

**Dynamical fluctuations of
classical and quantum
square dimer models**

by

Tom Oakes

Thesis submitted to
The University of Nottingham
for the degree of
Doctor of Philosophy

September 2018

“It’s still magic even if you know how it’s done.”

Terry Pratchett

Abstract

The work in this thesis is split into two main parts; the dynamical study of the fully packed classical dimer model, and the development of advanced large deviation techniques applied to the classical dimer model, highlighting the connection to the quantum dimer model.

The classical dimer model on the square lattice is a paradigmatic example of a system subject to strong local constraints. We study its behavior under local stochastic dynamics, by means of Monte Carlo simulations and theoretical arguments. We observe clear signatures of correlated dynamics in both global and local observables and over a broad range of time scales, indicating a breakdown of the simple continuum description that approximates well the statics. We show that this collective dynamics can be understood in terms of one-dimensional “strings” of high mobility. We introduce a coarse-grained description of the strings, which leads to exact results in the limit of low string density and provides a detailed qualitative understanding of the dynamics in all flux sectors.

We then go on, still using the fully packed dimers on the square lattice as a paradigmatic system, to study the connection between the phase behaviour of the ground state of quantum dimers and the dynamics of classical stochastic dimers. At the so-called Rokhsar-Kivelson (RK) point, a quantum dimer Hamiltonian is equivalent to the Markov generator of the dynamics of classical dimers. A less well understood fact is that away from the RK point the quantum-classical connection persists: in this case the quantum Hamiltonian maps to a non-stochastic “tilted” master operator of the classical stochastic problem. This implies a direct relation between the phase behaviour of quantum dimers and properties of ensembles of stochastic trajectories of classical dimers. Using transition path sampling (TPS) – supplemented by trajectory umbrella sampling – we obtain the large deviation statistics of dynamical activity in the classical problem, and show the correspondence between the phase behaviour of the classical and quantum systems.

Finally, we take these ideas further by developing a feedback-based modified dynamics used for umbrella sampling alongside TPS. This feedback approach was

first used in the cloning (or population dynamics) method. We compare different modified dynamics within TPS and show how much more efficient the feedback-modified dynamics can be than the original dynamics, for the purpose of rare event sampling. We also discuss how our approach may generalise to other problems with slow complex dynamics.

List of Publications

The work detailed in this thesis has produced a number of publications and preprints. Those directly related to the work in this thesis are:

Chapters 2:

1. Tom Oakes, Juan P. Garrahan and Stephen Powell, Emergence of cooperative dynamics in fully packed classical dimers, *Physical Review E*, 93, 032129, 2016.

Chapters 3 and 4:

2. Tom Oakes, Stephen Powell, Claudio Castelnovo, Austen Lamacraft, and Juan P. Garrahan, Phases of quantum dimers from ensembles of classical stochastic trajectories, *Physical Review B*, 98, 064302, 2018.

Acknowledgements

For the four years of my post-graduate study I have had the pleasure of learning and growing with some great people. They have shaped me as a person and I would like to acknowledge those that helped me the most.

Firstly, I would like to dedicate the work I have presented here in this thesis to my mom and grandparents. They never understood what I was talking about when I would tell them about my work, but they would always like to hear about it anyway.

I would also like to thank my supervisors Prof. Juan p. Garrahan and Prof. Stephen Powell. Both offered me support throughout my PhD and taught me how to approach scientific research.

I must mention my friends and colleagues that helped keep me sane during the initial stages of my PhD. Ben Everest, who gave me many hours of enlightening conversation and potential comedy skits. And Emanuele Levi, who's guiding wisdom helped me through tough times.

I am also incredibly grateful to everyone who help proof read the many drafts of this thesis, which I can only imagine would have been like trying to decode a previously unknown language. So thank you to Jemma Needham, Miake Ostmann and Dominic Rose.

I can not fail mention my close friends that followed me along this ride. Georgia Bishop and Ellie Wilson who allowed me to power through the last few year with glee. Conor Smith, Jim Smith and Toby Wilson all have been great sources of conversation and much needed down time from work.

Contents

1	Introduction	1
1.1	From Statics to Dynamics	4
1.2	Thesis Outline	7
2	Emergence of Cooperative Dynamics	10
2.1	Dimer Background	11
2.2	Model	12
2.2.1	Staggered Configurations and Strings	13
2.2.2	Height Mapping	14
2.2.3	Dynamics	14
2.2.4	An Outline of Continuous Time Monte Carlo	15
2.2.5	Continuous Height-Field Theory	17
2.3	String dynamics	18
2.3.1	Transfer-Matrix Calculation of String Configurations	18
2.3.2	Flippability	21
2.3.3	String Width	22
2.3.4	Edwards–Wilkinson Equation	23
2.4	Dynamical Tests	27

2.4.1	Height Correlation Function	28
2.4.2	Persistence and Persistence Time	30
2.5	Conclusions	33
3	Large Deviation Methods for Stochastic Dynamics	35
3.1	Large Deviation Theory	36
3.2	s -ensemble	38
3.3	Evidence of Dynamical Behaviours	43
3.4	Numerical Techniques and Sampling Methods	46
3.4.1	Transition Path Sampling	46
3.4.2	Numerical Implementations: Shooting and Shifting	48
3.4.3	Umbrella Sampling with Associated Errors	51
3.4.4	Multistate Bennett Acceptance Ratio	53
3.5	Perfect Sampling Dynamics: Doob Transformation	55
3.6	Effective Doob Dynamics	58
3.6.1	Open 2×2 Approximation	59
3.6.2	Optimising the Effective Doob Dynamics	62
3.6.3	Comparison of Methods	64
4	Using Large Deviations to Study Classical and Quantum Dimers	67
4.1	Rokhsar-Kivelson Hamiltonians	68
4.1.1	s -ensemble Link to Quantum Ground State Properties	70
4.2	The Quantum Phase Diagram	71
4.3	Learning About the Classical Dynamical Phase Diagram from the Quantum Phase Diagram	75
4.3.1	The Minimal Activity State	79
4.4	The Way Classical Dynamics Gives Instruction About the Quantum Ground State	82

4.5	Conclusions	85
5	Trajectory Sampling with Feedback-Augmented TPS	88
5.1	Review of the Cloning Method: Population Dynamics with a Feedback Doob Approximation	89
5.2	Feedback Dynamics for Cloning	91
5.2.1	Cloning Approximations	95
5.3	Feedback Process in TPS	97
5.3.1	Local Approximation for the Dimer Model Transitions	99
5.3.2	Algorithm: Feedback TPS	103
5.4	Numerical Comparisons Between Standard TPS and the Modified Dy- namics Methods	103
5.4.1	A Property of the Feedback Method	110
5.5	Conclusion	112
6	Summary	114

List of Figures

1.1	Examples of a fully packed dimer configuration on the square lattice and the dimer phenomenon arctic circle.	3
2.1	An illustration of the flux of a configuration and a possible height map of that configuration.	13
2.2	An illustration of the local flip dynamics undertaken by a capable plaquette.	15
2.3	example of string structure along with time slices from a trajectory to indicate the movement of the strings and the relaxation effect. . .	18
2.4	The components that are used to make up a string using a transfer matrix.	19
2.5	The results for flippability as a function of flux.	21
2.6	The static results of the string width as a function of displacement along the string.	23
2.7	The results that track the width of the string as it evolves in time. . .	27
2.8	The time evolution of the height field correlation function for various flux sectors and various wavevectors.	29

LIST OF FIGURES

2.9	Logarithm plot of the time evolution of the persistence, including various fits to illustrate time scales.	31
2.10	Logarithm plot of the persistence time as a function of flux.	32
3.1	plot of distributions of activity of trajectories for various starting conditions.	44
3.2	Sample trajectories from the average, high, and low flippable distributions.	45
3.3	Illustration of the transition path sampling method shooting	49
3.4	Illustration of the transition path sampling method shifting	50
3.5	Illustration of all possible dimer occupation of the open 2×2 , excluding rotations.	59
3.6	Comparison of TPS acceptance rates for the modified dynamics and the original dynamics.	63
3.7	Comparison of the various methods to explore the activity against exact results.	65
4.1	Illustration of the different possible phases for the ground state of the quantum dimer model.	73
4.2	Illustration of proposed phase diagrams of the quantum dimer model.	75
4.3	Plot of the activity as a function of s with a subplot showing the true distribution for largest systems size.	77
4.4	Plot of the susceptibility as a function of s	78
4.5	Illustration of the least flippable state.	80
4.6	Illustration of the possible structure of the distribution of magnetisation for different phases of the quantum dimer model.	82
4.7	Results of the distribution of the magnetisation for various values of s	84

4.8	Various plots top indicate the phases of the quantum dimer model for $s < 0$	85
5.1	Illustraiton of trajectories produced using various large deviation techniques.	93
5.2	Illustraion of all possible configurations of the next nearest neighbours of the 2×2 with two dimers in the center plaquette, excluding rotations.	102
5.3	Plot to compare modified dynamical methods.	104
5.4	plot of TPS acceptance rates for comparison of modified dynamics to the original one.	105
5.5	Plots of the distributions of accepted time cuts used in the TPS shifting algorithm for the various dynamical methods.	107
5.6	Comparison of trajectory update lengths for different dynamics system sizes and values of s	108
5.7	Example TPS trajectories for the different dynamical methods.	109
5.8	The evolution of the approximate probability distribution for various values of s	111
6.1	Plot to show how the free energy function of the largest and smallest flux sectors behave.	117

1

Introduction

Equilibrium statistical mechanics is an incredibly successful theory [1, 2]. Physicists often come across problems where the mechanical description becomes too complex to adequately detail how a system behaves. In order to construct a theory to describe how particles behave as part of a larger system, it is typical to start with one particle then introduce another and so until the full system is represented. The problem with studying such many body systems is the increasing difficulty as the particle number increases. This is where purely mechanical approaches can start to struggle to be useful in understanding larger systems. With statistical mechanics techniques this problem can be approached from a different direction, including through the

use of computational methods. This then often reveals the collective behaviours of the system as a whole. These types of collective behaviours can be seen in a variety of fields, from molecular dynamics of fluid systems, to those that are used to understand reaction processes between chemicals [3–5], and even to active matter systems, currently being used to understand animal or bacterial behaviour [6–8].

Cooperative behaviour gives many interesting properties to systems. Emergence of greater structure from the interactions of the smaller components is a fascinating phenomenon. Taking active matter as an example, it can be seen that the development and evolution of a flock or swarm formation is a process that, while relying on the individual components interactions, is truly its own entity and can be understood only through collective behaviours [9]. There is an inherent link between how the microscopic components interact and the overall configuration behaviours. That is to say, the type of collective processes that can be observed for a particular realisation might vary depending on the initial setup of the individual components. This is where statistical mechanics plays an important role. In order to study the effects of the cooperative dynamics that emerges, multiple realizations need to be generated and their statistical properties analysed.

Now is a good point to introduce the fully packed dimer model, as it will be the main model of interest in this thesis. The dimer model used in this thesis will be on a square lattice, $L \times L$, where L is the number of lattice sites along one dimension. A dimer is a rod or a two point particle. Each dimer has one end placed on a lattice site \mathbf{r} such that the other end occupies a neighbouring site $\mathbf{r} + \boldsymbol{\delta}_\mu$, where $\mu \in \{x, y\}$ and $\boldsymbol{\delta}_\mu$ is a lattice vector. The fully packed model only allows for configurations where each site is occupied by exactly one end of a dimer. It is similar to the Ising model except rather than single particles filling the lattice site the dimer occupy two lattice sites or one lattice link, leaving no site unoccupied. An example of a $L = 4$ configuration is shown in Fig. 1.1.a.

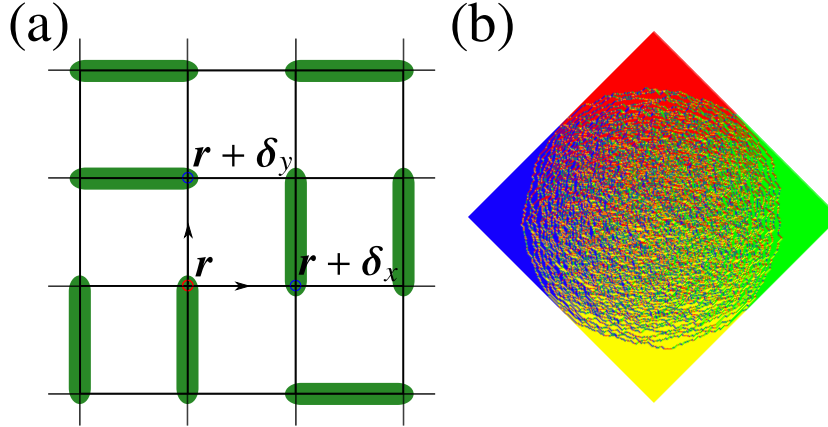


Figure 1.1: (a) is an example of an $L = 4$ configuration of fully packed dimers on the square lattice. We have a total of $N = L^2$ lattice sites, and so there are $N/2$ dimers. (b) is an example of the fully packed dimer model on the Aztec Diamond lattice, showing the emergence of the arctic circle phenomenon. Taken from [13], where a general method of producing configurations of the dimers on the Aztec diamond is outlined.

The methods used in statistical mechanics have allowed for understanding of the statics of many systems. They have often revealed more interesting properties than expected. An example is the fully packed dimer model on an Aztec diamond [10,11]. The Aztec diamond can be created from four staircase shapes and gluing them together at the straight edges. The effect is to appear like a rotated square lattice as seen in Fig. 1.1.b. When configurations are placed on the Aztec diamond a strange phenomenon forms. The areas near the four corners of the lattice appear to be in a staggered formation, which will be detailed latter. Favouring this particular configuration structure gave rise to the appearance of a “frozen” region around a widely varied or “noisy” central circle. This leads to a circular structure as can be seen Fig. 1.1.b, which is dubbed the “arctic circle” [12]. The arctic circle shows how the dimer model, which does not present this behaviour on the normal square lattice, forms a particular structure due to the geometry it has. This is an indicator of why models like the dimer model are worth studying, because of the surprising phenomena that can emerge.

Dimers are quite “versatile” objects to study; they arise in many different fields. From a mathematical point of view, studies such as the arctic circle phenomenon are interesting along with a combinatorics question of how many defect free configurations can exist on the square lattice of a given size [14]. Typically mathematical studies are done with hardcore rectangular dimers which occupy lattices [15]. From a chemical physics angle, the tiling of rigid dimer-shaped molecules, TCPC, onto a graphene surface has been of interest recently [16] due to its formation structure and relaxation properties. In some physical studies, however, a softer version of the dimer is of interest, for example in a packing problem [17].

1.1 From Statics to Dynamics

From a physical point of view, more can be explored than simply the statics. If such an interesting structure, like the arctic circle, can emerge from a static description, it is possible the dimer model is hiding a wealth of behaviours in the dynamics. Glassy systems have been studied in such dynamical contexts for the last few decades [18,19]. For example, kinetically constrained models (KCMs) [20] are simple lattice systems with local constraints in their dynamics that mimic steric restrictions, and are one of the paradigms for the slow relaxation characteristic of glassy systems [21]. It can be seen that even the simplest constraint of hard-core repulsion as in simple exclusion processes can give rise to interesting nonequilibrium dynamics [22]. The classical dimer model (CDM) is a more complex example of a system with such rich dynamical behaviour [23,24]. The configuration dependence on the relaxation time becomes apparent when studying such systems. This is an important theme throughout this thesis: whether it is in understanding how the CDM model uses cooperative behaviour to relax, or when trying to explore the link from the CDM to the ground state of the quantum dimer model (QDM).

The application of thermodynamic methodology to dynamical systems has had

many successes, and is a well studied method. Ruelle paved the way for the adaptation by writing the formalism outlined in [25]; specifically when applying statistical mechanics methods to trajectories of the dynamics. The best way to practically implement this formalism is still being determined [26, 27]. This application has given rise to many studies [28–31] that look into the thermodynamics of trajectories. This work is built on large deviation theory, where trajectory ensembles can be realized. Trajectory ensembles, a collection of trajectories, are similar to configuration ensembles but for the relevant objects for the dynamics. They are used to statistically map out the dynamical behaviours that would otherwise be hidden to thermodynamic studies of configurations. Due to the formalism of this theory being derived from usual statistical mechanics applied to thermodynamics, there are many analogies that can be drawn between the two. These analogies make it easier to understand how and why the trajectory space can be explored. For example, in a system such as the Ising model, a magnetic field can be applied, allowing us to determine how the system moves between different phases, ferromagnetic and paramagnetic. In the trajectory ensemble an analogous field can be considered to control some properties of the trajectories. Typically this dynamical field is not a physically tunable one in the way the magnetic field is for the Ising model, but it allows the same type of analysis of systems that might have different dynamical ordering and could experience a dynamical phase transition as it passes between the two.

The concept of a dynamical phase transition can be quite intuitive when compared to standard thermodynamics. When, for the Ising model, the two phases are explored the configurations are part of ensembles that give information of the static properties on either side of the transition. Quantities such as persistence, which is a way of tracking local relaxation within a system, often used in glassy studies, tells of how the dynamics appear to be affected by a given field being applied. Ultimately, the statistics that can be gathered, builds up a picture of what the different sides of the

phase transition look like. For dynamical transitions we have all of this information, in dynamical terminology, to build up a description of the model using dynamical observables. An important part of this for the dynamical study is information about how, for a given configuration, the local constituents interact with each other. This is fundamentally how the system wants to walk through its configuration space. Depending on the applied field there is a preferred path through configurations. The presence of a dynamical phase transition tells us that there are multiple competing regimes of paths to be taken. This information is not accessible through traditional thermodynamic studies. So unknown, possibly highly influential, behaviours are overshadowed when viewing unbiased samples taken from the original model.

It can be useful to understand the effects of various physical fields in order to see what behaviours they cause and to find how to perturb a system such that it would exhibit desired behaviours in such a field. This useful output of statistical mechanical studies can be extended to the dynamical ones. If we want a system to behave in a particular way during the life time of its trajectory then we could do so by extracting the type of dynamics that the system undergoes when we see the behaviours we desire. This has currently not been possible, as the majority of techniques are limited to exploring the dynamical observable rather than constructing the dynamics used to produce these observables.

So the key questions this thesis will try to ask are as follows: What is the true dynamical behaviour of the CDM? This is an interesting question, as previously it had only been suggested what the full dynamical behaviour was, it had not been rigorously numerical verified. Another question this thesis will try to answer is, can a dynamical study of a classical model be used to explore the ground state of a corresponding quantum model? and subsequently, what is the behaviour of the ground state of that model? It is known that certain models have a linking point between the classical and quantum version. So using techniques from large

deviation studies of the dynamics of the classical model, this thesis will try to show how information about the ground state of the quantum model can be acquired. A natural extension that arose from that question was, how can the dynamical study be optimised? This can be done with priory knowledge of the system in combination with taking the form of a hard to extract large deviation technique known as the Doob transformation, which provides the ideal dynamic. The goal in this thesis was in making it more general, in the sense it can take the form of various less optimal dynamics, which can be arrived at in an easier more automated manor.

1.2 Thesis Outline

This thesis will explore many aspects of the dimer model, both classical and quantum. In Chapter 2 work will be presented that improves on previous attempts made to explore the dynamics of the fully packed dimer model on a square lattice [24]. This was based on a previous static analysis [32], and was a reasonable extension of the static theory to the dynamics. However due to the lack of numerical capability at the time the breakdown of the approach was not apparent, an issue we now have the ability to address and do so in Chapter 2. The description was, therefore, considered a solid one until it could be numerically checked exhaustively. As will be shown in this thesis, it turns out that a better description can be used that breaks down to a lesser degree in a certain parameter regime.

Once the work in Chapter 2 was completed there was a natural direction to take the research. Expanding on the ideas of large deviation methods, it was to explore the various dynamical phases of the CDM, which can be taken further and used to map the CDM to the quantum dimer model (QDM). Chapter 3 explores the background involved in the work of Rokhsar and Kivelson [33] that identified the ground state of the quantum dimer model to the left eigenstate of the master operator for the CDM, and how to move beyond this point, referred to as the Rokhsar-Kivelson (RK) point.

In order for this to be done, large deviation theory and implementations [28] are also presented. These allow for a deformation of the classical master operator in such a way as to permit the mapping at the RK point to continue away from this point.

Chapter 4 will, utilizing the ideas of Chapter 3, show how the large deviation method applied to a RK Hamiltonian can be used to learn about both the classical and the quantum models by studying corresponding properties of each other. Knowledge of the quantum model will provide insight as to how we expect the dynamical properties of the classical model to behave. Equally, the more controversial topic of the (configuration) phases that occupy one side of the quantum phase diagram, can be answered by looking at configurations taken from sample trajectories from the ensemble of trajectories of the classical model [28, 34].

Furthermore Chapter 4 will outline how we dealt with the problem that arose from using current techniques used in large deviation studies [31]. In the process of applying these methods we identified that efficient sampling of trajectories became too difficult. It is known that trying to explore a phase transition by only ever trying to approach it from one direction is difficult. In traditional thermodynamic studies this is not really an issue as the conjugate field to the order parameter is a physical parameter and can simply be added as an extra term in the dynamics. In large deviation sampling techniques the dynamics is typically the unaltered dynamics. Since the basis of these methods is to identify a dynamical phase transition but the approach to exploring this is to do so with dynamics taken from only one side of the transition, it is maybe not surprising that in some hard cases this approach falls short.

In Chapter 5 we explore in detail how we can use alternative dynamics in order to explore the phase diagram much more efficiently. This was clearly a large problem that many people in the community were facing, as many similar solutions have appeared in recent years to try and solve it [26, 27, 30, 35]. We originally designed a

process that, while only governed by a single parameter, would need to be optimized before large data collection could begin. This process worked well and was used to collect the data presented in Chapter 4. Building on Ref. [30] we managed to adapt it into a feedback method, which removed the need to optimize beforehand, though it can be used in conjunction with it.

The final chapter discusses the future avenues to pursue with this method. It has the potential to be wide reaching, with the ability to provide the dynamics that is desired for a given task. We will also outline how we could further explore the dimer model, for example confirming the theoretical prediction for non-bipartite lattices such as the triangular lattice. This could also be extended to dimers in dimensions larger than two.

2

Emergence of Cooperative Dynamics

This Chapter will discuss the classical dimer model (CDM), starting by outlining the numerical technique of Markov chain Monte Carlo (MCMC) which is used to compare to the analytic description of the dynamics of the CDM. The numerical study we undertook was an extensive one, in order to have a detailed description of the CDM. The theory outlined in this chapter tries to go beyond the work of Henley, who took the first steps towards extending a static description of the dimer model to a dynamic one [24]. In order to compare with the dynamic description outlined by Henley, we use a description of the CDM known as height interpretation. We go on to study other dynamical properties such as the persistence, all with the aim of describing the relaxation of the CDM. In the CDM, there emerges a cooperative mechanism, through which the system undergoes relaxation, which the dynamic picture extracted

from the statics misses. We propose a different analytical description of the dynamics that agrees with our numerical observations to a greater degree than expected. The work from this Chapter was published in [23]

2.1 Dimer Background

Dimer models are archetypal systems for the study of the effects of strong local constraints [36,37]. Despite their simplicity, classical dimer models on bipartite lattices, such as the square or bi-bone lattices, exhibit a number of interesting phenomena, such as macroscopic ground-state degeneracy, topological order, and deconfinement of monomers [38]. Their static properties are well understood, and are captured by an effective coarse-grained theory, involving a height field [32] in two dimensions (2D) or an effective gauge field in higher dimensions [39]. In either case, the result is a critical equilibrium phase, with power-law correlations between local degrees of freedom. This class of systems provides the simplest examples of “exotic” thermodynamic behavior purely determined by entropy [38].

Even when the thermodynamic properties of a system have a simple effective description, its dynamics can be more intricate and interesting [40]. It is natural to ask whether this is the case for dynamical extensions of the classical dimer model, about which much less is known. In particular, Ref. [24] considered the simplest extension of the coarse-grained description to dynamics, predicting simple relaxational decay of correlations, while Ref. [41] considered the dimer model with nonlocal loop dynamics. These works should be contrasted with studies of defect-driven dynamics in dimer models [42] and of monopole dynamics in spin ice [43,44].

The first main contribution of this chapter is to demonstrate, using simulations, significant deviations from exponential relaxation in global and local observables over a broad range of time scales. We will argue below that a simple continuum description fails because the dynamics is *facilitated* by local objects—in this case, one-

dimensional *strings* [45–49]—and is hence highly heterogeneous. The understanding of the importance of these objects, which has broad implications for the study of cooperative dynamical phenomena, is our second main contribution.

In this chapter, we consider local stochastic dynamics in the defect-free square-lattice dimer model, using both simulations and theoretical arguments. We show that the simple continuum picture can fail to describe the true physics even over long time scales and that the phenomenology is, in fact, far richer than the simplicity of the model would suggest.

2.2 Model

We study a dimer model on an $L \times L$ square lattice with periodic boundaries. The occupation variable $d_\mu(\mathbf{r})$ gives the number, 0 or 1, of dimers on the link joining sites \mathbf{r} and $\mathbf{r} + \boldsymbol{\delta}_\mu$, where $\mu \in \{x, y\}$ and $\boldsymbol{\delta}_\mu$ is a lattice vector. A configuration is allowed only if every site is occupied by a single dimer. We refer to a plaquette as *flippable* when it contains two parallel dimers. The *flippability* f of a configuration is defined as the proportion of plaquettes that are flippable,

$$f = \frac{1}{L^2} \sum_{\mathbf{r}} \sum_{\mu, \nu \neq \mu} d_\mu(\mathbf{r}) d_\nu(\mathbf{r} + \boldsymbol{\delta}_\nu). \quad (2.1)$$

We define the effective “magnetic field” $B_\mu(\mathbf{r}) = \varepsilon_{\mathbf{r}} [d_\mu(\mathbf{r}) - \frac{1}{4}]$, where $\varepsilon_{\mathbf{r}} = \pm 1$ on the two sublattices [38]. The constraint on dimer configurations then becomes Gauss’ law, $\text{div}_{\mathbf{r}} B = 0$, where

$$\text{div}_{\mathbf{r}} B = \sum_{\mu} [B_\mu(\mathbf{r}) - B_\mu(\mathbf{r} - \boldsymbol{\delta}_\mu)] \quad (2.2)$$

is the lattice divergence. The flux $\boldsymbol{\Phi}$ corresponding to B can be defined by $\Phi_\mu = \sum_{\mathbf{r}} B_\mu(\mathbf{r}) = \sum_{\mathbf{r}} \varepsilon_{\mathbf{r}} d_\mu(\mathbf{r})$, an example of the flux for a specific configuration can be

seen in Fig. 2.1(a). Because $\varepsilon_{\mathbf{r}+\delta_\nu} = -\varepsilon_{\mathbf{r}}$, a pair of neighboring parallel dimers, of either orientation, gives zero net contribution to Φ , so plaquette-flip dynamics, which can be seen in Fig. 2.2, conserves the flux.




Figure 2.1: a) Example of the flux of a configuration. The effective magnetic field is shown by the black arrows. To work out the flux for direction ϕ_x total the number of red arrows pointing left and subtract the number of red arrows pointing right. Similarly for the flux in the ϕ_y total the number of upward pointing blue arrows and subtract the downward blue arrows. A indicates a flippable plaquette which is a plaquette containing 2 dimers, where a plaquette is the square made of 4 lattice sites, and B indicates an unflippable plaquette. b) Example of a possible height map for a configuration. the central plaquette is the origin of the map and is set to zero. The rest of the plaquettes height values are determined by rotating around a lattice site, when crossing a lattice link with no dimer the height is increased by 1, if the link is occupied then the height is reduced by 3. The direction of rotation is determined by the sub-lattice (blue circles or red crosses) being rotated around.

2.2.1 Staggered Configurations and Strings

The flux is maximized by a staggered dimer configuration. For example, if $d_y(\mathbf{r}) = 0$ for all \mathbf{r} , $d_x(\mathbf{r}) = 1$ for $\varepsilon_{\mathbf{r}} = +1$, and $d_x(\mathbf{r}) = 0$ otherwise, then $\Phi = +\Phi_{\max}\hat{x}$, where $\Phi_{\max} = \frac{1}{2}L^2$. The other three staggered configurations, related by symmetry, have flux of the same magnitude, $|\Phi| = \Phi_{\max}$, along other lattice directions. We define

the flux relative to its maximum by $\phi = \Phi/\Phi_{\max}$ and, for flux along the x direction, the deviation from maximum $\vartheta = 1 - \phi_x$.

To reduce the flux from maximum, one can shift a row of dimers spanning the system, which changes Φ_x by $-L$. We refer to such a set of shifted dimers as a string [48, 50]. After the shift, plaquettes along its length become flippable; flipping these deforms the string but conserves the flux. A possible path for a single string and the resulting configuration are shown in Fig. 2.3. With N_s strings introduced into a staggered configuration, we get $\vartheta = 2N_s/L$. The linear density of strings is therefore $\frac{1}{2}\vartheta$.

2.2.2 Height Mapping

The constraint $\text{div}_{\mathbf{r}} B = 0$ can be resolved by defining the *height* z on each plaquette [24], in terms of which $B_{\mu}(\mathbf{r}) = -\frac{1}{4} \text{curl}_{(\mathbf{r}, \mu)} z$, where the curl is the difference between the plaquettes on each side of a link. The example of a height map can be seen in Fig. 2.1(b). Global shifts of z do not affect B , corresponding to the gauge redundancy in 3D [38]. With an appropriate gauge choice, flipping a plaquette modifies z only on that plaquette [24]. If B has periodic boundary conditions, $z(\mathbf{r} + L\boldsymbol{\delta}_{\mu}) = z(\mathbf{r}) + 4L^{-1} \sum_{\nu} \epsilon_{\mu\nu} \Phi_{\nu}$, where ϵ is the Levi-Civita tensor. This is due to a tilt in the height of the configuration, so we need to be aware the height is not necessarily periodic even if the boundaries are. The spatial average of the derivative of the height (the *tilt*) is therefore intensive and proportional to ϕ . We define $\zeta(\mathbf{r}) = z(\mathbf{r}) - 2 \sum_{\mu\nu} \epsilon_{\mu\nu} r_{\mu} \phi_{\nu}$, with periodic boundary conditions. So we have now mapped to a height scale that is periodic.

2.2.3 Dynamics

The most natural dynamics for the dimer system is one where individual plaquettes flip randomly, see Fig. 2.2. This dynamics is efficiently implemented numerically via

continuous time Monte Carlo, in which, when flippable, plaquettes flip according to a Poisson process with rate constant γ . Dynamics at equilibrium within a sector of fixed flux ϕ can be studied by starting from a fixed configuration and equilibrating using plaquette-flip dynamics. We will denote by $\langle \dots \rangle$ an average both over the equilibrium ensemble (where all allowed states with flux ϕ have equal weight) and, where applicable, over subsequent trajectories.




Figure 2.2: Example of the configuration change to occur when a flip happens. The rest of the configuration, not shown here, is unaffected by this process.

The dynamical correlation function of the height is defined by $G_\zeta(\mathbf{q}, t) = \langle \tilde{\zeta}(\mathbf{q}, t) \tilde{\zeta}(-\mathbf{q}, 0) \rangle$ where $\tilde{\zeta}$ is the Fourier transform of ζ . We also consider the persistence $p(t)$, defined as the proportion of plaquettes that have not flipped at any point up to time t , which provides a local probe of the evolution.

2.2.4 An Outline of Continuous Time Monte Carlo

A stochastic process can be thought of as a series of configurations, \mathcal{C} , each surviving for a certain amount of time until a transition occurs to move to a new configuration. A transition between configurations is randomly chosen from a selection of accessible configurations \mathcal{C}' compiling a probability distribution, $P(\mathcal{C}'|\mathcal{C})$, each with its own weighting depending on the current configuration \mathcal{C} . These sequences of configurations are often referred to as *trajectories*. The transition times and transitions are random, so a trajectory is representative of the set of transitions that could have occurred. The trajectories represent the exploration of configuration space of a particular model. Multiple trajectories need to be sampled in order to suppress the

noise produced from the random transitions and extract the governing physical behaviour. The typical way to formulate such a stochastic process is through a master operator

$$\mathbb{W} \equiv \sum_{\mathcal{C}, \mathcal{C}' \neq \mathcal{C}} W(\mathcal{C} \rightarrow \mathcal{C}') |\mathcal{C}'\rangle \langle \mathcal{C}| - \sum_{\mathcal{C}} R(\mathcal{C}) |\mathcal{C}\rangle \langle \mathcal{C}|, \quad (2.3)$$

Where the transitions are given by $W(\mathcal{C} \rightarrow \mathcal{C}')$ and $R(\mathcal{C})$ is the escape rate from configuration \mathcal{C} . The escape rate is also described by the transition rates $R(\mathcal{C}) = \sum_{\mathcal{C}' \neq \mathcal{C}} W(\mathcal{C} \rightarrow \mathcal{C}')$

Our model fits into the definition of a Markov process, with the transitions being those of plaquette flips and the escape rates given by the number of configurations that can be accessed via a flip (i.e. the flippability of the configuration). These processes are only dependent on the configuration the trajectory is currently in and the transitions have no memory of the trajectory's past. Continuous Markov processes can be numerically simulated by continuous time Monte Carlo (CTMC). CTMC works as follows [51];

1) The system is in configuration \mathcal{C}_i . A configuration change to \mathcal{C}_{i+1} is selected from the list of configurations that can be accessed. In a poissonian process each configuration change has equal probability of occurrence, $\omega_{\mathcal{C}_i \rightarrow \mathcal{C}_{i+1}}$.

2) Based on the escape rate, $R(\mathcal{C}_i) = \sum_j \omega_{\mathcal{C}_i \rightarrow \mathcal{C}_j}$, of the current configuration, \mathcal{C}_i , produce a random number $r_t \in [0, 1]$ and calculate the time until the next jump using $\Delta t_i = -\ln(r_t) \frac{1}{R(\mathcal{C}_i)}$.

3) Update the current configuration, $\mathcal{C}_i \rightarrow \mathcal{C}_{i+1}$, and increase the time, $t_i = t_{i+1} + \Delta t_i$. Then repeat from step 1.

The CTMC algorithm allows us to create trajectories that are representative

of the dynamics of the system. Taking sufficient samples of such trajectories will provide us with an apt description of the distribution of our system dynamics.

2.2.5 Continuous Height-Field Theory

The static properties of the dimer model can be described by a continuum theory in terms of the coarse-grained height $h(\mathbf{r})$ resulting from averaging $\zeta(\mathbf{r})$ over short length scales [38]. Neglecting terms irrelevant at long distances, the effective dimensionless free energy is

$$\mathcal{F} = \frac{1}{2} \int d^2\mathbf{r} [K_y(\nabla_x h)^2 + K_x(\nabla_y h)^2] \quad (2.4)$$

for ϕ along $\hat{\mathbf{x}}$, implying correlations $\langle \tilde{h}(\mathbf{q})\tilde{h}(-\mathbf{q}) \rangle = [\omega(\mathbf{q})]^{-1}$ for the Fourier transform \tilde{h} , where $\omega(\mathbf{q}) = K_y q_x^2 + K_x q_y^2$. K_μ is the stiffness constant in direction μ .

The simplest extension of the continuum description to dynamical properties is the Langevin equation [24]

$$\frac{\partial}{\partial t} h(\mathbf{r}, t) = -\Gamma \frac{\delta \mathcal{F}}{\delta h(\mathbf{r}, t)} + \eta_h(\mathbf{r}, t), \quad (2.5)$$

where the noise has correlations $\langle \eta_h(\mathbf{r}, t)\eta_h(\mathbf{r}', t') \rangle = 2\Gamma\delta(\mathbf{r} - \mathbf{r}')\delta(t - t')$, and Γ is a damping coefficient controlling diffusion. The resulting two-time correlations are [24]

$$G_h(\mathbf{q}, t) \equiv \langle \tilde{h}(\mathbf{q}, t)\tilde{h}(-\mathbf{q}, 0) \rangle = \frac{e^{-\Gamma\omega(\mathbf{q})t}}{\omega(\mathbf{q})}, \quad (2.6)$$

implying exponential decay at long time scales. We show below that this prediction can break down, even when the height-field approach is accurate for the statics, due to cooperative effects that dominate the dynamics.

2.3 String dynamics

The collective character of the dynamics can be uncovered by considering the behavior near maximal flux (i.e., for small θ), where most of the system is unflippable. This regime can be understood in terms of a low density of well-separated strings. We first consider the dynamics of a single string, using a continuum description based on the Edwards–Wilkinson equation, before turning to the consequences for the two classes of observables, correlation functions and the persistence.




Figure 2.3: Left: A dimer configuration with a single string, relative to a fully staggered configuration with maximal flux $\Phi_x = \frac{1}{2}L^2$ along the horizontal direction. A single string spanning the system once reduces the flux Φ_x by $-L$, the smallest possible amount, irrespective of its path. Flippable plaquettes, which appear when the dimers are shifted, are marked with stars (\star); the fully staggered configuration has none. Center and right: Configurations with $\phi_x = \frac{1}{4}$ evolved for time t . Persistent plaquettes are white, while those that have flipped are blue; strings are yellow. The dynamics is spatially heterogeneous: Even at relatively long times, extended regions are unvisited by strings and hence, they are persistent.

2.3.1 Transfer-Matrix Calculation of String Configurations

A string can be divided into four types of segment, illustrated in Fig. 2.4; the ensemble \mathcal{C}_L^s of configurations for a single string is given by the set of ways in which these segments can be combined to produce a closed path of length L . One can write a




Figure 2.4: Four types of segment that can be combined to form a string. The first two, type α , are steps at which the vertical position of the string changes by ± 1 , while the last two, type β , are the two segments of a horizontal step. (This step is split so that all four segments involve the same horizontal displacement.)

transfer matrix

$$\mathbf{T}(k, \mu) = \begin{pmatrix} e^{-ik} & e^{-ik}e^{-\mu} & 0 & e^{-ik} \\ e^{ik}e^{-\mu} & e^{ik} & 0 & e^{ik} \\ 1 & 1 & 0 & 1 \\ 0 & 0 & e^{-2\mu} & 0 \end{pmatrix}, \quad (2.7)$$

such that $T_{\beta\alpha}(k, \mu)$ is nonzero only when a segment of type α can be followed by one of type β (labeled according to the order in Fig. 2.4). Each successive pair of segments is weighted by $e^{-\mu}$ for every flippable plaquette it produces and by e^{-ik} for every step in the positive y direction. We define the partition function \mathcal{Z}_L° as the weighted sum over all *open* paths of length L with *any* net vertical displacement y ,

$$\mathcal{Z}_L^\circ(k, \mu) = \sum_{\mathcal{C}_L^\circ} e^{-iky} e^{-\mu N_f}, \quad (2.8)$$

where \mathcal{C}_L° denotes the ensemble of such paths and N_f is the number of flippable plaquettes in the resulting dimer configuration. Summing over all sequences of path segments and (vertical) starting positions gives ¹

$$\mathcal{Z}_L^\circ(k, \mu) = \frac{L}{2} \text{Tr}[\mathbf{T}(k, \mu)]^L = \frac{L}{2} \sum_{\lambda \in \sigma_{\mathbf{T}(k, \mu)}} \lambda^L, \quad (2.9)$$

¹Summing over all starting positions for paths results in a factor of L , but this effectively includes both (horizontal) fully staggered configurations; the factor of $\frac{1}{2}$ corrects the double counting.

where the sum is over the set $\sigma_{\mathbf{T}(k,\mu)}$ of eigenvalues λ of $\mathbf{T}(k,\mu)$. Since this is a partition function we require it to be real. To check that it is we can compute the expectation $E_c = \langle e^{-iky} \rangle = \int dy e^{-y^2/2\sigma} e^{-iky}$ and compare it to the complex conjugate of this expectation, $E_{c^*} = \langle e^{iky} \rangle = \int dy e^{-y^2/2\sigma} e^{iky}$. Due to the symmetry of the paths, the number of paths are equal if we take $y \rightarrow -y$, we can therefore see that $E_c = E_{c^*}$. This is only true if the eigenvalues λ are real.

The allowed paths for a single string are those that return to their starting point after winding once around the system, and hence have net displacement $y = 0$. The partition function for such paths is

$$\mathcal{Z}_L^s(\mu) = \sum_{\mathcal{C}_L^s} e^{-\mu N_f} = \sum_{\mathcal{C}_L^o} \delta_{y,0} e^{-\mu N_f} = \int_{-\pi}^{\pi} \frac{dk}{2\pi} \sum_{\mathcal{C}_L^o} e^{-iky} e^{-\mu N_f} = \frac{L}{2} \int_{-\pi}^{\pi} \frac{dk}{2\pi} \sum_{\lambda \in \sigma_{\mathbf{T}(k,\mu)}} \lambda^L. \quad (2.10)$$

For large L , the saddle-point approximation gives

$$\ln \mathcal{Z}_L^s(\mu) = L \ln |\lambda_{\max}(\mu)| + \mathcal{O}(\ln L), \quad (2.11)$$

where

$$|\lambda_{\max}(\mu)| = \max_k |\lambda_{\max}(k, \mu)| = \max_k \max\{|\lambda| : \lambda \in \sigma_{\mathbf{T}(k,\mu)}\} \quad (2.12)$$

is the largest eigenvalue (by magnitude) of $\mathbf{T}(k,\mu)$ for any k . The maxima are $\lambda_{\max}(\mu) = 1 + \sqrt{2} - \sqrt{2}\mu + \mathcal{O}(\mu^2)$, occurring at the points $k = 0$ and $\pm\pi$, and hence

$$\frac{1}{L} \ln \mathcal{Z}_L^s(\mu) \simeq \ln(1 + \sqrt{2}) - (2 - \sqrt{2})\mu. \quad (2.13)$$

Setting $\mu = 0$ gives the entropy of a single string,

$$S_L^s = \ln \mathcal{Z}_L^s(0) = L \ln(1 + \sqrt{2}) + \mathcal{O}(\ln L). \quad (2.14)$$

2.3.2 Flippability

The mean number of flippable plaquettes in the presence of a single string is given by

$$N_f^s = -\frac{d}{d\mu} \ln \mathcal{Z}_L^s(\mu) \Big|_{\mu=0} = (2 - \sqrt{2})L + \mathcal{O}(\ln L). \quad (2.15)$$

At flux $\phi = \{1 - \vartheta, 0\}$, the number of strings is $N_s = \frac{1}{2}L\theta$. If the strings can be treated as approximately independent, as expected for sufficiently small θ , then the number of flippable plaquettes is simply $N_s N_f^s$, and so the mean flippability in equilibrium is

$$\langle f \rangle = \left(1 - \frac{1}{\sqrt{2}}\right)\theta + \mathcal{O}\left(\frac{\ln L}{L}\right). \quad (2.16)$$




Figure 2.5: Mean flippability in equilibrium $\langle f \rangle$ as a function of deviation from maximum flux, $\vartheta = 1 - \phi_x$. The solid (green) line shows Monte Carlo results for system size $L = 256$ (error bars are smaller than symbols), while the dashed (black) line shows the analytical prediction from Eq. (2.16).

Numerical results, shown in Fig. 2.5, confirm Eq. (2.16) in the limit of small ϑ and are in approximate agreement even for fairly large ϑ , suggesting that the independent-string picture is reasonable, (the logarithmic corrections modify the coefficient of ϑ , and are of the order of a few percent for $L = 256$).

2.3.3 String Width

The width distribution of a single string can be determined by a similar approach: The probability distribution for the net vertical displacement Y of a section of string with horizontal extent X is ²

$$P^s(X, Y) = \frac{\sum_{\mathcal{C}_X^o} \delta_{y, Y}}{\mathcal{Z}_X^o(0, 0)} = \int_{-\pi}^{\pi} \frac{dk}{2\pi} e^{ikY} \frac{\sum_{\lambda \in \sigma_{\mathbf{T}(k, 0)}} \lambda^X}{\sum_{\lambda \in \sigma_{\mathbf{T}(0, 0)}} \lambda^X}. \quad (2.17)$$

For large X , the ratio of sums can be found by Taylor expanding $\lambda_{\max}(k, 0)$ around its maxima, giving a pair of Gaussians of variance $\frac{\sqrt{2}}{X}$ centred at $k = 0$ and $\pm\pi$. Taking the Fourier transform, one finds that $P^s(X, Y)$ is given by a normal distribution of variance $\frac{1}{\sqrt{2}}X$ when X and Y have the same parity, and vanishes otherwise (as required by the structure of a string).

If, for a single string traversing the system in the horizontal direction, we denote by $y(x)$ the vertical position at horizontal position x , this result can be restated as follows: At length scales much larger than the lattice scale, but smaller than the system size, $1 \ll |x - x'| \ll L$, the vertical displacement $y(x) - y(x')$ is normally distributed with zero mean and

$$\langle [y(x) - y(x')]^2 \rangle = \frac{1}{\sqrt{2}} |x - x'|. \quad (2.18)$$

This result is confirmed by the small-displacement limit in Fig. 2.6.

²The ensemble \mathcal{C}_X^o has the additional constraint that the last and first segments of the open path are compatible, which allows \mathcal{Z}_X^o to be expressed as a trace. This has no effect on the large- X behavior of $P^s(X, Y)$.




Figure 2.6: Mean-square transverse displacement of a string $y(x, t)$ in equilibrium, given by the equal-time correlation $\langle [y(x, t) - y(x', t)]^2 \rangle$, as a function of displacement along the string, $|x - x'|$. The solid (green) line shows Monte Carlo results using a dimer configuration containing a single string in a system of size $L = 1024$. (Error bars are smaller than symbols.) The dashed (black) line shows the analytical result of Eq. (2.29), with $D = 1/\sqrt{2}$ determined using Eq. (2.18).

2.3.4 Edwards–Wilkinson Equation

We are still interested in the perturbations in the “surface” of the systems, as in [24]. At large lengths and time scales we expect $y(x, t)$, which can be considered as a description of the surface, to obey the Edwards–Wilkinson equation [52, 53],

$$\frac{\partial}{\partial t} y(x, t) = \frac{1}{2} \Lambda \frac{\partial^2}{\partial x^2} y(x, t) + \eta(x, t), \quad (2.19)$$

where $\langle \eta(x, t) \eta(x', t') \rangle_0 = D \Lambda \delta(x - x') \delta(t - t')$, and Λ and D parameterize, respectively, the stiffness of the string and the strength of the noise. The average $\langle \dots \rangle_0$ is taken over trajectories starting from a given initial configuration $y(x, 0)$.

The Green function $G(x, s)$, which is defined as any solution of $L_d G(x, s) = \delta(x - s)$ with L_d being a linear differential operator, can be used to solve differential equations of the form $L_d \Delta(x) = f(x)$. Accounting for the periodicity in the x

direction (but not in y), the Greens function for Eq. (2.19) is

$$\Delta(x, t) = \frac{1}{L} \sum_k e^{ikx} e^{-\frac{1}{2}\Lambda tk^2}, \quad (2.20)$$

where $kL/2\pi \in \mathbb{Z}$. The Edwards–Wilkinson equation, Eq. (2.19), has a general solution (for $t \geq 0$)

$$y(x, t) = \int_0^L dx' \Delta(x - x', t) y(x', 0) \quad (2.21)$$

$$+ \int_0^L dx' \int_0^\infty dt' \Delta(x - x', t - t') \eta(x', t'). \quad (2.22)$$

The two-time correlation function, within an ensemble of trajectories with fixed initial configuration, is therefore given by

$$\begin{aligned} \langle [y(x, t) - y(x', t')]^2 \rangle_0 &= \left\{ \int_0^L dx'' [\Delta(x - x'', t) - \Delta(x' - x'', t')] y(x'', 0) \right\}^2 \\ &+ D\Lambda \int_0^L dx'' \int_0^\infty dt'' [\Delta(x - x'', t - t'') - \Delta(x' - x'', t' - t'')]^2. \end{aligned} \quad (2.23)$$

To calculate the equivalent correlation function in an equilibrium ensemble, we take both times to infinity while keeping their difference finite:

$$\langle [y(x, t) - y(x', t')]^2 \rangle = \lim_{t_0 \rightarrow \infty} \langle [y(x, t_0 + t) - y(x', t_0 + t')]^2 \rangle_0. \quad (2.24)$$

In this limit, $\Delta(x, t_0 + t) = L^{-1}$, and so the first term in Eq. (2.23), which depends on the initial configuration, vanishes. The integrals in the second term can be performed to give

$$\langle [y(x, t) - y(x', t')]^2 \rangle = V_L(|x - x'|, |t - t'|), \quad (2.25)$$

where V_L is the two-time correlation function in the equilibrium ensemble, where we

take both times to infinity with their difference finite,

$$V_L(x, t) \equiv \langle [y(x, t) - y(0, 0)]^2 \rangle \quad (2.26)$$

$$\begin{aligned} &= \lim_{t_0 \rightarrow \infty} \langle [y(x, t_0 + t) - y(0, t_0)]^2 \rangle_0 \\ &= \frac{D\Lambda t}{L} + \frac{2D}{L} \sum_{k \neq 0} \frac{1 - e^{ikx} e^{-\frac{1}{2}\Lambda t k^2}}{k^2}. \end{aligned} \quad (2.27)$$

For $t = 0$, the sum in Eq. (2.27) can be evaluated exactly, using

$$\sum_{n=1}^{\infty} \frac{1 - \cos n\vartheta}{n^2} = \frac{\vartheta(2\pi - \vartheta)}{4}, \quad (2.28)$$

for $0 \leq \vartheta \leq 2\pi$, to give

$$\langle [y(x, t) - y(0, t)]^2 \rangle = V_L(x, 0) = Dx \left(1 - \frac{x}{L}\right), \quad (2.29)$$

for $0 \leq x \leq L$. Comparison with the microscopic result of Eq. (2.18) fixes $D = 1/\sqrt{2}$. In Fig. 2.6, both Eq. (2.29) and the value of D are confirmed using results of Monte Carlo (MC) simulations of the dimer model.

The typical width of a string in equilibrium can be characterized by the mean-square displacement between the points x and 0 at equal time, which is given, for $0 \leq x \leq L$, by $V_L(x, 0) = Dx(1 - x/L)$. The time scales for dynamics can similarly be understood through $V_L(0, t)$, which is shown in the following section to obey

$$\langle [y(0, t) - y(0, 0)]^2 \rangle \approx \begin{cases} D\sqrt{\frac{2\Lambda t}{\pi}} & \text{for } \Lambda t \ll L^2 \\ \frac{D}{L}\Lambda t & \text{for } \Lambda t \gg L^2. \end{cases} \quad (2.30)$$

The short-time result, $\Lambda t \ll L^2$, gives the dynamical scaling relation between the characteristic length in the y direction and time through the ‘‘growth exponent’’

β [53],

$$l_y \sim t^\beta, \quad \beta = \frac{1}{4}. \quad (2.31)$$

At long times, $\Lambda t \gg L^2$, the whole string can be treated as a random walker, with effective diffusion constant $D\Lambda/L$.

These results, along with those from the string microscopics, fix the values of D and Λ . Comparison of $V_L(0, t)$ with the exact result for the equal-time displacement variance gives $D = 1/\sqrt{2}$. The mean rate of plaquette flips in equilibrium is γN_f . Since each flip changes the mean vertical position by $\pm L^{-1}$, the variance of the total shift is $\gamma t N_f / L^2$ in the long-time limit. Comparison of Eqs. (2.15) and (2.30) therefore gives $\Lambda = 2(\sqrt{2} - 1)\gamma$.

In calculating $V_L(0, t)$ for small t/L^2 , the sum can be replaced by an integral,

$$\sum_{n=1}^{\infty} \frac{1 - e^{-\frac{1}{2}n^2\tau}}{n^2} \simeq \int_0^{\infty} dn \frac{1 - e^{-\frac{1}{2}n^2\tau}}{n^2} = \sqrt{\frac{\pi}{2}} \tau, \quad (2.32)$$

giving

$$\langle [y(x, t) - y(x, 0)]^2 \rangle = V_L(0, t) = D \sqrt{\frac{2\Lambda t}{\pi}}, \quad (2.33)$$

for $\Lambda t \ll L^2$. Eq. (2.16) can be used to fix $\Lambda = 2(\sqrt{2} - 1)\gamma$. This result, including the value of Λ , is confirmed using MC results in Fig. 2.7.

At large t and $x = 0$, the first term in Eq. (2.27) dominates, and so

$$\langle [y(x, t) - y(x, 0)]^2 \rangle = V_L(0, t) = \frac{D}{L} \Lambda t, \quad (2.34)$$

for $\Lambda t \gg L^2$.




Figure 2.7: Growth of mean-square transverse displacement of a string, $\langle [y(x, t) - y(x, 0)]^2 \rangle$, at a fixed position x , as a function of time t . The solid (green) line shows Monte Carlo results using a dimer configuration containing a single string in a system of size $L = 1024$. (Error bars are smaller than symbols.) The dashed (black) line shows the analytical result of Eq. (2.33), using the values of D and Λ fixed using Eqs. (2.18) and (2.16).

2.4 Dynamical Tests

Having probed how the strings should behave according to the coarse-grained theory and identified that these behaviors can be observed through numerical simulations, we can move on to test the dynamical properties of the dimer model, rather than the static or string based observables as we discussed in 2.2. We start with the height correlation function, as this was used primarily in the extension of the static description [24]. Then we will show the predictions and results for the persistence, which is an observable more commonly used when studying the relaxation of glassy type systems.

2.4.1 Height Correlation Function

To calculate height correlations based on the coarse-grained string description, we write

$$\nabla_y h(\mathbf{r}, t) \propto \int_0^L dx (\delta^2(\mathbf{r} - \{x, y(x, t)\}) - \frac{1}{L}), \quad (2.35)$$

which treats the string as a step in the height plus a uniform gradient to preserve the boundary conditions. The height correlation function for $\mathbf{q} \neq \mathbf{0}$ can then be written, for a single string, as $G_h(\mathbf{q}, t) \propto (Lq_y^2)^{-1} C_s(\mathbf{q}, t)$, where

$$C_s(\mathbf{q}, t) = \frac{1}{L} \int_0^L dx \int_0^L dx' e^{-iq_x(x-x')} \left\langle e^{-iq_y[y(x,t)-y(x',0)]} \right\rangle.$$

For small-enough density $\frac{1}{2}\vartheta$, string contributions add incoherently, resulting in $G_h(\mathbf{q}, t) \propto \vartheta q_y^{-2} C_s(\mathbf{q}, t)$. Since $y(x, t) - y(x', 0)$ is Gaussian distributed with zero mean, we get

$$C_s(\mathbf{q}, t) = \int_{-\frac{L}{2}}^{\frac{L}{2}} dx e^{-iq_x x} e^{-\frac{1}{2}q_y^2 V_L(x, t)}, \quad (2.36)$$

where the periodicity of $V_L(x, t)$ under $x \rightarrow x \pm L$ has been used to shift the limits of integration.

Asymptotic expressions for the correlations can be found in various limits. For the static correlations, $G_h(\mathbf{q}, 0) \propto \vartheta/\omega_s(\mathbf{q})$, where $\omega_s(\mathbf{q}) = q_x^2 + \frac{1}{8}q_y^4$ in the thermodynamic limit. For $t \ll L^2$, we find time dependence

$$\frac{G_h(\mathbf{q}, t)}{G_h(\mathbf{q}, 0)} \propto \begin{cases} \exp[-\frac{1}{2}\Lambda t \omega_s(\mathbf{q})] & \text{for } \Lambda t \ll q_y^{-1/\beta} \\ t^\beta \exp[-\sqrt{\Lambda t \tilde{\omega}_s(\mathbf{q})}] & \text{for } q_y^{-1/\beta} \ll \Lambda t, \end{cases} \quad (2.37)$$

where $\tilde{\omega}_s(\mathbf{q}) = (4\pi)^{-1}q_y^4 \exp[2(\operatorname{erfi}^{-1} \sqrt{8}q_x/q_y^2)^2]$. A crossover from simple to stretched exponential therefore occurs at $\Lambda t \sim q_y^{-1/\beta}$, which are the time scale corresponding to the wavelength $\sim q_y^{-1}$, according to Eq. (2.31). The stretching is the result of contributions from the continuum of modes of the string. These expressions, along

with the full result found by numerical integration of Eq. (2.36), are compared with simulations in Fig. 2.8. We find close agreement, with no adjustable parameters, at large flux and small wavevector, and qualitative agreement at smaller ϕ_x and larger q .

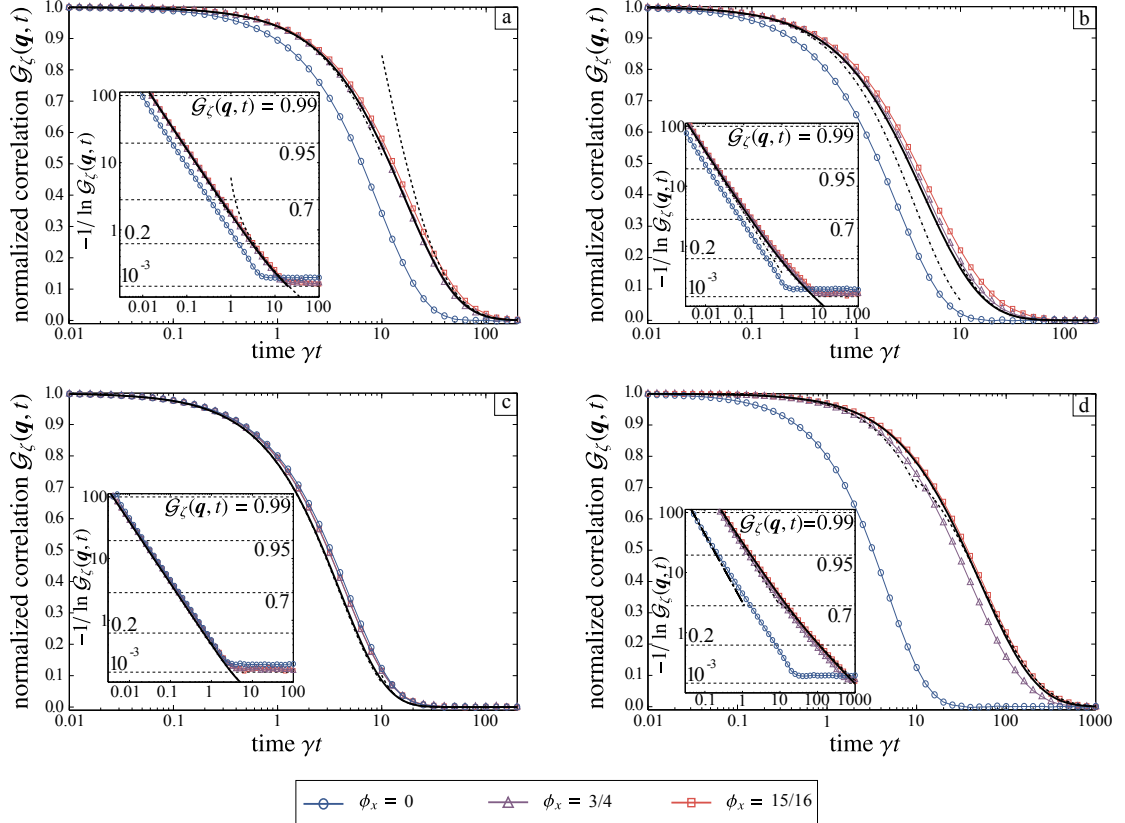


Figure 2.8: Normalized height correlations $\mathcal{G}_\zeta(\mathbf{q}, t) = G_\zeta(\mathbf{q}, t)/G_\zeta(\mathbf{q}, 0)$ at flux $\phi = \{\phi_x, 0\}$ and wavevectors (a) $\mathbf{q} = \{\frac{\pi}{8}, \frac{\pi}{8}\}$, (b) $\mathbf{q} = \{\frac{\pi}{4}, \frac{\pi}{4}\}$, (c) $\mathbf{q} = \{\frac{\pi}{4}, \frac{\pi}{16}\}$, and (d) $\mathbf{q} = \{\frac{\pi}{16}, \frac{\pi}{4}\}$. In each case, the symbols show MC results with system size $L = 256$ (error bars are smaller than symbols), while the thick (black) line shows the theoretical prediction for the large-flux limit (i.e., close to the maximum $\phi_x = 1$). The short- and long-time limits are shown with dot-dashed and dashed lines, respectively. Insets: Same data with double-logarithmic vertical scale; stretched exponentials appear as straight lines. The saturation at long times is an artifact resulting from the statistical uncertainty.

The independent-string approximation should be valid for $\Lambda t \ll \vartheta^{-1/\beta}$, where the time corresponding to a y displacement is equal to the mean string separation. The stretched-exponential form therefore applies up to a time that diverges at $\theta = 0$.

For larger t , string interactions are important for the dynamics, and Eq. (2.37) is no longer valid. When $\Lambda t \gg \vartheta^{-1/\beta}$ many strings contribute, their discreteness becomes unimportant, so we expect a crossover to the Coulomb-phase behavior of Eq. (2.6), which is described through the random noise in the height field [24] rather than the string description.

Our analytical results for the correlations are based on a coarse-grained description of the strings, and so are expected to be quantitatively accurate only for small \mathbf{q} . The values of \mathbf{q} used in Fig. 2.8 are chosen to show the stretching most clearly on time scales accessible in the CTMC simulations. We also require $Lq_y \gg 1$, because we have neglected the periodicity in the y direction.

Results for other values of \mathbf{q} are shown in Fig. 2.8. As expected, the quantitative accuracy of the analytical results decreases as $|\mathbf{q}|$ is increased. Clear evidence of stretched-exponential decay is visible, as a decreased slope on a double-logarithmic scale, only for the larger values of q_y .

2.4.2 Persistence and Persistence Time

Given a trajectory, a plaquette is referred to as “persistent” if it has not flipped at any point during the trajectory. At each time t during the trajectory, the persistence is defined as the proportion of plaquettes that are persistent, i.e.,

$$p(t) = \frac{1}{N} \sum_{\mathbf{r}} \begin{cases} 1 & \text{if plaquette } \mathbf{r} \text{ is persistent} \\ 0 & \text{otherwise.} \end{cases} \quad (2.38)$$

The persistence time τ_p is the average, over starting configurations and trajectories, of the integral of the persistence,

$$\tau_p = \left\langle \int_0^\infty dt p(t) \right\rangle. \quad (2.39)$$




Figure 2.9: MC results for persistence $p(t)$ at flux ϕ_x labeled as in Fig. 2.8, except where indicated. (Error bars are smaller than symbols.) The scale chosen for the y-axis is unusual, but allows for exponentials, stretched or not, to be seen as straight lines. System size $L = 4096$ is required to reach flux $\phi_x = 2047/2048$ (see Section 2.2.1); for the other flux values, $L = 256$. As expected based on the string picture, an initial exponential decay (straight line with slope 1) is followed by a stretched exponential (slope < 1), lasting until $\Lambda t \sim (1 - \phi_x)^{-1/\beta}$. The dashed line at the right, with slope $1/4$, is the function $p(t) \propto e^{-ct^{1/4}}$, in agreement with Eq. (2.40) for the limit $\phi_x \rightarrow 1$.

The persistence $p(t)$ can similarly be understood in terms of the behavior of strings. At very short times, $p(t) = e^{-\langle f \rangle \gamma t}$: plaquettes which are flippable at $t = 0$ will flip independently with rate γ . For times $\gamma t \gg 1$, each point x on a string performs a subdiffusive random walk, according to Eq. (2.30). As only plaquettes adjacent to strings are flippable, the mean persistence is equal to the probability that a plaquette is not reached by any string up to time t . A plaquette's persistence is therefore given by the survival probability for a stationary target in the presence of a density $\frac{1}{2}\vartheta$ of subdiffusive traps. Using the results of [54], relating the dynamic exponent to the persistence, we get

$$\langle p(t) \rangle \propto \exp \left[-\frac{\vartheta}{\Gamma(5/4)} (\Lambda t)^\beta \right] \text{ for } \gamma^{-1} \ll t \ll \Lambda^{-1} \vartheta^{-1/\beta}. \quad (2.40)$$

This form ceases to apply for $\Lambda t \sim \vartheta^{-1/\beta}$, when string interactions become important, or for $\Lambda t \sim L^2$, beyond which the long-time behavior in Eq. (2.30) applies and $\langle p(t) \rangle \propto e^{-ct^{1/2}}$.




Figure 2.10: Inverse persistence time τ_p^{-1} as a function of deviation from maximum flux, $\vartheta = 1 - \phi_x$. The symbols show Monte Carlo results for system sizes $L = 64$ (red), 128 (blue), and 256 (green). (Error bars are smaller than symbols.) The dashed (black) line shows the analytical prediction $\tau_p \sim \vartheta^{-4}$, which applies for $0.05 \lesssim \vartheta \lesssim 0.2$ (close, but not too close, to maximum flux). For larger θ , the density of strings is sufficiently high that their interactions become important, and the picture of independent strings breaks down. When $\vartheta \lesssim L^{-1}$, finite-size effects become important.

Our simulation results, Fig. 2.9, are in qualitative agreement with these arguments, showing an initial exponential followed by a stretched exponential, for all fluxes. The stretching exponent decreases continuously with ϕ_x , approaching $\beta = 1/4$, in agreement with Eq. (2.40), as $\phi_x \rightarrow 1$. At longer times, a faster decay, consistent with a single exponential, is observed. According to Eq. (2.30), in the thermodynamic limit, the typical spread of $y(x, t)$ is proportional to $t^{1/4}$. The typical time to reach a plaquette at a distance ℓ from the starting position of the string is therefore ℓ^4 . Since a plaquette can only flip when a string is nearby, its persistence time is given by the time at which a string first reaches it. The linear density of

strings is $\sim \vartheta$, and so the typical distance from a plaquette to the nearest string is $\sim \vartheta^{-1}$. The typical persistence time is therefore $\sim \vartheta^{-4}$.

In Fig. 2.10, this prediction is confirmed using MC simulations for intermediate values of ϑ . For larger ϑ , the string density is sufficiently high that interactions between strings become important, while for the smallest values of ϑ , the distance from a plaquette to its nearest string is bounded by the system size L .

2.5 Conclusions

We have shown that the close-packed square-lattice dimer model, subject to local, plaquette-flip dynamics, displays emergent collective relaxation that is not anticipated by simple extensions of its static properties. Approximations to the dynamics based on free-energy gradients plus noise, such as Eq. (2.5), fail to capture the intrinsic heterogeneity: Due to the constrained nature of the system, motion is only allowed in the vicinity of strings, and relaxation is dominated by spatial fluctuations. In a sense, the noise that triggers rearrangements is not uniform in space and time; rather, its strength depends sensitively on the local configuration. Strings *facilitate* local rearrangements, dynamics is heterogeneous and collective (see Fig. 2.3), and relaxation functions are nonexponential.

This situation is reminiscent of glass-forming systems [18]: In a slowly relaxing material such as a glass former, “facilitation” indicates the fact that local relaxation can occur only near an already locally relaxing region [55]. Similarly, in the dimer model, plaquette moves are only possible in the vicinity of a string. This is the reason that the Langevin dynamics of Eq. (2.5) is not accurate for relaxation in regions where string density is low. The additive noise assumed in that approximation would allow rearrangements to occur anywhere in space. However if dynamics is facilitated, the noise that drives local rearrangements is not uniformly distributed, but rather concentrated near already mobile regions. A Langevin description, along

the lines of Eq. (2.5), would therefore require a form of noise that is multiplicative and whose magnitude is strongly dependent on the local flippability.

Instead we have developed a string description of the dynamics, which directly incorporates the local nature of the relaxation. At low string density (high flux), we are able to make exact theoretical predictions for the correlations and persistence, which are confirmed by our simulations. We in fact find that much of the qualitative behavior is unchanged at smaller flux, where interactions between strings are certainly important. This indicates that the usefulness of the string picture, as well as the concept of facilitated and heterogeneous dynamics, extends well beyond the regime of high flux.

Besides their fundamental importance, our results are likely to be of relevance to spin ice, where closely analogous string excitations have been evidenced directly using neutron scattering [56], and where correlations with stretched-exponential decay have been noted [57]. Dynamical results for classical dimers are also relevant to the corresponding quantum dimer model at its Rokhsar–Kivelson point [58, 59].

3

Large Deviation Methods for Stochastic Dynamics

Equilibrium studies using statistical mechanics have illuminated many interesting phenomena. Chapter 2 is a prime example of how seemingly simple models can contain a multitude of unusual behaviours. These studies are often restricted to static or time integrated properties of the system, which means that certain influential behaviours of the non-equilibrium or rare aspects of the system are neglected. As was shown in the previous chapter the full picture could not be clearly seen by only considering the static description. In order to get a more complete description of the non-equilibrium behaviour of the system, the full dynamics has to be taken into account.

The mathematical formalism developed by Ruelle [25] allows for the application of large deviation theory [60] to the dynamical behaviours of many systems, allowing the tools of statistical mechanics to shine light on the dynamics of many body systems. This kind of analysis is often referred to as the “thermodynamics of trajectories” [28, 29, 34, 61–65].

This chapter will outline many of the techniques that will be used in the remainder of this thesis. Starting with a summary of large deviation theory, we can motivate the connection between the work in the previous chapter to using large deviation techniques to link the classical and quantum dimer models. This will be done by numerically exploring the so called s -ensemble [28, 29]. As the model we study was already introduced in great detail in the previous chapter, we will mainly discuss numerical tools necessary for studying rare fluctuations such as transition path sampling (TPS), multistate Bennett acceptance ratio (MBAR) and Umbrella sampling.

3.1 Large Deviation Theory

Large deviation theory (LDT) [60] is a way of exploring the exponentially small tails of a distribution. Lets consider n independent samples X_0, X_1, \dots, X_n which have all been sampled from the distribution $P(X_i)$. If we take the example where $P(X_i)$ has the form of a Gaussian distribution we have

$$P(X_i = x_i) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x_i - \mu)^2}{2\sigma^2}}, \quad (3.1)$$

where μ and σ^2 are the mean and variance of the distribution. If we are interested in the probability distribution of the sample mean $J_n = \frac{1}{n} \sum_i^n X_i$ we can write its as

$$P(J_n = j) = \sqrt{\frac{n}{2\pi\sigma^2}} e^{-\frac{n(j - \mu)^2}{2\sigma^2}}. \quad (3.2)$$

If we now take the $n \rightarrow \infty$ limit then we can see the largest contributing factor is that of the exponential, which lets us write

$$P(J_n = j) \asymp e^{-n\varphi(j)}, \quad (3.3)$$

where \asymp denotes the large n limit. We have defined $\varphi(j) = \frac{(j-\mu)^2}{2\sigma^2}$, which we call the large deviation rate function (LDF). With this we can use the Gärtner-Ellis theorem, which says if a LDF exists it can be related to the scaled cumulant generating function (SCGF) [60]. This is one of the most useful properties of LDT for work done in this thesis. The SCGF is the natural logarithm of the moment generating function (MGF), which provides a different description of the probability function. The MGF for an arbitrary observable (lets use the sample mean from above J_n) is given by

$$Z(s) = \langle e^{-nsJ_n} \rangle = \int dJ_n e^{-nsJ_n} P(J_n). \quad (3.4)$$

s can be understood as a real parameter that is the conjugate to J_n . In the case where the probability density is continuous, the MGF is simply a Laplace transformation of the probability distribution.

When the Gärtner-Ellis theorem is applied to a set of independent and identically distributed random variables, a well known result from probability theory, known as Cramér's theorem, can be used which allows us to take the large n limit for the SCGF

$$\theta(s) = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \int dJ_n e^{-nsJ_n} P(J_n) \asymp \frac{1}{n} \ln \int dJ_n e^{-n(sJ_n + \varphi(J_n))}, \quad (3.5)$$

from which we can use a saddle point approximation to say

$$\theta(s) = -\min_J [\varphi(J) + sJ], \quad (3.6)$$

and

$$\varphi(J) = -\min_s [\theta(s) + sJ]. \quad (3.7)$$

Continuing our example of the independent and identically distributed (IID) samples from a Gaussian we are allowed to write

$$\theta(s) = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \left\langle e^{-s \sum_{i=1}^n X_i} \right\rangle = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \prod_{i=1}^n \langle e^{-s X_i} \rangle = \ln \langle e^{-s X} \rangle, \quad (3.8)$$

which allows us to compute directly the form of the SCGF

$$\theta(s) = \ln \int dX e^{-(sX + \varphi(X))} = \ln \int dX e^{-(sX + \frac{(J-\mu)^2}{2\sigma^2})}, \quad (3.9)$$

which gives a result of

$$\theta(s) = \frac{1}{2} s^2 \sigma^2 - s\mu. \quad (3.10)$$

From this it is clear that the first derivative with respect to s , setting s to zero, would give us the mean and the second would give the variance, as expected.

3.2 s -ensemble

For the study we wish to do on the dimer model, understanding the LDT is important. Extending the ideas from generic large numbers, to large time ($n \rightarrow \tau$) we see how useful LDT is in the context of dynamical studies, where we look at the activity ($X \rightarrow K$). Consider a more generic master equation, [1, 66]

$$\partial_t |P(t)\rangle = \mathbb{W} |P(t)\rangle \quad (3.11)$$

where the probability distribution, $|P(t)\rangle$, at time t is given by

$$|P(t)\rangle \equiv \sum_{\mathcal{C}} P(\mathcal{C}; t) |\mathcal{C}\rangle, \quad (3.12)$$

with the probability, $P(\mathcal{C}; t)$, of being in configuration \mathcal{C} at time t . $|\mathcal{C}\rangle$ is an orthonormal configuration basis, such that $\langle \mathcal{C} | \mathcal{C}' \rangle = \delta_{\mathcal{C}, \mathcal{C}'}$.

The master operator is defined as

$$\mathbb{W} \equiv \sum_{\mathcal{C}, \mathcal{C}' \neq \mathcal{C}} W(\mathcal{C} \rightarrow \mathcal{C}') |\mathcal{C}'\rangle \langle \mathcal{C}| - \sum_{\mathcal{C}} R(\mathcal{C}) |\mathcal{C}\rangle \langle \mathcal{C}|, \quad (3.13)$$

where $W(\mathcal{C} \rightarrow \mathcal{C}')$ is the rate of transition from configuration \mathcal{C} to \mathcal{C}' (also called a jump), and $R(\mathcal{C}) = \sum_{\mathcal{C}'} W(\mathcal{C} \rightarrow \mathcal{C}')$, the escape rate from configuration \mathcal{C} . We also introduce the flat state

$$\langle - | \equiv \sum_{\mathcal{C}} \langle \mathcal{C} |, \quad (3.14)$$

so we can also write $\langle - | P(t) \rangle = 1$ due to probability conservation. The flat state is the left eigenvector of the master operator \mathbb{W} corresponding to the largest eigenvalue, which is zero, such that

$$\langle - | \mathbb{W} = 0. \quad (3.15)$$

An expectation value of an operator A can be written by $\langle A(t) \rangle = \langle - | A | P(t) \rangle$.

The dynamics described by Eq. (3.11)–Eq. (3.13), is realised in terms of stochastic trajectories. A trajectory, \mathbf{X}_τ , is just the chronological record of configurations and times at which jumps occur up to a finite time τ which is the overall length of the trajectory. From this the probability density of observing a trajectory, $P(\mathbf{X}_\tau)$ can be constructed as

$$P(\mathbf{X}_\tau) = p_0(\mathcal{C}_0) \left[\prod_{i=1}^n e^{-(t_i - t_{i-1})R(\mathcal{C}_{t_{i-1}})} W(\mathcal{C}_{t_{i-1}} \rightarrow \mathcal{C}_{t_i}) \right] e^{-(\tau - t_n)R(\mathcal{C}_{t_n})}. \quad (3.16)$$

Starting with the probability of the initial configuration, $p_0(\mathcal{C}_0)$, the trajectory is evolved in time, $t_i - t_{i-1}$, between configurations $\mathcal{C}_{i-1} \rightarrow \mathcal{C}_i$. The final factor in Eq. (3.16) comes from the difference in time of the final jump to the fixed end of the trajectory, $\tau - t_n$. Dynamical properties can be studied using time extensive observables. Here we will focus in particular on the activity K [25,34,63,67], defined as the total number of jumps that occur in a trajectory \mathbf{X}_τ , [28,34]. The distribution of the activity over all trajectories is given by

$$P_\tau(K) = \sum_{\mathbf{X}_\tau} \delta(K - \hat{K}[\mathbf{X}_\tau]) P(\mathbf{X}_\tau), \quad (3.17)$$

where $\hat{K}[\mathbf{X}_\tau]$ counts the number of jumps in a trajectory. The generating function, which also contains the counting information

$$Z_\tau(s) \equiv \sum_K e^{-sK} P_\tau(K) = \sum_{\mathbf{X}_\tau} e^{-s\hat{K}[\mathbf{X}_\tau]} P(\mathbf{X}_\tau), \quad (3.18)$$

has the advantage that its derivatives define the moments of the activity,

$$\langle K^n \rangle = (-)^n \partial_s^n Z_\tau(s)|_{s=0}. \quad (3.19)$$

We can take the large time limit of the trajectories, $\lim_{\tau \rightarrow \infty}$. In this limit, the intensive version of the activity, $k = K/\tau$, makes more sense to study. From [34,60], it is known that both the probability of observing k and the corresponding generating function have large deviation forms

$$P(k) \sim e^{-\tau\varphi(k)}, \quad (3.20)$$

$$Z_\tau(s) \sim e^{\tau\theta(s)}. \quad (3.21)$$

To draw parallels to standard thermodynamics, increasing the length of a trajectory is equivalent to increasing the system volume to ∞ , the trajectories can be thought of as configurations, acting as microstates in trajectory space, and the time extensive variables like the activity become the order parameters. The large deviation functions $\varphi(k)$ and $\theta(s)$, which describe the order parameter distribution and generating function respectively, are similar to the entropy and the free energy of the system. $\theta(s)$ has a dependence on the counting field s which is equivalent to the chemical potential for the activity. Though it is a fair analogy for the counting field, unlike the chemical potential, s is not a physical field as it is not a parameter that can be altered during an experiment. From [34, 60] we also know that $\varphi(k)$ and $\theta(s)$ are related by a Legendre-Fenchel transformation

$$\varphi(k) = -\min_s [\theta(s) + ks], \quad (3.22)$$

$$\theta(s) = -\min_k [\varphi(k) + ks]. \quad (3.23)$$

Due to the Legendre transform, the functional form of θ is convex and piecewise differentiable. In the large time limit we get the scaled cumulant generating function (SCGF) for the activity

$$\frac{\langle\langle K^m \rangle\rangle}{\tau} = (-)^m \partial_s^m \theta(s)|_{s=0} \quad (3.24)$$

where $\langle\langle K^m \rangle\rangle$ is the m th cumulant. Drawing the connection to the free energy, we can also see that θ contains all the statistical information about the order parameter in question, K , at long times. With that information it is possible to construct a phase diagram of the activity, but as this is a dynamical observable, any singularities

in θ would indicate a “dynamical phase transition”. Considering the area around $\theta(s = 0)$, this has information pertaining to all trajectories, not just the typical behavior of trajectories. It makes sense that if the counting field is turned off, the trajectories produced will be the most likely ones. If we look at the functional form of $\theta(s)$ at $s \neq 0$, rare trajectories are taken into account, which are produced due to unusual behaviours. This means we can interpret s as defining the ensemble of trajectories

$$P_s(X_\tau) \equiv Z_\tau^{-1}(s) e^{-s\hat{K}[X_\tau]} P(X_\tau). \quad (3.25)$$

This ensemble is typically called the s -ensemble [28,34]. It contains trajectories made from the standard dynamics outlined in Eq. (3.11)-Eq. (3.13), but the counting field biases the probability of trajectory occurrence. In order to obtain $\theta(s)$ we can work with a deformed master operator, \mathbb{W}_s [60,68], often referred to as the tilted operator, which is defined as [28,34]

$$\mathbb{W}_s \equiv \sum_{\mathcal{C}, \mathcal{C}' \neq \mathcal{C}} e^{-s} W(\mathcal{C} \rightarrow \mathcal{C}') |\mathcal{C}'\rangle \langle \mathcal{C}| - \sum_{\mathcal{C}} R(\mathcal{C}) |\mathcal{C}\rangle \langle \mathcal{C}|. \quad (3.26)$$

The largest eigenvalue of this operator is $\theta(s)$ [60]. The original master operator obeys Eq. (3.15) but, as we deform it with s , probability is no longer conserved, thus \mathbb{W}_s no longer obeys Eq. (3.15). The deformed master operator is the transfer matrix for the generating function

$$Z_\tau(s) = \langle - | e^{-\tau \mathbb{W}_s} | p_0 \rangle, \quad (3.27)$$

where p_0 is the vector of all initial states, $|p_0\rangle \equiv \sum_{\mathcal{C}} p_0(\mathcal{C}) |\mathcal{C}\rangle$. From all this we can write

$$\langle K \rangle_s = \frac{\langle K e^{-sK} \rangle}{Z_\tau(s)}. \quad (3.28)$$

So we have a connection between a dynamical dependent observable of the system and the distributions at $s \neq 0$.

3.3 Evidence of Dynamical Behaviours

Chapter 2 studied the fully packed dimer model with periodic boundary conditions. It looked at properties such as the flippability and the persistence. These are properties that indicated how the dynamics of the system behaves. The local flip dynamics is isolated to the region around the strings that span the system. It is also true that the fewer strings in the system then the lower the flippability. If there are fewer strings in the system it takes longer to relax. Therefore it is not unreasonable to identify a link between how quickly the systems evolves, and the type of configuration it is in. The more flippable the configuration the more likely the dynamics will be fast and vice versa.

In order to explore the s -ensemble, keeping track of the activity will indicate how much the system evolves during the length of the trajectory. Fig. 3.1 describes the dependence of a trajectory on the initial configuration.

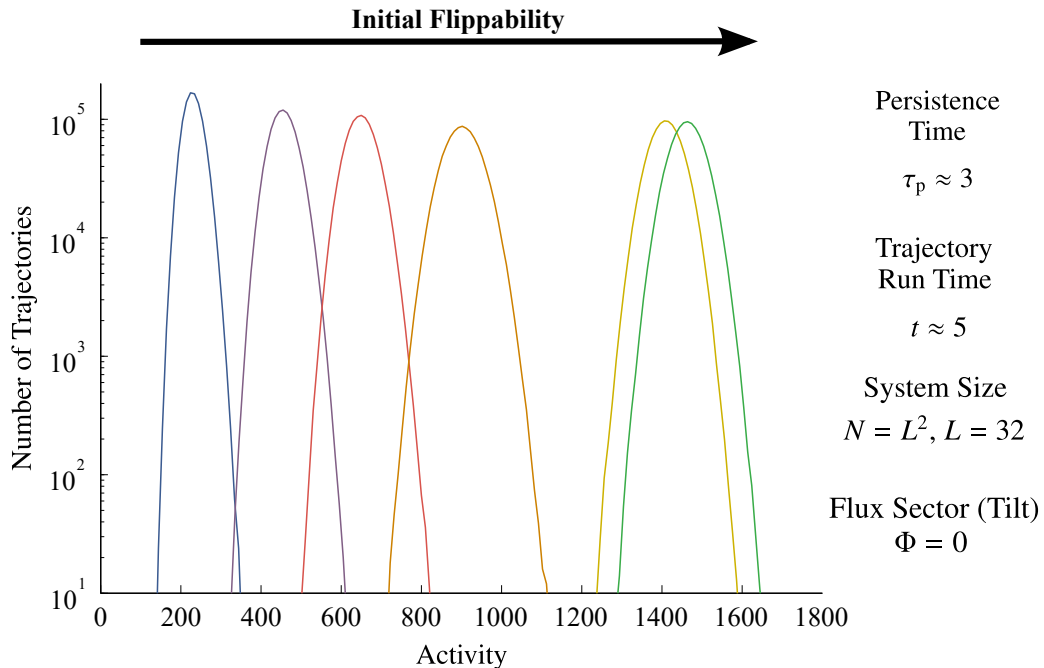


Figure 3.1: The distributions of recorded activity for various initial conditions. All configurations exist inside the same flux sector, $\phi = 0$. The initial conditions range from the least flippable (blue), $f = \frac{1}{32}$ moving to the most flippable (green), $f = \frac{1}{2}$, with purple $f = \frac{1}{16}$, red $f = \frac{3}{32}$, orange $f = \frac{1}{8}$, and yellow $f = \frac{1}{4}$. Configurations starting with the average flippability are represented by the yellow curve. The time of the trajectories is greater than the persistence time for a trajectory taken from equilibrium results, Fig. 2.10.

All of the initial configurations in Fig. 3.1 are taken from $\phi = 0$ flux sector. It can be seen that when starting in the very flippable configurations (green), an activity similar to that of the trajectories that start in equilibrium (yellow) is recovered. This is explained by considering the sample trajectories taken from each distribution, Fig. 3.2, which shows how a more flippable configuration relaxes quickly to equilibrium.




Figure 3.2: Sample trajectories taken from the distributions in Fig. 3.1, parameters are the same as in that figure. Blue is a trajectory starting from a less flippable state, green from the most flippable state and yellow, starts from an average flippability. The black line represents the average flippability. Each dot is the recording of a flip.

None of the behaviours exhibited by the highly flippable trajectories are seen in the least flippable trajectories. The activity distribution differs greatly from both the highly flippable and the equilibrium distributions, Fig. 3.1. The sample trajectory does not relax to the equilibrium value, Fig. 3.2. Moreover the trajectory itself shows flips occurring at a much slower rate than the trajectories starting in more flippable configurations.

All of this indicates that the model is one that has active and inactive dynamical behaviours. These dynamical behaviours seem to be linked to the configuration structure. Furthermore, the equilibrium state in the zero flux sector is predominantly governed by the more active behaviour.

3.4 Numerical Techniques and Sampling Methods

3.4.1 Transition Path Sampling

To explore the ensemble of trajectories outlined in section 3.2, a more complex approach than standard trajectory sampling is needed. While standard sampling has been sufficient to indicate the possibility of multiple dynamical behaviours (Fig. 3.1) in order to see the features of the dynamical phases more information about the trajectories needs to be extracted. In the cases where the master operator is exactly diagonalizable, extracting the SCGF is trivial, but in our case only small system sizes are diagonalizable. With an increase in system size, it quickly becomes apparent how hard it is to calculate the SCGF. Since the Eq. (3.26) is not stochastic, there is no simple way to generate the relevant trajectories directly. As in a static context (think for example of sampling the equilibrium of a spin system at finite temperature), this is resolved by importance sampling. In the case of trajectories one such importance sampling scheme is Transition Path Sampling (TPS) [2,31], which allows the SCGF to be determined.

The concept of TPS is similar to that of standard Monte Carlo (MC), except on the level of trajectories, rather than configurations. Where typical Markov chain Monte Carlo (MCMC) algorithms have configurations take random walks to favored configuration space, TPS will see trajectories undergo a random walk in the space of trajectories. In the original use of TPS, which was to simulate rare chemical reactions [4,31], the trajectories would start in a region of phase space dominated by reactants moving to one of the products, then move toward the phase space dominated by the desired reactions. For the s -ensemble the trajectories would start from the steady state of the system with no counting field in effect, and move towards a region dictated by the weighting factor e^{-sK} .

There are a few different types of these methods that have been shown to work

[29], such as shooting (forwards and backwards) and shifting (forwards and backwards). While shifting can only be used if both forwards and backwards updates are used, the shooting method can utilize both or only the forward updates. If the forward and backward processes are to be used in conjunction then the system being studied needs to obey detailed balance for the trajectories

$$P(X'_\tau|X_\tau)P(X_\tau) = P(X_\tau|X'_\tau)P(X'_\tau). \quad (3.29)$$

where $P(X'_\tau|X_\tau)$ is the probability of producing trajectory X'_τ given trajectory X_τ and $P(X_\tau)$ is the probability of occurrence of trajectory X_τ . $P(X_\tau|X'_\tau)$ can be broken down into more useful terms,

$$P(X'_\tau|X_\tau) = \pi(X'_\tau|X_\tau)r_a(X'_\tau|X_\tau), \quad (3.30)$$

where $\pi(X'_\tau|X_\tau)$ is the proposal distribution, with the probability of proposing X'_τ given the current trajectory X_τ , and $r_a(X'_\tau|X_\tau)$ is the probability of accepting this update to the trajectory. This means we can now write

$$\frac{r_a(X'_\tau|X_\tau)}{r_a(X_\tau|X'_\tau)} = \frac{\pi(X_\tau|X'_\tau)P(X'_\tau)}{\pi(X'_\tau|X_\tau)P(X_\tau)}. \quad (3.31)$$

Typically the proposal is given by a method that sets

$$\frac{\pi(X_\tau|X'_\tau)}{\pi(X'_\tau|X_\tau)} = 1, \quad (3.32)$$

but this is not necessarily the most efficient choice. We also have that $P(X_\tau) = e^{-sK}P_\tau(K)$, which is the distribution we are trying to sample from. Using this, we can construct a Metropolis acceptance rate of

$$r_a(X'_\tau|X_\tau) = \min(1, e^{-s\Delta K}), \quad (3.33)$$

where ΔK is the difference in activity between the new and old trajectories.

As has been emphasized before, the counting field is not physical, therefore when applying the weights in TPS, it is done in a layer of MC above that of the system dynamics. The dynamics that is used does not need to be that of the original system, and we explore later how this fact can be used to optimize the sampling process. The processes that occur in the random walk are edits to the current trajectory. The edit can be a small one that leaves as much of the trajectory intact as it can, while randomly modifying a part near the edges of the trajectory. That would have a high acceptance rate, which is computationally more favorable, but would then have a large correlation between trajectories, i.e. the majority of the trajectory would look the same. Alternatively the edit can be a large one where much larger sections are replaced. The edit can still be small but if it is near the bulk of the trajectory and the evolution of the trajectory maintains the same noise, the propagated change is large. In that case the, likelihood of producing a trajectory with favored properties, while still using the original dynamics, is low, thus giving a low chance of accepting the change. If accepted, a larger section change would then lead to the correlation between trajectories decreasing.

3.4.2 Numerical Implementations: Shooting and Shifting

The two types of methods mentioned above [31] are the key to understand how data is collected in what follows, so it will be schematically outlined how they both work.

Shooting moves occur as follows:

- 1) Randomly select a shooting time, t_{cut} , along the trajectory
- 2) Determine the direction for the shoot; randomly select forwards or backwards in time with equal probability
- 3) Generate a new part of the trajectory starting from t_{cut} going to τ for forward, or going to 0 using reverse dynamics in the backwards case.




Figure 3.3: An illustration of how the TPS shooting works. Starting from a trajectory, black, a cut time, t_{cut} , is selected uniformly across the trajectory. The direction of the shoot is chosen equally between forwards and backwards. If forwards is selected then a new part of the trajectory is created starting from the configuration at t_{cut} , and evolved until time τ (green). If backwards is selected, the starting configuration is the one at t_{cut} . When producing the rest of the trajectory (which is done using forward dynamics) it is done for a time $t_{cut} - t_0$, then reversed so that it appears to go backwards in time (blue). In either case, both the old and new trajectories have a portion in common. Once a new trajectory is created, the old (black) and new (green or blue) trajectories are compared using the e^{-sK} weighting for the Metropolis method of accept/reject the update.

Note: the shooting method suffers from correlation problems if the edits of the trajectories only occur at the ends. If the acceptance rate is low then typically smaller proposed changes are accepted, these might be in the first and last 10% of the trajectory. In that case there is 80% of the trajectory that is not following the correlation times that the acceptance rate might indicate.

Shifting follows a very similar procedure:

- 1) Randomly select a shifting time, t_{cut} , along the trajectory
- 2) Determine which part to keep by randomly selecting forwards or backwards with equal probability
- 3) From either the start (backwards) or end (forwards) produce a new part of

the trajectory such that the length of the new trajectory is equal to $\tau - t_{cut}$

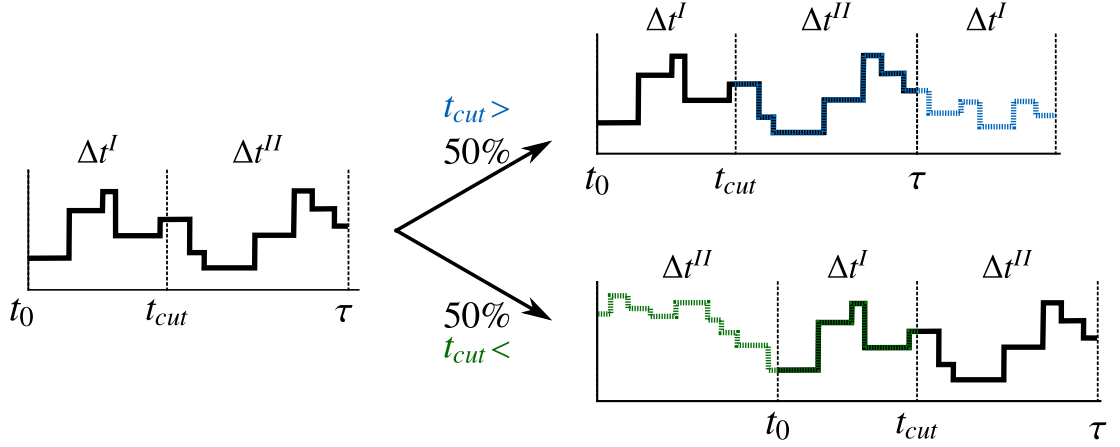


Figure 3.4: The Shifting method is similar to the shooting one. The process starts with an initial trajectory (black) and then a time is chosen uniformly along the trajectory, t_{cut} . The direction of the new trajectory is decided between forwards and backwards, equally. If forwards is chosen then the starting point is the final configuration in the trajectory, at time τ , which is then evolved for a time of $t_{cut} - t_0$, see the blue curve. If backwards is selected then the starting point is the initial configuration at time t_0 . As the trajectory producing dynamics is forward in time, part of a trajectory of length $\tau - t_{cut}$ is produced and then reversed, as shown by the green curve. Just as in the shooting case both the green and blue trajectories share a section with the black one. The new trajectory (green or blue) is then compared to the old one (black), using the weighting e^{-sK} for the metropolis method to accept or reject the proposed update to the trajectory.

Both methods suffer from a non-uniform acceptance of the proposed sizes of the trajectory to be edited. However, for shifting, as the center of the trajectory is moved by shifting of the “saved” part of the old trajectory, it acts a little like a random walker. This means its has a slight advantage over the shooting method. Ultimately both give an exponential improvement on the “brute force” method of generating a completely new trajectory at each step [29, 69, 70], which suffers from the exponential suppression that comes from Eq. (3.28). These alternative methods still suffer from an exponential suppression but to a lesser degree as explained below.

3.4.3 Umbrella Sampling with Associated Errors

When studying rare events of certain physical processes, it often becomes very difficult to properly sample these events using standard techniques. A “brute force” approach may work for some simple models, but in many cases this approach quickly starts to become too computationally expensive. There are some techniques that allow for the exploration of these rare events through manipulation of the data production - such as TPS (outlined above) or Cloning (to be discussed below) - although these methods tend to fail for many systems, albeit not as quickly as the “brute force” method.

There is, however, a possibly more intuitive approach for gaining samples away from the original point of $s = 0$, in a more achievable manner. Studying rare events of a system can be thought of as studying typical events of a modified system. From such a dynamics, data could be gathered with comparative ease, and appropriate weighting can be used to produce the desired distribution.

In order to explore the form of the mean activity as a function of s , we are interested in the following distribution

$$\langle K(\mathbf{X}_\tau) \rangle_s = \frac{\langle K(\mathbf{X}_\tau) e^{-sK(\mathbf{X}_\tau)} \rangle}{\langle e^{-sK(\mathbf{X}_\tau)} \rangle}. \quad (3.34)$$

When using the original dynamics to produce trajectories, \mathbf{X}_τ , even with a variety of methods, sampling inherently becomes burdened by the presence of the $e^{-sK(\mathbf{X}_\tau)}$, which is to say the typical trajectories have a small contribution to the sum at a given s . The distribution in (3.34), can be rewritten as if each trajectory was produced using a different \mathbb{W}_{ref} in the following way:

$$\langle K(\mathbf{X}_\tau)e^{-sK(\mathbf{X}_\tau)} \rangle = \sum_{\mathbf{X}_\tau} P(\mathbf{X}_\tau)K(\mathbf{X}_\tau)e^{-sK(\mathbf{X}_\tau)} \quad (3.35)$$

$$= \sum_{\mathbf{X}_\tau} P^{\text{ref}}(\mathbf{X}_\tau) \left(\frac{P(\mathbf{X}_\tau)}{P^{\text{ref}}(\mathbf{X}_\tau)} \right) K(\mathbf{X}_\tau)e^{-sK(\mathbf{X}_\tau)} \quad (3.36)$$

$$= \langle g(\mathbf{X}_\tau)K(\mathbf{X}_\tau)e^{-sK(\mathbf{X}_\tau)} \rangle_{\text{ref}} \quad (3.37)$$

where

$$g(\mathbf{X}_\tau) = \frac{P(\mathbf{X}_\tau)}{P^{\text{ref}}(\mathbf{X}_\tau)}. \quad (3.38)$$

This now equates the averaging of the observable e^{-sK} over the original dynamics to the averaging of observable $g(\mathbf{X}_\tau)e^{-sK(\mathbf{X}_\tau)}$ over the reference dynamics. This can be useful in the following way: given a reference dynamics we would estimate Eq. (3.35) by an empirical average over N_{sp} sample trajectories,

$$\langle e^{-s\hat{K}} \rangle = \langle e^{-sK} \rangle_{\text{ref}} \approx \frac{1}{N_{\text{sp}}} \sum_{\alpha=1}^{N_{\text{sp}}} g(\mathbf{X}_\tau^\alpha)e^{-sK(\mathbf{X}_\tau^\alpha)}. \quad (3.39)$$

The sampling error is given by the variance of the average squared of the empirical average,

$$\begin{aligned} \varepsilon^2 &= \frac{\text{var}_{\text{ref}} \left(\frac{1}{N_{\text{sp}}} \sum_{\alpha=1}^{N_{\text{sp}}} g(\mathbf{X}_\tau^\alpha)e^{-sK(\mathbf{X}_\tau^\alpha)} \right)}{\left\langle \frac{1}{N_{\text{sp}}} \sum_{\alpha=1}^{N_{\text{sp}}} g(\mathbf{X}_\tau^\alpha)e^{-sK(\mathbf{X}_\tau^\alpha)} \right\rangle_{\text{ref}}^2} \\ &= \frac{1}{N_{\text{sp}}} \left(\frac{\langle g^2 e^{-2sK} \rangle_{\text{ref}}}{\langle g e^{-sK} \rangle_{\text{ref}}^2} - 1 \right) \\ &= \frac{1}{N_{\text{sp}}} \left(\frac{\langle g e^{-2sK} \rangle}{\langle e^{-sK} \rangle^2} - 1 \right), \end{aligned} \quad (3.40)$$

where we have used Eq. (3.35) between the second and third lines to recast ε^2 in terms of the original averages. Consider the case where the reference dynamics is

just the original one, $g = 1$, the sampling error reads,

$$\varepsilon^2 = \frac{1}{N_{\text{sp}}} \left(\frac{\langle e^{-2sK} \rangle}{\langle e^{-sK} \rangle^2} - 1 \right) \approx \frac{e^{[\theta(2s) - 2\theta(s)]}}{N_{\text{sp}}}, \quad (3.41)$$

where in the last line Eq. (3.21) is used for long times. The convexity of the SCGF function implies that $\theta(2s) \geq 2\theta(s)$ for all s , and the error diverges exponentially with time. This is why sampling with the original dynamics is inefficient, and accurate estimation requires exponentially many trajectories N_{sp} . The aim is therefore to find an alternative reference dynamics which makes convergence of Eq. (3.39) more efficient.

3.4.4 Multistate Bennett Acceptance Ratio

The Multistate Bennett acceptance ratio (MBAR) method [71] is built from the Bennett acceptance ratio (BAR) [72], which is simply a way to approximate the free energy difference between two similar but distinct systems, or super states. MBAR expands this idea to multiple states, essentially giving an estimation of the free energy function over states that vary for a given parameter.

Working in standard thermodynamics we can consider two super states \mathbf{A} and \mathbf{B} . These two states have an overlap in configuration space. Each is sampled using dynamical simulation from the canonical ensemble. Observations are taken from along the trajectories and averaged. While \mathbf{A} and \mathbf{B} share some microstates, the associated energies of these micro states are not necessarily the same, $U_{\mathbf{A}}(x_n) \neq U_{\mathbf{B}}(x_n)$. For example the kinetic terms may be similar, but, due to different interaction strengths in the super states, the potential energies will differ. The goal is to find a form of the free energy difference, $\Delta F = F_{\mathbf{B}} - F_{\mathbf{A}}$, when moving from state \mathbf{A} to state \mathbf{B} . Adaptations of work from Bennett [72] give us a relationship of

$$\Delta F_{\mathbf{AB}} = -\beta^{-1} \ln \frac{Z_{\mathbf{B}}}{Z_{\mathbf{A}}}, \quad (3.42)$$

where $Z_{\mathbf{i}}$ is the canonical partition function for super state \mathbf{i} , and $\beta = 1/k_b T$. A similar relation can be found in [71]

$$Z_{\mathbf{A}} \langle \Gamma_{\mathbf{AB}} e^{-\beta U_{\mathbf{B}}(x_n)} \rangle_{\mathbf{A}} = Z_{\mathbf{B}} \langle \Gamma_{\mathbf{BA}} e^{-\beta U_{\mathbf{A}}(x_n)} \rangle_{\mathbf{B}}, \quad (3.43)$$

where $\Gamma_{\mathbf{ij}}$ is the weighting factor that links the two super states, \mathbf{i} and \mathbf{j} , similar to umbrella sampling. A neat form for these weighting factors has been shown to be [73]

$$\Gamma_{\mathbf{AB}} = \frac{N_{\mathbf{B}} Z_{\mathbf{B}}^{-1}}{N_{\mathbf{A}} Z_{\mathbf{A}}^{-1} e^{-\beta U_{\mathbf{A}}(x_n)} + N_{\mathbf{B}} Z_{\mathbf{B}}^{-1} e^{-\beta U_{\mathbf{B}}(x_n)}}, \quad (3.44)$$

where $N_{\mathbf{i}}$ is the number of samples taken from a super state. From these results, we can expand from two states to \mathbf{K} states and write an iterative formula for the free energy difference

$$\Delta F_i = -\beta^{-1} \ln \sum_{j=1}^{\mathbf{K}} \sum_{n=1}^{N_j} \frac{e^{-\beta U_i(x_n)}}{\sum_{k=1}^{\mathbf{K}} N_k e^{\beta F_k - \beta U_k(x_n)}}. \quad (3.45)$$

This can be simplified further by realising that the sum over n is running over the samples from state j , but in this formula it does not matter which state each sample comes from, so we can say $N = \sum_{j=1}^{\mathbf{K}} N_j$, giving us

$$F_i = -\beta^{-1} \ln \left\langle \frac{e^{-\beta U_i(x_n)}}{\sum_{k=1}^{\mathbf{K}} \frac{N_k}{N} e^{\beta F_k - \beta U_k(x_n)}} \right\rangle. \quad (3.46)$$

Note that this is no longer in terms of the free energy difference ΔF_i , but is instead a free energy. This is because for a given set of states, the free energy is only uniquely determined up to an additive constant. So we set one of the free energies as a reference value. Starting the iterations optimally, it is best to choose a set of free

energies close the true value. However if this value is unknown (which is commonly the case) any guess will work as this is a convergent formula.

This was all done using standard thermodynamic free energies and potentials but as has been indicated, objects like the free energy exist in the s -ensemble. The goal is to extract the “free energy” in the form of the large deviation function θ . Rather than running over a potential field, we consider various states by looking at trajectories produced from a variety of different strengths of the counting field s .

3.5 Perfect Sampling Dynamics: Doob Transformation

The main problem associated with this LD sampling is that convergence is quick for the system when the dynamics field is turned off ($s = 0$) but is exponentially slower when moving away from this point ($s \neq 0$). Given the slow convergence that occurred when applying the current widely used methods of TPS to the CDM a different angle is required to address the problem. If we consider our system and how easy it is to sample at $s = 0$, we can think about a similar model with slightly different dynamics, that is just as easy to study but corresponds to $s \neq 0$. This idea is the basis of umbrella sampling method explained in the previous section. In the case where the new model is found to have the same sampling efficiency as the original one at $s = 0$, the new one can be used to “fill out” the tails of the probability distribution of the original one. Influenced by the properties of the original system such optimal dynamics can in principle be constructed.

Exploring aspects of the original dynamics gives us the starting point we required to construct the new master operator, \mathbb{W}_{Doob} , which is stochastic and whose natural trajectories are those of the s -ensemble. The transformation from the original dynamics to this optimal dynamics is known as the generalised Doob trans-

form [65, 74, 75]. From our original dynamics we know

$$\langle - | \mathbb{W} = 0, \quad (3.47)$$

$$\mathbb{W} = \mathbb{W}^\dagger, \quad (3.48)$$

$$\mathbb{W} | - \rangle = 0, \quad (3.49)$$

where $| - \rangle$ is the flat state with an equal probability for each configuration. These properties tell us the following: (i) probability is conserved, (ii) the model obeys detailed balance and it is at an infinite temperature (i.e. only entropy plays a role), and (iii) the equilibrium state is the flat state. When we deform the master operator (3.26), the corresponding properties are

$$\langle l_s | \mathbb{W}_s = \langle l_s | \theta_s, \quad (3.50)$$

$$\mathbb{W}_s = \mathbb{W}_s^\dagger, \quad (3.51)$$

$$\mathbb{W}_s | r_s \rangle = \theta_s | r_s \rangle, \quad (3.52)$$

where $| r_s \rangle$ and $\langle l_s |$ are the right and left eigenstates. Since \mathbb{W}_s is Hermitian we could equally write $| r_s \rangle = | l_s \rangle$. Setting $s = 0$, Equations Eq. (3.49) and Eq. (3.47) are recovered and $\theta_s = 0$. This information can be used to construct a reference dynamics such that we satisfy the properties of the original dynamics but for $s \neq 0$. The properties required from the new reference dynamics are

$$\mathbb{W}_{\text{ref}} | - \rangle = 0, \quad (3.53)$$

$$\langle - | \mathbb{W}_{\text{ref}} = 0; \quad (3.54)$$

This requires the definition of an operator L_s , such that

$$\langle - | L_s = \langle l_s |, \quad (3.55)$$

$$|r_s\rangle = L_s^{-1} |-\rangle, \quad (3.56)$$

making L_s take the form of a diagonal matrix whose entries are those of the vector $\langle l_s |$. Now the Doob master operator can be written in terms of the deformed master operator and L_s [65, 74],

$$\mathbb{W}_{\text{Doob}} = L_s (\mathbb{W}_s - \theta_s \mathbb{I}) L_s^{-1}. \quad (3.57)$$

If the flat state is acted upon by this operator then the result is as desired,

$$\langle - | \mathbb{W}_{\text{Doob}} = 0, \quad (3.58)$$

$$\mathbb{W}_{\text{Doob}} = \mathbb{W}_{\text{Doob}}^\dagger, \quad (3.59)$$

$$\mathbb{W}_{\text{Doob}} |-\rangle = 0. \quad (3.60)$$

It is worth noting that the desired properties of Eq. (3.57) holds only for long times. For arbitrary times (i.e. not sufficiently long) the transformation is time dependent and complicated [74, 75]. Doob dynamics can now be used to sample at any value of s with the same efficiency as if we were sampling the original problem at $s = 0$. Expanding this, the new Doob master operator is

$$\mathbb{W}_{\text{Doob}} = \sum_{c, c' \neq c} \frac{l_{c'}}{l_c} e^{-s} W_{c \rightarrow c'} |c'\rangle \langle c| - \sum_c R_c^{\text{Doob}} |c\rangle \langle c|, \quad (3.61)$$

where the escape rate $R_c^{\text{Doob}} = R_c + \theta$ and the values of l_c are elements of the matrix L_s .

The only real disadvantage to this new dynamics is that in order to construct the Doob master operator, one requires the solution of the problem trying to be studied, i.e. $|r_s\rangle$, $\langle l_s|$ and θ . This means that any numerical results are an unnecessary self confirmation of the theory, which is therefore redundant. This explains why the transformation is not commonly used in the literature. The transformation can, however, be utilized as a first step towards a technique that can be useful.

3.6 Effective Doob Dynamics

Starting with \mathbb{W}_{Doob} from (3.61), the equation tells us two things: First, if the left eigenstate, depending on s , is not the flat state, the transition rates from $\mathcal{C} \rightarrow \mathcal{C}'$ are no longer uniform. Furthermore, the escape rates are defined by the corresponding θ . Addressing the second point, we know how to construct the escape rates manually, without the knowledge provided by θ , as long as L_s is given,

$$R_{\mathcal{C}}^{\text{ref}} = \sum_{\mathcal{C}'} \frac{l_{\mathcal{C}'}}{l_{\mathcal{C}}} e^{-s} W_{\mathcal{C} \rightarrow \mathcal{C}'}. \quad (3.62)$$

This means that the dynamics is still stochastic, but the alteration to the master operator can be written only in terms of the components of L_s . At this point we are free to change the form of L_s to an approximation and the dynamics will remain stochastic. This defines an approximation, \mathbb{W}_{ref} , to the true Doob dynamics \mathbb{W}_{Doob} .

The eigenvectors used in the Doob dynamics tell us where each configuration is located in the whole configuration space of interest. This idea can be taken further by simplifying the problem such that the focus is on a region of configuration space that is a local approximation to the true problem; this provides a starting point to find a suitable form for L_s .

3.6.1 Open 2×2 Approximation

We can obtain an approximation to the Doob transformed dynamics of a large system by focusing on the properties of a local region, such as 2×2 region shown in Fig. 3.5. This corresponds to a CDM with $L = 2$ and open boundary conditions: while the four sites in this region are covered by a dimer, such dimers may or may not be contained within the 2×2 box (a dimer may connect a site inside with one outside the region). Thus we can think of a site in the 2×2 as occupied by either a dimer or a monomer, a single site occupying particle. Looking at Fig. 3.5, and rotating the plaquettes where necessary, we see that in total there are seven configurations in this open $L = 2$ problem: two configurations with two dimers within the box, Fig. 3.5 (a), four configurations with one dimer and two monomers, Fig. 3.5 (b), and a single configuration with four monomers, Fig. 3.5 (c). The dynamics of the larger CDM, induces a dynamics between these seven states of the local 2×2 region.




Figure 3.5: Dimer configurations of a 2×2 region with open boundary conditions. There are seven configurations, which can be divided into three classes; one example of each is shown. (a) One of the two configurations with two dimers inside the included region. (b) One of the four configurations with a single dimer inside. (c) The single configuration with no dimers inside. In (b) and (c), the red circles represent sites whose dimers point out of the region.

The dynamical generator in the reduced system has the form

$$\mathbb{W}^{2 \times 2} = \begin{bmatrix} -4b & b & b & b & b & 0 & 0 \\ b & -(a+b) & 0 & 0 & 0 & a & 0 \\ b & 0 & -(a+b) & 0 & 0 & 0 & a \\ b & 0 & 0 & -(a+b) & 0 & 0 & a \\ b & 0 & 0 & 0 & -(a+b) & a & 0 \\ 0 & a & 0 & 0 & a & -(2a+c) & c \\ 0 & 0 & a & a & 0 & c & -(2a+c) \end{bmatrix}, \quad (3.63)$$

where the components correspond to the all monomer configuration, the four single dimer, and the two double dimers. The above generator has three kinds of transitions: between the two double dimer states at rate c , between single and double dimer configurations at rate a , and between the no dimer and single dimer configurations at rate b . The former kind of transition corresponds to a plaquette flip within the 2×2 region, while the latter two represent a plaquette flip at the boundary. The values of the rates depend on the size of the system in which the smaller region is embedded, and can be obtained numerically.

We can deform $\mathbb{W}^{2 \times 2}$ to obtain a scaled cumulant generating function (CGF) for the number of flips from the largest eigenvalue of the deformed operator

$$\mathbb{W}_s^{2 \times 2} = \begin{bmatrix} -4b & b & b & b & b & 0 & 0 \\ b & -(a+b) & 0 & 0 & 0 & a & 0 \\ b & 0 & -(a+b) & 0 & 0 & 0 & a \\ b & 0 & 0 & -(a+b) & 0 & 0 & a \\ b & 0 & 0 & 0 & -(a+b) & a & 0 \\ 0 & a & 0 & 0 & a & -(2a+c) & e^{-s}c \\ 0 & 0 & a & a & 0 & e^{-s}c & -(2a+c) \end{bmatrix}. \quad (3.64)$$

We are only counting the changes between the two dimer occupied states in the 2×2 region to avoid over counting when we reconstruct the large system by overlaying 2×2 regions, as the other transitions correspond to flips in neighbouring regions.

To obtain an approximation to the Doob transform of the larger system, we use the left eigenstate of 3.64. Approximating the vector $\langle l_c |$ of the whole system as a product of the exact left eigenvector of 3.64 for each 2×2 region, and using conservation of number of dimers, we find that we can write ℓ_c in the form

$$l_c = e^{DN_f^c}, \quad (3.65)$$

where N_f is the number of flippable plaquettes and the constant D is a sum of the constants in the eigenvector which is obtained by diagonalising $\mathbb{W}_s^{2 \times 2}$. Using this, the escape rates are defined as,

$$R_c^{\text{ref}} = \sum_{c'} e^{D(N_f^{c'} - N_f^c) - s} w_{c \rightarrow c'} \quad (3.66)$$

The implementation of this method is now a case of building the model as before, a simple Poissonian process, but editing the escape and transition rates to take the new dynamics into account. The data produced in this method needs to be reweighted in order to be used as a sample of the original dynamics. This is achieved through the acceptance ratio used in the TPS Metropolis algorithm. This and similar approaches amounting to umbrella sampling trajectory space have been applied recently both in cloning dynamics and in TPS [26, 27, 76]

We are interested in the distribution outlined in (3.34), which we can write as

$$\langle k_s(X_\tau) \rangle = \frac{\langle k(X_\tau)g(X_\tau) \rangle_{\text{ref}}}{\langle g(X_\tau) \rangle_{\text{ref}}}. \quad (3.67)$$

where $g(X_\tau)$ is the reweighting value used in the TPS Metropolis acceptance rates,

$$r_a(X'_\tau|X_\tau) = \min\left(1, \frac{g(X'_\tau)}{g(X_\tau)}\right). \quad (3.68)$$

The reference master operator can be written as,

$$\mathbb{W}_{\text{ref}} = \sum_{\mathcal{C}, \mathcal{C}' \neq \mathcal{C}} \frac{l_{\mathcal{C}'}}{l_{\mathcal{C}}} e^{-s} W_{\mathcal{C} \rightarrow \mathcal{C}'} |\mathcal{C}'\rangle \langle \mathcal{C}| - \sum_{\mathcal{C}} R_{\mathcal{C}}^{\text{ref}}, \quad (3.69)$$

where $l_{\mathcal{C}}$ is defined by Eq. (3.65) and the escape rates are given by Eq. (3.66). Using this with Eq. (3.38) and Eq. (3.16), the form of $g(X_\tau)$ can be calculated,

$$g(X_\tau) = e^{D(N_f^k - N_f^0)} \prod_{\mathcal{C}=0}^{k-1} e^{-\delta t_{\mathcal{C}}} \sum_{\mathcal{C}'} e^{D(N_f^{\mathcal{C}'} - N_f^{\mathcal{C}})}, \quad (3.70)$$

where $\delta t_{\mathcal{C}}$ is survival time of configuration \mathcal{C}

3.6.2 Optimising the Effective Doob Dynamics

The new dynamics can now be simplified to being governed by a single s -dependent constant, $D = D_s$. Even though we obtained the form of $l_{\mathcal{C}}$ by studying the open 2×2 case we are still free to change this form. Hence, there is no obvious reason that the obtained value is the optimal one. Therefore the properties dictating how optimal this new dynamics need to be considered. The most obvious property of the original dynamics, which the true Doob dynamics captures for non zero s , is the acceptance rate in the TPS Monte Carlo. The acceptance rate is unity when using the optimal dynamics. This makes sense because it is equivalent to stating samples produced via the optimal dynamics are typical samples for the behaviour of interest. The acceptance rate of new trajectories, therefore, seems to be a good indicator of how close the samples are to the characteristic behavior of interest.

A simple way of optimizing the sampling process is to start by collecting data from a stochastic version of the tilted master operator




Figure 3.6: Comparison of TPS acceptance rates for $L = 6$ as a function of s . The orange curve corresponds to the TPS acceptance when using the original dynamics. The green curve is for the reference dynamics with $D = D_s$ from the 2×2 approximation of the previous section. The blue curve corresponds to the optimal value of D found from exploring the acceptance rate landscape. The data is for trajectories of length $\tau_{obs} = 50$ for s positive, and $\tau_{obs} = 5$ for s negative (convergence in time is much faster on the active side $s < 0$). Each point corresponds to 5×10^6 attempted TPS moves. The optimized D_s values are $D_s = 0.25s$ for $s < 0$, and $D_s = s$ for $s \geq 0$. Inset: acceptance rates for $L = 12$ in the region where the active-inactive transition occurs for this system size.

$$w_{c \rightarrow c'}^s = e^{-s} w_{c \rightarrow c'}, \quad (3.71)$$

with escape rates

$$R_c^s = \sum_{c'} e^{-s} w_{c \rightarrow c'}, \quad (3.72)$$

so that the overall dynamics is stochastic. This is equivalent to setting $D = 0$. This will only change the time between flips, as the transition rates are uniformly affected by this new dynamics which means there is no biasing between the configurations

themselves. This is not necessarily the optimal choice for reference dynamics, which naturally leads to the question of how to optimise the selection of D in order to achieve the closest approximation to Doob dynamics. There are a couple of ways in which this can be done, depending on the parameter we choose to optimise. If we look at the acceptance rate of the TPS Metropolis, we know that for Doob dynamics we should have an acceptance rate of one which is equivalent to that of the original dynamics at $s = 0$. For a given value of s we can explore the landscape of the acceptance rate as we vary D . As written in Eq. (3.65), D is a function of s . To sum up we optimise the dynamics by varying the value of D and selecting the value which provides the highest acceptance rate. The results of the optimisation can be seen in Fig. 3.6

It is worth noting that exploring the acceptance rate is not always a good indicator that “good” dynamics have been reached. As previously discussed, if all updates are small and near the ends of the trajectory then there will be a high correlation in the data thus making the sampling process poor even though the acceptance rate would be high. We will see in Chapter 5 more evidence shows that in our case we do not have this issue, but it is always something to be weary of for this kind of study.

3.6.3 Comparison of Methods

A good comparison of the proposed technique to the current methods is to compare them at a system size where the exact result is known. The features of the first derivative of the CGF are difficult to pick out for some methods and so indicate how well the methods perform. This is illustrated in Fig 3.7.

The first method is a brute-force sampling, running the original dynamics to produce independent trajectories. This can either be achieved by producing one long trajectory and slicing it up, or by running many independent trajectories. The advantage of the large single trajectory is that one can slice it into subsets of trajectories of




Figure 3.7: The intensive activity as a function of s . A comparison of different methods against the first derivative of the exact CGF for $L=6$, black using 10^5 samples from each method: Sampling the original dynamics (cyan), applying TPS on top of original dynamics (green), sampling of Doob dynamics (purple) and sampling in effective Doob dynamics utilizing TPS (orange).

arbitrary size. This provides an easy way of testing if the length of the trajectory is sufficient in order to see if the desired long time behaviour has been reached. However, for producing 100 trajectories independent trajectories, the single long trajectory will take 100 times as long as producing the smaller ones as they can be computed in parallel. So, with a sufficient number of processes, the sample time can be reduced several orders of magnitude by producing the trajectories independently.

As expected, the brute-force method fails relatively close to $s = 0$, but it seems to agree quite well with the exact result before $s = 0.1$. This corresponds to the linear response regime of small s . The reweighting required to sample appropriately in this regime, is still exponentially hard but with a sufficiently small coefficient. As soon as this coefficient is large enough to no longer suppress the exponential behaviour to sample we see this method disagrees with the exact result, and direct sampling becomes inefficient as one moves away from $s = 0$. To improve the convergence would

require a substantially larger sample size, which starts to become computationally impractical relatively quickly.

The standard TPS method solves this exponentially bad sampling problem to some extent. However, the cost in the second method still scales exponentially, just in a more manageable way. This is best understood by looking at the error function related to the sampling methods, outlined in section 3.4.3. Looking at Eq. (3.40), we have already seen that the original dynamics, which sets $g = 1$, yields an exponential error given by Eq. (3.41). In the Doob dynamics however we can see that the function g takes a different form,

$$g = e^{sK}. \tag{3.73}$$

Inserting this into Eq. (3.40), we obtain

$$\varepsilon^2 = \frac{1}{N_{\text{sp}}} \left(\frac{\langle g^2 e^{-2sK} \rangle_{\text{ref}}}{\langle g e^{-sK} \rangle_{\text{ref}}^2} - 1 \right) = 0. \tag{3.74}$$

This tells us that the error associated with the perfect dynamics is one that removes the exponential factor. There will still be a standard sampling error that will reduce with the number of samples. So the best we can hope for by using the effective dynamics is reducing this to a non exponential problem. However, being a bit more realistic, the most likely thing that can be achieved through the effective dynamics will be a reduced exponential error. In other words, we hope the function g will be close to the form of e^{sK} but any deviations from this, which is highly likely given our approximation to the exact Doob form, will lead to an exponential build up in the error.

4

Using Large Deviations to Study Classical and Quantum Dimers

In previous Chapters we have discussed the effect of configuration on the evolution of the system for the classical dimer model (CDM) (i.e. trajectories starting with more flippable configurations relax faster than trajectories starting with less flippable configurations). Coupled to this are the properties of a Rokhsar-Kivelson Hamiltonian, where at a specific point in parameter space the quantum model maps to the classical stochastic one. With large deviation (LD) techniques detailed in Chapter 3, we can extend the mapping to a much larger parameter range in the quantum model, ultimately allowing for the exploration of the ground state of the quantum model, by a stochastic study of the classical one. In turn, knowledge of the

quantum model will improve our understanding of the dynamical phases of the the classical system. Due to the underlying complex nature of the fully packed dimer model, the problem is too computationally taxing to be explored using the standard techniques, such as the transition path sampling (TPS) techniques outlined in the previous chapter [28, 31, 77].

In this Chapter, the mechanism through which one can overcome the difficulties regarding sampling will be described. With the new method the dynamical phases of the classical model can be explored and an ambiguity with regards to the quantum phase diagram can be settled. The work from this chapter was published in [78]

4.1 Rokhsar-Kivelson Hamiltonians

Quantum dimer models (QDMs) were initially conceived by Anderson and collaborators [79, 80] as representations of singlet pairings of quantum spins. The simplest QDMs [33] correspond to systems where dimers are the basic degrees of freedom, fully pack a lattice, and have a Hamiltonian with kinetic terms flipping groups of neighbouring dimers – for example pairs of parallel dimers on the square lattice – and potential terms counting the number of such flippable clusters. In spite of their apparent simplicity – and of how old these models are – we lack a full understanding the ground state phase behaviour of such QDMs [81, 82].

As with the CDM, for the QDM the elementary degrees of freedom are dimers, which occupy the links of a lattice. We consider an $L \times L$ square lattice with periodic boundary conditions and define $d_\mu(\mathbf{r}) \in 0, 1$ as the dimer occupation number on the link joining sites \mathbf{r} and $\mathbf{r} + \boldsymbol{\delta}_\mu$, where $\boldsymbol{\delta}_\mu$ is a lattice vector and $\mu \in x, y$. The allowed configurations are those where the dimers are fully packed, i.e., where every site of the lattice is touched by precisely one dimer,

$$\sum_{\mu} [d_{\mu}(\mathbf{r}) + d_{\mu}(\mathbf{r} - \boldsymbol{\delta})] = 1. \quad (4.1)$$

At zero temperature, the phase diagram of a fully packed QDM is controlled by the ratio of the energy per plaquette to the flipping frequency. The case where the ratio is 1, usually called the Rokhsar–Kivelson (RK) point, is of special importance. Here the ground state is given by an equal superposition of all dimer configurations. The first example of a Hamiltonian with an RK point was the dimer model on the square lattice, [33]. In that work Rokhsar and Kivelson identify the quantum Hamiltonian of the fully packed dimer model,

$$\mathbb{H}_{t,v} = -t \sum (|\text{II}\rangle \langle \text{I}|\text{I} + |\text{I}\text{I}\rangle \langle \text{II}|) + v \sum (|\text{II}\rangle \langle \text{II}| + |\text{I}\text{I}\rangle \langle \text{I}\text{I}|), \quad (4.2)$$

where the kinetic terms flips the dimers with a prefactor of t and the potential terms try to either align the dimers in a parallel or staggered structure depending on the sign of v . This is equivalent to the classical master operator,

$$\mathbb{W} = \gamma \sum (|\text{II}\rangle \langle \text{I}|\text{I} + |\text{I}\text{I}\rangle \langle \text{II}|) - \gamma \sum (|\text{II}\rangle \langle \text{II}| + |\text{I}\text{I}\rangle \langle \text{I}\text{I}|), \quad (4.3)$$

where the transition rates $w_{c \rightarrow c'} = \gamma$ are uniform. The mapping between Eq. (4.3) and Eq. (4.2) is uniquely located in the quantum model when $v = t$ and so

$$\mathbb{W} = -\mathbb{H}_{\gamma,\gamma}. \quad (4.4)$$

This is known as the RK point. This raises the question of whether it is possible to use this relation to extract information about the ground state of the quantum model from the classical stochastic model, which is easier to simulate. The reverse may also be true: utilizing knowledge of the quantum model may enlighten our understanding of the classical.

4.1.1 s -ensemble Link to Quantum Ground State Properties

What we have from Section 3.2 is a deformed master operator, \mathbb{W}_s that recovers the original \mathbb{W} when $s = 0$. We know that the point $s = 0$ is the point in s that is equivalent to setting the potential and kinetic terms to be equal in the quantum model; which gives us the RK point. From the deformation Eq. (3.26) we can write

$$\mathbb{W}_s = \gamma e^{-s} \sum (|\text{II}\rangle \langle \text{II}| + |\text{I}\rangle \langle \text{I}|) - \gamma \sum (|\text{II}\rangle \langle \text{II}| + |\text{I}\rangle \langle \text{I}|). \quad (4.5)$$

So we can now take an equivalence of

$$e^{-s} = v/t, \quad (4.6)$$

which links Eq. (4.2) and Eq. (4.5). This means the s -ensemble of the classical model contains the information of the ground state of the quantum model away from the RK point. There are limitations on how far one can explore in the v/t axis using s , since when $s \rightarrow \infty$ we are asymptotically heading towards $v/t \rightarrow 0$. By this it is clear that $v/t < 0$ is unreachable.

To be specific, consider the ground-state, $|\text{gs}\rangle$, expectation value of a quantum observable represented by an operator \mathbb{Q} that is diagonal in the basis of dimer configurations. To find this, we evaluate \mathbb{Q} in the configuration at the midpoint of each trajectory, and average over trajectories with s weighting. The latter gives $\langle \mathbb{Q}(\frac{\tau}{2}) \rangle_s = \sum_c P_s(c, \frac{\tau}{2}) \langle c | \mathbb{Q} | c \rangle$, where P_s is given by $P(X_\tau)$ Eq. (3.16) with the replacement $p_0 \rightarrow p_0^s$ and $W \rightarrow W_s$. Using Eq. (3.34), and applying the same steps that led from Eq. (3.18) to Eq. (3.27), one finds

$$\langle \mathbb{Q}(\frac{\tau}{2}) \rangle_s = \frac{\langle - | e^{\frac{\tau}{2} \mathbb{W}_s} \mathbb{Q} e^{\frac{\tau}{2} \mathbb{W}_s} | \text{i} \rangle}{\langle - | e^{\tau \mathbb{W}_s} | \text{i} \rangle}, \quad (4.7)$$

where $|\text{i}\rangle$ is the initial distribution. The expectation value of \mathbb{Q} in $|\text{gs}\rangle$, the ground

state of $\mathbb{H}_{e^{-s}, \gamma}$, is therefore given by the limit of large τ ,

$$\langle \text{gs} | \mathbb{Q} | \text{gs} \rangle = \lim_{\tau \rightarrow \infty} \langle \mathbb{Q}(\frac{\tau}{2}) \rangle_s. \quad (4.8)$$

Note that this is true for any arbitrary initial distribution $|i\rangle$, as long as the overlap $\langle \text{gs} | i \rangle$ is nonzero.

In the next chapter we exploit the relation between the quantum and classical models, and connect the ground state phase diagram of the QDM to the properties of CDM trajectories explored numerically via path-sampling methods.

4.2 The Quantum Phase Diagram

Having explained our approach to sampling, we now provide some detail on what is already known about the phases of the dimer problem. Possible ground states of the square-lattice QDM can be divided into two kinds: dimer liquids, which are topologically ordered phases that break no symmetries [83], and conventional ordered phases (“dimer solids”), which break lattice symmetries. The ordered states, illustrated in Fig. 4.1, can be further divided according to their value of the flux, which vanishes in the columnar, Fig. 4.1(a), plaquette and mixed phases Fig. 4.1(b), and is maximized in the staggered phase Fig. 4.1(c).

In the quantum case a lot of work explores the ground state of the fully packed quantum dimer model [84, 85]. The study is usually done with respect to $\lambda = v/t$. There are a handful of features of the model, such as the existence of a phase transition at the RK point, which are among the reasons it is of interest to study the classical model from a LD point of view. If the mapping is correct and there exists a phase transition in the ground state at the RK point, then this corresponds to a dynamical phase transition at the same point in the classical model. According to a result of Polyakov [86, 87] a quantum U(1) liquid phase, i.e., a phase whose

long-wavelength description is a $U(1)$ gauge theory, is not possible in 2D. It can be shown, however, that the ground state at the RK point, $v/t = 1 (v, t > 0)$, can be found exactly, and is a $U(1)$ quantum liquid [87]. This fine-tuned liquid, existing only an isolated point in the phase structure, is consistent with Polyakov's argument.

It is also known that to the right of the RK point, i.e. when $\lambda > 1$, the configurations that are favoured are the staggered ones. The staggered phase can be seen in Fig. 4.1.c. Note that, the staggered phase has no flippable plaquettes. The effect of that on the classical approach can be thought of in two different ways; firstly it means that once the system is in this configuration there is no dynamical way out. So if this is the starting configuration the trajectory is frozen. Secondly it implies there is a flux dependence on this problem. If a configuration, such as the staggered one, does not exist inside the flux sector of the current configuration, then it cannot be reached under local dynamics. This will be discussed in more detail later.

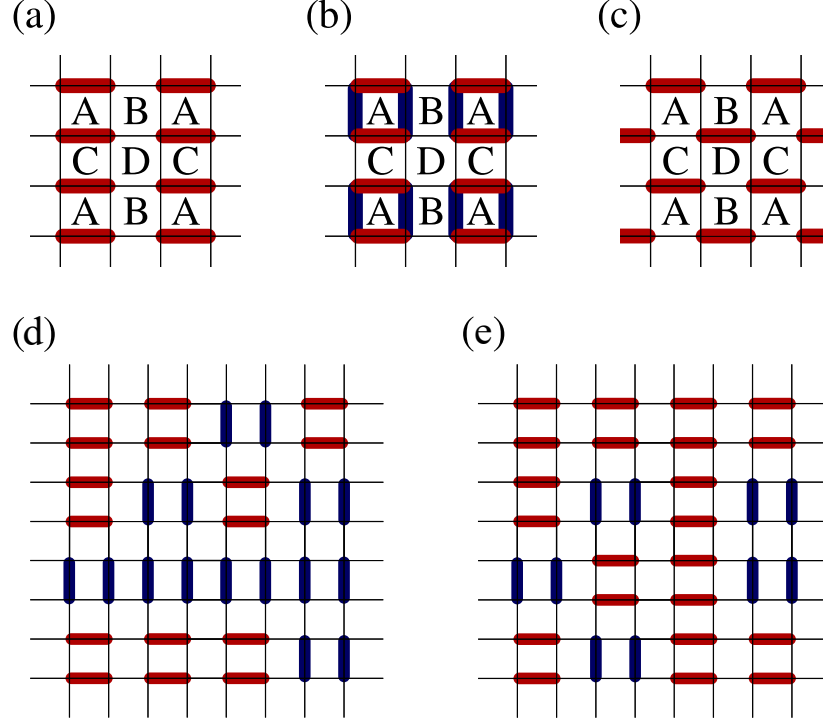


Figure 4.1: Four proposed phases [84] of the quantum dimer model. (a) shows the different types of the columnar phase. Each of the different cases are simply rotations of the others. Each of the sub-lattices has either two dimers or none. This is the phase that is represented by the highest flippable state. (b) is the plaquette phase. This is a phase where the flipping is local to certain plaquettes as indicated by the red and blue dimers. The flippable plaquettes are located on one of the four sub-lattices. In this phase the blue and the red orientations have equal probability so the state is $|\text{II}\rangle + |\text{III}\rangle$. (b) is also the case where there is some favored orientation, such as the columnar phase, there is a disparity between the blue and red weightings, $\alpha |\text{II}\rangle + \beta |\text{III}\rangle$, where $\alpha \neq \beta$. This is referred to as the mixed phase, as it is a mixture of the columnar and plaquette phases. To better understand the difference between the plaquette and mixed phases, look at the examples of configurations shown in (d) for the plaquette phase and (e) for the mixed phase. The final phase is the staggered phase (C). Just like the columnar phase it has four different orientations, but in this case each of the sublattices always has only one dimer. The staggered phase is also the maximal flux phase and has zero flippability.

Another feature of the phase diagram is that far to the left of the RK point, $\lambda \rightarrow -\infty$, the most flippable phase exists. This can be seen in Fig. 4.1(a), and is referred to as the columnar phase. It has a flippability of $N_f = \frac{1}{2}N$; this implies it has the potential to be very active as it has the maximum value of N_f .

The region of interest lies between the RK point and the columnar phase. The phase diagram for $v < t$ has been extensively studied using exact diagonalization [88, 89] and quantum Monte Carlo (QMC) [84, 85, 90], leading to a variety of contradictory conclusions. There are three theories describing this region, as shown in Fig. 4.2. Each has a different suggestion regarding the region to the left of the RK point; the first is straight to columnar, i.e. the RK point is a phase transition from staggered to columnar. Another theory identifies a new phase referred to as the plaquette phase, which exists for a small region with a phase transition on either side, columnar for $s < 0$ and staggered when $s > 0$. The final theory is similar to the second, with a third phase in this region, where this phase called the mixed phase is a hybrid of the columnar and plaquette phases.

The plaquette phase is more similar to the columnar phase than the staggered phase, see Fig. 4.1.(b). This phase consists of parallel plaquettes, just like the columnar, however, the orientations of all the plaquettes in the configuration do not need to be aligned. If the lattice is broken down to four sub-lattices, the plaquettes that are flippable in this phase are all on one sublattice. A plaquette on the sublattice of interest then has a superposition of $\alpha |II\rangle + \beta |\overline{II}\rangle$, where $\alpha = \beta$. For this to become the mixed phase then rather than having an equal superposition of the two orientations, $\alpha = \beta$, the columnar phase imposes some more structure which favours a direction, $\alpha \neq \beta$.

Ref. [84] provides different predictions from the literature and obtained the most likely phase diagram. That work concluded that the transition was a simple columnar to staggered one, Fig. 4.2.(1). This was confirmed with a numerical technique that was different to ours, as they utilized quantum Monte Carlo methods in a height field representation of the model.

4.3. LEARNING ABOUT THE CLASSICAL DYNAMICAL PHASE DIAGRAM FROM THE QUANTUM PHASE DIAGRAM

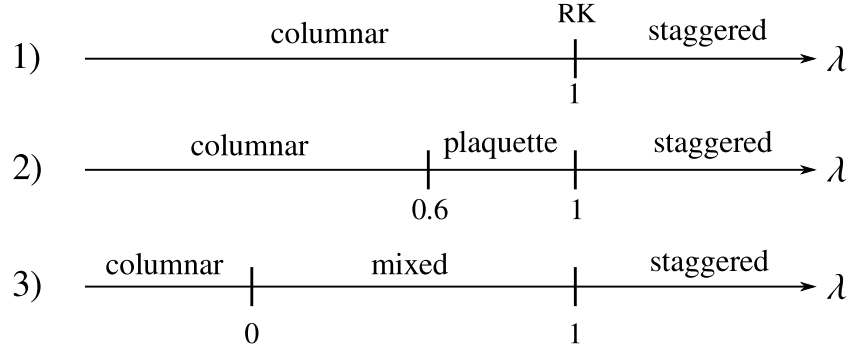


Figure 4.2: This set of predictions was collected in [84], which is where the figure comes from. It shows various possible behaviours of ground state of the QDM in the thermodynamic limit. In the first case (1), only one phase transition occurs and that is at the RK point between the staggered and the columnar phases. (2) and (3) both have multiple potential phase transitions. (2) echoes the known facts of the diagram shown in (1) of the columnar phase for $s \ll 0$, the RK point is a phase transition to the staggered phase for $s > 0$. The addition is the presence of the plaquette phase on the left of the RK point, that also has another phase transition to the columnar phase at approximately $\lambda \approx 0.6$ (in terms of the counting field, $s \approx -0.5$). The final possibility (3) follows a similar structure to (2), but the extra phase indicated is a mixed one, as apposed to the plaquette, where the transition to the columnar phase is at $\lambda = 0$.

4.3 Learning About the Classical Dynamical Phase Diagram from the Quantum Phase Diagram

From the understanding of the quantum phase diagram, we expect there to be a phase transition in the activity of the system at the RK point, $s = 0$. This is exactly what we found. Phase transitions are predicted in the thermodynamic limit but might become distorted at small system size, so we expect we may observe some systems size effect. The system starts to show interesting features at $L = 6$, see Fig. 3.7. This is the largest size that we can exactly diagonalize. The number of configurations scales as 2^N , [91], and when $L > 6$ the transition matrix is far too large to enumerate. See for example Ref. [92] for a study trying to enumerate various system sizes for the QDM, as an indication of the difficulty. For $L = 6$, we do not

4.3. LEARNING ABOUT THE CLASSICAL DYNAMICAL PHASE DIAGRAM FROM THE QUANTUM PHASE DIAGRAM

yet see a phase transition at the assumed transition point of $s = 0$ as expected for the thermodynamic limit [84].

Instead we observe two features of interest for $s > 0$. Fig. 4.3 shows the activity vs field s for several system sizes, obtained by sampling the s -ensemble with TPS and umbrella sampling as described above. Looking at Fig. 4.3, it is clear that with an increase in the system size the features apparent at $L = 6$ simultaneously reduce and collapse to the RK point. This gives clear indication that there is a dynamical phase transition at the RK point. The collapse to the single feature at $s = 0$ can also be seen in the higher order cumulants such as the susceptibility which is defined as

$$\chi = \frac{\langle K^2 \rangle - \langle K \rangle^2}{t_{\text{obs}}}. \quad (4.9)$$

Fig. 4.4 shows that the peak of the susceptibility sharpens with an increase of system size at the RK point as the secondary feature collapses down.

This confirms that extending the mapping beyond the RK point between the classical and quantum models holds, so we can turn our attention to the probability distribution. As outlined in previous chapters, in order to extract the different levels of cumulants of the system, we also collect the structural form of the LDF θ . We did this by utilizing the MBAR technique. Based on Eq. (3.22) and Eq. (3.20) we can reconstruct the probability distribution.

In Chapter 2, we found that the equilibrium flippability is, in the thermodynamic limit, $f = \frac{1}{4}$. The flippability is the number of flippable plaquettes, scaled by system size, of a typical configuration from the ensemble of configurations. The flippability and the average of the intensive activity are equivalent. The intensive activity identifies the typical number of possible flips a configuration can make during the length of the trajectory from the ensemble of trajectories. This shows that without any biasing the trajectory ensemble has the property of $\langle k \rangle = \langle f \rangle = \frac{1}{4}$.

4.3. LEARNING ABOUT THE CLASSICAL DYNAMICAL PHASE DIAGRAM FROM THE QUANTUM PHASE DIAGRAM




Figure 4.3: The main plot: activity $\langle k \rangle$ vs s , shows that $\langle k \rangle$ has a function of s is a first order phase transition, at the RK point, between inactive and active phases. For $L = 6$ it was possible to show the agreement between the exact diagonalization (black) and the numerical method (orange). Larger system sizes show that the transition region collapses towards zero. The subplot shows the true probability distribution, for $L = 24$, when the tails have been explored to a greater extent (green), in comparison with a Gaussian with the mean and standard deviation taken from exploration of the distribution at the RK point (dotted black). All data points shown are 10^6 independent samples which have been processed through the MBAR method and the first order cumulant is extracted.

Taking the largest system size, $L = 24$, from Fig. 4.3, it can be seen that this value corresponds to the curve coming from the left hand side, or the active side of the phase diagram. At a dynamical phase transition the dynamics is dominated by one of the two phases. The point is that the original dynamics are apparently governed predominantly by the active dynamical phase, rather than an equal combination of the inactive and active sides. This is due to the fact that most initial configurations lead to active trajectories.

4.3. LEARNING ABOUT THE CLASSICAL DYNAMICAL PHASE DIAGRAM FROM THE QUANTUM PHASE DIAGRAM

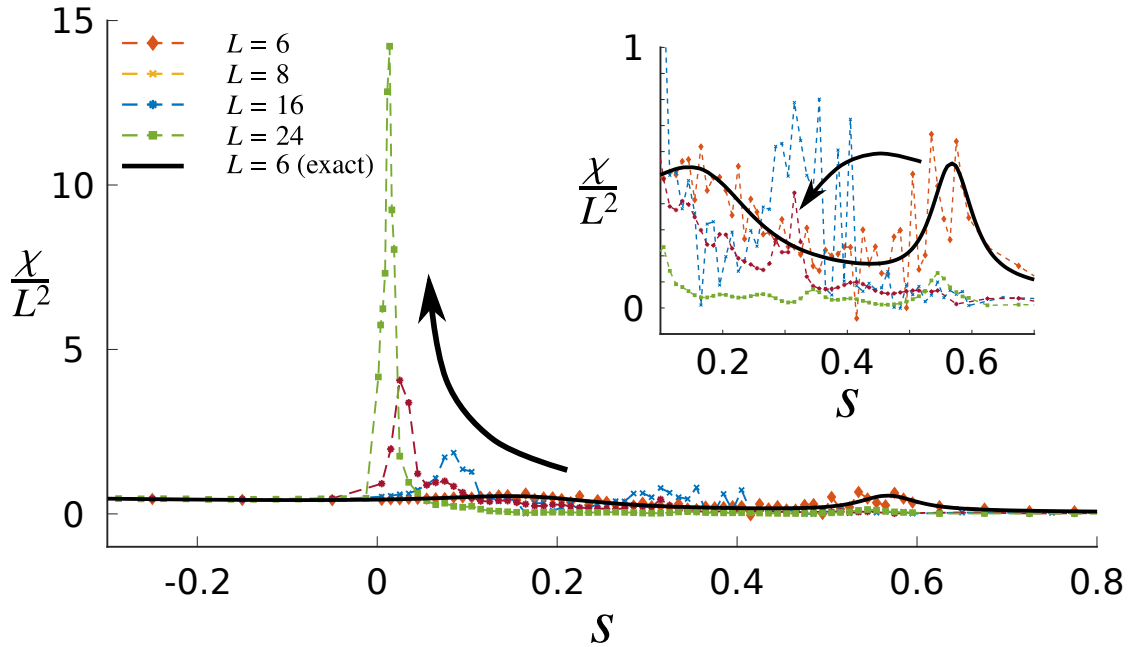


Figure 4.4: The susceptibility (the second order cumulant of the scaled CGF) as a function of s . The phase transition is just as clear in this cumulant. What becomes more apparent is the behaviour of another feature further from $s = 0$, identified most easily in the subplot. Just as the increase in system size shows the appearance of the first order transition, it also seems to show the disappearance of this feature, indicating this is a system size effect. All data points shown are 10^6 independent samples which have been processed through the MBAR method and the second order cumulant is extracted.

In the case where one dynamics (either active or inactive) is more dominant than the other at $s = 0$, the sampling would not truly represent the actual distribution therefore the true distribution will appear skewed in the direction of the less favoured phase as there is more importance in the “weaker” dynamics than the sampling indicates. If the true distribution is skewed for this reason, a normal Gaussian distribution should not fit when provided with statistics gathered from $s = 0$. However it would most likely lie on, or close to, top of the dominant side part of the distribution. We see in the inset of Fig. 4.3, that no side can be fitted by the Gaussian distribution. There is predominantly more stretching towards the less typically behaviour of inactivity,

however there is also stretching toward the more active side. Stretching on both sides implies that there is rare behaviour on either side of the transition that normal dynamical studies would miss.

With regard to comparing results from Fig. 4.3 to that of the predictions shown in Fig. 4.2, we are comparing finite system size results with the thermodynamic predictions. We believe this is acceptable, as for the largest system size we use, $L = 24$, system size effects have become negligible. The results of Fig. 4.4 indicate that this is true, since a prominent feature in the subplot is shown to reduce with system size and no longer be present for $L = 24$. This means we can suitably compare our results to the predictions in Fig. 4.2.

4.3.1 The Minimal Activity State

The theory from the QDM tells us that to the right of the RK point the configuration favours the staggered arrangement. The staggered configuration, as is outlined in Chapter 2, is part of the maximal flux sector. Under the local flip dynamics, flux is a conserved quantity. A flux sector is the set of configurations that contain the same value of flux, and are therefore the only configurations that can be reached by a trajectory that starts in such a sector. A limiting case to the flow between configurations is the maximal flux sector; this sector contain the staggered configuration which has no parallel dimers, and thus no way to undergo the flip dynamics.

The main debated area of the QDM is to the left of the RK point. It is known that the far left corresponds to columnar order. The perfect columnar state is located in the zeroth flux sector. In our simulations we restrict samples to the zeroth flux sector. There is a clear problem with the ability to explore both sides of the RK point while isolated in this sector; mainly that the true staggered configuration is unreachable. However we find that there are configurations in the zeroth flux sector which try to behave like the staggered configuration. With communication with

4.3. LEARNING ABOUT THE CLASSICAL DYNAMICAL PHASE DIAGRAM FROM THE QUANTUM PHASE DIAGRAM

Z.Lan (private comm.) and our own study we find that, up to the system sizes that we can study numerically, for any system including or larger than $L = 4$ there is a minimally flippable state of 4 plaquettes, an example of which is shown in Fig. 4.5. While there is no proof that this holds for all system sizes, the following argument suggests that this generalises to all L . For $L = 4$ the minimum number of flippable plaquettes, N_f , is 4 and when the system size is increased to $L = 6$ and again to $L = 8$ this value holds. The process of increasing system size is to add two more rows and columns of dimers to the configuration.

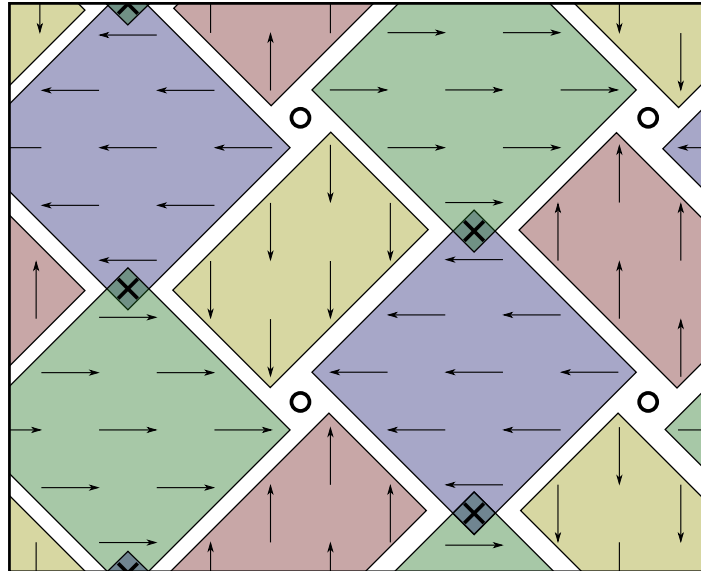


Figure 4.5: The minimal state of 4 flippable plaquettes is shown in this example configuration. What can be identified by observing various configuration of the minimal state is that there exists within it regions of staggered dimers. Given that the flux has to sum to zero for configuration in the zeroth flux sector, each staggered region has a “counter” staggered region that contains dimers that provide counter flux. Take for example the blue and green regions; both are the same size, but they contain dimers that occupy lattice links that provide the opposite flux, identified by the direction of the arrows. The same is true for the red and yellow regions. The only time a flippable plaquette is observed is when opposing flux regions intersect, identified by the \times in the green-blue overlap. In addition to this we can identify “vortices” that either have flow in or out where no regions extend to, these are identified by the \circ .

This is equivalent to wrapping the open boundary version of the previously system

4.3. LEARNING ABOUT THE CLASSICAL DYNAMICAL PHASE DIAGRAM FROM THE QUANTUM PHASE DIAGRAM

size in a way to ensure the new system size has periodic boundary conditions. This process cannot result in the loss of a flippable plaquette. Either there are no flippable plaquettes on the boundary in which case the additional rows have no decreasing impact on N_f . Alternatively, there are flippable plaquettes on the boundary (which do not appear as flippable in the open boundary conditions) which become flippable plaquettes to make sure the boundaries become wrapped.

The true ground state for positive s ($v/t > 1$) is fully staggered, but our simulations are restricted to the zeroth flux sector, and this has important consequences for the dynamics. For $s > 0$, the activity decreases in a sequence of rounded steps, which are strongly system-size dependent, which is the second feature identifiable in Fig. 4.3. This is a consequence of commensuration effects, as different low-flippability configurations are favoured depending on the precise value of s . For large positive s , the system is mostly restricted to the minimally flippable zero flux configurations, which, as illustrated in Fig. 4.5, have the 4 flippable plaquettes for any system size. Second-order perturbation theory gives the ground-state energy as $-\theta(s) = -2e^{-2s}$, which leads, using Eq. (3.24), to a mean activity of $\langle K \rangle / t_{\text{obs}} = 4e^{-2s}$ in this limit. As the system size is increased the number of minimal states increases rapidly, this is possibly the reason for the collapse of the second feature in Fig. 4.3.

It is also worth pointing out that in statistical physics the principle of equivalence means that there must be a configuration of the whole system that consists of tilings of sub regions of thermodynamically stable phases. This is similar to that of the phase coexistence observed in a ferromagnet with zero magnetisation below the critical temperature. The main difference being that in the ferromagnet case the domain walls dictate the tiling whereas in the dimer case it is the vortices identified by the \times and \circ in Fig. 4.5.

4.4 The Way Classical Dynamics Gives Instruction About the Quantum Ground State

Having shown how well we can explore trajectories for the dynamical phases of the classical model, we turn our attentions to the configurations themselves to study the ground state quantum phases. A well established order parameter in the context of classical dimer models [88, 93] is provided by the magnetisation \vec{N} . The magnetisation \vec{N} is,

$$N_\mu = \sum_{\mathbf{r}} (-1)^{r_\mu} d_\mu(\mathbf{r}). \quad (4.10)$$




Figure 4.6: Example of what the different phases would look like in terms of the order parameter, \vec{N} , when in the thermodynamic limit. The columnar phase lies on the corners as maximum possible values along a given axis. The staggered phase is isolated to a single delta function at the origin. Finally the plaquette phase lies along the edge of possible values but diagonal to the axis. The design is based on figures from [84]. Linking to Fig. 4.2 the various phases would exist as differently positioned peaks located as shown in this diagram. Where the mixed phase would be a peak that lies somewhere along the line between the columnar and plaquette peaks.

This points along a lattice direction $\pm\delta_\mu$ in the columnar phase and along the four

4.4. THE WAY CLASSICAL DYNAMICS GIVES INSTRUCTION ABOUT THE QUANTUM GROUND STATE

diagonal directions $\pm\delta_x \pm \delta_y$ in the plaquette phase, while interpolating between the two in the mixed phase, which can be seen in Fig. 4.6. It vanishes by symmetry at the RK point and also in the staggered states. The main results regarding the phase structure of the QDM are displayed in Fig. 4.7, which shows the distribution of the order parameter \vec{N} in the quantum ground state, for various values of $s \leq 0$. For all negative s , ($\lambda < 1$), the maxima of the distribution of the magnetisation occur for \vec{N} aligned with the square axes, indicating that the ground state has columnar order. Particularly for small $|s|$ the selection of this order is weak, and the distribution has approximate SO(2) symmetry under continuous rotations with the largest probabilities occurring on a ring of fixed $|\vec{N}|$.

To characterize quantitatively the degree of selection of columnar order, we consider the quantities $\langle \text{gs} | |\vec{N}|^2 | \text{gs} \rangle$ and $\langle \text{gs} | \cos(4\phi) | \text{gs} \rangle$, where $\tan \phi = N_y/N_x$, evaluated in the ground state $|\text{gs}\rangle$ of the QDM. As Fig. 4.8 shows, both quantities decrease as the RK point, $s = 0$, is approached. This is in qualitative agreement with the results of Banerjee et al. [84].

The microscopic model is only symmetric under the discrete rotations of the lattice, and so the approximate SO(2) symmetry of the order parameter \vec{N} is emergent. As argued by Fradkin et al. [87], this can be understood qualitatively by considering the renormalization group (RG) flow structure. Near the RK point, an effective action can be written in terms of a continuum height field h as

$$\mathcal{S} = \frac{1}{2}(\partial_0 h)^2 + \frac{1}{2}\rho_2(\nabla h)^2 + \frac{1}{2}\rho_4(\nabla^2 h)^2 + \lambda \cos(2\pi h) + \dots, \quad (4.11)$$

where ρ_2 , ρ_4 , and λ are real parameters, and ∇ and ∂_0 denote the space and imaginary-time derivatives, respectively. The coefficient $\rho_2 \sim t - v \sim -s$ is tuned through zero at the RK point.

4.4. THE WAY CLASSICAL DYNAMICS GIVES INSTRUCTION ABOUT THE QUANTUM GROUND STATE

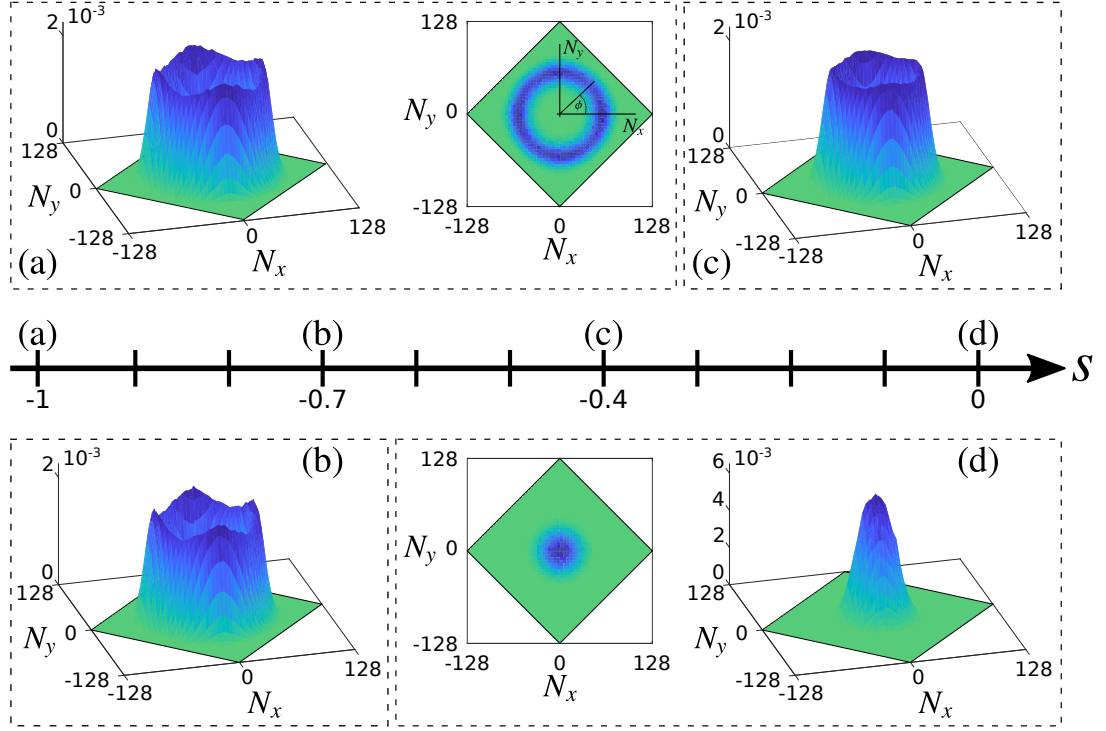


Figure 4.7: (a–d) Distribution of magnetisation \vec{N} in the quantum dimer model, for various values of the parameter $s = \ln(v/t)$, as indicated on the central s axis, or a system size $L = 16$. For all $s < 0$, the distribution has approximate circular symmetry, but with peaks along the square axes, corresponding to columnar order. The prominence of the peaks and the magnitude of $|\vec{N}|$ of the ring decrease as one approaches the RK point, at $s = 0$, where the distribution is Gaussian around $\vec{N} = \vec{0}$.

While the magnetisation \vec{N} does not appear explicitly in Eq. (4.11), it is related to the coarse-grained height by $N_x + iN_y \sim \exp[-i\frac{\pi}{2}(h + \frac{1}{2})]$. The λ term therefore breaks the $SO(2)$ symmetry and obtains the discrete subgroup of lattice symmetries determining the ultimate direction of the RG flow towards columnar order for positive λ . At the RG fixed point, corresponding to the RK point, λ is strongly irrelevant, with RG eigenvalue $y_\lambda = -12$. Standard scaling arguments in the presence of a irrelevant perturbation [94, 95] imply the existence of an additional length scale, $\propto |s|^{-3}$ for small negative s . The selection between columnar and plaquette order occurring only beyond this large scale [87]. This is qualitatively consistent with the

weak columnar ordering observed for small $|s|$ at the system sizes accessible in our MC simulations. Which is why, even though the results of Fig. 4.7 differ from the sharp peaked behaviour expected of the prediction from Fig. 4.2.1 and Fig. 4.6, we still identify this prediction to be the most likely.

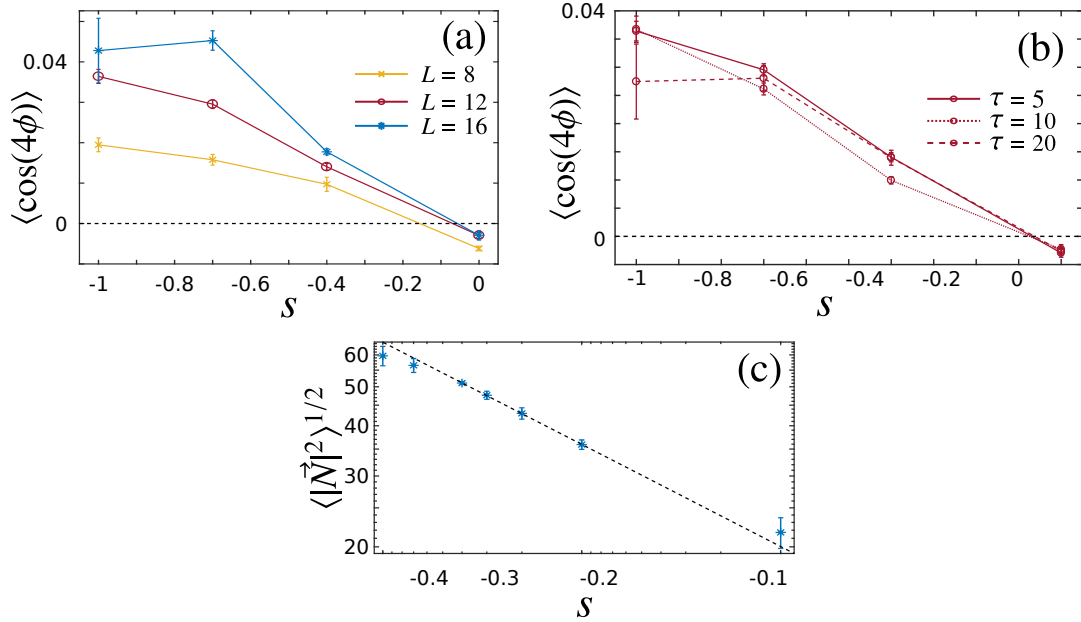


Figure 4.8: (a–b) Anisotropy measure $\langle \cos(4\phi) \rangle$, with $\tan \phi = N_y/N_x$ versus s for various system sizes L and, for $L = 12$, as a function of the trajectory time used in the simulations. Positive values correspond to a distribution peaked along the square axes confirming that the ordering is columnar, and that it becomes more pronounced as $|s|$ increases. (The small negative value at $s = 0$ is, presumably, a consequence of the discrete values taken by \vec{N} .) (c) Root-mean-square magnetisation magnitude $\langle |\vec{N}|^2 \rangle^{1/2}$ as a function of s for $L = 16$, corresponding roughly to the radius of the ring in the distribution of \vec{N} . The slope of the dashed line is $\frac{1}{4}$.

4.5 Conclusions

The results of this chapter provide an example of the connection between the statistical properties of long-time trajectories of a classical system and the properties of the low-lying spectrum of a related quantum system. We have focused on the fully packed CDM on the square lattice and, correspondingly, the QDM. However

this process could be expanded to other models with ease. As already illustrated this connection works both ways: from the known existence of a quantum phase transition in the QDM at the RK point we infer the existence of a transition - which we confirm numerically - between active and inactive dynamical phases in the CDM. Conversely, from the statistics of atypical trajectories of the CDM, we learn about the ground state properties of the QDM away from the RK point. Other examples of this classical-quantum connection include classical self-exclusion processes and XXZ chains [96–98], and the one-dimensional Ising model with Glauber dynamics, and the transverse field Ising chain [65].

For the QDM, our main results, shown in Figs. 4.7 and 4.8, are that the ground state is the columnar phase for all $v/t < 1$. Our results in this regime also show an approximate emergent $SO(2)$ symmetry of the order parameter. Both of these findings agree with the observations of Ref. [84], and contradict earlier results [85,90]. We note that the method we use here is closer in spirit to that of the earlier work. For other numerical approaches to the QDM problem see also [85,90]. For the CDM, the main result is the non-trivial structure of fluctuations in the dynamics away from typical behaviour. The first-order transition at $s = 0$, see Fig. 4.3, implies a coexistence in the equilibrium dynamics of space-time regions of high and low activity, and therefore, a broad distribution of the dynamical order parameter, (see Fig. 4.3 subplot). Furthermore, the two competing phases display different kinds of structural order: while the inactive phase, $s > 0$, is staggered, the active phase ($s < 0$) is columnar. This change in the nature of configurations, in order to optimize large dynamical fluctuations, is reminiscent of what occurs in other systems. For example simple exclusion processes where - even its state of typical activity and current are featureless - rare inactive trajectories are associated with phase separated states, and atypical large currents to hyperuniform (super-homogeneous) states [98–100].

A consequence of the large fluctuations in the dynamics is that sampling rare

trajectories is difficult. This is more apparent in a system like the CDM with periodic boundary conditions due to the constrained nature of configuration space and the conservation of the flux. To sample trajectory space, we used transition path sampling [a Monte Carlo meta-dynamics in the space of trajectories guaranteed to converge to the s -ensemble Eq. (3.34)]. In order to overcome the numerical difficulty of accessing exponentially suppressed trajectories we supplemented TPS with a version of umbrella sampling in trajectory space [26, 27, 76, 101]. TPS is well suited to our problem as the CDM dynamics obeys detailed balance. Our umbrella sampling could be improved by obtaining the reference dynamics in an adaptive manner, as is done in [27] for cloning dynamics. This is in fact the topic of the following chapter. Other interesting avenues to pursue include considering open boundary conditions (where we expect exploration of dynamics to be easier due to the absence of flux conservation), and to study, in a similar manner as here, dimer coverings in other geometries including higher dimensions.

5

Trajectory Sampling with Feedback-Augmented TPS

The previous Chapters have discussed the mapping between the quantum dimer model (QDM) to the classical dimer model (CDM) using large deviation (LD) techniques. There we enhanced our approach to the transition path sampling (TPS), by supplementing it with the modified dynamics and umbrella sampling in trajectory space [26, 27, 76]. In work done by Nemoto et al, [27, 30], “cloning”, a common form of LD numerical analysis, different from TPS, is adapted to include a modified dynamics that works through a feedback procedure.

In this Chapter we will explore improvements to the TPS method by combining the modified dynamics similar to the one used in Chapter 4 and the feedback ideas

introduced in [27]. The goal of the improvement is to provide a method of probing previously hard to explore dynamical phase diagrams. The feedback approach comes from work based on the cloning algorithms, so this Chapter will begin by explaining this alternative method to TPS. While different, TPS and cloning give rise to the same sampling error, calculated in section 3.4.3, as both deal with the weighting factor of e^{-s} in similar ways. Although the two LD numerical methods have no formal mapping, we show how the feedback algorithm can be adapted to TPS. In fact, due to the structure of trajectories in cloning, it is conceptually easier to extract the desired object out of TPS trajectories than cloned ones.

We use the CDM to compare dynamics of a feedback variant of an approximation to the Doob transformation to that of the original dynamics method, as an illustrative example. The results shown in the previous Chapter act as a good comparison to show that the feedback method works. While we do not expect a large numerical improvement on the previous method (which was in effect a non-systematic realisation of the feedback approach presented below), the feedback process offers a more autonomous and generally applicable approach to data collection. Furthermore, it naturally provides another way to confirm whether the most efficient, modified dynamics used in umbrella sampling has been reached.

5.1 Review of the Cloning Method: Population Dynamics with a Feedback Doob Approximation

In order to explain the feedback process as it is presented in the literature [27, 30], we must introduce another numerical techniques the so-called cloning method [64, 77, 102, 103]. Cloning approaches the sampling problem in a different way to TPS. Where TPS typically confines itself to a two trajectory comparison, cloning has many

parallel trajectories continuously running at the same time. There are a few different types of implementation of cloning, but we will focus on the population-dynamics method presented here [77]. The dynamics used on the various different trajectories is that of the stochastic master operator

$$\mathbb{W}_{\text{mod}}^{\mathcal{C}\mathcal{C}'} = W_{\text{mod}}(\mathcal{C} \rightarrow \mathcal{C}') - R_{\text{mod}}(\mathcal{C}), \quad (5.1)$$

where the escape rate is

$$R_{\text{mod}}(\mathcal{C}) = \sum_{\mathcal{C}'} W_{\text{mod}}(\mathcal{C} \rightarrow \mathcal{C}'), \quad (5.2)$$

and $W_{\text{mod}} = e^{-s}W$. This will neither give the correct distribution of the original dynamics, nor that of the tilted operator. The object of interest is the distribution of configurations $P(\mathcal{C}, s, t)$ when acted upon by the master operator. The non-stochastic tilted master operator can be written in terms of the tilted stochastic one (5.1),

$$\delta_t P(\mathcal{C}, s, t) = \sum_{\mathcal{C}'} \mathbb{W}_{\text{mod}}^{\mathcal{C}\mathcal{C}'} P(\mathcal{C}', s, t) + [R_{\text{mod}}(\mathcal{C}) - R(\mathcal{C})] P(\mathcal{C}, s, t). \quad (5.3)$$

There is a breakdown in the conservation of probability due to the $R_{\text{mod}}(\mathcal{C}) - R(\mathcal{C})$ term. The dynamics provided from Eq. (5.1) allow for the basic numerics of the cloning algorithm. Consider running N_c trajectories in parallel. A trajectory is created according to continuous time Monte Carlo, where the next configuration is chosen from the transitions rates $W_{\text{mod}}(\mathcal{C} \rightarrow \mathcal{C}')$, and the time until the next jump Δt is governed by the escape rate $R_{\text{mod}}(\mathcal{C})$. However when a jump is made there is a possibility that the α th trajectory (or “clone”) will be cloned or pruned, which means that it is replicated or destroyed. These happen with a probability

$$\mathcal{Y}(\mathcal{C}') = e^{\Delta t(R_{\text{mod}}(\mathcal{C}') - R(\mathcal{C}'))}. \quad (5.4)$$

This probability is then combined with some random noise to determine $y = \mathcal{Y}(\mathcal{C}') + \epsilon$, where the random number $\epsilon \in [0, 1]$. For $y = 0$, the clone is erased. However, if $y > 0$ then $y - 1$ new clones are created. It is easy to see how such an algorithm could lead to the complete decay or exponential growth of N_c by a factor of $X = \frac{N_c + y - 1}{N_c}$. Such a change to N_c would be either useless for sampling or too computationally expensive. This leads to a correction to the number of clones. Every time a change to N_c is made, a counter measure is taken. If $y > 0$ then $y - 1$ random clones are selected and destroyed, and if $y = 0$ then a clone is chosen to be duplicated at random.

With the build up of clones surviving or dying, the cloning factor X , which represents the exponential evolution of $P(\mathcal{C}, s, t)$, can be used to construct the LD function θ_A in the long time limit;

$$\frac{1}{t} \ln(X_1 X_2 \dots X_h) = \frac{1}{t} \ln \langle e^{sA} \rangle \sim \theta_A, \quad (5.5)$$

where h is the number of events occurring during the entirety of all trajectories. From the LD function, it is then possible to extract a cumulant of any order.

This is not the only form of cloning, since it can be done with discrete time, or by using standard dynamics rather than tilted. From this description of the cloning process, it is quite natural to see that altering the dynamics simply changes the cloning function, which in turn might arrive at the LD function in a different way. It can be assumed that, given the best possible dynamics, the LD function could be obtained more efficiently than using original dynamics or even the tilted stochastic one.

5.2 Feedback Dynamics for Cloning

Cloning is a method of LD sampling that naturally lends itself to sampling with an idea of alternative (or reference) dynamics, [27, 30]. We know that the best possible

dynamics comes from the Doob transformation, Section 3.5. But since we know that using the Doob dynamics is not practical for reasons previously mentioned, we explore an alternative dynamics constructed, similar to section 3.6, using the path probability of the original dynamics $P(X_\tau)$, and the path probability of some modified dynamics $P^U(X_\tau)$. Here, the ‘‘potential’’ $U(\mathcal{C})$ is the configuration-dependent weighting giving rise to the modified dynamics. The two dynamics are linked by

$$P(X_\tau)e^{-s\tau K(\tau)} = e^{U(\mathcal{C}(\tau))-U(\mathcal{C}(0))}P^U(X_\tau)e^{\mathbb{K}_{\text{mod}}(\tau)}, \quad (5.6)$$

where $\mathbb{K}_{\text{mod}} = \int_0^\tau dt[R_{\text{mod}}(\mathcal{C}_t) - R(\mathcal{C}_t)]$, with R being defined in (5.2). For the modified dynamics, the transition rates feel the influence of the potential as it affects configurations in a non-uniform way. In LD studies the limit of $\tau \rightarrow \infty$ is taken. In this limit the difference in potential between the initial and final configurations is negligible. Looking at Eq. (5.6), the weighting of the LHS is $e^{-s\tau K(\tau)}$ and the RHS has a biasing of $e^{\mathbb{K}_{\text{mod}}}$ and a negligible factor. The challenge is choosing a good form of $U(\mathcal{C})$ to obtain the best results.

Here we introduce two new probability distributions; the final time distribution

$$p_{\text{end}}(\mathcal{C}) = \lim_{\tau, N_c \rightarrow \infty} \frac{1}{\tau N_c} \sum_{\alpha=1}^{N_c} \int_0^\tau dt \delta_{\mathcal{C}, \mathcal{C}_t^\alpha} \quad (5.7)$$

where \mathcal{C}_t^α is the trajectory of the α th clone, and the intermediate time distribution

$$p_{\text{ave}}(\mathcal{C}) = \lim_{\tau, N_c \rightarrow \infty} \frac{1}{\tau N_c} \sum_{a=1}^{N_c} \int_0^\tau dt \delta_{\mathcal{C}, \bar{\mathcal{C}}_t^a} \quad (5.8)$$

where $\bar{\mathcal{C}}_t^a$ is the trajectory of the a th clone that has survived until the final time τ .




Figure 5.1: [Not real data] Comparison of cartoon trajectories from the equilibrium original dynamics (no biasing) (green solid), a standard LD numerical technique (TPS) (blue dashed), a modified dynamical process coupled to LD numerical technique (red dot-dashed), and an optimal dynamics (Doob) (black dotted). For green and blue lines, the ends align because both are sampled by the original dynamics. In the middle, the two differ because the weighting from the TPS algorithm drives the bulk of the trajectory towards the equilibrium for a given value of s . The red and black lines have ends that do not match with other lines due to the differing dynamics used in sampling. The bulk of the trajectories, however, do coincide with the blue dashed line, as this is the equilibrium for the given value of s .

These objects can be understood equally well in TPS language. Take Fig. 5.1, where the green curve shows equilibrium sampling ($s = 0$). The structure is relatively flat because the configurations at the ends, from $|R\rangle$, have the same probability distribution as the configurations sampled in the bulk of the trajectory, $|LR\rangle$ (or $|R^2\rangle$). This is due to $|R\rangle$, which is equivalent to $\sum_C p(C) |C\rangle$, taking a normalised form of the flat state $|-\rangle$, which is the equal distribution of states. Moving to rare event trajectories ($s \neq 0$), the bulk is sampled from a different distribution to that of the ends of the trajectories, resulting in the dip feature, as is seen in the blue curve. This is due to the fact that $|R\rangle$ is no longer proportional to the flat state.

So the expectation value of observable A at the ends of the trajectory looks like $\langle -|A|R\rangle$ while in the bulk of the trajectory they are $\langle L|A|R\rangle$. This is equivalent to saying that the configurations at the end of the trajectory are sampled from the flat state while in the bulk they are sampled from $\langle L|$. This is noteworthy as this

description only applies to unaltered dynamics. When the dynamics is “improved” the sampling at the ends of the trajectory is done from a different $\langle L_{\text{ref}} |$ which is closer than the flat state to the desired $\langle L |$, red curve. If the perfect dynamics are used then the configurations are sampled from $\langle L |$ and we get the black curve.

The reason to outline these properties is that the language used with regards to p_{end} and p_{ave} are equivalent to referring to $|R\rangle$ and $|LR\rangle$. We now have a link between the TPS terminology for the distributions, in the bulk and at the ends of the trajectory, and terminology used in cloning. In terms of the potential U , the ratio of p_{end} and p_{ave} gives rise to the relationship

$$U^*(\mathcal{C}) = -\log \frac{p_{\text{ave}}(\mathcal{C})}{p_{\text{end}}(\mathcal{C})}, \quad (5.9)$$

where U^* is the optimal potential (or the Doob potential).

Instead of finding U^* , the goal is to find an appropriate approximation to it. Starting from the optimal potential and relate it to the final and intermediate probability distributions of an arbitrary potential, U , can be represented by p_{end}^U and p_{ave}^U . They have a relation to the original distributions

$$p_{\text{end}}^U(\mathcal{C}) = p_{\text{end}}(\mathcal{C}), \quad (5.10)$$

$$p_{\text{ave}}^U(\mathcal{C}) = p_{\text{ave}}(\mathcal{C})e^{-U(\mathcal{C})}, \quad (5.11)$$

which allows Eq. (5.9) to be written as

$$U^*(\mathcal{C}) = U(\mathcal{C}) - \log \frac{p_{\text{ave}}^U(\mathcal{C})}{p_{\text{end}}^U(\mathcal{C})}. \quad (5.12)$$

If the system being studied has detailed balance, the equilibrium (original dynamics) distribution can be utilized. Detailed balance comes from $p_{\text{ave}}(\mathcal{C})w_{\mathcal{C} \rightarrow \mathcal{C}'}^{\text{mod}} = p_{\text{ave}}(\mathcal{C}')w_{\mathcal{C}' \rightarrow \mathcal{C}}^{\text{mod}}$ and $p_{\text{eq}}(\mathcal{C})w_{\mathcal{C} \rightarrow \mathcal{C}'} = p_{\text{eq}}(\mathcal{C}')w_{\mathcal{C}' \rightarrow \mathcal{C}}$ which leads to

$$p_{\text{ave}}^U(\mathcal{C}) = \frac{p_{\text{end}}^U(\mathcal{C})^2}{p_{\text{eq}}(\mathcal{C})} e^{-U(\mathcal{C})}. \quad (5.13)$$

The enhancement in sampling the equilibrium of the original dynamics in comparison to the rare events Eq. (5.13), simplifies the problem of finding the desired potential.

5.2.1 Cloning Approximations

So far the method entails using the whole probability distributions for the entire configuration space. These distributions can be difficult to construct analytically as well as numerically when the system size increases. The feedback process has the goal of iteratively adapting a potential to be close to the ideal potential. In order to achieve this an approximation has to be applied. Rather than considering the whole configuration space one can make a small system size approximation, as we did in previous Chapters; for example only consider an area around the site which undergoes the local update. The local update has no structural effect on the configuration nearby (in our system), but it does alter the configurations that can be accessed. Therefore an approximation that considers an area that is large enough to encompass this neighborhood effect would be a reasonable one. Formally this is considered as the difference between the potentials of configurations before and after an update,

$$U(\mathcal{F}_i[\mathcal{C}]) - U(\mathcal{C}) = u_d(n_{i-d} \dots n_i \dots n_{i+d}). \quad (5.14)$$

The potential of the approximation u_d covers an area of the total configuration, d sites, that has undergone an update by the operator \mathcal{F}_i at site i . This approximation is valid if the neighborhood effect is a small tiling of the whole system; if so, then it is possible to recreate every configuration from tiling different configurations of the

small area, and the approximation holds. The new potential has different distributions associated with it, due to a much smaller number of configurations in the small area being considered,

$$p_{\text{end}}(\mathcal{C}) = \lim_{\tau, N_c \rightarrow \infty} \frac{1}{\tau N_c} \sum_{a=1}^{N_c} \int_0^\tau dt \prod_{j=i-d}^{i+d} \delta_{n_j, (n_j(t))_a}, \quad (5.15)$$

where $(n_j(t))_a$ is the j th site of the a th clone at time t ; and the intermediate time distribution

$$p_{\text{ave}}(\mathcal{C}) = \lim_{\tau, N_c \rightarrow \infty} \frac{1}{\tau N_c} \sum_{a=1}^{N_c} \int_0^\tau dt \prod_{j=i-d}^{i+d} \delta_{n_j, (\bar{n}_j(t))_a}, \quad (5.16)$$

where $(\bar{n}_j(t))_a$ is the configuration of the j th spin at time t for the a th clone that has survived until the final time τ .

From these new distributions and the approximate potentials, we can rewrite Eq. (5.12) as

$$u_d^{l+1}(n_{i-d} \dots n_i \dots n_{i+d}) = u_d^l(n_{i-d} \dots n_i \dots n_{i+d}) - a \log \frac{p_{\text{ave}}^d(\mathcal{F}_i[\mathcal{C}], i) p_{\text{end}}^d(\mathcal{C}, i)}{p_{\text{end}}^d(\mathcal{F}_i[\mathcal{C}], i) p_{\text{ave}}^d(\mathcal{C}, i)}. \quad (5.17)$$

An initial guess is made for the potential u_d^1 and then after data has been collected to construct more accurate forms of p_{end}^d and p_{ave}^d , the iterative formula Eq. (5.17) is applied to alter the potential. The value of d at this point is arbitrary, but what is shown in [27] is that once the size of d is sufficiently large, little benefit is gained by increasing it further. In fact, increasing d simply increases the number of approximate configurations that require sampling, making it harder to construct p_{end}^d and p_{ave}^d .

The variable a is a way to control the approach to the true potential. If a is small, the updates to the potential will be minor but stable. If a is large, then the updates will suffer from large fluctuations. In practice a is chosen through a manual

process of exploring various values and determining which yields the best results, similarly to how we approached the evaluation of D_s in Chapter 3.

It is worth noting that the construction of the probability distributions p_{end}^d and p_{ave}^d , can be achieved by either tracking a single site i , or by studying the whole system. In the case where the whole system is being recorded, the distributions build up quicker. However there may be some correlations affecting the distributions. The correlations would possibly identify how sound the approximation is, so it is worth identifying if the two approaches yield the same distributions before deciding which method to use.

Taking into account the conditions of a detailed balance system from Eq. (5.13) we can rewrite Eq. (5.17) as

$$u_d^{l+1}(n_{i-d}\dots n_i\dots n_{i+d}) = a \log \frac{p_{\text{end}}^d(\mathcal{F}_i[\mathcal{C}], i) p_{\text{eq}}^d(\mathcal{C}, i)}{p_{\text{eq}}^d(\mathcal{F}_i[\mathcal{C}], i) p_{\text{end}}^d(\mathcal{C}, i)}. \quad (5.18)$$

This form of the feedback process has the benefit of not having to keep track of the potential that was used before. In each iteration a new potential is created using only the distributions of the original dynamics and the previous iteration of modified dynamics.

5.3 Feedback Process in TPS

While in the previous Chapter the modified dynamics we employed was sufficient to sample the dynamics of the CDM, it requires a significant amount of time to optimize the dynamics. The feedback method outlined in the previous section suggests us that it is possible to automate this optimisation in terms of probability distributions based on the configurations of the system visited in the sampled trajectories. Therefore, this method has the benefit of self-optimisation.

The exact feedback equation for a system with detailed balance is written as

$$U^*(\mathcal{C}) = -2 \log \frac{P_{\text{mod}}(\mathcal{C})}{P_{\text{eq}}(\mathcal{C})}, \quad (5.19)$$

where the probability distribution $P_{\text{mod}}(\mathcal{C})$, is constructed from the dynamics produced as a result of using the potential U .

In order to create a feedback process for the TPS algorithm we look at the final form of the exact feedback equation in the cloning methodology. From there we determine which assumptions and properties have been used to derive the expression Eq. (5.19) and applying these to our system, and TPS, to see if we recreate a similar procedure.

Eq. (5.19) is a powerful equation, as the only requirement during a simulation is to record the probability of a configuration. Assuming that you have access to the equilibrium probability distribution, eventually the optimal dynamics will be achieved. In the language we have been using for the TPS modified dynamics from Chapter 4, e^{U^*} is the object L_s defined in Eq. (3.55).

We can construct Eq. (5.19) with a few simple statements from our previous study:

- 1) $w_{\mathcal{C} \rightarrow \mathcal{C}'}^{\text{mod}} = \frac{l_{\mathcal{C}'}}{l_{\mathcal{C}}} e^{-S} w_{\mathcal{C} \rightarrow \mathcal{C}'}$
- 2) $P^{\text{eq}}(\mathcal{C}) w_{\mathcal{C} \rightarrow \mathcal{C}'} = P^{\text{eq}}(\mathcal{C}') w_{\mathcal{C}' \rightarrow \mathcal{C}}$
- 3) $P^{\text{mod}}(\mathcal{C}) w_{\mathcal{C} \rightarrow \mathcal{C}'}^{\text{mod}} = P^{\text{mod}}(\mathcal{C}') w_{\mathcal{C}' \rightarrow \mathcal{C}}^{\text{mod}}$

The Doob transformation tell us the form the optimal dynamics should take. Even if we do not have access to the true form of l_c , a good assumption is 1). 2) states that the original dynamics obeys detailed balance. The modified dynamics also requires detailed balance, which is given by 3).

Inserting 1) in 3), and then combining it with the detailed balanced condition of

the original dynamics 2), we obtain

$$\frac{P^{\text{mod}}(\mathcal{C})}{P^{\text{mod}}(\mathcal{C}')} = \left(\frac{l_{\mathcal{C}}}{l_{\mathcal{C}'}}\right)^2 \frac{w_{\mathcal{C}' \rightarrow \mathcal{C}}}{w_{\mathcal{C} \rightarrow \mathcal{C}'}} = \left(\frac{l_{\mathcal{C}}}{l_{\mathcal{C}'}}\right)^2 \frac{P^{\text{eq}}(\mathcal{C})}{P^{\text{eq}}(\mathcal{C}')}. \quad (5.20)$$

This can be rewritten in terms of the full probability vectors P^{eq} and P^{mod} as

$$L_s = \sqrt{\frac{P^{\text{mod}}}{P^{\text{eq}}}}. \quad (5.21)$$

To justify the feedback we can consider using the tilted dynamics, as outlined in Eq. (3.71); this has uniform values for $l_{\mathcal{C}}$. The result will provide a P^{mod} based on the tilted dynamics, which will then dictate a form of the $l_{\mathcal{C}}$; that will differ from the uniform initial values as $P^{\text{mod}} \neq P^{\text{eq}}$ due to the influence of e^{-s} . This process will repeat until the optimal dynamics is reached. At which point the $l_{\mathcal{C}}$ would be the Doob values. Thus, $P^{\text{mod}} = P^{\text{Doob}}$ so the L_s constructed by the ratio in Eq. (5.21) will correspond to the $l_{\mathcal{C}}$ already used in the dynamics.

5.3.1 Local Approximation for the Dimer Model Transitions

As already discussed, the main issue with realising Eq. (5.19) is the difficulty of constructing the P^{mod} from the numerics. In the dimer model increasing the system size and trying to enumerate all of the configurations is an expensive task. Gathering enough statistics to construct the probabilities for each configuration is numerically unfeasible. In order to utilize the feedback method we cannot use Eq. (5.21) in the exact form, just as the cloning method could not.

In a similar way to the modified dynamics used in Chapter 4, which we have extracted from the 2×2 open system, an approximation is needed to construct the $P(\mathcal{C})$. In order to obtain an approximation we use the open 2×2 problem, which can be used to tile a whole configuration of arbitrary size as a starting point. As can be seen in Fig. 3.5, we can approximate the entirety of a configuration as products

of the smaller system. We can write

$$P(\mathcal{C}) = \prod_{i,j=1}^L \sigma^{i,j}, \quad (5.22)$$

where $\sigma^{i,j}$ corresponds to zero (σ_0), single (σ_1), and double (σ_2) occupation of dimers within the plaquette located at plaquette site (i, j) . It is much simpler to numerically record the probabilities for a plaquette with the three occupation numbers than to store each total configuration.

There are definite issues associated with the approximation. We know that the system has 2^N possible tilings [104] while the open 2×2 , which contains seven states, has 7^N tilings. Consequently the mapping between the full configuration space and the constructed configuration space is clearly not one to one. Logically this can be understood because once a plaquette has a given occupation of dimers, the neighboring plaquettes are not allowed to take on all seven possible values. Due to the inconsistent mapping it is possible that, within the approximation, the dynamics that the system settles on is not the optimal one. However, the dynamics obtained with this approximation is likely to be close to the desired one due to the following: When a transition between configurations occurs, we can see that there is a small neighborhood that actually undergoes some changes. Given that we intend to construct the matrix L_s in order to compute the ratio $\frac{l_{\mathcal{C}'}}{l_{\mathcal{C}}}$, we can make some simplifications. Rather than writing the probabilities in terms of the whole configuration, using the approximation from Eq. (5.22), we can write

$$\frac{P^{\text{mod}}(\mathcal{C})}{P^{\text{mod}}(\mathcal{C}')} = \frac{\sigma_{\mathcal{C}}^{(i+1,j)} \sigma_{\mathcal{C}}^{(i-1,j)} \sigma_{\mathcal{C}}^{(i,j+1)} \sigma_{\mathcal{C}}^{(i,j-1)}}{\sigma_{\mathcal{C}'}^{(i+1,j)} \sigma_{\mathcal{C}'}^{(i-1,j)} \sigma_{\mathcal{C}'}^{(i,j+1)} \sigma_{\mathcal{C}'}^{(i,j-1)}}. \quad (5.23)$$

This shows that the majority of the configuration can be ignored when considering a transition, including the plaquette at site (i, j) , since the transition has no effect on the flippability of this plaquette. Ignoring the majority of the configuration

gives rise to a constraint on d (used to identify the size of the approximation) which provides an upper limit on the effectiveness, as the approximation only cares about the transitions, or the changes in the configuration. Given that the transition rates can be constructed as Eq. (5.23), we are only interested in small areas around flip-pable plaquettes. As the form of the building blocks to these areas is known Fig. 3.5, we can construct all of the different configurations allowed. We also therefore, lose a large amount of the error expected to be picked up in mapping the configurations into the tiled configuration space. So our approximation should yield dynamics closer than may have been initially expected to the desired one.

Fig. 5.2 shows all different possible areas of a configuration that contain a flippable plaquette at the center. Each can undergo rotations to provide a different configuration. From the configurations in Fig. 5.2, we can construct $P(\mathcal{C})_{\text{apx}}$, which is the product of the four nearest neighbour plaquettes. We can also determine which of the nine configurations can transition to any of the others. Using the ratio Eq. (5.23) we can then determine the nine associated transition rates. When written out we realised that each transition can be written in terms of powers of

$$\Omega = \frac{(\sigma_1)^2}{\sigma_0\sigma_2}. \quad (5.24)$$

This means that we can write all possible transition as

$$\frac{l_{\mathcal{C}'}}{l_{\mathcal{C}}} = L_r = \sqrt{\frac{\Omega_r^{\text{mod}}}{\Omega_r^{\text{eq}}}}, \quad (5.25)$$

where L_r is the r^{th} component of the transition weighting vector.

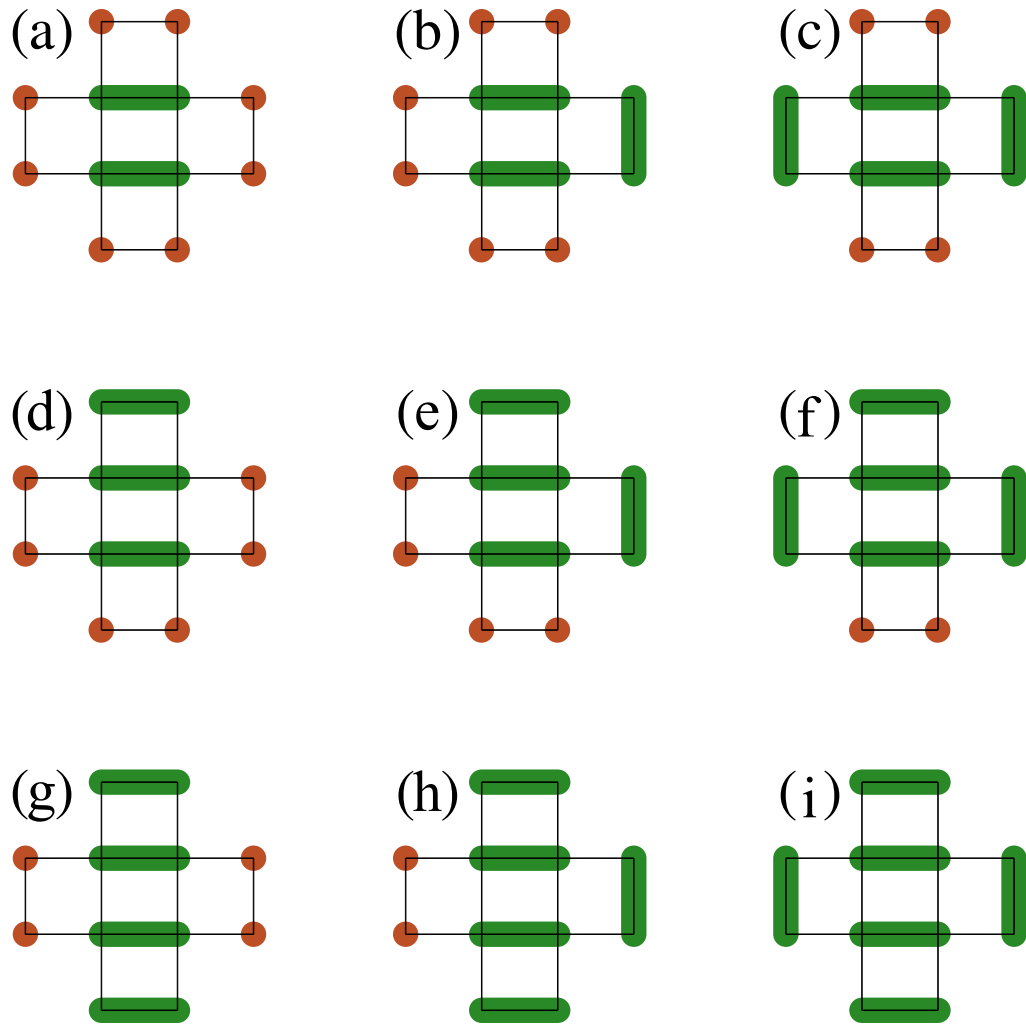


Figure 5.2: All configurations of the open 2×2 , with a flippable plaquette in the middle, when taking into account the next nearest neighbours from the central plaquette. Each row contains a different flippability, (a–c) have a single flippable plaquette, (d–f) have two flippable plaquettes and (g–i) have three, the maximum number of flippable plaquettes. Each configuration can undergo one or more rotations of 45° to obtain configurations of similar construction in terms of σ_0 , σ_1 and σ_2 .

5.3.2 Algorithm: Feedback TPS

With the approximation outlined, we can now begin to implement the feedback dynamics using the following algorithm:

- 1) Record the probability distribution of the seven 2×2 states for the original dynamics P^{eq} (TPS is not required) and proceed to using TPS from an alternative dynamics, e.g. the stochastic form of the tilted MO as a starting point.
- 2) Using the alternative dynamics record the probability distribution of the seven 2×2 states, P^{mod} .
- 3) Construction of the values of the transition rates using

$$\frac{l_{c'}}{\bar{l}_c} = a \sqrt{\frac{\Omega_r^{\text{mod}}}{\Omega_r^{\text{eq}}}}, \quad (5.26)$$

where we have introduced a parameter, a , that can be tuned, in order to stabilize the convergence.

- 4) Repeat step 2 and onwards using the new rates.

Ideally the first step is not required and the distribution is known exactly. However if this is not the case then recording the data at the start of the simulation works well, assuming the error is sufficiently low.

5.4 Numerical Comparisons Between Standard TPS and the Modified Dynamics Methods

The next task is implementing the feedback method in TPS Eq. (5.21) and outlining the approximation we needed for the CDM Eq. (5.25). A good place to start is to make sure that the method holds up in the case where we already have results. We know the behavior for several system sizes from our work in the previous Chapter. Fig. 5.3 shows the results from both modified dynamics methods. The re-

5.4. NUMERICAL COMPARISONS BETWEEN STANDARD TPS AND THE MODIFIED DYNAMICS METHODS

sults coincide (within error) indicating that the feedback method recovers the correct behaviour.

The difference in resources between the methods are as follows: Note for all cases that there was an equilibration period of 10^3 , starting with the modified dynamics method from Chapter 4. For all system sizes, the computational requirements were similar; 100 independent iterations of the simulations were run, each having attempted 10^5 trajectory updates. Once collected, to remove any correlation effects, a random subset of 10^3 was taken from each of the independent iterations, producing a total of 10^5 independent samples. The smaller system sizes, $L = 6$ and $L = 8$, did not require this level of detail.




Figure 5.3: The comparison of the activity as a function of s for the two alternative dynamics methods. The dashed lines are the modified dynamics taken from the previous Chapter, and the markers are the feedback results. The two methods coincide with each other.

In the case of the feedback method, each iteration had 10 “cycles”, updates to the dynamics. There were 10 independent iterations of the smaller system sizes, $L = 6$ & 8, and 100 used for the larger sizes, $L = 16$ & 24. This is typically achieved after the first two cycles. Each cycle contains 10^4 attempted trajectory updates. All cycles produce useful results, as each trajectory has its own re-weighting factor.

5.4. NUMERICAL COMPARISONS BETWEEN STANDARD TPS AND THE MODIFIED DYNAMICS METHODS

Therefore we again have 10^5 samples, from which we take a sub-set of 10^3 . This makes the larger system sizes equivalent in both cases. The smaller systems then have an order of magnitude fewer samples.




Figure 5.4: Comparison of the acceptance rate (of the different TPS methods) as a function of s , all scaled by the acceptance rate of the feedback method. The blue line shows the standard TPS efficiency. The orange dashed line is the improved efficiency of the modified dynamics. The yellow line is the efficiency of the feedback method, which lies very close to that of the modified dynamics. These results are for a system size of $L = 12$. The insert is the same data not scaled by the acceptance rate of the feedback method.

The acceptance rates of the feedback algorithm are shown in comparison to the modified dynamics and the standard TPS rates in Fig. 5.4. It is clear that both modified dynamics have a higher acceptance rate than the standard TPS counterpart. The comparison between the two new types of dynamics is close, which makes sense when considering that both are built of similar approximations. The feedback process does appear to be slightly more efficient as the value of s moves further from the transition region.

The acceptance rate does inform nicely as to which method is more efficient (the

modified dynamics over the original), but does not reveal the whole picture. More information can be gleaned from inspecting the accepted values of τ_{cut} . Recall τ_{cut} is the length of the edit to the trajectory in the shifting algorithm. In Fig. 5.5 the values of τ_{cut} that are accepted are shown for various system sizes and values of s . What can be seen immediately is the values of $s < 0$ have a much higher acceptance rate than, and are more uniformly spread, than $s > 0$.

However, in both cases it is clear that as system size increases the uniform acceptance of τ_{cut} is lost. The loss of the uniform distribution of τ_{cut} acceptance comes from the trajectories containing too much information to edit easily; that is to say, because the Doob dynamics is not being used, the trajectories being produced are not sampled from equilibrium from the point of view of that value of s . This makes it harder to propose large parts of trajectory to replace the current more favorable one. This problem has been identified multiple times throughout this thesis. The point of showing these distributions is to see the improvement on the uniformity of the distributions of τ_{cut} acceptance for the feedback method over the original dynamics. The best place to see this is the larger values of s in Fig. 5.5. The feedback distribution (green), has many more accepted τ_{cut} values in the center than the TPS distributions (blue), implying larger trajectory updates are accepted more readily with the feedback method.




Figure 5.5: Comparison of the times selected to be cut τ_{cut} and updated, for various system sizes and values of s . These plots include samples of forward and backward shifting. There is a favour to editing the ends of the trajectories, which shows a similarity to the cartoon Fig. 5.1. The blue histogram represents the standard TPS method, while the green histogram represents the modified, feedback dynamics. The red curve corresponds to the method outlined in Chapter 4. From top to bottom each row corresponds to $L = 6$, $L = 8$ and $L = 12$. Column (a) is for $s = -0.2$, column (b) is for $s = 0.2$ and column (c) is for $s = 0.6$. Each set of data consists of 10^5 independent samples.

The results shown in Fig. 5.5 do not indicate the length of the trajectory update; they simply show the distribution of τ_{cut} . This does not take into account the

5.4. NUMERICAL COMPARISONS BETWEEN STANDARD TPS AND THE MODIFIED DYNAMICS METHODS

direction of the shifting algorithm. To confirm that the trajectory updates are in fact small, and near the ends of the trajectories, Fig. 5.6 shows how the length of the update changes as a function of both system size and s for the different dynamics.

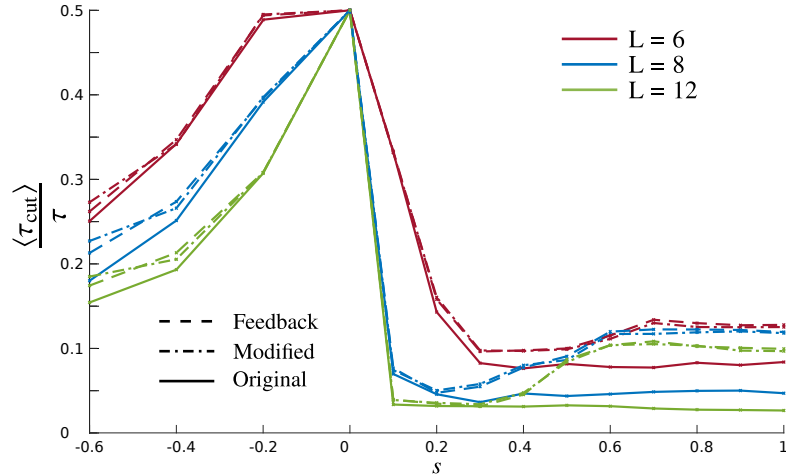


Figure 5.6: Comparison of the modified dynamics of the CDM against the original dynamics of the CDM. This plot shows the average length of a updated trajectory, $\langle \tau_{\text{cut}} \rangle$, (scaled by the total length of the trajectory, τ) to be accepted as an update for various values of s and for different system sizes.

If we look at an example of the TPS trajectory, which is to say the value of the intensive activity of the current trajectory as a function of TPS updates, we can see again how the modified dynamics improve upon the standard TPS. In the case of studying dynamical phase transitions, it can also be hard to sample away from the phase transition on the opposite side to the standard dynamics. This is mainly due to trajectories, in standard TPS, being produced using the dynamics from one side of the phase transition. So one would expect sampling on the same side of the transition to be “easy”, and that is precisely what we see from Fig. 5.5 and Fig. 5.7. What we also see from Fig. 5.7 is that the modified dynamics (yellow-feedback and orange-Chapter 4) typically explore the same region of values for the activity as the standard dynamics (blue).

5.4. NUMERICAL COMPARISONS BETWEEN STANDARD TPS AND THE MODIFIED DYNAMICS METHODS

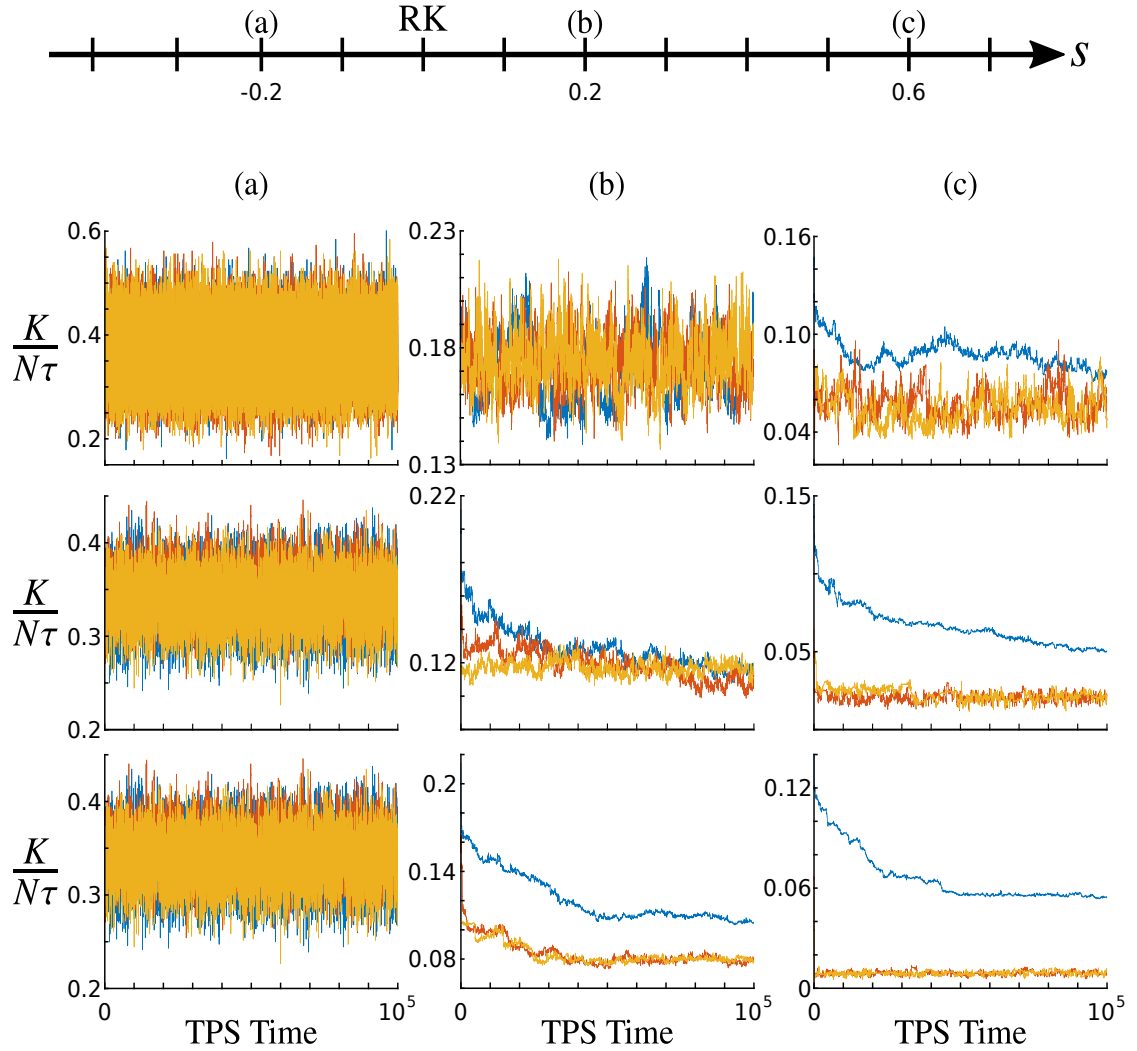


Figure 5.7: Examples of a comparison between the TPS trajectories for various system sizes and values of s . The blue curve is the standard TPS method, the orange is the modified dynamics used in Chapter 4, and the yellow is the modified, feedback dynamics. From top to bottom each row corresponds to $L = 6$, $L = 8$ and $L = 12$. Column (a) is for $s = -0.2$, column (b) is for $s = 0.2$ and column (c) is for $s = 0.6$.

However when the other side of the transition is explored, this coincidence ceases. For $L = 6$ we know that the transition region is not particularly clear, but based on Fig. 5.3 we can assume that $s = 0.6$ is on the less active side of the transition, and this is where Fig. 5.7 shows a difference in results between the modified dynamics and the

standard one. For the larger system size, $s = 0.2$ starts to show this disagreement too, as the transition moves towards $s = 0$, but the TPS trajectories for $s = 0.6$ clearly show the standard TPS fails to equilibrate (or appears to equilibrate incorrectly) whereas this is not true of the modified dynamics methods.

5.4.1 A Property of the Feedback Method

Having explored the differences between the methods' efficiencies, we can now highlight a feature that is specific to the feedback method. We outlined the construction of the transitions for the approximation by recording the probability distribution of the 2×2 configurations. There are seven states outlined in Fig. 3.5 that contribute to this distribution, but due to a rotational symmetry averaging out over enough realizations, the contributions reduce to the three values outlined in section 3.6.1. This can clearly be seen in the subplots of Fig. 5.8. In each subplot it is clear that the three different states (none, one or two dimers) settle to a steady value. The way this value changes as a function of s can be seen in the main plot.

A striking feature that is also noted in [27], is that on average, it takes around two to three updates for the dynamics to settle into a stable probability distribution. This is ideal because while all the data collected can be used as samples, the quicker the distributions settle the quicker the samples will approximate the average behavior for a given value of s . They are all initially set to a uniform value then allowed to evolve according to Eq. (5.26). As can be seen from the $s = 0$ case, the dynamics immediately behave as intended, albeit this is due to the influence of the transitions is affected by the value of s . Even when $s = 1$, we can see the stability is reached within 3 iterations.

5.4. NUMERICAL COMPARISONS BETWEEN STANDARD TPS AND THE MODIFIED DYNAMICS METHODS

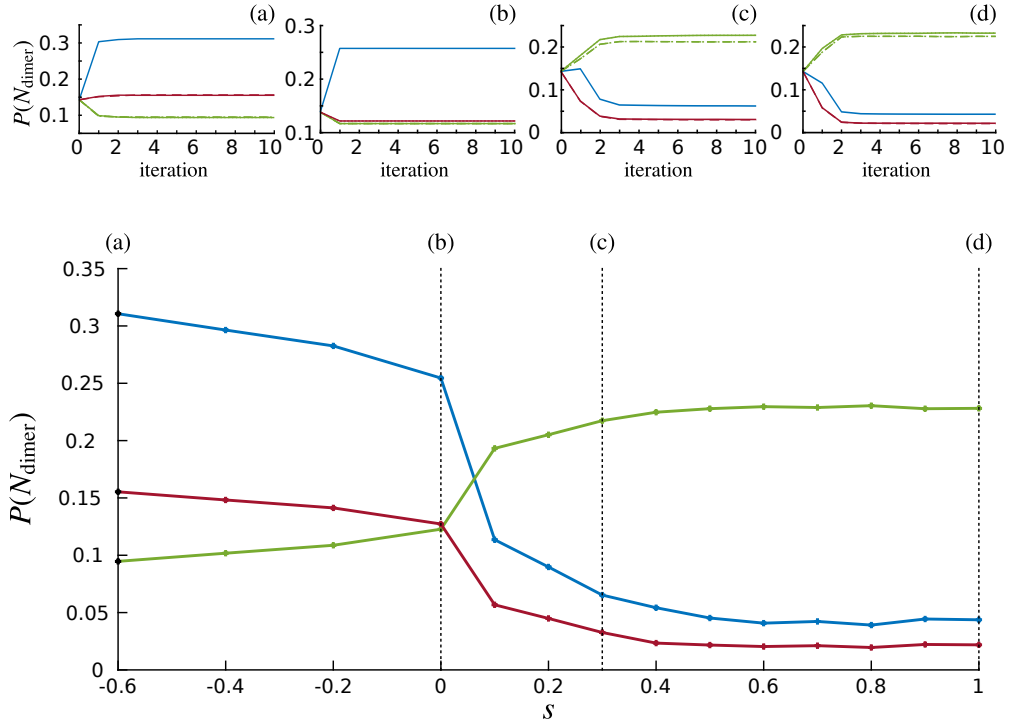


Figure 5.8: The probabilities of zero (blue), one (green, averaged over the four states) and two (red, averaged over the two states) dimers as a function of s . In each case the value is an average over all orientations, see Fig. 3.5. In each subplot the evolution of seven different states (with different orientations being shown by different line styles), discussed in 3.6.1, are shown as a function of iteration. Each subplot is taken to show a different feature. (a) is $s = -0.6$ is on the left of the transition region. (b) is $s = 0$ show the equilibrium distribution of the standard dynamics. (c) is $s = 0.3$ which is in the center of the transition region. (d) is $s = 1$ to show the distributions away from the transition region.

For the system size $L = 12$, $s = 0.3$ is the closest value to the area of the phase transition shown in Fig. 5.8(c). Naturally this required more samples to be taken in order to match accuracy of those away from the transition. For both of the cases $s = 0$ and $s = 1$ there were 10 independent samples, each of which contained 10 updates. Whereas, the distributions required 100 samples for $s = 0.3$, each of which contained 10 updates. An update was done after 10^5 attempted trajectory changes were made. For each attempt made to change a trajectory, the 2×2 probability distribution was added to by sampling the configuration at time $\frac{\tau}{2}$. Each site was used in collecting

the statistics. The possibility of a correlation problem, mentioned in section 5.2.1, was considered. After checking that the values of the distributions agreed, the distribution took less time to construct when using the whole configuration rather than the single site.

5.5 Conclusion

The motivation of this Chapter was to show that the feedback algorithm, first utilized for the cloning method [27], is applicable to the TPS method. In previous work (see Chapter 4) we had developed a similar, but less systematic, modified dynamics approach. We have shown how the feedback method agrees with the results collected by the method outlined in Chapter 4 for the CDM. Looking at $s > 0$ which, from the previous study of the CDM, we know is a difficult region to explore, we can see the improvement in efficiency for the modified dynamics over the original ones. The difficulty on the right of the RK point, comes from the phase transition becoming a crossover in a region of $s > 0$ for the smaller system sizes. The transition region has a notable effect on sampling. A good example of this issue is shown by the acceptance rates. The acceptance rate from samples collected in and around the transition region drop in comparison to those further from it. This feature is echoed in all methods.

With the feedback method shown to work, we compared the efficiency of the technique to the more commonly used TPS method utilising original dynamics. The motivation to adapt the feedback process to TPS was due to the poor ability to explore the dominant dynamics away from $s = 0$. We have shown, in this case, that the dominant dynamics at $s = 0$ is from the active side; this has lead to TPS using original dynamics failing to correctly characterise the inactive side of the phase diagram. Switching the sampling to a more appropriate modified dynamics allows for the true behaviours of the inactive phase to be revealed. To enforce this point, we

see that all dynamics converge to the left of the RK point, the active phase, which is dominant at $s = 0$.

Furthermore we have shown how the efficiency is increased using modified dynamics. The work in Chapter 4 alluded to this, but this Chapter shows the gain in sampling efficiency alongside the faster decorrelation of trajectories due to the ability to make large edits to the trajectories. The feedback method adds further benefits of being able to see what the “best” dynamics are, given no initial knowledge of those dynamics, and the ability to see when these dynamics have been reached. As the feedback dynamics are based around a probability distribution of configurations, the form of such a distribution is likely to indicate some structure to the configurations of different values of s .

The feedback method is somewhat more useful than the dynamics used in Chapter 4, as there is no need to do an exploration of the landscape for the variable D_s . Along side that, once the value for D_s is chosen then it is fixed, whereas the feedback process allows for more exploration of the dynamics space.

6

Summary

The basic outline of this thesis can be broken down into three sections: First, the dynamical description of the classical dimer model (CDM), using a collective understanding of how the local dynamics come together to act as larger objects, which we call strings, as the system relaxes. This study gave rise to the second section which links the dynamical behavior of the CDM to the ground state of the quantum dimer model (QDM).

The different types of behaviours exhibited by trajectories can be seen even from simple trajectory production. Using the Rokhsar-Kivelson mapping, the information known about the QDM identified the existence of a possible dynamical phase transition in the CDM. Through the use of LD theory the RK mapping could be extended beyond the RK point, allowing for much more of the QDM ground state to be made

accessible through numerical study of the CDM. This yielded a detailed picture of the phase diagram of the QDM. The picture was in agreement with other numerical studies that the disputed region (to the left of the RK point) is dominated by the columnar phase. Although it does not act simply as the columnar phase due to some emergent $SO(2)$ symmetry effects.

This gave rise to the third idea explored in this thesis, utilizing alternative dynamics in trajectory ensemble studies. The difficulty of trying to explore even small system sizes of the CDM using the standard TPS techniques proved that this algorithm is too inefficient. Only through the use of an alternative dynamics did the study become possible. Utilising the Doob transformation allowed us to identify a way to modify the dynamics used in sampling, which gave us the required efficiency. Due to this being an area of interest to the LD community, we improved on our method by adapting a feedback process initially created for a different LD numerical technique, cloning. In comparison to the original TPS method we saw a marked improvement in efficiency, alongside some major benefits including, but not limited to, the ability to extract the form of the dynamics that is close to the Doob dynamics. With this knowledge it would be possible to explore different systems with the dynamics now known to represent the rare trajectories of the system of interest.

A similar study to the one done in Chapter 2 could be made on the CDM for different lattices, or even non-fully packed models. The mapping of the classical dynamical phases to the ground state of the appropriate quantum model, Chapter 4, could either be taken to a different lattice structure or dimension, or even to the non-fully packed CDM. Sticking with the fully packed square 2D CDM, the features of the LD function could be interesting. Due to the fact that the large deviation function θ is required to be a convex function, we expect the function to follow a shape presented in Fig. 6.1.

Given the results of the study in Chapter 4 we know the QDM is dominated by

the staggered phase for $s > 0$, which exists in the maximal flux sector, while $s < 0$ is dominated by the columnar ordering, which exists in the zeroth flux sector. This tells us that for $s > 0$ the LDF should be zero, due to it being incapable of undergoing the local dynamics, while for $s < 0$ the LDF is described by the one identified in this thesis, shown in blue in Fig. 6.1. The way to test this theory is to do the same study with supplemented TPS methods only on the open boundary fully packed CDM. In this model dimers are allowed to “point out” of the edges of the lattice, which is equivalent to introducing monomers isolated to sites at the boundaries of the lattice. This requires new dynamics that can be thought of as adsorption and desorption of dimers the sites around the edge, similar to the works of Stannard et al [16]. The open boundary conditions then allow for the constraints of the flux sector to be relaxed as movement between them would be possible through the new dynamics. With the new model and using the same methods as before, it would be simple to see whether the LDF does indeed exhibit the predicted behaviour shown in Fig. 6.1.




Figure 6.1: The LD function, or free energy as a function of s . The blue curves show results from the zero flux sector MC method using TPS feedback, where the system size increases the convexity of the curve from $L = 8$ (dot-dash), and $L = 16$ (dash) to $L = 24$ (solid). Given the maximal flux sector has a known configuration of zero flippable plaquettes, it will have zero activity for all values of s . Therefore it has a LDF represented by the black line.

The methods outlined in Chapters 4 and 5 are compatible with each other. Taking this further there are works outlining an error guiding function [105], which could be incorporated into the dynamics during the simulations to help the efficiency of the modified dynamics.

Having the ability to access dynamics that are not properly represented by the typical dynamics of the problem is quite powerful. One of the major struggles that arises during dynamical ensemble studies is caused by sampling trajectories without using the dynamics associated with them. While the structure of LD theory means that eventually all numerical methods will produce the desired results, the timescales for convergence are often infeasible. The alternative dynamics approach can be used to speed up the numerics.

Another feature of the feedback method is the gain of information about the dynamics. The method does not require much knowledge of the true alternative dynamics, and so any guess can be made at the start of the simulations and the feedback nature will optimize until the best dynamics is achieved. This is especially useful in cases where the goal is not only to study the dynamical behaviours, but to reproduce them. Once the dynamics has been determined in the simulation then they can be implemented in real world situations. A prime example for an application of this method is a set of simple small drones used to explore rubble in order to find survivors after a disaster. Determining the optimal dynamical behavior to find more survivors in the quickest time can be simulated using the feedback TPS methodology would not only allow for the identification of such dynamics but would provide the basic structure required to replicate it. This example may seem a little bit of a stretch from the dimer case outlined in this thesis. None the less it is a dynamical system that would benefit from this type of work.

Also an adaption of the feedback or effective dynamics method could be applied to alternative ensembles; namely the x -ensemble, outlined in [106], where it is shown this ensemble is equivalent to the s -ensemble and yet numerically can be less computationally expensive. Utilizing the alternative dynamics here might see a greater improvement still.

The main message of this thesis is the wealth of information that can be accessed when alternative dynamics are considered in exploring dynamical ensembles. Using the dimer model as a case study, we have shown that a full description of a dynamical model is not only more complex than simply extending a static description, but it has a large amount of untapped information when the dynamics are truly explored. While each level of description provides information it usually misses some understanding. The static description is a valid approximation in certain limits, however the dynamical one we outlined extends beyond even its expected limit.

Going further than that, the dynamical description did indicate the existence of multiple dynamical phases, but offered no way with which to explore them. Using modified dynamics to detail the rare events of such a model enlightened the behaviors that dominate in equilibrium. Identifying a first order dynamical phase transition in the case of the CDM showed how the typical dynamic behavior was dominated by one of two different phases. To truly understand the model both phases should be understood. Furthermore in the case such as the CDM, where an RK point exists to link it to the QDM, we can now use relatively straightforward classical numerical methods to shed light on the, typically harder to simulate, quantum counterpart.

Bibliography

- [1] L. Peliti. *Statistical Mechanics in a Nutshell*. Princeton University Press, 2011.
- [2] D. Chandler. *Introduction to modern statistical mechanics*. Oxford University Press, 1987.
- [3] C. Dellago, P. Bolhuis, and D. Chandler, “Efficient transition path sampling: Application to Lennard-Jones cluster rearrangements”, *Journal of Chemical Physics*, vol. 108, p. 9236, 1998.
- [4] C. Bennett. *Algorithms for Chemical Computations (Series No. 46)*. ACS Publications, 1977.
- [5] A. Grossfield and D. M. Zuckerman, “Quantifying uncertainty and sampling quality in biomolecular simulations”, *Annual reports in computational chemistry*, vol. 5, p. 23, 2009.
- [6] N. Kumar, H. Soni, S. Ramaswamy, and A. Sood, “Flocking at a distance in active granular matter”, *Nature Communications*, vol. 5, p. 4688, 2014.
- [7] J. Toner and Y. Tu, “Flocks, herds, and schools: A quantitative theory of flocking”, *Physical Review E*, vol. 58, p. 4828, 1998.
- [8] S. Ramaswamy, “The mechanics and statistics of active matter”, *Annual Review of Condensed Matter Physics*, vol. 1, p. 323, 2010.
- [9] R. Bouffanais. *Design and Control of Swarm Dynamics*. Springer, 2016.

- [10] N. Elkies, G. Kuperberg, M. Larsen, and J. Propp, “Alternating-sign matrices and domino tilings (part i)”, *Journal of Algebraic Combinatorics*, vol. 1, no. 2, p. 111, 1992.
- [11] H. Cohn, J. Propp, and N. Elkies, “Local statistics for random domino tilings of the aztec diamond”, *Duke Mathematics Journal*, vol. 85, no. 0008243, p. 117, 1996.
- [12] W. Jockusch, J. Propp, and P. Shor, “Random domino tilings and the arctic circle theorem”, *arXiv preprint math/9801068*, 1998.
- [13] J. Propp, “Generalized domino-shuffling”, *Theoretical Computer Science*, vol. 303, p. 267, 2003.
- [14] P. W. Kasteleyn, “The statistics of dimers on a lattice: I. the number of dimer arrangements on a quadratic lattice”, *Physica*, vol. 27, no. 12, p. 1209, 1961.
- [15] H. N. Temperley and M. E. Fisher, “Dimer problem in statistical mechanics-an exact result”, *Philosophical Magazine*, vol. 6, no. 68, p. 1061, 1961.
- [16] A. Stannard, J. C. Russell, M. O. Blunt, C. Salesiotis, M. del Carmen Giménez-López, N. Taleb, M. Schröder, N. R. Champness, J. P. Garrahan, and P. H. Beton, “Broken symmetry and the variation of critical properties in the phase behaviour of supramolecular rhombus tilings”, *Nature Chemistry*, vol. 4, no. 2, p. 112, 2012.
- [17] A. Baule and H. A. Makse, “Fundamental challenges in packing problems: from spherical to non-spherical particles”, *Soft Matter*, vol. 10, no. 25, p. 4423, 2014.
- [18] J. Garrahan and D. Chandler, “Geometrical explanation and scaling of dynamical heterogeneities in glass forming systems”, *Physical Review Letters*, vol. 89, p. 35704, 2002.
- [19] L. Berthier and G. Biroli, “Theoretical perspective on the glass transition and amorphous materials”, *Reviews of Modern Physics*, vol. 83, p. 587, 2011.

- [20] F. Ritort and P. Sollich, “Glassy dynamics of kinetically constrained models”, *Advances in Physics*, vol. 52, no. 4, p. 219, 2003.
- [21] D. Chandler and J. P. Garrahan, “Dynamics on the way to forming glass: Bubbles in space-time”, *Annual Review of Physical Chemistry*, vol. 61, p. 191, 2010.
- [22] B. Derrida, “Non-equilibrium steady states: fluctuations and large deviations of the density and of the current”, *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2007, no. 07, p. P07023, 2007.
- [23] T. Oakes, J. P. Garrahan, and S. Powell, “Emergence of cooperative dynamics in fully packed classical dimers”, *Physical Review E*, vol. 93, p. 032129, 2016.
- [24] C. L. Henley, “Relaxation time for a dimer covering with height representation”, *Journal of Statistical Physics*, vol. 89, p. 483, 1997.
- [25] D. Ruelle. *Thermodynamic Formalism*. Cambridge University Press, 2004.
- [26] U. Ray, G. K.-L. Chan, and D. T. Limmer, “Importance sampling large deviations in nonequilibrium steady states. i”, *The Journal of Chemical Physics*, vol. 148, no. 12, p. 124120, 2018.
- [27] T. Nemoto, F. Bouchet, R. L. Jack, and V. Lecomte, “Population-dynamics method with a multicanonical feedback control”, *Physical Review E*, vol. 93, no. 6, p. 062123, 2016.
- [28] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, “Dynamical first-order phase transition in kinetically constrained models of glasses”, *Physical Review Letters*, vol. 98, p. 195702, 2007.
- [29] L. O. Hedges, R. L. Jack, J. P. Garrahan, and D. Chandler, “Dynamic order-disorder in atomistic models of structural glass formers”, *Science*, vol. 323, p. 1309, 2009.
- [30] T. Nemoto, R. L. Jack, and V. Lecomte, “Finite-size scaling of a first-order dynamical phase transition: Adaptive population dynamics and an effective

- model”, *Physical Review Letters*, vol. 118, no. 11, 2017.
- [31] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, “Transition path sampling: Throwing ropes over rough mountain passes, in the dark”, *Annual Review of Physical Chemistry*, vol. 53, p. 291, 2002.
- [32] H. W. J. Blöte and H. J. Hilhorst, “Roughening transitions and the zero-temperature triangular Ising antiferromagnet”, *Journal of Physics A*, vol. 15, 1982.
- [33] D. S. Rokhsar and S. A. Kivelson, “Superconductivity and the quantum hardcore dimer gas”, *Physical Review Letters*, vol. 61, p. 2376, 1988.
- [34] V. Lecomte, C. Appert-Rolland, and F. van Wijland, “Thermodynamic formalism for systems with Markov dynamics”, *Journal of Statistical Physics*, vol. 127, p. 51, 2007.
- [35] T. R. Gingrich and P. L. Geissler, “Preserving correlations between trajectories for efficient path sampling”, *Journal of Chemical Physics*, vol. 142, p. 234104, 2015.
- [36] M. E. Fisher, “Statistical mechanics of dimers on a plane lattice”, *Physical Review*, vol. 124, no. 6, p. 1664, 1961.
- [37] J. Jacobsen, S. Ouvry, V. Pasquier, D. Serban, and L. F. Cugliandolo. *Exact Methods in Low-dimensional Statistical Physics and Quantum Computing: Lecture Notes of the Les Houches Summer School*. Oxford University Press, 2010.
- [38] C. L. Henley, “The “coulomb phase” in frustrated systems”, *Annual Review of Condensed Matter Physics*, vol. 1, p. 179, 2010.
- [39] D. A. Huse, W. Krauth, R. Moessner, and S. Sondhi, “Coulomb and liquid dimer models in three dimensions”, *Physical Review Letter*, vol. 91, p. 167004, 2003.

- [40] G. Biroli and J. P. Garrahan, “Perspective: The glass transition”, *Journal of Chemical Physics*, vol. 138, p. 12A301, 2013.
- [41] D. Das, G. Farrell, J. Kondev, and B. Chakraborty, “Critical dynamics of dimers: Implications for the glass transition”, *Journal of Physical Chemistry B*, vol. 109, p. 21413, 2005.
- [42] J. P. Garrahan, A. Stannard, M. O. Blunt, and P. H. Beton, “Molecular random tilings as glasses”, *Proceedings of the National Academy of Sciences USA*, vol. 106, p. 15209, 2009.
- [43] L. D. C. Jaubert and P. C. W. Holdsworth, “Signature of magnetic monopole and Dirac string dynamics in spin ice”, *Nature Physics*, vol. 5, p. 258, 2009.
- [44] C. Castelnovo, R. Moessner, and S. L. Sondhi, “Spin ice, fractionalization, and topological order”, *Annual Review of Condensed Matter Physics*, vol. 3, p. 35, 2012.
- [45] S. M. Bhattacharjee and J. F. Nagle, “Finite-size effect for the critical point of an anisotropic dimer model of domain walls”, *Physical Review A*, vol. 31, p. 3199, 1985.
- [46] S. M. Bhattacharjee, “Crossover in an exactly solvable dimer model of domain walls with dislocations”, *Physical Review Letters*, vol. 53, p. 1161, 1984.
- [47] S. M. Bhattacharjee and J. J. Rajasekaran, “Absence of anomalous dimension in vertex models: Semidilute solution of directed polymers”, *Physical Review A*, vol. 44, p. 6202, 1991.
- [48] L. D. C. Jaubert, J. T. Chalker, P. C. W. Holdsworth, and R. Moessner, “Three-dimensional Kasteleyn transition: Spin ice in a [100] field”, *Physical Review Letters*, vol. 100, p. 067207, 2008.
- [49] H. Otsuka, “Classical dimer model with anisotropic interactions on the square lattice”, *Physical Review E*, vol. 80, p. 011140, 2009.

- [50] S. Powell and J. T. Chalker, “Classical to quantum mappings for geometrically frustrated systems: Spin-ice in a [100] field”, *Physical Review B*, vol. 78, p. 024422, 2008.
- [51] D. P. Landau and K. Binder. *A guide to Monte Carlo simulation in statistical physics*. Cambridge University Press, 2009.
- [52] S. F. Edwards and D. R. Wilkinson, “The surface statistics of a granular aggregate”, *Proceedings of the Royal Society of London A*, vol. 381, p. 17, 1982.
- [53] A. J. Bray, S. N. Majumdar, and G. Schehr, “Persistence and first-passage properties in nonequilibrium systems”, *Advances in Physics*, vol. 62, p. 225, 2013.
- [54] J. Franke and S. N. Majumdar, “Survival probability of an immobile target surrounded by mobile traps”, *Journal of Statistical Mechanics*, vol. 2012, p. 5024, 2012.
- [55] A. S. Keys, L. O. Hedges, J. P. Garrahan, S. C. Glotzer, and D. Chandler, “Excitations are localized and relaxation is hierarchical in glass-forming liquids”, *Physical Review X*, vol. 1, p. 21013, 2011.
- [56] D. J. P. Morris, D. A. Tennant, S. A. Grigera, B. Klemke, C. Castelnovo, R. Moessner, C. Czternasty, M. Meissner, K. C. Rule, J. U. Hoffmann, K. Kiefer, S. Gerischer, D. Slobinsky, and R. S. Perry, “Dirac strings and magnetic monopoles in the spin ice $\text{Dy}_2\text{Ti}_2\text{O}_7$ ”, *Science*, vol. 326, p. 411, 2009.
- [57] H. M. Revell, L. R. Yaraskavitch, J. D. Mason, K. A. Ross, H. M. L. Noad, H. A. Dabkowska, B. D. Gaulin, P. Henelius, and J. B. Kycia, “Evidence of impurity and boundary effects on magnetic monopole dynamics in spin ice”, *Nature Physics*, vol. 9, p. 34, 2013.
- [58] C. L. Henley, “From classical to quantum dynamics at Rokhsar–Kivelson points”, *Journal of Physics: Condensed Matter*, vol. 16, p. S891, 2004.

- [59] A. M. Läuchli, S. Capponi, and F. F. Assaad, “Dynamical dimer correlations at bipartite and non-bipartite Rokhsar–Kivelson points”, *Journal of Statistical Mechanics*, vol. 8, p. 153, 1973.
- [60] H. Touchette, “The large deviation approach to statistical mechanics”, *Physics Reports*, vol. 478, p. 1, 2009.
- [61] E. Pitard, V. Lecomte, , and F. van Wijland, “Dynamic transition in an atomic glass former: A molecular-dynamics evidence”, *Europhysics Letters*, vol. 96, no. 5, p. 56002, 2011.
- [62] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, and K. van Duijvendijk, “First-order dynamical phase transition in models of glasses: an approach based on ensembles of histories”, *Journal of Physics A: Mathematical and Theoretical*, vol. 42, no. 7, p. 075007, 2009.
- [63] M. Merolle, J. P. Garrahan, and D. Chandler, “Space–time thermodynamics of the glass transition”, *Proceedings of the National Academy of Sciences USA*, vol. 102, p. 10837, 2005.
- [64] C. Giardinà, V. Kurchan, J Lecomte, , and J. Tailleur, “Simulating rare events in dynamical processes”, *Journal of Statistical Physics*, vol. 145, no. 4, p. 787, 2011.
- [65] R. L. Jack and P. Sollich, “Large deviations and ensembles of trajectories in stochastic models”, *Progress of Theoretical Physics Supplement*, vol. 184, p. 304, 2010.
- [66] C. Gardiner. *A Handbook of Stochastic Methods*. Springer, 2004.
- [67] M. Baiesi, C. Maes, and B. Wynants, “Fluctuations and response of nonequilibrium states”, *Physical Review Letters*, vol. 103, no. 1, p. 010602, 2009.
- [68] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Springer, 1998.

- [69] Y. S. Elmatad, R. L. Jack, D. Chandler, and J. P. Garrahan, “Finite-temperature critical point of a glass transition”, *Proceedings of the National Academy of Sciences*, vol. 107, no. 29, p. 12793, 2010.
- [70] T. Speck, A. Malins, and C. P. Royall, “First-order phase transition in a model glass former: coupling of local structure and dynamics”, *Physical Review Letters*, vol. 109, no. 19, p. 195703, 2012.
- [71] M. R. Shirts and J. D. Chodera, “Statistically optimal analysis of samples from multiple equilibrium states”, *The Journal of Chemical Physics*, vol. 129, no. 12, p. 124105, 2008.
- [72] C. H. Bennett, “Efficient estimation of free energy differences from Monte Carlo data”, *Journal of Computational Physics*, vol. 22, p. 245, 1976.
- [73] T. Zhiqiang, “On a likelihood approach for Monte Carlo integration”, *Journal of the American Statistical Association*, vol. 99, p. 1027, 2004.
- [74] R. Chetrite and H. Touchette, “Nonequilibrium markov processes conditioned on large deviations”, in *Annales Henri Poincaré*, vol. 16, p. 2005, Springer, 2015.
- [75] J. P. Garrahan, “Classical stochastic dynamics and continuous matrix product states: gauge transformations, conditioned and driven processes, and equivalence of trajectory ensembles”, *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2016, no. 7, p. 073208, 2016.
- [76] K. Klymko, P. L. Geissler, J. P. Garrahan, and S. Whitelam, “Rare behavior of growth processes via umbrella sampling of trajectories”, *Physical Review E*, vol. 97, no. 3, p. 032123, 2018.
- [77] V. Lecomte and J. Tailleur, “A numerical approach to large deviations in continuous time”, *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2007, no. 03, p. P03004, 2007.

- [78] T. Oakes, S. Powell, C. Castelnovo, A. Lamacraft, and J. P. Garrahan, “Phases of quantum dimers from ensembles of classical stochastic trajectories”, *Physical Review B*, vol. 98, p. 064302, 2018.
- [79] P. W. Anderson, “Resonating valence bonds: A new kind of insulator?”, *Materials Research Bulletin*, vol. 8, no. 2, p. 153, 1973.
- [80] P. Fazekas and P. Anderson, “On the ground state properties of the anisotropic triangular antiferromagnet”, *Philosophical Magazine*, vol. 30, no. 2, p. 423, 1974.
- [81] R. Moessner, K. S. Raman, P. Mendels, and F. Mila. *Quantum dimer models in: Introduction to Frustrated Magnetism*. Springer, 2011.
- [82] J. T. Chalker, C. Chamon, M. O. Goerbig, R. Moessner, and L. F. Cugliandolo. *Spin liquids and frustrated magnetism: Lecture Notes of the Les Houches Summer School*. Oxford University Press, 2014.
- [83] L. Balents, “Spin liquids in frustrated magnets”, *Nature*, vol. 464, no. 7286, p. 199, 2010.
- [84] D. Banerjee, M. Bögli, C. P. Hofmann, F. J. Jiang, P. Widmer, and U. J. Wiese, “Interfaces, strings, and a soft mode in the square lattice quantum dimer model”, *Physical Review B*, vol. 90, p. 245143, 2014.
- [85] O. F. Syljuåsen, “Plaquette phase of the square-lattice quantum dimer model: Quantum Monte Carlo calculations”, *Physical Review B*, vol. 73, no. 24, p. 245105, 2006.
- [86] A. M. Polyakov, “Quark confinement and topology of gauge theories”, *Nuclear Physics B*, vol. 120, no. 3, p. 429, 1977.
- [87] E. Fradkin, D. A. Huse, R. Moessner, V. Oganesyan, and S. L. Sondhi, “Bipartite Rokhsar–Kivelson points and cantor deconfinement”, *Physical Review B*, vol. 69, p. 2244, 2004.

- [88] S. Sachdev, “Spin-Peierls ground states of the quantum dimer model: A finite-size study”, *Physical Review B*, vol. 40, no. 7, p. 5204, 1989.
- [89] P. W. Leung, K. C. Chiu, and K. J. Runge, “Columnar dimer and plaquette resonating-valence-bond orders in the quantum dimer model”, *Physical Review B*, vol. 54, no. 18, p. 12938, 1996.
- [90] F. Carollo, J. P. Garrahan, I. Lesanovsky, and C. Pérez-Espigares, “Generic mixed columnar-plaquette phases in Rokhsar–Kivelson models”, *Physical Review Letters*, vol. 100, p. 037201, 2008.
- [91] S. R. Finch. *Mathematical Constants*. Cambridge University Press, 2003.
- [92] Z. Lan and S. Powell, “Eigenstate thermalization hypothesis in quantum dimer models”, *Physical Review B*, vol. 96, no. 11, p. 115140, 2017.
- [93] F. Alet, Y. Ikhlef, J. L. Jacobsen, G. Misguich, and V. Pasquier, “Classical dimers with aligning interactions on the square lattice”, *Physical Review E*, vol. 74, no. 4, p. 041124, 2006.
- [94] T. Senthil, L. Balents, S. Sachdev, A. Vishwanath, and M. P. A. Fisher, “Quantum criticality beyond the Landau-Ginzburg-Wilson paradigm”, *Physical Review B*, vol. 70, p. 144407, 2004.
- [95] G. J. Sreejith and S. Powell, “Critical behavior in the cubic dimer model at nonzero monomer density”, *Physical Review B*, vol. 89, p. 014404, 2014.
- [96] C. Appert-Rolland, B. Derrida, V. Lecomte, and F. van Wijland, “Universal cumulants of the current in diffusive systems on a ring”, *Physical Review E*, vol. 78, p. 021122, 2008.
- [97] V. Lecomte, J. P. Garrahan, and F. van Wijland, “Inactive dynamical phase of a symmetric exclusion process on a ring”, *Journal of Physics A*, vol. 118, no. 3, p. 030601, 2012.
- [98] D. Karevski and G. M. Schutz, “Conformal invariance in driven diffusive systems at high currents”, *Physical Review Letters*, vol. 89, p. 014404, 2017.

- [99] R. L. Jack, I. R. Thompson, and P. Sollich, “Hyperuniformity and phase separation in biased ensembles of trajectories for diffusive systems”, *Physical Review Letters*, vol. 114, p. 060601, 2015.
- [100] F. Carollo, J. P. Garrahan, I. Lesanovsky, and C. Pérez-Espigares, “Fluctuating hydrodynamics, current fluctuations and hyperuniformity in boundary-driven open quantum chains”, *Physical Review E*, vol. 96, p. 052118, 2017.
- [101] U. Ray, G. K. Chan, and D. T. Limmer, “Exact fluctuations of nonequilibrium steady states from approximate auxiliary dynamics”, *Physical Review Letters*, vol. 120, no. 21, p. 210602, 2018.
- [102] C. Giardinà, J. Kurchan, and L. Peliti, “Direct evaluation of large-deviation functions”, *Physical Review Letters*, vol. 96, no. 12, p. 120603, 2006.
- [103] J. Tailleur and V. Lecomte, “Simulation of large deviation functions using population dynamics”, in *AIP Conference Proceedings*, vol. 1091, p. 212, AIP, 2009.
- [104] P. W. Kasteleyn, “The statistics of dimers on a lattice : I. the number of dimer arrangements on a quadratic lattice”, *Physica*, vol. 27, pp. 1209–1225, 1961.
- [105] T. R. Gingrich and P. L. Geissler, “Preserving correlations between trajectories for efficient path sampling”, *Journal of Chemical Physics*, vol. 142, p. 234104, 2015.
- [106] A. A. Budini, R. M. Turner, and J. P. Garrahan, “Fluctuating observation time ensembles in the thermodynamics of trajectories”, *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2014, no. 3, p. P03012, 2014.