3.74)	32	256	0.214	1.432	1.7120(86)
6.0	1.0	γ/ν	Eq. (3.73)	32	256	0.867	0.371	1.726(15)
6.0	1.0	$-\beta/\nu$	Eq. (3.72)					
6.0	1.0	$2 - \eta$	Eq. (3.74)	32	256	0.096	1.889	1.7346(92)

Table B.2 Here we fit the peak (or inflection point) locations to the form Eq. (3.78) to estimate the infinite-volume critical point β_c and the critical exponent ν . Empty cells indicate that it was not possible to obtain a trustworthy fit.

q	h_q	From	L_{min}	L_{max}	Q-value	χ^2/dof	ν	β_c
2.0	0.1	$dU_M/d\beta _{max}$	32	256	0.119	1.767	1.13(10)	0.8775(12)
2.0	0.1	$C_V _{max}$	32	256	0.016	2.816	1.03(11)	0.87726(38)
2.0	0.1	$\chi_M _{max}$	32	256	0.256	1.32	1.054(54)	0.87753(48)
2.0	0.1	$\langle M \rangle _{infl}$	32	256	0.172	1.561	1.014(80)	0.87734(49)
2.0	0.1	$F(\vec{q}) _{max}$	32	256	0.179	1.534	1.066(45)	0.87816(76)
2.0	1.0	$dU_M/d\beta _{max}$	32	256	0.123	1.752	0.979(73)	0.65440(53)
2.0	1.0	$C_V _{max}$	32	256	0.164	1.59	1.04(16)	0.65442(27)
2.0	1.0	$\chi_M _{max}$	32	256	0.842	0.407	0.991(55)	0.65453(28)
2.0	1.0	$\langle M \rangle _{infl}$	32	256	0.963	0.194	0.908(82)	0.65423(28)
2.0	1.0	$F(\vec{q}) _{max}$	32	256	0.778	0.497	1.012(42)	0.65482(40)
3.0	0.1	$dU_M/d\beta _{max}$	32	256	0.133	1.706	0.791(35)	0.99873(93)
3.0	0.1	$C_V _{max}$	32	256	0.924	0.276	0.863(41)	1.00073(74)
3.0	0.1	$\chi_M _{max}$	32	256	0.03	2.507	0.811(19)	0.99924(45)
3.0	0.1	$\langle M \rangle _{infl}$	32	256	0.176	1.545	0.756(29)	0.99878(54)
3.0	0.1	$F(\vec{q}) _{max}$	32	256	0.215	1.429	1.168(44)	1.0047(16)
3.0	1.0	$dU_M/d\beta _{max}$	32	256	0.22	1.415	0.808(31)	0.82568(18)
3.0	1.0	$C_V _{max}$	32	256	0.173	1.558	0.909(47)	0.82603(13)
3.0	1.0	$\chi_M _{max}$	32	256	0.093	1.905	0.760(21)	0.82574(10)
3.0	1.0	$\langle M \rangle _{infl}$	32	256	0.12	1.764	0.760(27)	0.82575(11)
3.0	1.0	$F(\vec{q}) _{max}$	32	256	0.015	2.855	0.823(18)	0.82603(19)
4.0	0.1	$dU_M/d\beta _{max}$	32	256	0.158	1.607	2.26(59)	1.098(27)
4.0	0.1	$C_V _{max}$						
4.0	0.1	$\chi_M _{max}$	32	256	0.229	1.387	1.96(12)	1.0797(54)
4.0	0.1	$\langle M \rangle _{infl}$	32	256	0.082	1.972	2.05(26)	1.082(12)
4.0	0.1	$F(\vec{q}) _{max}$	32	256	0.839	0.412	1.88(13)	1.0772(67)
4.0	1.0	$dU_M/d\beta _{max}$	32	256	0.954	0.216	1.34(17)	0.9918(47)
4.0	1.0	$C_V _{max}$						
4.0	1.0	$\chi_M _{max}$	32	256	0.053	2.21	1.251(47)	0.9883(13)
4.0	1.0	$\langle M \rangle _{infl}$	32	256	0.46	0.935	1.52(26)	0.9952(60)
4.0	1.0	$F(\vec{q}) _{max}$	32	256	0.791	0.479	1.334(67)	0.9912(20)

Table B.3 Here we fit the peak (or inflection point) locations to the form Eq. (3.84) to estimate the infinite-volume critical point β_c of the small- β phase transition.

q	h_q	From	L_{min}	L_{max}	Q-value	χ^2/dof	β_c
5.0	0.1	$\chi_M _{max}$	32	256	0.167	1.579	1.1308(77)
5.0	0.1	$\langle M \rangle _{infl}$	32	256	0.037	2.402	1.143(23)
5.0	0.1	$F(\vec{q}) _{max}$	32	256	0.26	1.312	1.133(13)
5.0	1.0	$\chi_M _{max}$	32	256	0.571	0.774	1.1267(96)
5.0	1.0	$\langle M \rangle _{infl}$	32	256	0.078	2.007	1.133(22)
5.0	1.0	$F(\vec{q}) _{max}$	32	256	0.939	0.246	1.126(15)
6.0	0.1	$\chi_M _{max}$	32	256	0.11	1.819	1.1237(76)
6.0	0.1	$\langle M \rangle _{infl}$	32	256	0.012	2.971	1.123(13)
6.0	0.1	$F(\vec{q}) _{max}$	32	256	0.348	1.127	1.128(12)
6.0	1.0	$\chi_M _{max}$	32	256	0.209	1.445	1.125(11)
6.0	1.0	$\langle M \rangle _{infl}$	32	256	0.339	1.143	1.099(16)
6.0	1.0	$F(\vec{q}) _{max}$	32	256	0.206	1.456	1.118(13)