
EFFICIENT PATH SAMPLING FOR
TRAJECTORY ENSEMBLES WITH
APPLICATIONS TO
NON-EQUILIBRIUM SYSTEMS

BY

ROBERT M TURNER, MSCI.

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

October 2016

ABSTRACT

This thesis utilises large deviation methods to study nonequilibrium phenomena in both quantum and classical systems. The dynamical analogues of the ensembles of statistical mechanics are used to explore dynamical phase spaces of systems, quantifying atypical fluctuations that can play a critical role in long term behaviour. A dynamical ensemble based on fixed numbers of dynamical events, allowing trajectory observation time to fluctuate, is introduced. This ensemble, denoted the x -ensemble, is found to be well suited to numerically simulate atypical fluctuations using transition path sampling (TPS). x -ensemble TPS schemes are analysed with reference to existing methods in both quantum and classical stochastic systems, and are found to offer more flexibility and efficiency in a variety of situations. The potential to develop this scheme into a self-optimizing algorithm is discussed with examples. The x -ensemble is then used in three non-equilibrium scenarios. Firstly in plaquette models of glass formers, in an effort to provide insight into the nature of the glass transition. It is shown that a two-dimensional triangular plaquette model (TPM) exhibits both a trajectory phase-transition between dynamical active and inactive phases, and when two replicas are coupled, a thermal phase transition between states of low and high overlap between the replicas. These two transitions are similar to those seen to occur in more realistic glass formers. When the TPM is generalised to a three-dimensional square pyramid plaquette model (SPyM) these dynamical and thermodynamic features of interest remain. It is argued that these models therefore provide an ideal test-bed for competing theories of the glass transition. Secondly the x -ensemble is used to define and analyse the dynamical analogue of the Jarzynski equality, allowing for the computation of dynamical free energy differences with, in principle, arbitrarily fast protocols linking two dynamical states. This relation is tested and found to hold in open quantum systems. Finally the partition sum zeros method of Lee and Yang is used to extract the location of dynamical phase transitions from the high-order, short-time cumulants of the x -ensemble. Results in both classical and open quantum systems are compared with previously studied dynamical ensembles, providing insight into the nature in which dynamical behaviours are encoded by these ensembles.

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 3 and 4:

1. Adrián A. Budini, **Robert M. Turner**, and Juan P. Garrahan, *Fluctuating observation time ensembles in the thermodynamics of trajectories*, J. Stat. Mech. P03012 (2014).

Chapter 5:

2. **Robert M. Turner**, Robert L. Jack and Juan P. Garrahan, *Overlap and activity glass transitions in plaquette spin models with hierarchical dynamics*, Phys. Rev. E. **92**, 022115 (2015).

Chapter 6:

3. **Robert M. Turner**, Thomas Speck and Juan P. Garrahan, *Meta-work and the analogous Jarzynski relation in ensembles of dynamical trajectories*, J. Stat. Mech. P09017 (2014).

Chapter 7:

4. **Robert M. Turner**, James M. Hickey, Christian Flindt, and Juan P. Garrahan, *Dynamic Lee-Yang zeros and phase transitions in Fluctuating time ensembles of trajectories*, To be published.

ACKNOWLEDGEMENTS

First and foremost I thank, and express my utmost respect for, my supervisor, Juan Garrahan for years of support and enjoyable discussion. I could not imagine a better fit for a supervisor, you went above and beyond all that I could have anticipated before starting. To Igor Lesanovsky, my secondary supervisor, I express my gratitude for providing a continuing source of knowledge and inspiration from my years as an undergraduate, and for first setting me on the path of non-equilibrium physics.

I also thank James Hickey for his tutelage, friendship and the example he set for me to follow, as well as my friends in the office, Ben, Tom, and Claire to name a few, for keeping work lively and entertaining. I further thank the post-docs in the group, Michael, Sam, Emanuele, Matteo and Robin for their insights and many interesting discussions.

To Jenna, I can not even begin to express my appreciation for the things you do every single day. You are the most kind, caring person I've ever encountered. To my brothers, who I have looked up to always, and finally to my closest friends, who have stuck with me through the turbulence, I am ever grateful.

CONTENTS

LIST OF PUBLICATIONS	1
ACKNOWLEDGEMENTS	2
CONTENTS	3
LIST OF FIGURES	7
1 INTRODUCTION	18
2 BACKGROUND THEORY: THERMODYNAMICS, LARGE DEVIATIONS AND STOCHASTIC PROCESSES	23
2.1 Large deviation theory	24
2.1.1 Connection to thermodynamics	27
2.2 Classical stochastic processes	32
2.3 Open quantum systems	34
2.4 Thermodynamics of trajectories - s -ensemble	35
2.5 Systems and relations of interest	38
2.5.1 Models of glass formers	38
2.5.2 Fluctuation theorems and the Jarzynski equality	42
2.5.3 Lee-Yang partition function zeros method	44
3 FLUCTUATING OBSERVATION TIME ENSEMBLES IN THE THERMODY- NAMICS OF TRAJECTORIES: THEORY	46
3.1 Fluctuating observation time: the x -ensemble	47

3.2	Ensemble correspondence	50
3.2.1	Example: classical two-level system	52
3.3	Generalisation of the x -ensemble to multiple observables	54
3.3.1	Example: classical three-level system	57
3.4	x -ensemble in open quantum systems	58
3.4.1	Example: quantum two-level system	59
3.5	Outlook	61
4	FLUCTUATING OBSERVATION TIME ENSEMBLES: TRANSITION PATH SAM- PLING APPLICATIONS	62
4.1	Transition path sampling for trajectory ensembles	63
4.1.1	Classical continuous time Monte Carlo	64
4.1.2	Quantum jump Monte Carlo	65
4.1.3	TPS and Metropolis-Hastings in trajectory space	66
4.1.4	Comments on s -ensemble efficiency	67
4.2	TPS with fluctuating observation time	70
4.3	Applications	71
4.3.1	$T \neq 0$ quantum two-level system	72
4.3.2	Micromaser	73
4.3.3	East facilitated spin model of glasses	76
4.4	Efficiency of the x -ensemble scheme	79
4.5	Beyond the standard TPS algorithm: a self-tuning approach	82
4.6	Outlook	89
5	DYNAMICAL AND STATIC TRANSITIONS IN PLAQUETTE MODELS OF GLASSES	90
5.1	Plaquette models, and coupled replicas	92
5.1.1	Models	92
5.1.2	Coupled replicas	94

5.1.3	Dualities and phase transitions	95
5.1.4	Other consequences of dualities and symmetries	96
5.2	Numerical results for the TPM in a field	102
5.3	Active-inactive dynamical transitions in the TPM	107
5.4	Overlap and activity transitions in a three-dimensional plaquette model	110
5.4.1	Model	110
5.4.2	Phase transition in (annealed) coupled replicas	114
5.4.3	Evidence for a dynamical (space-time) phase transition . . .	115
5.5	Discussion	117
5.5.1	Connection of phase transitions to long-lived metastable states	117
5.5.2	Connection between multiple coupled replicas and biased- activity ensembles	121
5.5.3	Outlook	123
6	META-WORK AND AN ANALOGOUS JARZYNSKI RELATION	125
6.1	Jarzynski relation in trajectory space	125
6.1.1	Meta-dynamics: Dynamics in the space of trajectories	126
6.1.2	Meta-work and the Jarzynski relation	126
6.1.3	Computing the meta-free energy $g(x)$	128
6.2	Application to open quantum systems	129
6.2.1	Two-Level System	130
6.2.2	Micromaser	132
6.2.3	Driving across a first-order phase transition	134
6.3	Outlook	137
7	DYNAMICAL LEE-YANG ZEROS IN THE x-ENSEMBLE	139
7.1	Dynamical Lee-Yang zeros of the x -ensemble	141
7.2	East Model	143

7.3	Dissipative Three-level system	146
7.4	Outlook	152
8	CONCLUSIONS AND FUTURE WORK	154
A	MAPPING COUPLED TPMS TO SINGLE SYSTEM IN FIELD	158
	BIBLIOGRAPHY	160

LIST OF FIGURES

2.1	(i), (ii), (v) Taken from Ref. [7]. (i) Super-Arrhenius relaxation of ortho-terphenyl (OTP) showing good agreement with both VFT (blue) and parabolic (red) laws. Also plotted is a mode-coupling theory (green) prediction, showing it break down at low temperature. (ii) Dynamical heterogenities in an equilibrium trajectory of two-dimensional mixture. Overlap of particles with their original position after one tenth of the systems relaxation time are plotted with high overlap (no movement) in blue and low overlap in red. (iii) Schematic illustration of temperature dependent entropy of a glass former. Extrapolating the slope of the liquid phase intersects the entropy of the crystal at a non-zero temperature. Theories must either account for the inability to crystallise, or uncover a phase transition to an ideal glass state that occurs before T_K . (iv) Taken from Ref. [111] Relaxation of a large number of different materials, all exhibiting super-Arrhenius relaxation. (v) The difference in specific heat of glass and liquid phases, C_p of OTP showing anomalous differences when the sample is cooled (black) and heated (green) through the same temperature range. This illustrates that the properties of the material are unusually dependent on its history.	39
-----	--	----

- 3.1 (a) The s -ensemble is the set of all trajectories that are possible with the dynamics (2.27)-(2.28), of fixed total time τ , and where the probability of each trajectory is weighed by the number of configuration changes. It is the ensemble defined by fixed (τ, s) . The ensemble of trajectories that corresponds to the actual dynamics (2.27)-(2.28) is given by $(\tau, 0)$. We sketch two trajectories, the squares indicate the times where trajectories begin and end, and the tick where jumps between configurations take place. (b) The x -ensemble is the set of all trajectories that are possible with the dynamics (2.27)-(2.28), of fixed number of configuration changes K , and where the probability of each trajectory is weighed by the trajectory length. It is the ensemble defined by fixed (x, K) 48
- 3.2 (a) Classical two-level system. (b) Classical three-level system. (c) Quantum $T = 0$ two-level system: the full line indicates a coherent transition of frequency Ω and the wavy line a dissipative quantum jump of rate γ associated to emission into the bath. (d) Quantum $T \neq 0$ two-level system: same as before, but now absorption from the bath leads to a second kind of quantum jump with rate λ 52
- 4.1 Illustration of typical s -ensemble TPS protocols with a generic evolution of “coordinates” of a system as a function of time. (a) Shooting moves: a point in time is selected (blue line) and based on a coin flip, either the past or future is discarded (pink) and replaced by a new segment (green). (b) Shifting moves: before generating a new segment the kept portion of a trajectory is first shifted forwards or backwards (backwards shift not shown). It should be noted that both protocols require the ability to produce time reversed dynamics. 68

- 4.2 $T \neq 0$ quantum two-level system. (a) x -ensemble: $\langle \tau \rangle / \langle K_1 \rangle$ as a function of x , along the curve $x^*(s_1)$ in the x -ensemble; symbols are from TPS simulations and the curve from exact diagonalisation of \mathcal{T}_{x,s_1} . (b) s -ensemble: the symbols are the s -ensemble expectation values as a function of s , as obtained from the x -ensemble TPS simulation, and the curve those from the exact diagonalisation of \mathcal{W}_{s_1} . The parameters here are $\gamma = 6\Omega$ and $\lambda = 2\Omega$. (c) and (d) Efficiency comparison of (c) x -ensemble TPS to (d) s -ensemble forwards/backwards shooting for parameters $\gamma = 4\Omega$ and $\lambda = 0$ (zero-temperature). The simplicity of the system represents the most efficient possible scenario for both approaches. Real time taken to converge to within 2% of the analytical rates showing the scaling with trajectory length. The typical trajectory lengths plotted are 400 times longer in (c) than (d). (Inset) As before but now the scaling is with x/s . The range of x and s used are equivalent. . . . 74
- 4.3 Micromaser: average number of L_1 jumps, $\langle M_1 \rangle$, as a function of s_1 and α/π . (a) Results from x -ensemble TPS of trajectories of fluctuating observation time τ , transformed to the s -ensemble. (b) Exact numerical diagonalisation of \mathcal{W}_{s_1} . Both plots at the same resolution. (c) Results from s -ensemble TPS. This small segment of the phase diagram took approximately 100 times the computation time of the x -ensemble plot. 76

- 4.4 (a) The raw x -dependent mean jump waiting time per site τ/NK explicitly converted to: (b) The s -dependent mean activity per site, $K/N\tau$ and (c) The s -dependent mean excitation density per site, ρ_n/N . Both plots are at temperature $T = 0.5$, for chain lengths $N = 15, 30, 60$, and the simulations were made using x -ensemble TPS. (d/e/f) Same as before but now for $N = 60$ and for temperatures $T = 0.91, 0.75, 0.5$. The change in relationship between x and s with a change in both chain length and temperature is demonstrated in the conversion from (a) \rightarrow (b) and from (d) \rightarrow (e) respectively . . . 77
- 4.5 Efficiency comparison of the (a) x -ensemble TPS scheme, and (b) a forwards/backwards shooting s -ensemble TPS scheme. Both plots are for $T=0.5$, $N=60$ and show the mean real time taken for the mean excitation density per site, ρ_n/N , of an initially infinite-temperature system to fall below some critical value ρ_{crit} for TPS with $s = 0.2$. Due to overall difference in efficiency we used $\rho_{crit} = 0.02$ for the x -ensemble in (a) and $\rho_{crit} = 0.05$ for the s -ensemble in (b). The trajectory lengths are over similar ranges for trajectories in the inactive phase with $s > 0$ 79
- 4.6 Efficiency comparison of the “small r ” scheme (red circles) with the original scheme (blue triangles) in the quantum two-level system with parameters $\gamma = 4\Omega$ and $\lambda = 0$ ($T=0$), $K = 10000$. (a) Mean real time taken to converge to within 0.5% of the analytical rates across a range of x values with their corresponding acceptance rates (b). 82

- 4.7 Efficiency of the initial relaxation of an East model, mean real time taken (blue triangles) for the excitation density to fall below $\rho_{crit} = 0.012$ with $x = -0.04$ with $N = 30$, $T = 0.5$, and $K = 7500$. The values of dr are integer i.e. changes to whole sets of random numbers. The corresponding acceptance rates (red circles) are shown, and the relationship between acceptance rate and time taken is non-linear (as expected). 84
- 4.8 Implementation of a self-tuning x -ensemble TPS algorithm in a $T = 0$ quantum two-level system. The real time taken to converge within 0.5% of the analytical rates across a range of values for a range of target acceptance rates. Also included is the original method with $dr = 1$ (maroon asterisks). For large x all schemes converge towards the same time taken as they consistently hit the lower dr bound of 0.03. The plateaus, and subsequent improvements in time taken, for the high acceptance rate targets (40 – 60%, teal triangles, black triangles and green crosses) are caused by encountering this lower dr bound much sooner, and the inability to maintain unnecessarily high acceptance rates. 87
- 5.1 (a) Illustration of spins and plaquettes in the TPM. The spins σ_i are located on the vertices of the lattice. The plaquette variables τ_μ are located on the upward pointing triangles (shaded). Each plaquette is associated with three spins and the variable τ_μ is given by the product of these spins. (b),(c) Geometrical illustration of the duality relation (A.9) in the TPM. Panel (b) shows two TPM systems, a and b , with coupling as in (5.3). Panel (c) shows the location of the sites of the dual problem, again two coupled TPMs a^* and b^* . The plaquettes in the dual system bisect the coupling interactions in the direct system, and vice versa. 93

- 5.2 Illustration of the relation between critical behaviour of the Ising model and the TPM. (a) Ising phase diagram. the $h_I = 0$ axis is a symmetry line, there is a critical point indicated by a circle, with a first-order (phase coexistence) line for large J_I , indicated by a solid line. Selected lines of constant J_I are indicated by dotted lines. (b) The corresponding situation for the TPM in a field. On the solid/dashed line (5.7), the system has a discrete (Z_2) symmetry: a critical point and phase coexistence both occur on this line, as indicated. The dotted lines are obtained from (5.13) for three different values of h_0 , and correspond to the lines of constant J_I in panel (a). Near the critical point, they indicate the direction of the most relevant renormalisation group flow. 98

- 5.3 Simulations of the TPM in a field. (a) Distribution of the magnetisation at various values of J for state points on the self-dual line (5.8), at system size $L = 128$. The bimodal distribution $P(m)$ indicates a first-order transition, which disappears on reducing J . From (A.2), the same distributions would be obtained when considering the overlap between two coupled TPMs, at appropriate state points. (b) Representative configuration at phase coexistence ($\beta J = 2.9$ and $L = 128$) showing interfaces between regions of small and large magnetisation (corresponding to regions of small and large overlap in the two-replica problem). (c) At our estimated critical point, ($J_c = 2.634$, $h_c = 0.072$) and for various system sizes, we show distributions of the variable x that is obtained by rescaling the order parameter \mathcal{M} to zero mean and unit variance. The full line is the corresponding result for the $2d$ Ising model at criticality, indicating that the critical point of the TPM in a field (and therefore of the two coupled TPMs) is in the $2d$ Ising universality class. 99
- 5.4 Phase diagrams of coupled plaquette models, the two-dimensional TPM (left) and the three-dimensional SPyM (right). The full line corresponds to a line of first-order transitions between a thermodynamic phase of small overlap and one of large overlap between the replicas. This curve is on the self-dual line (5.8) (dashed line). The first-order transition line ends at a critical point that is in the $2d$ Ising universality class for the TPM and the $3d$ Ising universality class for the SPyM. 104

- 5.5 Average activity $k(s)$, and the associated susceptibility $\chi(s)$ in the TPM at $T = 0.5$. The main panels show data for system size $L = 8$. These results were obtained by the x -ensemble method (see text), using trajectories with fixed numbers of events K , as shown. On increasing the trajectory length, the crossover from active to inactive behaviour becomes increasingly sharp and the susceptibility peak increases. The behaviour for smaller systems ($L = 4$) is shown in the insets, with both quantities normalised by the system size. In the absence of a phase transition, one expects both $k(s)/L^d$ and $\chi(s)/L^d$ to be independent of L , so the sharper crossover at $L = 8$ is again consistent with an underlying phase transition. 109
- 5.6 (a) The SPyM consists of spins (grey circles) on the sites of a BCC lattice, which interact in quintuplets at the vertices of upward pointing square pyramids. One such pyramid is indicated; the central spin also participates in four other upward pointing pyramids whose apexes are the four spins on the upper face of the cube. 111
- 5.7 Relaxation of the energy of the SPyM at low temperature starting from a random configuration (the system size is $L = 16$). The curve shows the characteristic plateaus indicative of hierarchical relaxation, as in the East model and the TPM. Inset: average relaxation time as a function of inverse temperature, showing super-Arrhenius behaviour. 114

- 5.8 (a) Activity as a function of s in the SPyM for system size $L = 4$, various trajectory lengths, and $T/J = 0.65$. The inset shows the corresponding susceptibility. (b) and (c) s -ensemble phase diagrams for the TPM and SPyM. The full curves are an estimate of the transition point from the simulations. The dashed lines are extrapolations in the low temperature regime inaccessible to numerics. . . . 116
- 5.9 Three dimensional stack of coupled two-dimensional TPMs. . . . 121
- 6.1 2-level system. (a) Comparison of the meta-free energy $g(x)$ obtained numerically via the trajectory Jarzynski relation (symbols) to the exact analytical result (6.9) (solid line) for a range of x , in the zero temperature case, with $\gamma = 4\Omega$. (b) Same as in (a), but now for the finite temperature case, with $\gamma = 6\Omega$ and $\lambda = 2\Omega$. The statistical error is smaller than the symbol sizes. Insets to (b): Sampled histograms for the meta-work distribution $P_{\uparrow}(W)$ for the forward (red) and $P_{\downarrow}(-W)$ for the backward process (blue), at the two final values of x shown. 131
- 6.2 Micromaser. (a) Comparison of the meta-free energy $g(x)$ obtained numerically via the trajectory Jarzynski relation (symbols) to results obtained by direct diagonalisation of the master operator (solid line) for $\alpha = 1.2\pi$, where the system is initially equilibrated to $x = 2$. (b) Same as in (a), but now for $\alpha = 4\pi$. Different simulations, equilibrated to different initial values of x are denoted by different symbols. 133

- 6.3 Micromaser with cross-phase Jarzynski protocol. (a) Comparison of the numerical meta free energy, $g(x)$, obtained numerically via the Jarzynski relation (symbols) to results obtained by direct diagonalisation (solid line) in a micromaser with pump parameter $\alpha = 1.2\pi$, at a finite temperature ($\gamma/\kappa = 0.15$). The second largest eigenvalue (dashed line) is plotted to illustrate the meta free energy calculation being locked to the metastable branch after the transition. Inset to (a): the expected waiting time per event showing the differing dynamic properties of the two phases. (b) Same as in (a), but now at zero temperature ($\gamma/\kappa = 0$). Insets to (b): sampled meta-work distributions for the forward (red) and backward (blue) process for the three points shown. 136
- 7.1 East Model results for $N = 60$ spins. (a) The observation time activity as a function of x displays a discontinuity at $x = 0$. This marks a first order transition from an inactive glassy phase, $x < 0$, to an active liquid phase. (b) The leading Lee-Yang zero pair extracted from the high-order cumulants. These converge to points close to the origin with increasing K , this highlights one infer the presence of a trajectory transition point at the origin in the limit of $K \rightarrow \infty$ from these small K cumulants. (c) The approximated cumulants (dashed lines) are found to match their numerically extracted values (full lines) for small K , confirming the validity of the approximation. 144

- 7.2 We consider $\kappa = 4\Omega_1$ and fix $\Omega_1 = 1$ and $\Omega_2 = 0.15$. (a) Schematic of a dissipative 3-level system. The system is driven by two lasers of frequency Ω_1 and Ω_2 which drive transitions from $|0\rangle \rightarrow |1\rangle$ and $|0\rangle \rightarrow |2\rangle$ respectively. There is a single decay channel $|1\rangle \rightarrow |0\rangle$ with an associated decay rate κ . (b) The large deviation function $\theta(s)$ and associated dynamical activity $k(s)$ as a function of s . A rounded crossover is seen at $s \sim 0$ where the system changes from behaving as an active two-level system comprised of the states $|1\rangle$ and $|0\rangle$, i. e. an active phase, to a photon inactive phase. (c) The large deviation function associated with the fluctuating time trajectory ensemble along with the observation time is plotted as a function of x . Although the system is intermittent this does not manifest as a crossover as in the s -ensemble case. (d) Dynamical Lee-Yang zeros of the fluctuating time ensemble MGF extracted from cumulants of order $m = 6, 7, 8, 9$. These are found to oscillate with increasing K and do not converge to a single point. This is indicative of a lack of a trajectory transition point existing in the complex x plane. 147
- 7.3 We consider $\kappa = 4\Omega_1$ and fix $\Omega_1 = 1$ and $\Omega_2 = 0.15$. Plots of the Mandel Q Parameter for (a) the s -ensemble and (b) the x -ensemble. Despite the information being encoded in different ways, the same physical properties are realised by both ensembles 152

1. INTRODUCTION

The theory of equilibrium statistical mechanics provides powerful tools for describing and predicting the static properties of many-body systems [1, 2]. In lieu of a deterministic microscopic approach, seemingly complex problems are well described by a small number of aggregate statistical measures. Many-body interactions with degrees of freedom on the order of Avogadro's number ($\sim 10^{24}$) are reduced to a small number of quantities exchanged with an environment that is too large to be influenced by the system under study. Few theories have enjoyed such broad success, from providing the underpinnings of thermodynamics, ideal gases, heat engines, and phase transitions, to applications in quantum systems such as Bose-Einstein condensation and quantum computing. Despite this success, there are areas of scientific interest where the static picture alone proves insufficient. This is most notably an issue when systems are driven out of equilibrium, where the breaking of time-reversal symmetry necessitates the inclusion of dynamic processes. This has long been a problem in the study of biophysics, where nearly every system of interest is driven in some manner, and indeed, many of the tools to be discussed in this thesis were developed to study such systems [3].

There is a more subtle area where the static picture also encounters problems - in systems where ergodicity is broken, such as glass. The nature of glass formation is a highly controversial topic, with no single theory fully describing the observed

phenomena [4–8]. Glasses have no apparent structural order, characteristic of a liquid state, but long structural relaxation times make them solids. Unlike most materials the specific protocol used in the formation of glass, such as initial temperature and rate of cooling, plays a critical role in the properties of the resulting solid - and this provides the first clue that the dynamics might provide insight into the static structure of the material. Furthermore glasses exhibit *dynamical heterogeneities* - space-time “bubbles” within the bulk material with vastly different dynamical activity [9, 10], which any theory of glass must necessarily account for. The controversy, then, is centred on whether there is an equilibrium phase transition to the glass state [8, 11], or if it is the result of unusually correlated dynamical behaviour [12, 13].

The thermodynamic formalism developed by Ruelle [14] provided an opportunity to adapt the language of statistical mechanics to dynamical systems, resulting in a “thermodynamics of trajectories” [12, 13, 15–22]. Using large deviation theory [23–25], the concept of an ensemble was generalised to trajectories, and the statistics of dynamical processes probed in a manner analogous to the configurations of equilibrium statistical mechanics. Fields conjugate to dynamical variables of interest play the role of free energies and entropies, and are used to explore a *dynamical phase space*. Rare, but important dynamical behaviours can be uncovered by varying these fields, including the dynamical equivalent of a phase transition - a singular change in dynamical properties of a system at specific points in phase space. The physical dynamics of the system take place at zero field while atypical but potentially important dynamics occur with non-zero field, in contrast to the ensembles of equilibrium statistical mechanics, where fields are often physically tunable parameters that can be precisely controlled. A good example of the importance of these rare behaviours occurs in chemical reactions - the physical interaction that causes two species to bond might be incredibly uncommon, especially if it requires more than a two-body interaction, but such events play an

obviously crucial role in the long term behaviour of the system.

With this new approach to non-equilibrium dynamics there comes a need for tools to efficiently explore the dynamical phase spaces of systems. A novel approach that is well suited to sampling rare dynamical behaviours is transition path sampling (TPS) [3]. TPS is a collection of computational algorithms originally developed to study so called *reaction trajectories* to gain insight into the physical mechanisms involved in a chemical reaction. It amounts to a biased random walk, coupled with efficient methods of proposing steps, through the space of all trajectories - in a similar vein to how techniques such as simulated annealing operate on configurations [26]. This approach revealed the existence of a zero field first order dynamical phase transition between a high activity liquid phase, and an inactive glass phase, in a simple glass forming spin model, the East Model [13, 27, 28]. Crucially this model has trivial thermodynamic properties, with no static phase transitions, adding weight to the dynamic approach to the glass transition. Glassy behaviour arises because of kinetically constrained dynamics - neighbouring spins compete dynamically and block each others motion. Despite the success of the approach, capturing all the properties of a real glass former in a three-dimensional setting remains difficult.

The purpose of this thesis is to adapt and develop efficient path sampling schemes into the thermodynamics of trajectories, with the ultimate aim of application to three-dimensional glass formers. More generally it aims to improve the flexibility of the tools used to study dynamical systems, by exploring alternative dynamical ensembles and adapting important relations in conventional thermodynamics to the dynamical case. The work is laid out as follows: Chapter 2 provides the background theory on large deviations and their connection to Ruelle's formalism for thermodynamics that underpins the thermodynamics of trajectories. The conventional presentation of the thermodynamics of trajectories, ensembles of fixed

trajectory observation time known as the s -ensemble, is also developed here, with brief introductions to the problems tackled later in this thesis. Chapter 3 introduces a new dynamical ensemble where observation time fluctuates and the number of dynamical events (e.g. photon emissions, spin flips) is instead kept constant - which we denote the x -ensemble. The relationship between s - and x -ensemble is explored with simple examples of both classical [29] and quantum [30,31] stochastic systems.

A path sampling scheme implemented by Crooks and Chandler [32] is adapted to the x -ensemble in Chapter 4. Efficiency is a key concern here, and the x -ensemble scheme is examined in classical, glassy, and quantum stochastic systems, with comparisons drawn to the relative efficiency of s -ensemble schemes. One of the main problems encountered in TPS is tuning the algorithms to the system under study, as efficiency, including algorithmic scaling, can vary substantially. As such, a future outlook on a self-tuning x -ensemble algorithm is also discussed here.

Chapter 5 focuses on kinetically constrained models of glasses. We examine an old model glass former with (effectively) kinetically constrained dynamics - the two-dimensional triangular plaquette model (TPM) [33–35]. This model has recently become of interest to the discussion on glasses, as two coupled replicas exhibit a thermodynamic phase transition with coupling strength as the order parameter. However this transition is not present at zero coupling strength, except in the limit of zero temperature [36]. Furthermore the model also undergoes a dynamical phase transition between an equilibrium liquid state and a glass state at low temperature, making it a useful test bed for arguments about thermal and dynamical transitions in glasses. We generalise this model to a three-dimensional equivalent, which we denote the square pyramid model (SPyM) and study its dynamical phase space in an attempt to bring arguments of glass formers onto solid three-dimensional footing. Developing three-dimensional glass forming models is an important step

in the study of glassiness, not only because of the three-dimensional nature of physical glass, but also crucially due to inherent dimensionality arguments in the debate on glasses. Arguments proposing an equilibrium phase transition to glass state are compelling in a mean-field (i.e. infinite-dimensional) setting, but the transitions vanish in two dimensions. Conversely, early studies utilising the dynamical approach saw success in one- and two-dimensional models, but took longer to establish a footing in three dimensions. The efficiency provided by x -ensemble TPS allows the exploration of the SPyM's dynamical phase space, which we find to be strikingly similar to that of the two-dimensional TPM. The interesting thermodynamic properties of coupled replicas of these models, as well as the rich dynamical phase space allude to the possibility of a connection between the (apparently quite different) competing theories of glass formers.

In the final two chapters, 6 and 7, we expand upon the thermodynamics of trajectories, developing equivalents to useful methods and relations from conventional thermodynamics. First we examine the Jarzynski equality [37, 38] - a revolutionary relation that allows for the computation of equilibrium free energy differences, from non-equilibrium (i.e. fast) methods. It adapts quite naturally to dynamical ensembles and we explore its use in open quantum systems - including attempts to compute dynamical “free energies” across a trajectory phase transition point. Chapter 7 looks at the Lee-Yang zeros method [39, 40] in the context of the x -ensemble. This method has previously been adapted to the s -ensemble [41–43] and we note important consequences of the manner in which dynamical ensembles encode information about physical dynamics, based on the results of examinations of open quantum systems.

Finally Chapter 8 provides a general discussion of the work in this thesis, with comments on potential future developments.

2. BACKGROUND THEORY: THERMODYNAMICS, LARGE DEVIATIONS AND STOCHASTIC PROCESSES

Few theories have enjoyed the broad applicability of statistical mechanics, capturing the phenomenology of a wide array of systems, both quantum and classical. This is in a large part due to the flexibility allowed by statistical **ensembles** - a mathematical idealisation of a large (or infinite) number of replicas of a system. The ease with which expectation values can be calculated using the approach, coupled with the large numbers of microscopic degrees of freedom present in any macroscopic system ensuring minimal deviation from expected aggregate values, and crucially, the ability to easily tailor an ensemble to describe any physical system has resulted in decades of successful exploration of physical systems. A mathematical formalism developed by Ruelle [14], based on the theory of large deviations [25] has facilitated the extension of the tools of statistical mechanics into the realm of non-equilibrium physics, capturing the *dynamical* behaviours of systems. This extension is sometimes called the *thermodynamics of trajectories* [12, 13, 15–22].

This chapter introduces the background theory necessary for the work presented in the remainder of the thesis. Firstly large deviation theory will be introduced with connections made to thermal ensembles. This leads quite naturally into the thermodynamics of trajectories itself, which will be presented in their typical manner with an ensemble of fixed trajectory observation time - the s -ensemble. A brief primer on master equations, both quantum and classical, will be provided to facilitate the description of dynamical processes captured by the s -ensemble. Furthermore, brief backgrounds on systems and relations of interest studied in this thesis, including glasses and fluctuation theorems will be provided

2.1 Large deviation theory

Large deviation (LD) theory [16, 25] can naively be thought of as an extension of the law of large numbers. It is the statement that fluctuations in the mean value of some parameter, n , become exponentially suppressed in the large size limit, and hence the probability distribution of the observed mean value takes on the generic form $P(n) \approx e^{-nI(n)}$, with all other terms being subdominant in this limit.

As a simple illustrative example, consider the case of n independent Gaussian random variables, X_0, X_1, \dots, X_n with mean μ and variance σ^2 . The probability distribution for these variables is given by

$$P(X_i = x_i) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x_i - \mu)^2}{2\sigma^2}}. \quad (2.1)$$

If we now wish to calculate the probability distribution for the sample mean S of our n variables, defined by $S_n = \frac{1}{n} \sum_i X_i$, using the fact that the probability of drawing a sequence of Gaussian variables, represented by the vector $x = (x_1, x_2, \dots, x_n)$, is simply the product of their individual probabilities

$$P(x) = P(x_1, x_2, \dots, x_n) = P(x_1)P(x_2)\dots P(x_n) \quad (2.2)$$

the probability that the sample mean takes on a value s is

$$\begin{aligned}
 P(S_n = s) &= \int dx \delta(S_n(x) - s) P(x) \\
 &= \langle \delta(S_n - s) \rangle \\
 &= \sqrt{\frac{n}{2\pi\sigma^2}} e^{-\frac{n(s-\mu)^2}{2\sigma^2}}.
 \end{aligned} \tag{2.3}$$

As we take $n \rightarrow \infty$ the \sqrt{n} term is subdominant to the exponential term and we recover the large deviation form

$$P(S_n = s) \asymp e^{-nI(s)} \tag{2.4}$$

where the function $I(s) = (s - \mu)^2/2\sigma^2$ is known as the large deviation **rate function**, and the symbol \asymp indicates that in the limit of large n the behaviour of the probability distribution is dominated by the exponential term given.

While the applicability of LD theory is broad, the most useful result for the purposes of this thesis is known as the Gärtner-Ellis theorem, which connects the LD form of a probability distribution to its (scaled) cumulant generating function (CGF). The CGF is closely related to the moment generating function (MGF) of a distribution, which presents an alternative description to the probability distribution. For a random variable, A , which is parameterised by a real positive integer n (i.e. we are taking n samples), the MGF is defined by

$$M(t) = \lim_{n \rightarrow \infty} \frac{1}{n} \langle e^{tnA_n} \rangle \tag{2.5}$$

where

$$\langle e^{-ntA_n} \rangle = \int dA_n e^{-tnA_n} P(A_n) \tag{2.6}$$

In the case of a continuous probability density, the MGF amounts to the Laplace transform of $P(A_n)$ and t can be understood as the real parameter conjugate to A . For independent and identically distributed random variables, X_i , Cramr's theorem can be used to simplify the calculation of the MGF

$$M(t) = \lim_{n \rightarrow \infty} \frac{1}{n} \langle e^{t \sum_{i=1}^n X_i} \rangle = \lim_{n \rightarrow \infty} \frac{1}{n} \prod_{i=1}^n \langle e^{tX_i} \rangle = \langle e^{tX} \rangle. \tag{2.7}$$

Thus for our previous example, the MGF of a Gaussian sample mean can be calculated by inserting Eq. 2.1 into Eq. 2.6 using Eq. 2.7

$$\int dX e^{-tX} P(X) = \int dX e^{-tX} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(X-\mu)^2}{2\sigma^2}}. \quad (2.8)$$

This integral can be explicitly calculated as $\exp[t\mu + \frac{1}{2}t^2\sigma^2]$. The purpose of defining the MGF, if it exists, is that one can easily extract the k th moments of the distribution by differentiating the MGF w.r.t. t , k times, and then setting $t = 0$. The CGF is then simply defined as the natural logarithm of the MGF,

$$\theta(t) = \ln M_A(t) = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \langle e^{-ntA_n} \rangle \quad (2.9)$$

and is thus simply $\theta(t) = t\mu + \frac{1}{2}t^2\sigma^2$ for our Gaussian sample mean. As with the MGF, the k th cumulant of the probability distribution is obtained by differentiating θ k times w.r.t t and then setting $t = 0$. It is straightforward to see that the CGF of our Gaussian sample mean yields μ as the first cumulant and σ^2 as the second cumulant, with all other moments being zero, as expected.

The Gärtner-Ellis theorem then states that A satisfies a large deviation principle with

$$P(A_n) \asymp e^{-nI(A_n)} \quad (2.10)$$

provided that the CGF exists and is differentiable at all $t \in \mathbb{R}$. The LD rate function and the CGF are related by a *Legendre-Fenchel* transform (an extension to the Legendre transform, which can be seen as a saddle point approximation solution to Eq. 2.6 above) with

$$\theta(t) = -\min_{A_n} [I(A_n) + tA_n] \quad (2.11)$$

and its inversion

$$I(A_n) = -\min_t [\theta(t) + tA_n]. \quad (2.12)$$

Thus we can recover the large deviation rate function for our Gaussian sample mean given previously, from the CGF of its probability distribution.

2.1.1 Connection to thermodynamics

Before the development of statistical mechanics, thermodynamic calculations in systems of interest proceeded through the use of free energies Φ [1,2]. Given a set of **extensive variables**, (X_0, X_1, \dots, X_n) - variables that are additive when two systems are combined, number of particles, N , or internal energy, E^1 - the free energies relevant to a problem of interest are defined by Legendre transforming the system's energy by

$$\Phi = E - \frac{\partial E}{\partial X_i} X_i = E - f_i X_i. \quad (2.13)$$

Here f_i is the **intensive** field conjugate to X_i , defined by $f_i = \partial E / \partial X_i$. These fields act as constraints on thermal systems by controlling fluctuations in the relevant extensive variables, for example the chemical potential μ , which controls fluctuations in N . The free energies then capture the amount of energy available in a system, subject to all its constraints, to perform work. Although it is not immediately obvious, that they are obtained by a Legendre transform of the systems energy is a subtle manifestation of the Gärtner-Ellis theorem.

2.1.1.1 Ensembles of statistical mechanics

With the development of statistical mechanics, the rigorous microscopic theory underpinning thermodynamics, Ruelle was able to cast thermodynamics in a formalism that makes this connection to large deviation theory explicit, and one can

¹Strictly speaking the internal energy is not extensive as the total energy of two combined systems contains a contribution from interactions: $E_{tot} = E_1 + E_2 + E_{int}$. However for short range interactions E_{int} grows with the size of the interface between the two systems and is therefore subextensive compared to E_{tot} , which grows with the combined volume. Long range interactions such as magnetic or electric fields need to be treated more carefully, but can often be considered as a separate mechanical (not internal) energy.

identify the partition function of statistical ensembles as the MGF of the systems energy. We illustrate this connection here, with a heuristic derivation of statistical ensembles [1, 14].

Consider a system comprised of N interacting particles. At any moment in time it is possible to fully describe the state of the system with the positions \mathbf{r} and momenta \mathbf{p} of the particles. In a system with D dimensions, this corresponds to $2DN$ scalar values, which uniquely define a **representative point**, x , of a system. This space of $2DN$ points in which the representative point lives is known as **phase space**. In principle it is possible to compute the full behaviour of the system from Hamiltonian dynamics, but the scale of the computation makes it impossible in practical situations. Instead we assume the **thermodynamic limit** of the system, that is that the system is sufficiently large so that the law of large numbers ensures fluctuations in the extensive macroscopic quantities are suppressed, (X_0, X_1, \dots, X_n) of the system. It is then clear that the systems representative point exists in some accessible region of phase space, denoted Γ , where each of these variables takes on its thermodynamic value, which we denote as X_i^* . The volume of this accessible region of phase space $|\Gamma|$ then corresponds to the number of microstates accessible to the system, allowing us to write Boltzmann's fundamental entropy postulate as

$$S = k_B \ln |\Gamma| \tag{2.14}$$

where k_B is the Boltzmann constant.

Using the above concept we can define a probability distribution for the representative point $P(x)$ within phase space, that is to say the probability the representative point takes on value in the infinitesimal range $x + dx$. Specific probability distributions define an ensemble of all the possible realisations of a system's state, with any constraints on the system encoded by $P(x)$. The power of this statistical approach is the ability to predict the expected value of the observables, which aren't

necessarily variables considered in thermodynamics, of a system. For a generic observable $A(x)$, its equilibrium value, which we denote a^* , can be calculated by averaging the value of $A(x)$ over all of the accessible region of phase space:

$$a^* = \langle A(x) \rangle = \int_{\Gamma} dx A(x) P(x). \quad (2.15)$$

The simplest choice for $P(x)$ - a flat probability distribution with $P(x) = 1/\langle \Gamma \rangle$ - defines the **microcanonical** ensemble. This is the maximally constrained ensemble, as all extensive variables, including internal energy are fixed. The broad success of this approach, however, lies with the flexibility offered by defining alternative ensembles where we allow certain extensive variables to fluctuate. Generically we can imagine connecting our system of interest to a much larger reservoir, R , and allowing for the exchange of the quantity of interest, X_i . We imagine the reservoir is sufficiently large so that its properties are unaffected by any fluctuations in the system. In essence this means that the exchange of X_i between the system and reservoir is governed by the intensive variable conjugate to X_i , which for convenience we redefine here

$$-\frac{f_i}{T} = \left. \frac{\partial S^R}{\partial X_i^R} \right)_E. \quad (2.16)$$

The superscript R denotes variables associated with the reservoir. Treating the representative points of the system and reservoir separately, observables of the system can then be calculated as follows. From Eq. 2.15 we have

$$\langle A \rangle = \frac{1}{|\Gamma|} \int_{\Gamma} dx_R dx_{sys} A(x_{sys})$$

which can be rewritten as

$$\begin{aligned} \langle A \rangle &= \frac{1}{|\Gamma|} \int_{\Gamma} dx_R dx_{sys} A(x_{sys}) \delta(X_i^R(x_R) + X_i^{sys}(x_{sys}) - X_i) \\ &= \frac{1}{|\Gamma|} \int dx_{sys} A(x_{sys}) \int dx_R \delta(X_i^R(x_R) - (X_i - X_i^{sys}(x_{sys}))). \end{aligned} \quad (2.17)$$

The integral over the reservoir dx_R is simply the volume of accessible phase space when the reservoirs extensive variable X_i^R takes on a value $X_i - X_i^{sys}(x_{sys})$ which

can be related to the entropy using the fundamental postulate (Eq. 2.14)

$$\int dx_R \delta (X_i^R(x_R) - (X_i - X_i^{sys}(x_{sys}))) = \exp \left[-\frac{1}{k_B} S^R (X_i - X_i^{sys}(x_{sys})) \right]. \quad (2.18)$$

Because the system is much larger than the reservoir, we can Taylor expand this exponential

$$\exp \left[-\frac{S^R}{k_B} (X_i - X_i^{sys}(x_{sys})) \right] \simeq \exp \left[\frac{S^R}{k_B} (X_i) \right] \exp \left[-\frac{1}{k_B} \frac{\partial S^R}{\partial X_i} X_i^{sys}(x_{sys}) \right]. \quad (2.19)$$

Remembering Eq. 2.16 we finally arrive at

$$\langle A \rangle \simeq \frac{1}{Z} \int dx_{sys} A(x_{sys}) \exp \left[\frac{f_i X_i^{sys}(x_{sys})}{k_B T} \right] \quad (2.20)$$

where Z is the normalising quantity known as the partition function

$$Z = \int dx_{sys} \exp \left[\frac{f_i X_i^{sys}(x_{sys})}{k_B T} \right]. \quad (2.21)$$

In essence we have defined a new probability distribution where we have weighted different points in phase space by the factor $\exp(f_i X_i/k_B T)$. Furthermore the Boltzmann constant defines an inverse temperature scale for thermodynamic problems, $\beta = 1/k_B T$.

We have now defined a flexible mathematical language where an ensemble can be chosen to treat any system of interest depending on which quantities we allow to fluctuate within the system, and which we keep fixed. Either an extensive quantity is fixed in a system, or it can fluctuate and the fluctuations are controlled by fixing the conjugate intensive variable. For example, if we imagine a system that is able to exchange energy with the reservoir we define the **canonical** ensemble, with $P(x) = e^{-\beta E}/Z$, while allowing the exchange of particles as well defines the **grand canonical** ensemble, with $P(x) = e^{-\beta(E-\mu N)}/Z$ (note we have omitted the *sys* index since we are always working with system quantities now). This flexibility allows us to easily define an ensemble appropriate for treating any physical system we are interested in studying.

The important connection to large deviation theory arises from studying the partition function whose large deviation form should be clear from Equation 2.21. It is possible to show that the mean values of observables in the thermodynamic limit are dominated by the region of phase space where they are close to their equilibrium value X_i^* . Away from this value they fall off on the order of $N^{1/2}$. Thus we have

$$\begin{aligned} Z &= \int dx e^{\beta f_i X_i(x)} = \int dX' \int dx \delta(X_i(x) - X') e^{\beta f_i X'} \\ &= \int dX' e^{\beta(TS(X') + f_i X')} \end{aligned} \quad (2.22)$$

where we are implicitly summing over all of the relevant observables X_i . The similarity between this form and Eq. 2.6 is no coincidence, nor is the existence of a saddle point method solution in the thermodynamic limit, yielding

$$Z \simeq \exp[\beta(TS(X_i^*) + f_i X_i^*)]. \quad (2.23)$$

Explicitly then, the partition function plays the role of the (unscaled) MGF of the systems energy. The quantity in this exponent, and therefore the (unscaled) CGF, is the free energy for whichever ensemble we are using, for example in the canonical ensemble we have $\exp[-\beta(E^* - TS(E^*))] = \exp[-\beta F]$ and we have recovered the Helmholtz free energy. This precisely explains why free energies are obtained Legendre transformations of the systems energy.

The important result of this section comes from the generic form of Equation 2.20. While thermodynamics is concerned with the thermal properties of systems in space, it is equally possible to define a probability distribution for the observation of some dynamical event *in time*. This allows the power of the tools of statistical mechanics to be brought into the dynamical realm, defining the thermodynamics of trajectories, and allowing the study of non-equilibrium behaviours. Before presenting the thermodynamics of trajectories however, it is necessary to introduce

some of the formalism used to describe the dynamics of the systems to be studied in this thesis.

2.2 Classical stochastic processes

The work presented in this thesis will focus on **stochastic** systems [1, 29]. A stochastic process can be seen as a sequence of configurations defined at discrete moments in time, $\{\mathcal{C}_1(t_1), \mathcal{C}_2(t_2), \dots, \mathcal{C}_n(t_n)\}$ where the transitions between configurations contain an element of randomness. Such processes are naturally a part of a statistical mechanical description of nature - while the motion of a particle is inherently deterministic in nature the complexity of a large system of particles makes a deterministic solution prohibitive, instead the aggregate effect of an environment is considered in a probabilistic fashion. The problem can thus be solved by assigning a probability to possible sequences of configurations, $P(\mathcal{C}_n, t_n; \mathcal{C}_{n-1}, t_{n-1}; \dots; \mathcal{C}_i, t_i)$, where \mathcal{C}_i, t_i are the initial configuration and time, and instead studying the time evolution of this probability. For simplicity we focus on a subcategory of stochastic processes known as **Markovian** processes, in continuous time. These are processes where the transition probability to a new configuration is dependent only on the current configuration, with no memory of previous configurations and thus

$$P(\mathcal{C}_n, t_n | \mathcal{C}_{n-1}, t_{n-1}; \dots; \mathcal{C}_i, t_i) = P(\mathcal{C}_n, t_n | \mathcal{C}_{n-1}, t_{n-1}) \quad (2.24)$$

where $P(\mathcal{C}_n, t_n | \mathcal{C}_{n-1}, t_{n-2})$ denotes the conditional probability of being in configuration \mathcal{C}_n at time t_n given that the configuration \mathcal{C}_{n-1} was realised at time t_{n-1} . The time evolution of these probabilities is governed by the *Chapman-Kolmogorov* equation, where the probability of being in configuration \mathcal{C} at a time t is calculated by integrating over all intermediate configurations \mathcal{C}' at times t' ,

$$P(\mathcal{C}, t | \mathcal{C}_i, t_i) = \int d\mathcal{C}' P(\mathcal{C}, t | \mathcal{C}', t') P(\mathcal{C}', t' | \mathcal{C}_i, t_i). \quad (2.25)$$

While the Chapman-Kolmogorov equation is useful, it is much more convenient to collect the probabilities of being in all possible states into a vector,

$$|P(t)\rangle \equiv \sum_{\mathcal{C}} P(\mathcal{C}, t) |\mathcal{C}\rangle \quad (2.26)$$

where $\{|\mathcal{C}\rangle\}$ is an orthonormal configuration basis, $\langle \mathcal{C} | \mathcal{C}' \rangle = \delta_{\mathcal{C}, \mathcal{C}'}$. The time evolution of this vector can then be captured by a matrix equation, called the **master equation** [1, 31]

$$\partial_t |P(t)\rangle = \mathbb{W} |P(t)\rangle. \quad (2.27)$$

While we focus on continuous time Markov chains, generalisations of what we describe below are straightforward. The master operator \mathbb{W} is the matrix

$$\mathbb{W} \equiv \sum_{\mathcal{C}' \neq \mathcal{C}} W(\mathcal{C} \rightarrow \mathcal{C}') |\mathcal{C}'\rangle \langle \mathcal{C}| - \sum_{\mathcal{C}} R(\mathcal{C}) |\mathcal{C}\rangle \langle \mathcal{C}|, \quad (2.28)$$

where $W(\mathcal{C} \rightarrow \mathcal{C}')$ is the transition rate from \mathcal{C} to \mathcal{C}' , and $R(\mathcal{C})$ the escape rate from \mathcal{C} (these transition rates can be explicitly calculated by Taylor expanding the Chapman-Kolmogorov equation about infinitesimal t'). In this description, the expectation value of an operator A is given by $\langle A(t) \rangle = \langle - | A | P(t) \rangle$, where $\langle - | \equiv \sum_{\mathcal{C}} \langle \mathcal{C} |$ (such that $\langle - | P(t) \rangle = 1$ due to probability conservation).

The dynamics described by Eqs. (2.27-2.28) is realised by stochastic *trajectories*. A trajectory of total time τ is a time record of configurations, and of waiting times for jumps between them, observed up to a time τ . That is, if we denote by \mathbf{X}_τ such a trajectory, then $\mathbf{X}_\tau = (\mathcal{C}_0 \rightarrow \mathcal{C}_{t_1} \rightarrow \dots \rightarrow \mathcal{C}_{t_n})$, where \mathcal{C}_0 is the initial configuration and t_i the time when the transition from $\mathcal{C}_{t_{i-1}}$ to \mathcal{C}_{t_i} occurs (so that the waiting time for the i -th jump is $t_i - t_{i-1}$). The trajectory \mathbf{X}_τ has a total of n configuration changes (and $t_n \leq \tau$, i.e., between t_n and τ no jump occurred). Eqs. (2.27-2.28) imply that the probability $P(\mathbf{X}_\tau)$ to observe this trajectory out of all the possible ones of total time τ is given by

$$P(\mathbf{X}_\tau) = p_0(\mathcal{C}_0) \prod_{i=1}^n e^{-(t_i - t_{i-1}) R(\mathcal{C}_{t_{i-1}})} W(\mathcal{C}_{t_{i-1}} \rightarrow \mathcal{C}_{t_i}) \times e^{-(\tau - t_n) R(\mathcal{C}_{t_n})} \quad (2.29)$$

with $t_0 = 0$. The last factor is the survival probability of the configuration \mathcal{C}_{t_n} between t_n and τ , and we have also included the probability p_0 of the initial configuration.

2.3 Open quantum systems

In the previous section we focused for simplicity on stochastic Markovian classical systems. The extension to Markovian open quantum systems [30, 44, 45] is straightforward. In this case, instead of a master equation for the probability distribution we have a master equation for the density matrix ρ

$$\partial_t \rho(t) = \mathcal{W}[\rho(t)], \quad (2.30)$$

where the *quantum master operator* is the super-operator [45],

$$\mathcal{W}(\cdot) \equiv -i[H, \cdot] + \sum_{i=1}^{N_L} L_i(\cdot)L_i^\dagger - \frac{1}{2}\{L_i^\dagger L_i, \cdot\}. \quad (2.31)$$

Here H is the Hamiltonian, which generates the coherent part of the evolution, and L_i ($i = 1, \dots, N_L$) are (bounded) quantum jump operators corresponding to the incoherent effect of the interaction with the environment [30, 44, 45]. The evolution described by (2.30)-(2.31) can be realised by an “unravelling” in terms of stochastic wave-functions [30, 45, 46]. This stochastic evolution is given by propagation of the wave-function under the action of the non-Hermitian operator $H_{\text{eff}} \equiv H - \frac{1}{2} \sum_i L_i^\dagger L_i$, punctuated at random times by “quantum jumps” due to the action of the jump operators L_i . That is, a quantum version of the continuous time Markov chains discussed above.

2.4 Thermodynamics of trajectories - s -ensemble

In this section we introduce the thermodynamics of trajectories in its usual form - the so-called s -ensemble (also known as the canonical dynamical or canonical path ensemble) [12,13,16]. The properties of the dynamics can be studied by considering the statistics of *time-extensive* observables [14–16]. This allows the formalism of statistical ensembles to be brought into the realm of dynamics. One such trajectory observable is the “dynamical activity”, defined as the total number of configuration changes in a trajectory [12,16,47]. Its distribution over all trajectories \mathbf{X}_τ of total time τ is

$$P_\tau(K) = \sum_{\mathbf{X}_\tau} \delta\left(K - \hat{K}[\mathbf{X}_\tau]\right) P(\mathbf{X}_\tau) \quad (2.32)$$

where the operator \hat{K} counts the number of jumps in a trajectory. Equivalent information is contained in the MGF,

$$Z_\tau(s) \equiv \sum_K e^{-sK} P_\tau(K) = \sum_{\mathbf{X}_\tau} e^{-s\hat{K}[\mathbf{X}_\tau]} P(\mathbf{X}_\tau), \quad (2.33)$$

whose derivatives give the moments of the activity, $\langle K^n \rangle = (-)^n \partial_s^n Z_\tau(s)|_{s=0}$.

For large τ , the probability of observing an *intensive* number of jumps K/τ acquires a LD form [16,25],

$$P(K/\tau) \sim e^{-\tau\varphi(K/\tau)}. \quad (2.34)$$

This in turn implies that for large τ the generating function also acquires a LD form [16,25],

$$Z_\tau(s) \sim e^{\tau\theta(s)}. \quad (2.35)$$

The analogy with equilibrium statistical mechanics is now evident from Eqs. (2.34,2.35). For the dynamics, the equivalent objects to configurations, or microstates, are trajectories. Order parameters are time-extensive observables, in

this case the activity K . The large-size limit becomes that of large observation time τ (more specifically the limit of large space-time volume for a many-body system), and in this limit the order parameter distribution $P_\tau(K)$ is described by the function $\varphi(k)$, which plays the role of, say, a Helmholtz free energy, which for constant “volume” τ is only a function of the intensive “density” (of the number of transitions) $k = K/\tau$. Similarly, $Z_\tau(s)$ is like a partition sum with an associated free energy $\theta(s)$ (which in this analogy would be like a grand-potential) dependent on the counting field s (akin to a chemical potential for the activity). Just like thermodynamic potentials, the LD functions $\varphi(k)$ and $\theta(s)$ are related by a Legendre-Fenchel transformation [16, 25]

$$\varphi(k) = - \min_s [\theta(s) + ks], \quad (2.36)$$

together with the inversion formula

$$\theta(s) = - \min_k [\varphi(k) + ks]. \quad (2.37)$$

[In Eq. (2.36) above, the function θ is convex and (piecewise) differentiable, thus giving rise to a convex φ . In Eq. (2.37), if the function φ is non-convex the resulting θ is convex but singular, and the inverse transform, via (2.36), returns the convex envelope of φ .]

The LD function $\theta(s)$ is the quantity of interest. It is the scaled cumulant generating function for the activity, i.e., the n -th cumulant of the activity (per unit time) is given by

$$\frac{\langle\langle K^n \rangle\rangle}{\tau} = (-)^n \left. \frac{\partial^n \theta(s)}{\partial s^n} \right|_{s=0}, \quad (2.38)$$

where $\langle\langle \cdot \rangle\rangle$ indicates cumulant (mean, variance, etc.). It thus contains the full statistical information about K . Furthermore, just like a free energy, its analytic properties encode the phase structure of the dynamics. In particular, singularities of $\theta(s)$ are indicative of dynamical, or trajectory, phase transitions.

It is useful to clarify the meaning of s at this point. While K is a physical observable, its conjugate field s is not in principle a physical field (such as pressure or magnetic field); it is a mathematical field which defines the generating function of K . Furthermore, it would appear that only $s = 0$ matters. But to recover all cumulants, derivatives to all orders at $s = 0$ are needed, see (2.38), so that the behaviour of $\theta(s)$ for all values of s is relevant to the full statistics of K . In the vicinity of $s = 0$ the LD function θ encodes statistical information about all trajectories (as one calculates averages by setting $s = 0$) and therefore provides information about *typical* dynamics. In contrast, θ in the vicinity of $s \neq 0$ carries information about *rare* trajectories and thus atypical dynamics. In particular, we can think of s as defining an ensemble of trajectories whose probability is given by

$$P_s(\mathbf{X}_\tau) \equiv Z_\tau^{-1}(s) e^{-s\hat{K}[\mathbf{X}_\tau]} P(\mathbf{X}_\tau), \quad (2.39)$$

where \mathbf{X}_τ are the same trajectories as the ones generated by the dynamics (2.27-2.28), but their probability of occurring is now biased by their activity. Of particular interest, singular features in $\theta(s)$ away from $s = 0$ imply the existence of “fat” tails in the distribution $P_s(\mathbf{X}_\tau)$. The ensemble given by (2.39) is often termed *s-ensemble* [12, 13, 16]. Note that this is an ensemble of trajectories defined by controlling the total time τ and the field s , so it can be thought of as a (τ, s) -ensemble for trajectories, analogous to a (V, μ) -ensemble for configurations. Figure 3.1(a) illustrates this ensemble.

The function $\theta(s)$ can be obtained from a deformation of the master operator \mathbb{W} [24, 25, 48]. Specifically, for the case of the activity in classical systems, this deformed operator is [12, 16]

$$\mathbb{W}_s \equiv \sum_{\mathcal{C}' \neq \mathcal{C}} e^{-s} W(\mathcal{C} \rightarrow \mathcal{C}') |\mathcal{C}'\rangle \langle \mathcal{C}| - \sum_{\mathcal{C}} R(\mathcal{C}) |\mathcal{C}\rangle \langle \mathcal{C}|, \quad (2.40)$$

and $\theta(s)$ is its largest eigenvalue. For general s the operator \mathbb{W}_s does not conserve probability and does not describe a proper stochastic evolution. It only does so

at $s = 0$, where it reverts to \mathbb{W} , and where $\theta(0) = 0$. The operator \mathbb{W}_s is the “transfer matrix” for the “partition sum” $Z_\tau(s)$, that is,

$$Z_\tau(s) = \langle - | e^{\tau \mathbb{W}_s} | p_0 \rangle, \quad (2.41)$$

where $|p_0\rangle$ is the vector for the initial state probability, $|p_0\rangle \equiv \sum_{\mathcal{C}} p_0(\mathcal{C}) |\mathcal{C}\rangle$. The above expression is easy to prove from the definitions (2.29) and (2.33). Just like in equilibrium statistical mechanics, this provides a simplification, as calculating $\theta(s)$ then becomes an eigenvalue problem. For quantum systems we instead have

$$\mathcal{W}_s(\cdot) \equiv -i[H, \cdot] + \sum_{i=1}^{N_L} e^{-s} L_i(\cdot) L_i^\dagger - \frac{1}{2} \{L_i^\dagger L_i, \cdot\}. \quad (2.42)$$

2.5 Systems and relations of interest

While the primer on the thermodynamics of trajectories provides the background theory for the tools and methods to be used in this thesis, the remainder of this chapter is dedicated to providing a brief introduction to the systems and relations to be studied in the context of dynamical ensembles. This starts with a presentation of glass formers, where the s -ensemble has already provided some fresh insight into the debate on the nature of glass formation [12, 13]. Following this we look at some of the useful methods of statistical ensembles, namely the Jarzynski equality and the Lee-Yang zeros method, both of which should be adaptable to the dynamical context.

2.5.1 Models of glass formers

As supercooled liquids approach their glass transitions, one observes a very sharp increase in their viscosities and structural relaxation times. Below their glass temperature, T_g , these materials have a non-zero static shear modulus - indicative of

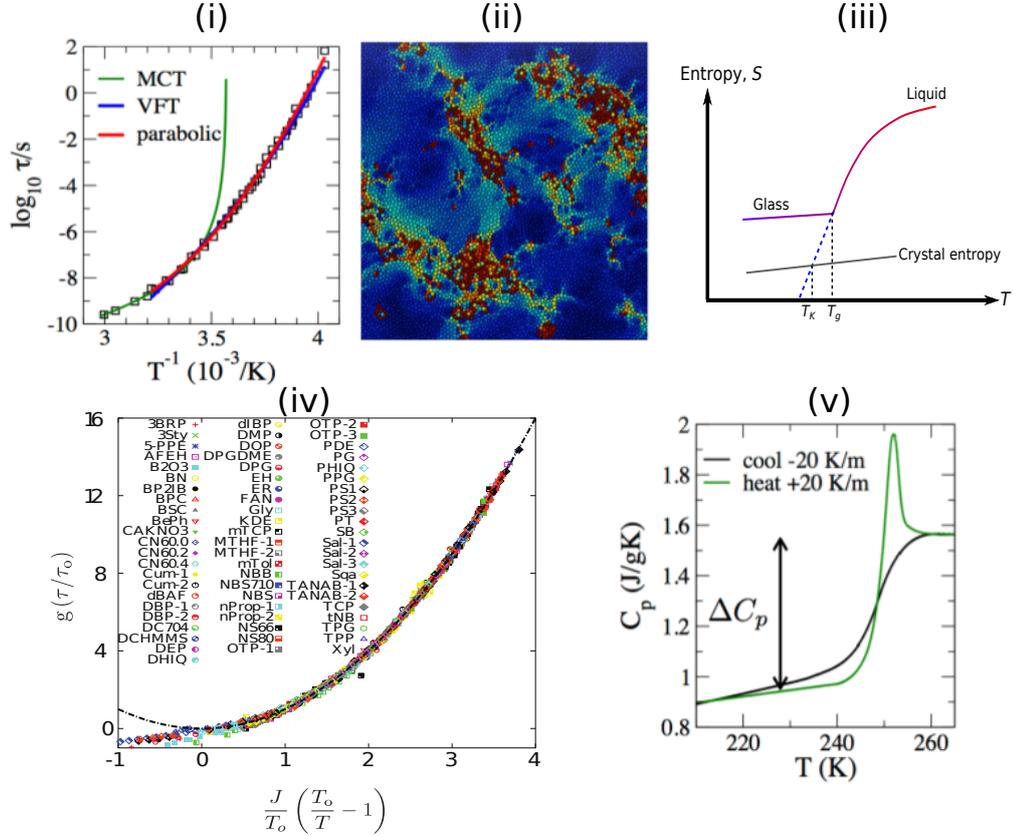


Figure 2.1: (i), (ii), (v) Taken from Ref. [7]. (i) Super-Arrhenius relaxation of orthoterphenyl (OTP) showing good agreement with both VFT (blue) and parabolic (red) laws. Also plotted is a mode-coupling theory (green) prediction, showing it break down at low temperature. (ii) Dynamical heterogeneities in an equilibrium trajectory of two-dimensional mixture. Overlap of particles with their original position after one tenth of the systems relaxation time are plotted with high overlap (no movement) in blue and low overlap in red. (iii) Schematic illustration of temperature dependent entropy of a glass former. Extrapolating the slope of the liquid phase intersects the entropy of the crystal at a non-zero temperature. Theories must either account for the inability to crystallise, or uncover a phase transition to an ideal glass state that occurs before T_K . (iv) Taken from Ref. [111] Relaxation of a large number of different materials, all exhibiting super-Arrhenius relaxation. (v) The difference in specific heat of glass and liquid phases, C_p of OTP showing anomalous differences when the sample is cooled (black) and heated (green) through the same temperature range. This illustrates that the properties of the material are unusually dependent on its history.

a solid - yet they lack any long range order. In essence glasses appear to be liquids that are unable to flow. Confusingly this glass temperature, and other properties of the resulting glass, are not only dependent in the material in question, but also the specific protocol used to cool the glass, with different rates resulting in different T_g . The viscosities of “strong” glass formers, η follow a characteristic *Arrhenius* law, $\eta \sim \exp(\beta\Delta)$ where Δ is the microscopic energy barrier to motion and β is the inverse temperature. Glassiness occurs at low temperature as the system is unable to overcome these barriers, and the dynamics become blocked. More confusing are so-called “fragile” glass formers, exhibit super-Arrhenius behaviour, the precise functional form of which is still debated (see Ref. [7] and Fig.2.1(i)). Glass formers also characteristically exhibit *dynamical heterogeneities* – space-time “bubbles” of differing activity within the bulk of a glass former (see Fig.2.1(ii)). This demonstrates that the relaxation does not occur evenly, in both space and time, and that the dynamics are unusually correlated beyond the apparent structure - which remains that of a liquid. Furthermore if one extrapolates the structural entropy of the liquid phase, ignoring the abrupt change in slope at T_g (see Fig. 2.1(iii)), one finds that the entropy of the liquid phase crosses that of an ideal crystal at a finite positive temperature, known as the Kauzmann temperature, T_k [49]. The argument is then that either a liquid must crystallise before T_k , or there must be a true thermodynamic phase transition to a glass state. The physical mechanism underlying this slow dynamics remains controversial [4–7]. Some theories, particularly the random first-order transition theory (RFOT) [8], propose that glassy systems are approaching some kind of thermodynamic phase transition, with associated collective (slow) dynamics. These approaches paint the picture of a “rugged” thermodynamic landscape with many deep, local minima in free energy functions. They argue that upon cooling the number of energy minima reduces, becoming sub-exponential in system size. Glassiness then arises because systems fall out of equilibrium as the liquid is cooled, breaking up into

small subsystems, each with its own amorphous order in a local energy minima, without ever reaching the global (liquid) minimum. A characteristic length scale, l^* can be defined, subsystems smaller than this scale behave as glasses and are unable to relax, while those on longer scales exist as supercooled liquids. It is argued that l^* is determined by the thermodynamic properties of the system, and that it diverges at T_g , ensuring a thermodynamic phase transition at some non-zero temperature T_c (which is often related to, but not necessarily exactly T_k), to an ideal glass state. The existence of such phase transitions can be probed by computing the free energy of a pair of coupled copies (or replicas) of the system, and searching for a transition as a function of both temperature and coupling strength. These transitions link an equilibrium-like phase where the replicas are different from each other (the liquid) to one where they become very similar (the glass) [11]. The similarity between the configurations is measured by an overlap variable, which is the order parameter for the transition. A functional form for super-Arrhenius relaxation, known as the Vogel-Fulcher-Tammann (VFT) law, accounts for the singularity at T_c in the RFOT description, with relaxation given by $\tau = \tau_0 \exp[A/(T - T_c)]$, where A is a system-dependent constant.

An alternative approach, that of dynamical facilitation (DF) [50], links glassy behaviour to a dynamical phase transitions that arise from particle motion being blocked (or facilitated) by neighbouring particles as in the East model, a one-dimensional lattice spin model where spinflips can only occur on sites whose neighbour to the left is in excited state. Such models are thermodynamically trivial, but with super-Arrhenius barriers to particle motion [12, 13]. By analysing their dynamical phase structure using the s -ensemble, one may infer the existence of transitions between an active dynamical phase (the equilibrium liquid) and an inactive phase (the non-equilibrium glass). This provides a natural explanation for the origins of dynamical heterogeneities, and accounts for the onset of glassiness without a thermodynamic transition. Proponents of this approach

argue that complex dynamical pathways alone are enough to account for the inability to crystallise, and the liquids invariably fall out of equilibrium on experimental timescales. DF predicts super-Arrhenius parabolic relaxation of the form $\tau = \tau_0 \exp[(J/T_o)^2(T_o/T - 1)^2]$ where T_o is a reference temperature for the onset of glassiness.

Frustratingly both of these very different forms for super-Arrhenius relaxation fit existing experimental data well (see Fig 2.1(iv)). While a more recent experiment on 20 million year old amber showing divergence with the VFT form and agreement with the parabolic form [51], there is still much debate over the correct functional form of super-Arrhenius relaxations - it is not clear whether all materials must necessarily follow parabolic relaxation.

The critical issue in resolving these differing approaches lies in the difficulty of studying models of real glass formers in three dimensions. RFOT based approaches are well studied and compelling in the mean field limit but symmetry considerations mean that the transitions they predict can not be present in two dimensions. On the other hand the dynamical approach has several compelling results in lower dimensional systems, but the complexity involved in three dimensional glass formers means that making connections to real glass is a difficult endeavour. As of yet, no single theory captures all of the properties of real glasses.

2.5.2 Fluctuation theorems and the Jarzynski equality

One of the more problematic questions in statistical physics, and indeed all of physics, is irreversibility, or more specifically how time-irreversible properties can emerge from fundamentally reversible dynamics. This question is, at its heart, a question of entropy - the second law of thermodynamics only holds on a statistical

average, and indeed, only for sufficiently large systems. With technology shrinking to smaller scales it became necessary to rigorously quantify the scales on which fully reversible descriptions of nature start to give way to the second law. The results of this investigation are known as *fluctuation theorems* [37,38,52–58]. While the second law ensures that, on average, entropy must increase for driven processes, this is not true for an individual realisation of a trajectory, where entropy may decrease. Fluctuation theorems encode the precise probability that such a drop in entropy occurs. Their study has led to one of the most celebrated recent developments in statistical physics - the Jarzynski equality [37], which connects equilibrium free energy differences with the work done in *non-equilibrium* (i.e. fast) processes joining the two states.

Early studies on fluctuation theorems by Evans et al. [59–61] considered the energy dissipated by some fixed time trajectory of a system, characterized by a dissipation function, Γ . This led to what is now known as the Fluctuation Theorem (FT), which relates the probability the dissipation function takes on a specific value, \mathcal{A} to its negative:

$$\frac{P(\Gamma = \mathcal{A})}{P(\Gamma = -\mathcal{A})} = e^{At}. \quad (2.43)$$

Negative values of Γ are thus suppressed exponentially, and it becomes clear how violations of the second law are only expected for small systems and short times.

The situation considered by the Jarzynski relation is slightly different - that of a system initially in thermal equilibrium with a heat reservoir, where the system is subsequently driven away from equilibrium by externally changing one or more parameters. The dynamics of the system obeys detailed balance with respect to the stationary distribution corresponding to the instantaneous values of the control parameters. Non-equilibrium can then be described as a “lag” between this stationary distribution and the actual distribution [62]. The Jarzynski equality relates the average work done over all trajectories with the change of free energy

between initial, \mathcal{C} and final, \mathcal{C}' state. From the second law one expects the free energy difference, $\Delta F = F_{\mathcal{C}'} - F_{\mathcal{C}}$ between the states to be related to the work done, W in moving from one to the other by

$$\Delta F \leq W. \quad (2.44)$$

Equation 2.44 is only an equality when the system is changed quasistatically (i.e. infinitely slowly). The Jarzynski equality, however, finds that

$$e^{-\beta\Delta F} = \langle e^{-\beta W} \rangle \quad (2.45)$$

where the brackets denotes the average over all possible trajectories connecting \mathcal{C} and \mathcal{C}' and $\beta = 1/k_B T$. This is closely related to (and can be derived from) the Crooks' Fluctuation Theorem [54,63], which relates the work done in an arbitrarily fast process connecting $\mathcal{C} \rightarrow \mathcal{C}'$, to the work done in the reverse process taking $\mathcal{C}' \rightarrow \mathcal{C}$:

$$\frac{P_{\mathcal{C} \rightarrow \mathcal{C}'}(W = \mathcal{A})}{P_{\mathcal{C}' \rightarrow \mathcal{C}}(W = -\mathcal{A})} = e^{\mathcal{A} - \Delta F} \quad (2.46)$$

Remarkably, these relations allow for the calculation of equilibrium free energy differences using non-equilibrium protocols.

2.5.3 Lee-Yang partition function zeros method

In 1952 Lee and Yang made an important advancement in our understanding of phase transitions [39,40]. By studying Ising models in a magnetic field, they noted the partition function could be expanded into the complex magnetic field plane, revealing that all of the zeros of the partition function lie on a ring in the complex plane. As the system is tuned towards a phase transition (including taking the system size to the thermodynamic limit), these zeros move towards the real axis, eventually intersecting it at a phase transition point. While there was limited

success in expanding this analysis of the zeros of the partition function in the thermodynamic case, it has been adapted to the dynamical case [41–43]. Through an analysis of high-order cumulants of the systems dynamics, the motion of the zeros of the MGF at short, but increasing, times can be calculated [64–66]. As in the thermodynamic case these zeros lie in the complex s plane, and move towards any dynamical phase transition points on the real s axis in the “thermodynamic limit” of large time and system size. The location of these dynamical phase transitions can thus be extrapolated from the short-time motion of the zeros. This has interesting implications for dynamical phase transitions, and particularly, instances where the zeros of the MGF lie near but not on, the real axis in the large time limit can heavily influence the dynamics of a system, causing a dynamical crossover without a true phase transition. As such the methods of Lee and Yang provide potentially useful tools for understanding a variety of dynamical behaviours.

3. FLUCTUATING OBSERVATION TIME ENSEMBLES IN THE THERMODYNAMICS OF TRAJECTORIES: THEORY

While the thermodynamics of trajectories has provided great insight into the dynamical properties of both classical and quantum systems [12, 13, 15–20, 42, 47, 48, 67–90], studies have thus far been mostly limited to either the “microcanonical” dynamical ensemble - where the set of trajectories is conditioned on all of them having a fixed value of a time-integrated quantity (for example fixed total activity, see Chapter 2), or the “canonical” dynamical ensemble - where the set of trajectories is such that a time-integrated quantity, while not strictly fixed, has a specific fixed average (as in the s -ensemble) [91]. A large part of the ubiquitous success of statistical mechanics, however, comes from the flexibility offered by the description of its ensembles. For any system of interest it is a simple matter of finding the ensemble best suited to treat it. As such, there is likely to be great benefit in exploring alternative dynamical ensembles as well, ensembles that use order parameters other than the number of counted dynamical events.

In this chapter we introduce a *fluctuating observation time* ensemble we denote the x -ensemble, which uses the total time of the trajectory, τ as the order parameter with fixed numbers of dynamical events, K . The correspondence between this and the s -ensemble is examined and discussed heuristically, with simple examples of both classical and quantum stochastic systems whose dynamics are Markovian, and described by master equations. This new ensemble will form the basis for the remainder of the thesis.

3.1 Fluctuating observation time: the x -ensemble

Consider the case where, instead of keeping fixed the total time τ of trajectories generated by (2.27)-(2.28), what is kept fixed is the total number of configuration changes K , i.e. the activity, in each trajectory. That is, if we denote by \mathbf{Y}_K such a trajectory, then $\mathbf{Y}_K = (\mathcal{C}_0 \rightarrow \mathcal{C}_{t_1} \rightarrow \dots \rightarrow \mathcal{C}_\tau)$, where the number of configuration changes is fixed to be K , but the time τ of the final K -th jump fluctuates from trajectory to trajectory. From (2.27)-(2.28) the probability of \mathbf{Y}_K is

$$P(\mathbf{Y}_K) = p_0(\mathcal{C}_0) \prod_{i=1}^K e^{-(t_i - t_{i-1})R(\mathcal{C}_{t_{i-1}})} W(\mathcal{C}_{t_{i-1}} \rightarrow \mathcal{C}_{t_i}), \quad (3.1)$$

where $t_0 = 0$ and $t_K = \tau$.

In analogy with Section 2.4, we ask the question: what is the distribution $P_K(\tau)$ of total trajectory length τ for fixed activity K . From the definitions above we have,

$$\begin{aligned} P_K(\tau) &= \sum_{\mathbf{Y}_\tau} \delta(\tau - \hat{\tau}[\mathbf{Y}_K]) P(\mathbf{Y}_K) \\ &= \sum_{\mathcal{C}_0 \dots \mathcal{C}_K} p_0(\mathcal{C}_0) \prod_{i=1}^{K-1} \int_0^{t_{i+1}} dt_i e^{-(t_i - t_{i-1})R(\mathcal{C}_{t_{i-1}})} \\ &\quad \times W(\mathcal{C}_{t_{i-1}} \rightarrow \mathcal{C}_{t_i}). \end{aligned} \quad (3.2)$$

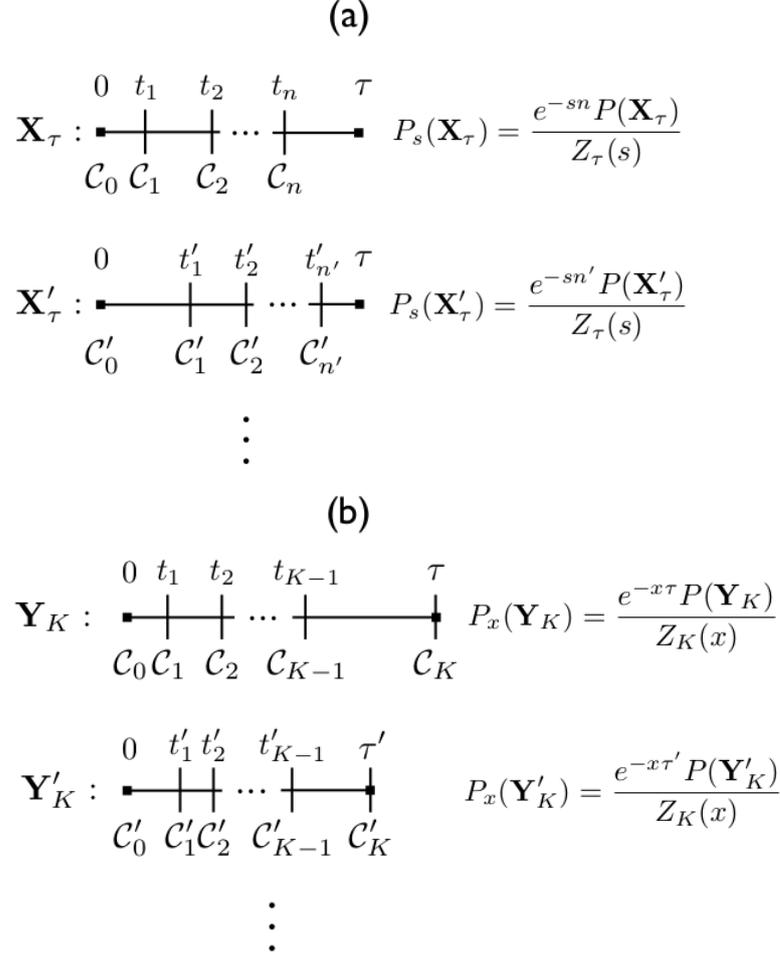


Figure 3.1: (a) The s -ensemble is the set of all trajectories that are possible with the dynamics (2.27)-(2.28), of fixed total time τ , and where the probability of each trajectory is weighed by the number of configuration changes. It is the ensemble defined by fixed (τ, s) . The ensemble of trajectories that corresponds to the actual dynamics (2.27)-(2.28) is given by $(\tau, 0)$. We sketch two trajectories, the squares indicate the times where trajectories begin and end, and the tick where jumps between configurations take place. (b) The x -ensemble is the set of all trajectories that are possible with the dynamics (2.27)-(2.28), of fixed number of configuration changes K , and where the probability of each trajectory is weighed by the trajectory length. It is the ensemble defined by fixed (x, K) .

The corresponding moment generating function for τ is

$$\begin{aligned} Z_K(x) &\equiv \int_0^\infty d\tau e^{-x\tau} P_K(\tau) \\ &= \sum_{\mathbf{Y}_K} e^{-x\hat{\tau}[\mathbf{Y}_K]} P(\mathbf{Y}_K), \end{aligned} \quad (3.3)$$

so that $\langle \tau^n \rangle = (-)^n \partial_x^n Z_K(x)|_{x=0}$.

Similarly to what we discussed for the s -ensemble, for large K , the probability of the *intensive* trajectory length, τ/K , has a LD form,

$$P(\tau/K) \sim e^{-K\phi(\tau/K)}, \quad (3.4)$$

and so does the generating function,

$$Z_K(x) \sim e^{Kg(x)}. \quad (3.5)$$

The definitions (3.2)-(3.5) are analogous to (2.32)-(2.35) above: all trajectories have fixed activity K (cf. τ above); the large limit is that of large K (cf. large τ); the LD function ϕ determines the probability of τ at large K (cf. φ for K at large τ); x is the conjugate field to τ (cf. s and K), and the LD function g is the cumulant generating function for τ at large K (cf. $\theta(s)$ for K at large τ). As before, the LD functions ϕ and g are related by Legendre-Fenchel transforms,

$$\phi(t) = -\min_x [g(x) + tx], \quad g(x) = -\min_t [\phi(t) + tx]. \quad (3.6)$$

Equation (3.3) is the ‘‘partition sum’’ for the ensemble of trajectories with probabilities

$$P_x(\mathbf{Y}_K) \equiv Z_K^{-1}(x) e^{-x\hat{\tau}[\mathbf{Y}_K]} P(\mathbf{Y}_K). \quad (3.7)$$

If the s -ensemble of the previous section, of fixed (τ, s) , is analogous to an equilibrium (V, μ) ensemble (since τ plays the role of volume and s of chemical potential for the activity), then this x -ensemble, of fixed (x, K) , can be thought of as analogous to an equilibrium (p, N) ensemble, as x is conjugate to the size trajectories τ (cf. p and V for configurations). The x -ensemble is sketched in Fig. 3.1(b).

The generating function $Z_K(x)$ can also be written in terms of a transfer matrix operator. This can be proved as follows. If we define the operator,

$$\mathbb{T}(t) \equiv \sum_{\mathcal{C}' \neq \mathcal{C}} W(\mathcal{C} \rightarrow \mathcal{C}') e^{-tR(\mathcal{C})}, \quad (3.8)$$

then the probability of a trajectory in Eq. (3.1) is given by the product of its matrix elements. This is because a fixed K trajectory can be thought of as a trajectory of a *discrete* Markov chain with transition operator $\mathbb{T}(t)$. The probability of observing a total time τ in K jumps therefore given by K convolutions of $\mathbb{T}(t)$. From Eq. (3.3) we have that $Z_K(x)$ is the Laplace transform of this convolution, and thus given by

$$Z_K(x) = \langle - | \mathbb{T}_x^K | p_0 \rangle, \quad (3.9)$$

where

$$\mathbb{T}_x \equiv \sum_{\mathcal{C}' \neq \mathcal{C}} \frac{W(\mathcal{C} \rightarrow \mathcal{C}')}{x + R(\mathcal{C})} |\mathcal{C}'\rangle \langle \mathcal{C}|, \quad (3.10)$$

is the Laplace transform of $\mathbb{T}(t)$. The LD function $g(x)$ then corresponds to the logarithm of the largest eigenvalue of \mathbb{T}_x .

3.2 Ensemble correspondence

Both the s -ensemble and the x -ensemble are different ways to consider the same underlying dynamics generated by Eqs. (2.27)-(2.28). Just like in the configurational equilibrium case, we expect that in the “thermodynamic limit” of large τ and large K the two ensembles will be equivalent (except perhaps at phase transitions [91]), and that the properties of one ensemble will be obtainable from those of the other. This correspondence can be proved directly from the spectral properties of the operators \mathbb{W}_s and \mathbb{T}_x .

The matrices \mathbb{W}_s and \mathbb{T}_x are directly related to each other. Specifically, from Eqs. (2.40) and (3.10) we have

$$e^{-s}\mathbb{T}_x = \mathbb{W}_s \cdot \mathbb{Q}_x + \mathbb{I} - x\mathbb{Q}_x, \quad (3.11)$$

where \mathbb{I} is the identity, $\mathbb{I} \equiv \sum_{\mathcal{C}} |\mathcal{C}\rangle\langle\mathcal{C}|$, and \mathbb{Q}_x the diagonal matrix,

$$\mathbb{Q}_x \equiv \sum_{\mathcal{C}} \frac{1}{x + R(\mathcal{C})} |\mathcal{C}\rangle\langle\mathcal{C}| \quad (3.12)$$

Consider now a left vector $\langle l|$ that is simultaneously an eigenvector of \mathbb{W}_s and \mathbb{T}_x with eigenvalues $\theta(s)$ and $e^{g(x)}$, respectively. If we left multiply (3.11) by this vector we obtain,

$$(e^{-s+g(x)} - 1) \langle l| = [\theta(s) - x] \langle l| \mathbb{Q}_x. \quad (3.13)$$

We see that our assumption (of $\langle l|$ being an eigenvector of both \mathbb{W}_s and \mathbb{T}_x) can only be satisfied if $g(x) = s$ and $\theta(s) = x$. That is, given the function g , the function θ is obtained from its inverse, and vice-versa,

$$\theta(s) = g^{-1}(s), \quad g(x) = \theta^{-1}(x). \quad (3.14)$$

Since the LD rate functions are convex the relation between g and θ is one-to-one (except perhaps at their boundaries, or at phase-transition points [91]). Equation (3.14) is the statement of the correspondence between the s -ensemble and the x -ensemble at the level of their respective “free-energies”.

It is illuminating at this point to expand on the analogy with thermodynamic ensembles. If one considers the trajectory observation time τ to be a effective volume, with the events K as particles that fill up this volume, then the field s controls the addition and removal of particles, making it a chemical potential, μ . The s -ensemble is then the analogue of fixed volume ensemble with fluctuating particle numbers - the grand canonical ensemble. The x field is then conjugate to the volume when particle numbers are fixed, i.e. a pressure, p . The x -ensemble is

then the analogue of the isobaric-isothermal ensemble, more commonly known as the (NpT) -ensemble. A summary of this analogy is presented below.

Dynamical property	Thermodynamic equiv.	s -ensemble	x -ensemble
τ (time)	V (Volume)	fixed	fluctuates
K	N (number of particles)	fluctuates	fixed
Counting fields s, x	μ, p	s fixed	x fixed
Moment generating function	Partition function	$Z_\tau(s)$	$Z_K(x)$
Cumulant generating function	Free energy density	$\theta(s)$	$g(x)$
LD rate function	Entropy density	$\psi(K/\tau)$	$\phi(\tau/K)$

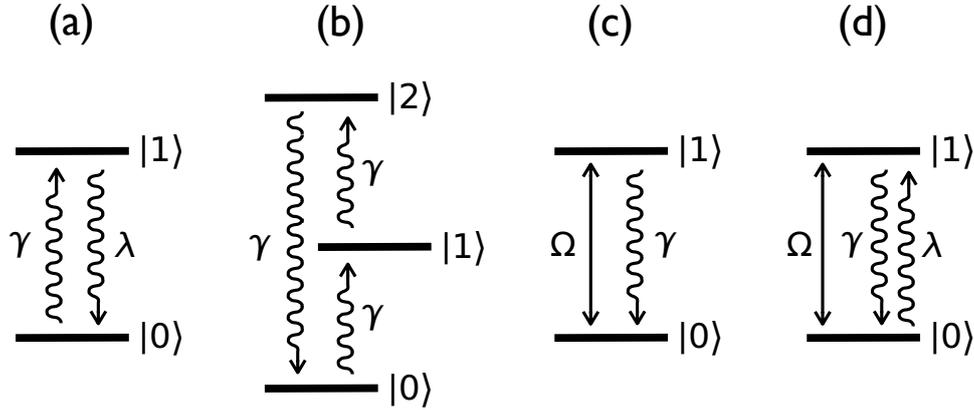


Figure 3.2: (a) Classical two-level system. (b) Classical three-level system. (c) Quantum $T = 0$ two-level system: the full line indicates a coherent transition of frequency Ω and the wavy line a dissipative quantum jump of rate γ associated to emission into the bath. (d) Quantum $T \neq 0$ two-level system: same as before, but now absorption from the bath leads to a second kind of quantum jump with rate λ .

3.2.1 Example: classical two-level system

As an elementary example consider the classical two level system of Fig. 3.2(a), where there are only two configurations, $\mathcal{C} \in \{0, 1\}$, and the transition rates are,

$W(0 \rightarrow 1) = \gamma$ and $W(1 \rightarrow 0) = \lambda$. With only one transition out of each state the escape rates are thus simply $R(0) = \gamma$ and $R(1) = \lambda$ respectively. The operator \mathbb{W}_s is,

$$\mathbb{W}_s = \begin{pmatrix} -\lambda & e^{-s}\gamma \\ e^{-s}\lambda & -\gamma \end{pmatrix}, \quad (3.15)$$

and $\theta(s)$ is given by its largest eigenvalue,

$$\theta(s) = \frac{1}{2}(\lambda + \gamma) \left(\sqrt{1 - \frac{4\lambda\gamma}{(\lambda + \gamma)^2}(1 - e^{-2s})} - 1 \right). \quad (3.16)$$

We can explicitly see from this that probability is only conserved at $s = 0$, as this is the only situation with a zero eigenvalue, corresponding to a steady state. From $\theta(s)$ we can extract the cumulants of the activity. For the average activity per unit time, i.e. the average transition rate between the two levels, we obtain

$$\frac{\langle K \rangle_\tau}{\tau} = -\theta'(0) = \frac{2\lambda\gamma}{\lambda + \gamma}, \quad (3.17)$$

as expected. For the case $\gamma = \lambda$ the LD function reduces to $\theta(s) = \lambda(e^{-s} - 1)$, which is the cumulant generating function for a Poisson process with rate λ .

Similarly, the operator \mathbb{T}_x for this problem reads,

$$\mathbb{T}_x = \begin{pmatrix} 0 & \frac{\gamma}{x+\gamma} \\ \frac{\lambda}{x+\lambda} & 0 \end{pmatrix}, \quad (3.18)$$

and from its largest eigenvalue we obtain the LD function $g(x)$

$$g(x) = \frac{1}{2} \log \left(\frac{\lambda\gamma}{(x + \lambda)(x + \gamma)} \right). \quad (3.19)$$

This function g is indeed the inverse of the function θ (3.16). The cumulants of the total time are obtained from $g(x)$. In particular, the average total time, scaled by the number of jumps, is

$$\frac{\langle \tau \rangle_K}{K} = -g'(0) = \frac{\lambda + \gamma}{2\lambda\gamma}, \quad (3.20)$$

which is the inverse of (3.17). Analogous relations between the moments of K in the fixed τ ensemble, and those of τ in the fixed K ensemble can be obtained by virtue of Eq. (3.14).

3.3 Generalisation of the x -ensemble to multiple observables

In the sections above we have proved the correspondence between the s -ensemble of fixed observation time τ , where s is conjugate to the overall activity, and the x -ensemble of fixed activity K , where x is conjugate to the total time. This correspondence can be extended to the case where one or more s fields couple to other time-extensive quantities.

Consider a setup like that of section 2.4, but now we are interested in the statistics of several different time-extensive quantities [12]. For example, one could think of counting, instead of the total activity, the total number of certain kind of transitions, or the time integral of a certain quantity such as the energy. Let us say that there are N different dynamical observables, which we denote collectively by the vector $\vec{M} \equiv (M_1, \dots, M_N)$. Under the dynamics Eqs. (2.27)-(2.28) there will be a joint probability for observing a combination of these M quantities, $P_\tau(\vec{M})$. For large τ this joint probability will have a LD form,

$$P(\vec{M}/\tau) \sim e^{-\tau\Phi(\vec{M}/\tau)}, \quad (3.21)$$

where the LD function now depends on the whole vector of intensive observables $(M_1/\tau, \dots, M_N/\tau)$. The corresponding moment generating function for \vec{M} also has a LD form at large τ [25],

$$Z_\tau(\vec{s}) \equiv \sum_{\vec{M}} e^{-\vec{s} \cdot \vec{M}} P_\tau(\vec{M}) \sim e^{\tau\Theta(\vec{s})}, \quad (3.22)$$

where for each observable M_n there is a counting field s_n , collected in the vector $\vec{s} \equiv (s_1, \dots, s_N)$, and where the LD function $\Theta(\vec{s})$ is a now function of this whole vector.

The partition function $Z_\tau(\vec{s})$ has a transfer matrix representation similar to (2.41)

in terms of an operator $\mathbb{W}_{\vec{s}}$. For simplicity we will assume that the time-extensive observables \vec{M} only change at jumps between configurations (extending to cases where observables accumulate in the periods between jumps is straightforward). In this case $\mathbb{W}_{\vec{s}}$ reads [16],

$$\begin{aligned} \mathbb{W}_{\vec{s}} \equiv & \sum_{\mathcal{C}' \neq \mathcal{C}} e^{-\vec{s} \cdot \vec{m}(\mathcal{C} \rightarrow \mathcal{C}')} W(\mathcal{C} \rightarrow \mathcal{C}') |\mathcal{C}'\rangle \langle \mathcal{C}| \\ & - \sum_{\mathcal{C}} R(\mathcal{C}) |\mathcal{C}\rangle \langle \mathcal{C}|, \end{aligned} \quad (3.23)$$

where $m_n(\mathcal{C} \rightarrow \mathcal{C}')$ is the change in M_n under the transition $\mathcal{C} \rightarrow \mathcal{C}'$ (for the activity this was just 1 for all $\mathcal{C}, \mathcal{C}'$ as it counted all transitions equally). $\Theta(\vec{s})$ is the largest eigenvalue of the operator (3.23). As before, LD functions are related by Legendre-Fenchel transforms

$$\Phi(\vec{m}) = - \min_{\vec{s}} [\Theta(\vec{s}) + \vec{m} \cdot \vec{s}], \quad \Theta(\vec{s}) = - \min_{\vec{m}} [\Phi(\vec{m}) + \vec{m} \cdot \vec{s}].$$

Eqs. (3.22),(3.23) define an (τ, \vec{s}) -ensemble for a general set of dynamical order parameters \vec{M} .

In analogy with section 3.1, there is a corresponding construct for studying the statistics of \vec{M} in trajectories where the total activity K is fixed. The probability of observing \vec{M} , together with a total time τ , for a fixed and large K has the form,

$$P(\tau/K, \vec{M}/K) \sim e^{-K\Phi(\tau/K, \vec{M}/K)}. \quad (3.24)$$

The corresponding moment generating function is

$$\begin{aligned} Z_K(x, \vec{s}) & \equiv \sum_{\vec{M}} \int_0^\infty d\tau e^{-x\tau - \vec{s} \cdot \vec{M}} P_K(\tau, \vec{M}) \\ & \sim e^{KG(x, \vec{s})}, \end{aligned} \quad (3.25)$$

with Φ and G related by

$$\begin{aligned} \Phi(t, \vec{m}) & = - \min_{x, \vec{s}} [G(x, \vec{s}) + tx + \vec{m} \cdot \vec{s}], \\ G(x, \vec{s}) & = - \min_{t, \vec{m}} [\Phi(t, \vec{m}) + tx + \vec{m} \cdot \vec{s}]. \end{aligned}$$

Equations (3.24),(3.25) define a generalised x -ensemble.

The partition sum $Z_K(x)$ can be written in terms of a transfer matrix, $Z_K(x) = \langle - | \mathbb{T}_{x,\vec{s}}^K | p_0 \rangle$, with

$$\mathbb{T}_{x,\vec{s}} \equiv \sum_{\mathcal{C}' \neq \mathcal{C}} \frac{W_{\vec{s}}(\mathcal{C} \rightarrow \mathcal{C}')}{x + R(\mathcal{C})} |\mathcal{C}'\rangle \langle \mathcal{C}|, \quad (3.26)$$

where $W_{\vec{s}}(\mathcal{C} \rightarrow \mathcal{C}')$ are the coefficients of the off-diagonal entries of (3.23). This allows us to prove the ensemble correspondence in this generalised case. From Eqs. (3.12),(3.23),(3.26) we have that

$$\mathbb{T}_{x,\vec{s}} = \mathbb{W}_{\vec{s}} \cdot \mathbb{Q}_x + \mathbb{I} - x\mathbb{Q}_x. \quad (3.27)$$

As in section 3.2 we search for conditions for which a left vector $\langle l |$ is simultaneously an eigenvector of both $\mathbb{W}_{\vec{s}}$ and $\mathbb{T}_{x,\vec{s}}$ with eigenvalues $\Theta(\vec{s})$ and $e^{G(x,\vec{s})}$, respectively. Multiplying $\langle l |$ into = (3.27) we get,

$$(e^{G(x,\vec{s})} - 1) \langle l | = [\Theta(\vec{s}) - x] \langle l | \mathbb{Q}_x. \quad (3.28)$$

It follows that Eq. (3.28) is satisfied when

$$\Theta(\vec{s}) = x_*(\vec{s}), \quad (3.29)$$

where $x_*(\vec{s})$ is the solution of

$$G(x_*(\vec{s}), \vec{s}) = 0. \quad (3.30)$$

Equations (3.29),(3.30) prove the correspondence between the general \vec{s} -ensemble and the general x -ensemble: they allow to obtain the “free energy” LD functions in one from those in the other, and thus encode the statistical properties of each other. In the case where \vec{M} corresponds only to the activity K , the function $G(x, s) = g(x) - s$, and Eqs. (3.29),(3.30) reduce to (3.14).

3.3.1 Example: classical three-level system

As a simple example of how the general s - and x - ensembles relate, consider the classical three-level system in a setup like that of Fig. 3.2(b). We have three allowed transitions each with the same rate, $W(0 \rightarrow 1) = W(1 \rightarrow 2) = W(2 \rightarrow 0) = \gamma$, with all other transitions prohibited $W(0 \rightarrow 2) = W(1 \rightarrow 0) = W(2 \rightarrow 1) = 0$, and the escape rate from each state is thus $R(0) = R(1) = R(2) = \gamma$. Suppose we only observe the jumps between configurations 2 and 0. In the notation above we have $N = 1$, and \vec{M} is just K_{20} , the total number of transitions between top and bottom levels. In the s -ensemble, the largest eigenvalue of the operator

$$\mathbb{W}_{s_{20}} = \gamma \begin{pmatrix} -1 & 1 & 0 \\ 0 & -1 & 1 \\ e^{-s_{20}} & 0 & -1 \end{pmatrix}, \quad (3.31)$$

(where s_{20} is the field conjugate to K_{20}) gives the LD function $\Theta(s_{20}) = \gamma(e^{-s/3} - 1)$, which is the cumulant generating function for the number of jumps K_{20} per unit time. In the x -ensemble context, the relevant operator is

$$\mathbb{T}_{x,s_{20}} = \frac{\gamma}{x + \gamma} \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ e^{-s_{20}} & 0 & 0 \end{pmatrix}. \quad (3.32)$$

From its largest eigenvalue we obtain the LD function $G(x, s_{20}) = -s_{20}/3 + \log \gamma - \log(x + \gamma)$. This is the generating function for cumulants of both τ/K and K_{20}/K . If we solve $G(x_*, s_{20}) = 0$ for x_* we get, $x_*(s_{20}) = \Theta(s_{20})$ above, in accordance with (3.29),(3.30).

3.4 x -ensemble in open quantum systems

We now turn to the quantum case, with dynamics described by the quantum master super-operator given in Eq. 2.31. We denote again by $\vec{M} \equiv (M_1, \dots, M_N)$ the time-integrated observables we wish to count, and by $\Theta(\vec{s})$ the large-deviation rate function corresponding to the cumulant generating function for \vec{M}/τ in the large τ limit. If under the action of the jump operator L_i the observable M_n is incremented by $m_n^{(i)}$, then the deformed quantum master operator for which $\Theta(\vec{s})$ is its largest eigenvalue reads [18],

$$\mathcal{W}_{\vec{s}}(\cdot) \equiv -i[H, \cdot] + \sum_{i=1}^{N_L} e^{-\vec{s} \cdot \vec{m}^{(i)}} L_i(\cdot) L_i^\dagger - \frac{1}{2} \{L_i^\dagger L_i, \cdot\}. \quad (3.33)$$

This is the open quantum equivalent s -ensemble operator to that of Eq. (3.23) for the classical case.

The generalised x -ensemble corresponds to controlling the fields x , conjugate to the total time τ , and \vec{s} , conjugate to \vec{M} , in quantum stochastic trajectories of total and fixed K quantum jumps. The corresponding LD function $G(x, \vec{s})$ is the largest eigenvalue of the super-operator,

$$\mathcal{T}_{x, \vec{s}}(\cdot) \equiv \sum_{i=1}^{N_L} e^{-\vec{s} \cdot \vec{m}^{(i)}} L_i [(x + \mathcal{R})^{-1}(\cdot)] L_i^\dagger, \quad (3.34)$$

where $(x + \mathcal{R})^{-1}$ is the inverse super-operator to $(x + \mathcal{R})$, i.e., $(x + \mathcal{R})^{-1}[(x + \mathcal{R})(\cdot)] = (\cdot)$. Here \mathcal{R} is the ‘‘escape’’ super-operator, $\mathcal{R}(\cdot) \equiv iH_{\text{eff}}(\cdot) - i(\cdot)H_{\text{eff}}^\dagger$, where $H_{\text{eff}}^\dagger \equiv H - \frac{1}{2} \sum_i L_i^\dagger L_i$. Equation (3.34) is, in the open quantum context, equivalent to Eq. (3.26) in the classical context. It is obtained in a similar manner as (3.26) by noting that the probability to observe a trajectory of K quantum jumps due to the action of operators $(L_{i_1}, \dots, L_{i_{K-1}}, L_{i_K})$ that occur at times $(t_1, \dots, t_{K-1}, \tau)$ is

$$\text{Tr} \left[L_{i_K} e^{-i(\tau - t_{K-1})H_{\text{eff}}} L_{i_{K-1}} \dots L_{i_1} e^{-it_1 H_{\text{eff}}} \rho_0 e^{it_1 H_{\text{eff}}^\dagger} L_{i_1}^\dagger \dots L_{i_{K-1}}^\dagger e^{i(\tau - t_{K-1})H_{\text{eff}}^\dagger} L_{i_K}^\dagger \right], \quad (3.35)$$

where ρ_0 is the initial density matrix.

Just like in the classical case, $\mathcal{W}_{\vec{s}}$ and $\mathcal{T}_{x,\vec{s}}$ are directly related,

$$\mathcal{T}_{x,\vec{s}}(\cdot) = \mathcal{W}_{\vec{s}}[\mathcal{Q}(\cdot)] + (\mathcal{I} - x\mathcal{Q})(\cdot), \quad (3.36)$$

where \mathcal{I} is the identity super-operator and $\mathcal{Q} \equiv (x + \mathcal{R})^{-1}$. The correspondence between the s and x -ensembles is proved in the same manner as before. We act to the *left* on a matrix λ which we ask to be simultaneously a left-eigenmatrix of $\mathcal{T}_{x,\vec{s}}$ and $\mathcal{W}_{\vec{s}}$ with eigenvalues $e^{G(x,\vec{s})}$ and $\Theta(\vec{s})$, respectively (where the left action of a super-operator is that of the adjoint), we get

$$\begin{aligned} (\lambda)\mathcal{T}_{x,\vec{s}} &= \mathcal{Q}[(\lambda)\mathcal{W}_{\vec{s}}] + (\mathcal{I} - x\mathcal{Q})(\lambda) \\ e^{G(x,\vec{s})}\lambda &= \Theta(\vec{s})\mathcal{Q}(\lambda) + \lambda - x\mathcal{Q}(\lambda) \\ \Rightarrow (e^{G(x,\vec{s})} - 1)\lambda &= [\Theta(\vec{s}) - x]\mathcal{Q}(\lambda), \end{aligned} \quad (3.37)$$

which has as solutions (3.29),(3.30) as in the classical case.

3.4.1 Example: quantum two-level system

As a simple example consider the quantum two-level system of Fig. 3.2(c), corresponding to a system of two quantum levels $|0\rangle, |1\rangle$ coherently driven on resonance at Rabi frequency Ω and coupled to a zero temperature bath [31]. The operators that enter in the definition of \mathcal{W} are the Hamiltonian,

$$H = \Omega (|0\rangle\langle 1| + |1\rangle\langle 0|), \quad (3.38)$$

and the single jump operator ($N_L = 1$)

$$L_1 = \sqrt{\gamma}|0\rangle\langle 1|. \quad (3.39)$$

We count the number of quantum jumps due to this operator, and consider for simplicity the case where $\gamma = 4\Omega$ (a particular parameter point where the algebra

is simple). We can write the super-operator \mathcal{W}_s as a matrix [18],

$$\mathcal{W}_s = \Omega \begin{pmatrix} 0 & 4e^{-s} & i & -i \\ 0 & -4 & -i & i \\ i & -i & -2 & 0 \\ -i & i & 0 & -2 \end{pmatrix}, \quad (3.40)$$

which acts on the 2×2 density matrix ρ which we write as the vector,

$$\begin{pmatrix} \rho_{00} \\ \rho_{11} \\ \rho_{01} \\ \rho_{10} \end{pmatrix}. \quad (3.41)$$

The largest eigenvalue of (3.40) is [18],

$$\theta(s) = 2\Omega(e^{-s/3} - 1). \quad (3.42)$$

In this matrix form the operator $\mathcal{T}_{x,s}$ reads,

$$\mathcal{T}_{x,s} = \frac{4\Omega e^{-s}}{(x + 2\Omega)^3} \begin{pmatrix} 2\Omega^2 & (x + \Omega)^2 + \Omega^2 & -ix\Omega & ix\Omega \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

From its largest eigenvalue we obtain,

$$G(x, s) = -3 \log \left(1 + \frac{x}{2\Omega} \right) - s, \quad (3.43)$$

and it is easy to see that $G[\theta(s), x] = 0$, as expected from Eqs. (3.29),(3.30). In particular, $g(x) = G(x, 0)$ is the cumulant generating function for the trajectory length τ . Using the transform (3.6) we obtain the distribution of total time τ in the large K limit,

$$P_K(\tau) \approx \left(\frac{2\Omega\tau}{3K} \right)^{3K} e^{-2\Omega\tau + 3K}, \quad (3.44)$$

which is the expected result given that the probability of waiting a time t_w between jumps (except for the first one if the initial condition is not $|0\rangle$) is $p(t_w) = 4(\Omega t_w)^2 e^{-2\Omega t_w}$.

3.5 Outlook

The x -ensemble introduced in this chapter provides an alternative means of studying atypical dynamical behaviours. While the analytical form of the s -ensemble, following naturally from common master equation descriptions of stochastic dynamics, may appear more intuitive than the transfer matrix prescription of the x -ensemble, the x -ensemble is not without benefit. The x -ensemble's utilisation of fixed K is well suited to simulations of stochastic dynamics - and its use in such situations will be the focus of the remainder of this thesis. One of the key aspects of this chapter that we again emphasise here is the flexibility offered by utilising alternative dynamical ensembles. We argue that the ability to work with the s -ensemble for more familiar analytics, and switch to the x -ensemble for numerical work, converting results between the two using the ensemble correspondence illustrated in section 3.2¹ is of significant benefit when exploring the dynamical properties of new systems.

¹This conversion between ensembles will be treated more explicitly in Chapter 4.

4. FLUCTUATING OBSERVATION TIME ENSEMBLES: TRANSITION PATH SAMPLING APPLICATIONS

While the thermodynamics of trajectories described in the previous chapter can unlock the full array of dynamical behaviours present in a system, it is not always a straightforward task to access this information. The examples presented all possessed master operators (for the classical case) or super-operators (i.e. operators with a larger number of degrees of freedom than number of states - for the quantum case) that are directly diagonalizable, but for more complex systems this is often not the case, making it difficult to obtain the moment generating function. The problem is compounded since the dynamical fields, s, x are not connected to physically tunable parameters as their thermodynamic counterparts are. While techniques such as thermodynamic integration and umbrella sampling [92] can be adapted to numerically recreate these dynamical free energies, they are computationally costly, and often the full LD function is not required. Transition path sampling (TPS) [3] offers an alternative method of efficiently sampling the low-order moments. Originally developed to generate so-called reaction trajectories that occur sufficiently infrequently that their direct observation from simulated dynamics is infeasible, it is well suited to sample the rare trajectories at the heart

of the s - and x -ensemble. It has already had success in numerically describing the s -ensemble [13, 72, 73].

In this chapter we develop a TPS scheme for the x -ensemble based on the method of Crooks and Chandler in Ref. [32], focusing on its efficiency relative to the s -ensemble. We generically introduce the simulation methodology in section 4.1, before detailing the operation of the x -ensemble scheme in section 4.2. This scheme is then applied to a variety of stochastic systems, both classical and quantum in section 4.3, with an analysis of the performance of the algorithms. At the end of the chapter the efficiency of x -ensemble TPS is discussed, with a look at the future potential of improving performance further through “self-tuning” TPS algorithms.

4.1 Transition path sampling for trajectory ensembles

TPS operates by performing a biased random walk through the space of all trajectories, towards a predefined region of the trajectory space, and is in essence a Markov chain Monte Carlo (MCMC) [92] sampling algorithm. In its original application to chemical reactions, the predefined region would be trajectories that start in the reactant region of phase space, and end in the product region. For the s - and x -ensembles, the regions defined by the weights e^{-sK} , $e^{-x\tau}$ are instead the target. It is convenient to view TPS as a wrapper, operating as a layer around the simulated dynamics of the system. While any method of generating dynamics can be used with TPS, for reasons of efficiency we focus on continuous time Monte Carlo methods in both the quantum and the classical case, which we present below.

The only mechanism for efficiency over raw sampling of the dynamics is clever

methods of proposing subsequent steps in the random walk, reusing much of the old trajectory without artificially altering the underlying dynamics. Broadly speaking there are two methods of proposing alterations; local and non-local changes. Local changes alter some small part of the system at a specific point in time, for example giving a single particle a small change in momentum, without drastically altering the trajectory as a whole. They tend to have high acceptance rates, but because of the limited scope of the change, decorrelation times are high at each stage of the chain, and they struggle to reach distant target distributions in a timely manner. Non-local changes alter entire sections of the current trajectory, and tend to suffer the opposite problems - low acceptance rates. Examples of these methods include the “forwards/backwards shooting” and “forwards/backwards shifting” techniques previously used with the s -ensemble [13, 72].

4.1.1 Classical continuous time Monte Carlo

Like in the previous chapter we consider stochastic dynamics produced with continuous time Markov chains, both classical or quantum. For systems subject to classical or quantum master equations, Eqs. (2.27)-(2.28) or (2.30)-(2.31), respectively, the standard way to simulate stochastic trajectories is by means of continuous time Monte Carlo [26] (often called quantum jump Monte Carlo for the case of open quantum systems [30]). For the classical case such a scheme amounts to the following [26]:

- i) Given the current configuration of the system \mathcal{C} , compute the time t_w to the next transition by solving $P_{\mathcal{C}}(t_w) = r_1$, with $P_{\mathcal{C}}(t_w) \equiv e^{-t_w R(\mathcal{C})}$ being the “survival probability”, $R(\mathcal{C})$ the escape rate from \mathcal{C} , and $r_1 \in [0, 1]$ a uniformly distributed random number

- ii) Choose a transition $\mathcal{C} \rightarrow \mathcal{C}'$ by drawing a second random number r_2 , where the probability to make the jump $\mathcal{C} \rightarrow \mathcal{C}'$ is given by $W(\mathcal{C} \rightarrow \mathcal{C}')/R(\mathcal{C})$
- iii) Change the current configuration from \mathcal{C} to \mathcal{C}' and repeat from (i).

4.1.2 Quantum jump Monte Carlo

In the open quantum case, quantum jump Monte Carlo amounts to an “unraveling” of the quantum master equation, Eqs. (2.30)-(2.31), that leads to a stochastic evolution for the wave function [30,45,46]. Again (quantum) jumps occur stochastically, but in contrast to the classical case the wave function also evolves between jumps through the action of the effective Hamiltonian H_{eff} . That is, if the wave function at t is $|\psi(t)\rangle$, and no quantum jumps occur between t and $t + t_w$, then $|\psi(t + t_w)\rangle = e^{-it_w H_{\text{eff}}} |\psi(t)\rangle$. Furthermore, the survival probability for the waiting time until the next quantum jump is given by $P_\psi(t_w) \equiv \|e^{-ihH_{\text{eff}}t_w} |\psi(t)\rangle\|^2$. A stochastic trajectory can then be generated in the following way [30]:

- i) Given the (normalised) state $|\psi(t)\rangle$, compute the time to the next jump by solving $P_\psi(t_w) = r_1$, where $r_1 \in [0, 1]$ is a uniformly distributed random number
- ii) Evolve the wave function by t_w , $|\psi(t + t_w)\rangle = e^{-it_w H_{\text{eff}}} |\psi(t)\rangle$
- iii) Draw a second random number r_2 to select which quantum jump to perform, where the probability to make the quantum jump i is proportional to $\langle \psi(t + t_w) | L_i^\dagger L_i | \psi(t + t_w) \rangle$
- iv) Make the selected quantum jump, $|\psi(t + t_w)\rangle \rightarrow L_i |\psi(t + t_w)\rangle$, normalise the resulting state, and repeat from (i).

4.1.3 TPS and Metropolis-Hastings in trajectory space

TPS then operates as a wrapper around the dynamics in the following manner:

- i) Generate a starter trajectory, \mathbf{X} , with continuous time Monte Carlo, computing observables of interest, τ, \vec{M}
- ii) Propose a new trajectory, \mathbf{X}' with observables τ', \vec{M}'
- iii) Accept or reject \mathbf{X}' with Metropolis acceptance probability $P_{\text{accept}} = \min\{1, e^{-(x\Delta\tau)}\}$ for the x -ensemble, $P_{\text{accept}} = \min\{1, e^{-(s\Delta K)}\}$ for the s -ensemble, or indeed $P_{\text{accept}} = \min\{1, e^{-(x\Delta\tau + \vec{s}\cdot\Delta\vec{M})}\}$ for a compound ensemble with many s, x fields
- iv) Repeat steps ii and iii until \mathbf{X} is sufficiently equilibrated to the target distribution

The Metropolis acceptance criteria in step iii) above ensures trajectory sampling obeys detailed balance with respect to the target distribution. It can be derived, starting from a statement of detailed balance, as follows. From

$$P(\mathbf{X} \rightarrow \mathbf{X}')P(\mathbf{X}) = P(\mathbf{X}' \rightarrow \mathbf{X})P(\mathbf{X}') \quad (4.1)$$

the transition probability is split into a proposal distribution (the probability of proposing \mathbf{X}' given \mathbf{X}), $\pi(\mathbf{X} \rightarrow \mathbf{X}')$, and an acceptance distribution (the probability of accepting the proposed change), $A(\mathbf{X} \rightarrow \mathbf{X}')$. Rearranging we then have

$$\frac{A(\mathbf{X} \rightarrow \mathbf{X}')}{A(\mathbf{X}' \rightarrow \mathbf{X})} = \frac{P(\mathbf{X}')\pi(\mathbf{X}' \rightarrow \mathbf{X})}{P(\mathbf{X})\pi(\mathbf{X} \rightarrow \mathbf{X}')} \quad (4.2)$$

Choosing a proposal method such that

$$\frac{\pi(\mathbf{X}' \rightarrow \mathbf{X})}{\pi(\mathbf{X} \rightarrow \mathbf{X}')} = 1, \quad (4.3)$$

and given, for example, that in the x -ensemble we have $P(\mathbf{X}) = e^{-x\tau} P_K(\tau)$ we arrive at

$$\frac{A(\mathbf{X} \rightarrow \mathbf{X}')}{A(\mathbf{X}' \rightarrow \mathbf{X})} = \frac{e^{-x\tau'}}{e^{-x\tau}} \quad (4.4)$$

which for convenience we write

$$A(\mathbf{X} \rightarrow \mathbf{X}') = \min\{1, e^{-x\Delta\tau}\}. \quad (4.5)$$

It should be noted that the Metropolis-Hastings algorithm approaches the target distribution asymptotically, so it is a somewhat subjective matter to decide when a trajectory is “equilibrated.” In practice one waits until the acceptance ratio (the fraction of proposed changes that get accepted) is stable and fluctuations in the mean value of observables have ceased. Furthermore because old trajectories are adapted to propose new ones, subsequent steps in the Markov chain are correlated, and so to sample statistical information fairly one must wait multiple steps before drawing samples. This, again, is a subjective matter that depends on the specific method of proposal. The key to successful TPS, then, lies in balancing acceptance rates vs decorrelation times. It can be difficult to know what the “ideal” acceptance rate should be tuned to, though it has been shown to be 50% for a one-dimensional Gaussian target distribution, dropping to 23% for an infinite-dimensional Gaussian distribution [93]. As such acceptance ratios in the region of 25 – 40% are a good target for most systems.

4.1.4 Comments on s -ensemble efficiency

The s -ensemble TPS schemes have thus far relied on two kinds of moves, *shooting* and *shifting*. Shooting moves operate as follows

- i) Randomly select a time along the trajectory, t'

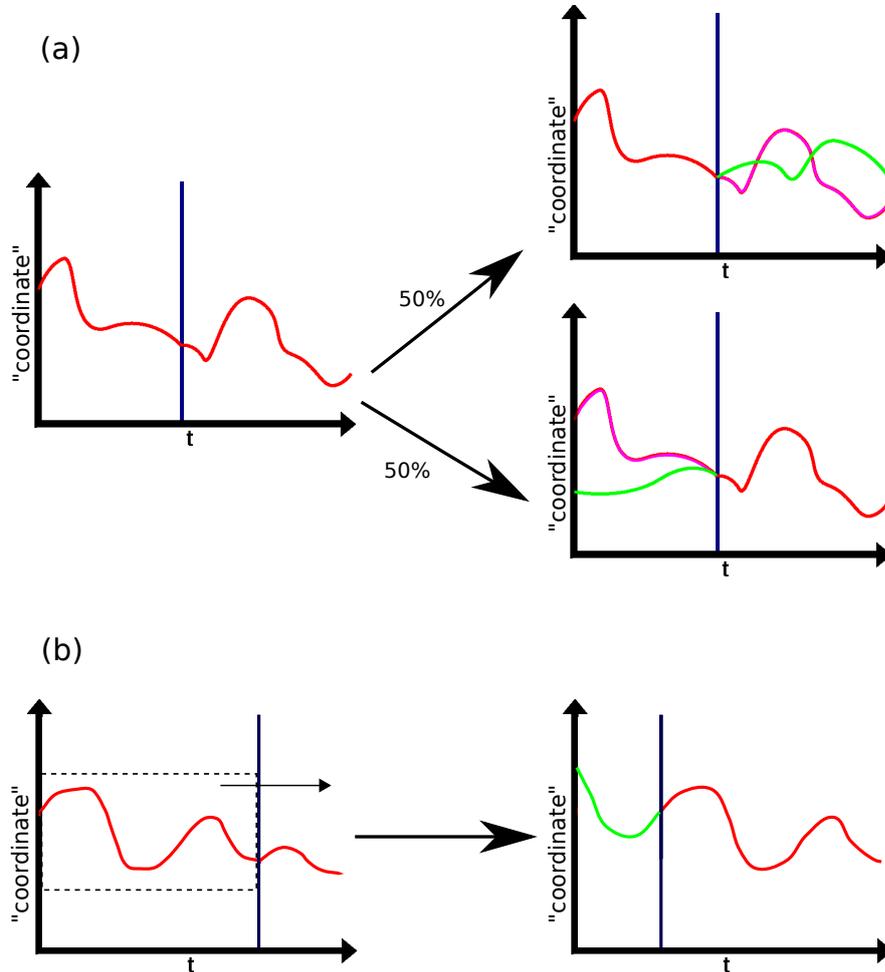


Figure 4.1: Illustration of typical s -ensemble TPS protocols with a generic evolution of “coordinates” of a system as a function of time. (a) Shooting moves: a point in time is selected (blue line) and based on a coin flip, either the past or future is discarded (pink) and replaced by a new segment (green). (b) Shifting moves: before generating a new segment the kept portion of a trajectory is first shifted forwards or backwards (backwards shift not shown). It should be noted that both protocols require the ability to produce time reversed dynamics.

- ii) Generate a random number $r \in [0, 1]$ to determine whether to replace the future or past from t'
- iii) If $r > \frac{1}{2}$ discard and regenerate the trajectory segment $t' \rightarrow t_f$
- iv) If $r < \frac{1}{2}$ discard and regenerate, using time reversed dynamics, the trajectory segment $t' \rightarrow t_i$

Shifting moves operate in the same manner, but add the additional step of shifting the section of trajectory that is not discarded backwards/forwards in time before replacing the missing segment (see Fig. 4.1 for an illustration). Since the sampling is done from an exponentially suppressed distribution, raw sampling of trajectories must scale as $O(e^s)$ and $O(e^\tau)$. Shooting moves are fundamentally limited by the need to accept alterations to the middle of the trajectory, and so scale similarly to raw sampling with $\tau/2$. This is a significant improvement, but ultimately they still suffer from exponential scaling. Shifting moves can function better, with multiple small shifts in the same direction working collaboratively to alter the middle portion of trajectories. However they suffer from bottlenecks as plateaus of “desired” dynamical behaviour, which are unlikely to be modified, form around “bad” regions in the trajectory. If one imagines an “optimal” shifting length, l that has an ideal acceptance rate, trajectories that are very much longer than this, $\tau \gg l$ are likely to suffer from such bottlenecks. s -ensemble schemes are thus a delicate balance act between the optimal trajectory length for efficient sampling, and achieving the “thermodynamic limit” of large τ . This is further confounded by variations in dynamical activity in a dynamical phase space, which alter l . Because of these issues, as well as the overall degree to which a trajectory is changed in a proposed move, s -ensemble schemes often have acceptance rates that are fractions of a percent in practical applications, well below the desired range. Both schemes are also critically dependent on the availability of time-reversed dynamics for efficient sampling, making their application to quantum

systems problematic. It is these problems we aim to address with the x -ensemble.

4.2 TPS with fluctuating observation time

The x -ensemble is well suited to describe systems whose dynamics is generated by such classical or quantum continuous time Monte Carlo. A trajectory with a fixed total number K of jumps is fully determined by K pairs of random numbers, $\{r_1, r_2\}^K$ - the first to determine how long until a configuration or quantum jump occurs, the second to determine which configuration change or quantum jump occurs.¹ A TPS scheme for the x -ensemble can then be devised based on the method of Ref. [32]. Since a trajectory is fully encoded in predetermined set of K pairs of random numbers, randomly selecting one of these pairs and modifying it is an efficient way to generate a new trajectory from an old one. The change in the extensive quantities of interest, $\Delta\tau$ and $\Delta\vec{M}$, where \vec{M} again denotes the counted observables, can be calculated and the new trajectory accepted or rejected based on the Metropolis acceptance criterion $P_{\text{accept}} = \min\{1, e^{-(x\Delta\tau + \vec{s}\cdot\Delta\vec{M})}\}$. After a sufficiently large number of accepted moves, a trajectory typified by the fields (x, \vec{s}) is generated. It is often impossible to implement local changes in the context of trajectories due to the exponential growth of differences upon making a change, however, this x -ensemble scheme is, in essence, a combination of local and non-local changes. In the context of the random numbers only a small, local change is made, but due to the highly non-linear mapping from random numbers to trajectory, and the exponential growth of differences, this can have non-local effects throughout the trajectory. Furthermore the overall extent of the change is small

¹This pair of random numbers could, of course, be combined into a single random number that determines both the waiting time and the configuration change, but it is far simpler just to generate a pair of random numbers.

compared to s -ensemble schemes (one can think of s -ensemble schemes as altering half of the random numbers associated with a trajectory on average), bolstering acceptance rates.

Due to the ensemble correspondence demonstrated in chapter 3, it is possible to convert x -ensemble TPS results, where the values of the fields (x, \vec{s}) are fixed, to their equivalent s -ensemble described by \vec{s} only. In order to do this it is necessary to find the curve $x^*(\vec{s})$, see Eq. (3.30). This curve passes through the origin of the space spanned by (x, \vec{s}) , since $G(0, 0) = 0$ trivially. It is then possible to move along the $x^*(\vec{s})$ curve by expanding (3.30) for small increments $\delta\vec{s}$ in \vec{s} which allow to relate the required change in δx to the current averages of τ and \vec{M} :

$$\delta x = \frac{\delta\vec{s} \cdot \langle \vec{M} \rangle_{x, \vec{s}}}{\langle \tau \rangle_{x, \vec{s}}} + o(\delta s^2). \quad (4.6)$$

In order to compute the s -ensemble we can therefore start at the point $(x = 0, \vec{s} = 0)$, i.e. unbiased dynamics, and progress towards \vec{s} by adjusting x according to (4.6), using an x -ensemble TPS algorithm for each value of (x, \vec{s}) . In this way we recover the properties of ensembles of trajectories with fixed total observation time from simulations of trajectories with fixed number of transitions or jumps.

4.3 Applications

We now apply the x -ensemble TPS scheme described above to a number of systems. We start with some simple open quantum systems because they are easy to simulate and the structure of their trajectory phase space can be analysed analytically, but inherent issues with time reversal make an s -ensemble exploration difficult. This provides an ideal test bed for the x -ensemble, with known results with which to compare the numerics generated. We then move onto a classical, kinetically constrained glass forming model, the East model, to test the performance

of x -ensemble TPS in an area the s -ensemble has already seen success.

4.3.1 $T \neq 0$ quantum two-level system

We again consider the quantum two-level system, but now for the case of $T \neq 0$, Fig. 3.2(d) [30]. In most open quantum systems, generating time-reversed dynamics is non-trivial, whether due to the combination of coherent dynamics interspersed with stochastic jumps in state, or the possibility of an unpaired Lindblad term for which there is no reverse process (that is a Lindblad term describing a transition between states $|i\rangle \rightarrow |k\rangle$ with no corresponding Lindblad term for a transition $|k\rangle \rightarrow |i\rangle$ – the population of state $|i\rangle$ might instead be brought about by the coherent evolution). In such systems conventional TPS is limited to forward shooting only, limiting its efficiency. We use the simple case of the $T \neq 0$ two-level problem to illustrate how the x -ensemble TPS can efficiently sample such systems.

Compared to the $T = 0$ case of Sect. 3.4.1, when $T \neq 0$ there is a second jump operator, $L_2 \equiv \sqrt{\lambda}|1\rangle\langle 0|$, associated to the absorption of a quanta from the bath, which leads to a projection to the $|1\rangle$ state at a rate λ (where the ratio λ/γ is determined by the temperature T). Let us say we are interested in the statistics of the number of jumps K_1 due to L_1 (3.39). Using the x -ensemble TPS scheme we can compute the average total time, $\langle \tau \rangle(x, s_1)$, and the average number of 1-jumps, $\langle K_1 \rangle(x, s_1)$, for fixed total jumps K (i.e. due to *both* L_1 and L_2) as a function of the fields x and s_1 . Figure 4.2(a) shows the ratio $\langle \tau \rangle / \langle K_1 \rangle$ as a function of x , along the curve $x^*(s_1)$ (3.29)-(3.30), comparing the TPS simulation to the exact result. Figure 4.2(b) shows the activity associated to K_1 in the s -ensemble, $\langle K_1 \rangle_s = -\theta'(s_1)\tau$, both the exact result from diagonalisation of \mathcal{W}_{s_1} , and the numerical estimation from the conversion from the x -ensemble TPS simulation, $\langle K_1 \rangle_s = \langle K_1 \rangle_{x^*(s_1)} \langle \tau \rangle_{x^*(s_1)}$. A rudimentary efficiency comparison of the x -ensemble

TPS scheme and a forwards/backwards shooting s -ensemble TPS scheme is shown for zero-temperature in Figure 4.2(c) and (d). The x -ensemble scheme scales linearly with trajectory length, and as $O(x^2)$, while the s -ensemble scheme scales exponentially with both trajectory length, and s .

4.3.2 Micromaser

Next we consider a micromaser [94], a single-mode cavity coupled to a finite-temperature bath, and pumped by sending excited two-level atoms through the cavity at a constant rate. There are four quantum jump operators associated with the system, two for the atom-cavity interaction $L_1 = \sqrt{r} \frac{\sin(\lambda\sqrt{aa^\dagger})}{\sqrt{aa^\dagger}} a$, $L_2 = \sqrt{r} \cos(\lambda\sqrt{aa^\dagger})$, and two for the cavity-bath interaction, $L_3 = \sqrt{\kappa} a$, $L_4 = \sqrt{\gamma} a^\dagger$. Here the a, a^\dagger are the raising/lowering operators of the cavity mode, r is the atom beam rate, and λ encodes the time of flight of atoms through the cavity. For simplicity, the system can be parameterised by a single ‘‘pump parameter’’ $\alpha = \lambda\sqrt{r/(\kappa - \gamma)}$. It can be shown that if the system is initiated in a state with diagonal density matrix, the system stays in a diagonal state. It should be noted that while this effectively reduces the problem to a classical one, generating time-reversed dynamics is still problematic as there is no clear reverse process for action under L_1 .

The micromaser has a rich trajectory phase diagram with many dynamical phase transitions between states characterised by different photon occupations, $\langle N \rangle$, of the cavity [95]. This complex dynamical phase structure is made manifest by coupling to the number of events M_1 under the action of L_1 , i.e. measurements on the atoms leaving the cavity where the atom is in its ground state. Below we will reproduce this behaviour by converting results from the x -ensemble, described by fields (x, s_1) , to the equivalent s -ensemble described by s_1 .

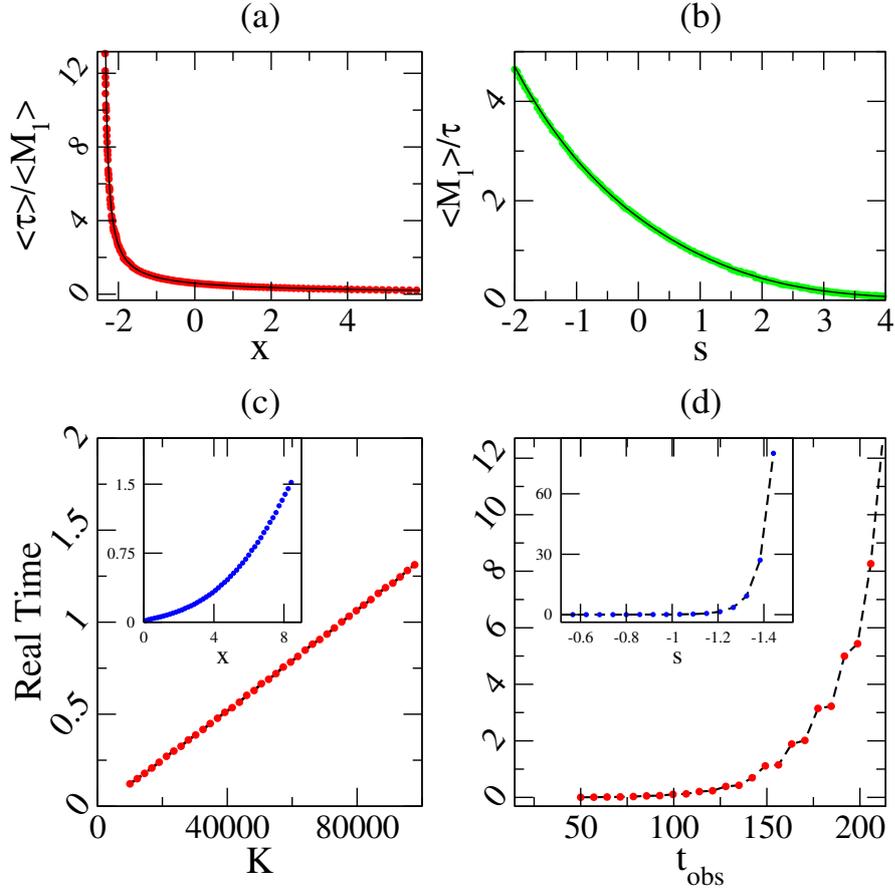


Figure 4.2: $T \neq 0$ quantum two-level system. (a) x -ensemble: $\langle \tau \rangle / \langle K_1 \rangle$ as a function of x , along the curve $x^*(s_1)$ in the x -ensemble; symbols are from TPS simulations and the curve from exact diagonalisation of \mathcal{T}_{x,s_1} . (b) s -ensemble: the symbols are the s -ensemble expectation values as a function of s , as obtained from the x -ensemble TPS simulation, and the curve those from the exact diagonalisation of \mathcal{W}_{s_1} . The parameters here are $\gamma = 6\Omega$ and $\lambda = 2\Omega$. (c) and (d) Efficiency comparison of (c) x -ensemble TPS to (d) s -ensemble forwards/backwards shooting for parameters $\gamma = 4\Omega$ and $\lambda = 0$ (zero-temperature). The simplicity of the system represents the most efficient possible scenario for both approaches. Real time taken to converge to within 2% of the analytical rates showing the scaling with trajectory length. The typical trajectory lengths plotted are 400 times longer in (c) than (d). (Inset) As before but now the scaling is with x/s . The range of x and s used are equivalent.

Transitions from states with high to low photon occupation occur far more readily than the reverse. This can present a problem when trying to numerically recreate a transition away from $x = 0$ from a state with low photon occupation to a state with a higher photon occupation. This is similar to the problem of a thermal system becoming stuck in a potential well that is not the global minimum. While there are many techniques to deal with such an issue, such as replica-exchange, there is a more novel approach in the micromaser. Since suitable time-reversed dynamics are not available, it is simple to fix the initial state of the trajectory to a large photon occupation, N , and set the trajectory length (defined by K) large enough that the initial conditions do not have significant impact on the latter stages of the trajectory. Since the micromaser frequently returns to the same state under a TPS algorithm, the x -ensemble TPS approach is not punished by large trajectory lengths in the same way the s -ensemble is (by returning to a previous state, in effect a constant amount of computation is needed to propose a new trajectory under the x -ensemble, where an amount of time of $O(\tau)$ is needed in the s -ensemble, since on average half the trajectory is always recomputed; see next section).

Figure 4.3(c) provides a comparison between the efficiency of the x -ensemble to that of the s -ensemble. It shows the same quantity as in (a) in the range $s \in [0, 0.1]$ and $\alpha \in [\pi, 4\pi]$ but generated using an s -ensemble forward shooting algorithm. It took about 100 times more computational effort using the s -ensemble TPS to generate the data in (c), over a fraction of the parameter range of (a), and clearly the convergence to the exact result is still poor, cf. panel (b). Furthermore, no useful data could be generated in that time using an s -ensemble TPS approach for $s_1 < 0$. This illustrates the efficiency of the x -ensemble TPS scheme as compared to the standard fixed observation time TPS.

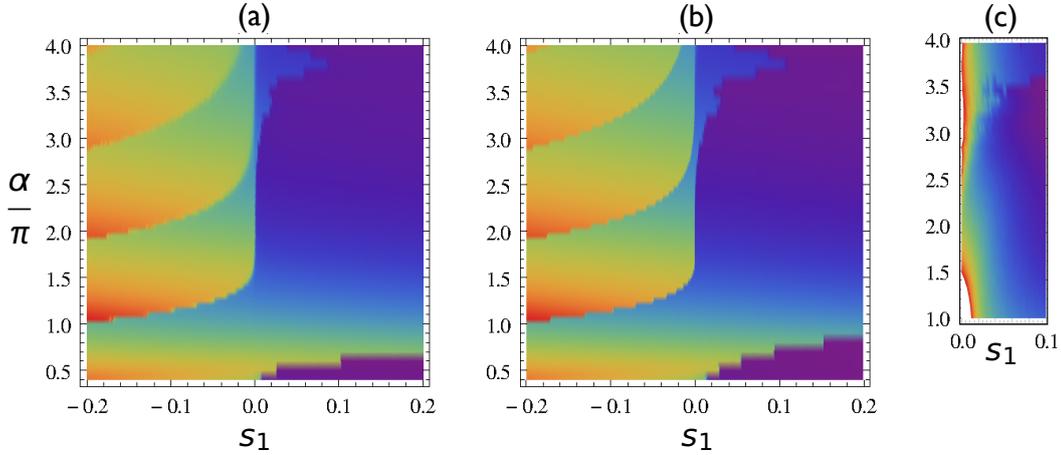


Figure 4.3: Micromaser: average number of L_1 jumps, $\langle M_1 \rangle$, as a function of s_1 and α/π . (a) Results from x -ensemble TPS of trajectories of fluctuating observation time τ , transformed to the s -ensemble. (b) Exact numerical diagonalisation of \mathcal{W}_{s_1} . Both plots at the same resolution. (c) Results from s -ensemble TPS. This small segment of the phase diagram took approximately 100 times the computation time of the x -ensemble plot.

4.3.3 East facilitated spin model of glasses

The s -ensemble method was first applied [12, 15] to uncover the dynamical phase structure of kinetically constrained models [28] of glassy systems. Such systems are thermodynamically simple but dynamically complex, and this can be traced back to a singularity in ensembles of trajectories, between “active” (equilibrium) and “inactive” (non-equilibrium) dynamical phases. These two phases are stabilised by negative or positive s , respectively, with a first-order transition between them at $s = 0$ (in the limit of large system size, and for the case where the kinetic constraints are “hard”, i.e. cannot be violated). A more in depth discussion on the nature of glasses can be found in chapter 5.

For a demonstration of the functionality of the x -ensemble approach in a many-body glassy systems, we consider the East model in one-dimension [28], defined

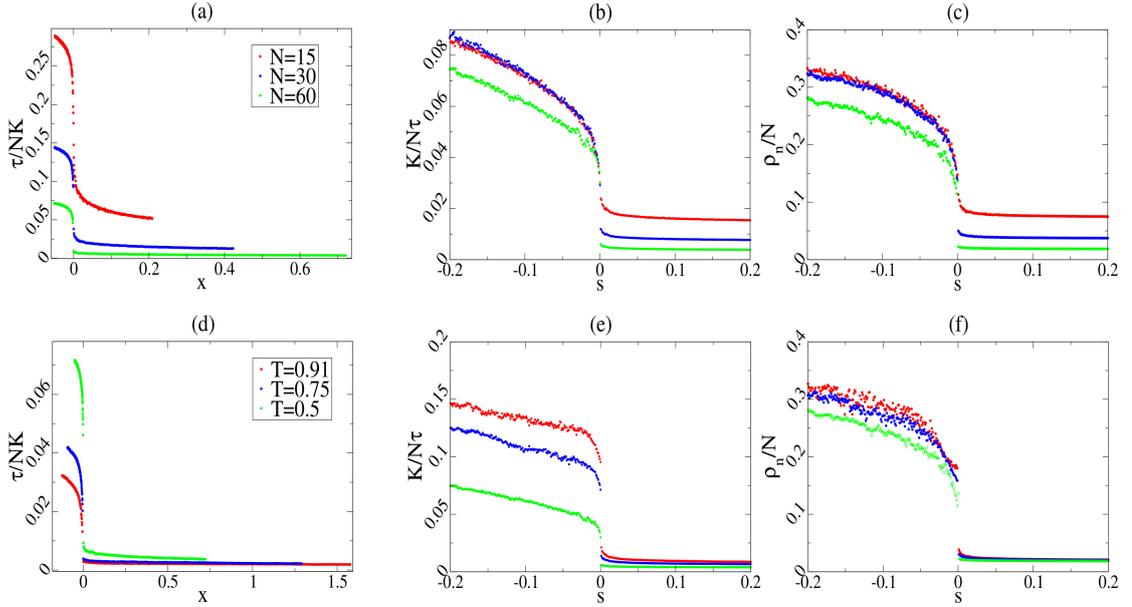


Figure 4.4: (a) The raw x -dependent mean jump waiting time per site τ/NK explicitly converted to: (b) The s -dependent mean activity per site, $K/N\tau$ and (c) The s -dependent mean excitation density per site, ρ_n/N . Both plots are at temperature $T = 0.5$, for chain lengths $N = 15, 30, 60$, and the simulations were made using x -ensemble TPS. (d/e/f) Same as before but now for $N = 60$ and for temperatures $T = 0.91, 0.75, 0.5$. The change in relationship between x and s with a change in both chain length and temperature is demonstrated in the conversion from (a) \rightarrow (b) and from (d) \rightarrow (e) respectively

on a lattice of N sites with a binary variable at each site, $n_i = 0, 1$ ($i = 1, \dots, N$), and with energy function $E = \sum_i n_i$. A transition at site i , from $0 \rightarrow 1$ with rate c , and from $1 \rightarrow 0$ with rate $(1 - c)$, can occur only if the neighbouring site to the left is excited, $n_{i-1} = 1$. This latter condition on the rates is the kinetic constraint. The transition rates are temperature dependent with $c = (1 + e^{1/T})^{-1}$. Note that with these definitions detailed balance is obeyed with respect to the Boltzmann equilibrium with energy E at temperature T , and since E is non-interacting, despite the strong dynamical interactions, the system evolves towards a non-interacting equilibrium state.

The x -ensemble TPS approach is able to efficiently recover the results obtained through s -ensemble TPS. Figure 4.4 shows that we recover the active-inactive crossover, which is seen to get sharper with increasing lattice size, and is present at all temperatures, as expected [12]. Again we have converted from an x -ensemble described by field x to an s -ensemble described by field s .

Finally we perform another efficiency comparison shown in Figure 4.5. An x -ensemble TPS scheme, and a forwards/backwards shooting s -ensemble TPS scheme were run from an initially infinite temperature system, and we measure the real time taken for the mean excitation density per site, ρ_n/N to fall below some critical value ρ_{crit} for $s = 0.2$. This corresponds to taking the system from a highly active state ($s < 0$) to the inactive state ($s > 0$). The lengths of the trajectories were set to be over a similar range for a trajectory in the inactive state. The linear dependence of the x -ensemble scheme with system size, K continues to hold in this system. The s -ensemble scheme on the other hand is significantly punished by trajectories that are too short or too long. Furthermore, due to a large overall difference in efficiency, a higher value of $\rho_{crit} = 0.05$ had to be used for the s -ensemble scheme, compared to $\rho_{crit} = 0.02$ for the x -ensemble.

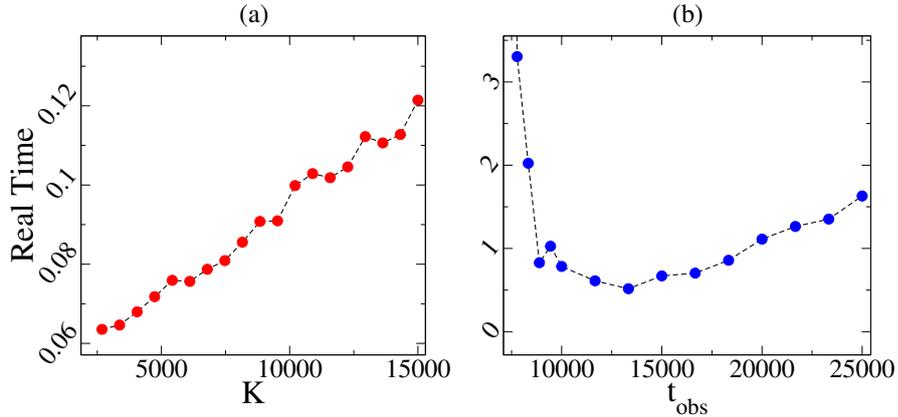


Figure 4.5: Efficiency comparison of the (a) x -ensemble TPS scheme, and (b) a forwards/backwards shooting s -ensemble TPS scheme. Both plots are for $T=0.5$, $N=60$ and show the mean real time taken for the mean excitation density per site, ρ_n/N , of an initially infinite-temperature system to fall below some critical value ρ_{crit} for TPS with $s = 0.2$. Due to overall difference in efficiency we used $\rho_{crit} = 0.02$ for the x -ensemble in (a) and $\rho_{crit} = 0.05$ for the s -ensemble in (b). The trajectory lengths are over similar ranges for trajectories in the inactive phase with $s > 0$

4.4 Efficiency of the x -ensemble scheme

While we have focused on systems simulated with continuous-time Monte Carlo (both classical and quantum), it should be noted that x -ensemble TPS is applicable to any system whose trajectories can be fully described by sets of random numbers. The efficiency of the approach will be dependent on the nature of the system under study, but in general there are several factors that should contribute to a greater efficiency in the x -ensemble than the equivalent s -ensemble TPS schemes. These factors are the following:

- (i) A new trajectory \mathbf{Y}' is generated by altering the i -th random number that defines the current trajectory \mathbf{Y} leaving the other random numbers, before and

after, unchanged. If at any later stage $j > i$ the trajectory visits the same state as in \mathbf{Y} then no further computation is required to generate \mathbf{Y}' . This drastically reduces the computation overhead required to propose a new trajectory. While this recurrence is unlikely in large systems, it will be the main source of efficiency in few-body problems, and perhaps also in systems with large state spaces but with limited dynamical pathways (cf. the micromaser).

(ii) Smoothness of the acceptance criteria. s -ensemble TPS has a Metropolis acceptance probability of $P_{\text{accept}} = \min(1, e^{-s\Delta K})$. Since the number of events is necessarily an integer, a change in activity is only seen if K changes by at least 1. This can lead to low acceptance probabilities, particularly for large s . The x -ensemble, on the other hand, has a metropolis acceptance probability of $P_{\text{accept}} = \min(1, e^{-x\Delta\tau})$ [or more generally for multiple observables, $P_{\text{accept}} = \min(1, e^{-x\Delta\tau - \vec{s}\cdot\Delta\vec{M}})$]. Since the trajectory length τ is continuous small incremental improvements, towards a trajectory typical of the desired value of (x, \vec{s}) , are more likely to be accepted.

(iii) Trajectories are not altered as drastically. Working in such a manner where trajectories are fully described by a sequence of sets of random numbers, an s -ensemble forwards-backwards shooting approach is equivalent to a replacement, on average, of half of the random numbers. While the relationship between the random numbers used to describe a trajectory and the activity of that trajectory is obviously highly non-trivial for a complex system, it is nevertheless to be expected that smaller changes to the random numbers defining the trajectories will have a smaller impact on the activity - although this is unlikely to be a factor in highly chaotic systems. Furthermore, while the mapping from random numbers to trajectories may be non-linear, it is still systematic, for example small random numbers might correspond to long waiting times in general. By modifying one pair of random numbers only, the x -ensemble approach makes smaller incremental

improvements facilitating faster convergence. This is particularly important as s -ensemble schemes often only manage acceptance rates that are fractions of one percent.

(iv) Alterations can be made at any point along the trajectory with comparative ease. This avoids the bottlenecks in s -ensemble schemes where the middle portions of a trajectory are difficult to alter.

We expect the x -ensemble to be particularly useful in systems that display an active-inactive dynamical phase coexistence, such as glassy systems [12,13]. Using a fixed-time trajectory, as in the s -ensemble, the length of the trajectory needs to be set long enough that any interesting behaviour in the inactive phase is captured. But this can lead to an unnecessarily large amount of information on the active phase being recorded. This is the primary reason for the poor performance of the s -ensemble in the micromaser (see Fig. 4.3) - the scheme fails whenever the system encounters a dynamical phase boundary because of the difference in activity in either side. Optimising τ for the inactive phase leads to poor sampling in the active phase and vice versa. It is also the reason for the computational cost when τ is too small seen in the East model (see Fig. 4.5). In contrast, in the x -ensemble little computation time is wasted on the active phase: by fixing total event numbers, the same quality of statistics is generated for both active and inactive dynamics. It should also be noted that at worst, s -ensemble shooting/shifting schemes can easily be adapted to the x -ensemble. This flexibility is of great importance, as in general with TPS, there is no a priori way of knowing for certain which schemes will be effective.

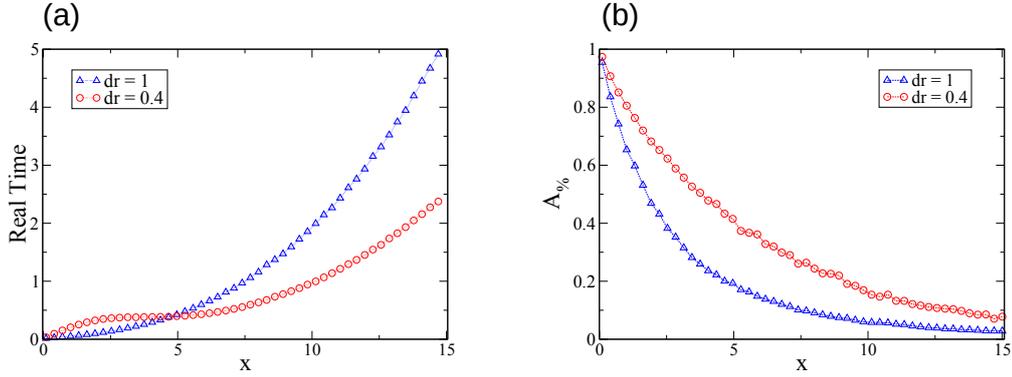


Figure 4.6: Efficiency comparison of the “small r ” scheme (red circles) with the original scheme (blue triangles) in the quantum two-level system with parameters $\gamma = 4\Omega$ and $\lambda = 0$ ($T=0$), $K = 10000$. (a) Mean real time taken to converge to within 0.5% of the analytical rates across a range of x values with their corresponding acceptance rates (b).

4.5 Beyond the standard TPS algorithm: a self-tuning approach

In this section we discuss the future outlook of work in progress on x -ensemble TPS schemes. While the x -ensemble has been shown to be more efficient than the s -ensemble in a wide variety of systems, its key advantage is its flexibility which can be exploited further. While in the previous section we utilised a scheme of replacing a single set of random numbers, it is not necessary to be limited in this manner. Any modifications that preserve the symmetry of the proposed move distributions (see Eq. 4.3) can be used. This includes alterations that change a set of random numbers by a small amount instead of replacing them completely, provided the flatness of the random number distribution is preserved. One possibility is to set a maximum increment, dr , and modify a random set of random numbers, $\{r_1, r_2\} \rightarrow \{r_1 + dr_1, r_2 + dr_2\}$ with $dr_1, dr_2 \in [-dr, dr]$. In order

to preserve the distribution of random numbers, we impose periodic boundary conditions on $[0, 1]$. It is straightforward to see that any proposed move occurs with the same probability as its reverse, i.e. $P(\{r_1, r_2\} \rightarrow \{r_3, r_4\}) = P(\{r_3, r_4\} \rightarrow \{r_1, r_2\})$ and Eq. 4.3 is satisfied.

The primary advantage of this “small r ” method is that it allows for more precise control of acceptance rates, something that has been lacking in TPS applications to the thermodynamics of trajectories. We apply this method, again to the $T = 0$ quantum two-level system of Fig. 3.2 (c). Figure 4.6 (a) shows the benefits obtained. We compare the performance of the original method ($dr = 1$) to the new one with a value of $dr = 0.4$. Initially, for small x where acceptance rates are high, there is some linear overhead where further bolstering of the acceptance rates is a hindrance. However as we increase x the small r method begins to outperform the original, with a 50% reduction in computation time for the $x > 10$ region. The corresponding acceptance ratios (defined as the ratio of accepted moves to total proposed moves $A\% = N_{accepted}/N_{proposed}$) in Figure 4.6 (b) illustrate this point. A crude analysis of the peak relative performance of the two methods suggests the optimum acceptance ratio for this system occurs at approximately $A\% = 0.3$ - in line with other applications of the Metropolis-Hastings algorithm [93].

There is still further flexibility that can be explored with the x -ensemble. Situations where acceptance rates of the standard scheme are higher than optimal are not overly common in practical situations, particularly not in difficult to sample regions of a trajectory phase space. The $x < 5$ region of Fig. 4.6, for example, is a manifestation of this. Improvements could potentially be made to lower the acceptance rates in such a case, and speed up sampling, but there is little need given the overall rapidity with which sampling occurs. However, often in the initial “burn in” phase of a complex system, particularly when attempting to sample close to a dynamical phase transition, there is a need for large overall changes to a trajectory,

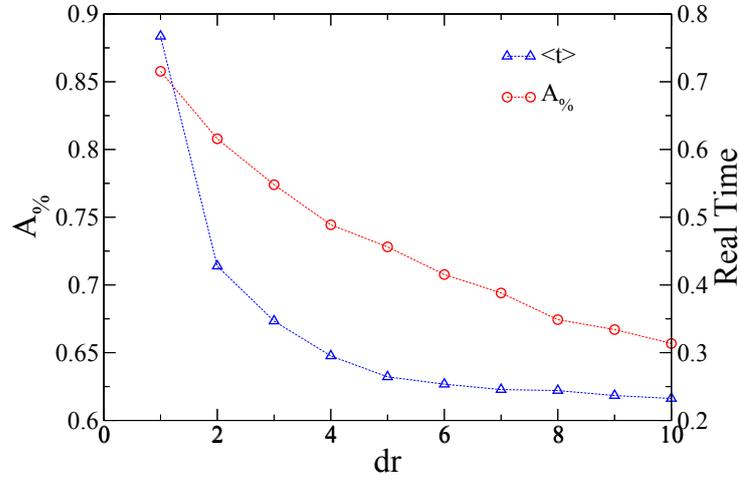


Figure 4.7: Efficiency of the initial relaxation of an East model, mean real time taken (blue triangles) for the excitation density to fall below $\rho_{crit} = 0.012$ with $x = -0.04$ with $N = 30$, $T = 0.5$, and $K = 7500$. The values of dr are integer i.e. changes to whole sets of random numbers. The corresponding acceptance rates (red circles) are shown, and the relationship between acceptance rate and time taken is non-linear (as expected).

to quickly decorrelate from an initial chosen trajectory. Such a situation occurs for example in the East model. If we want to fairly sample the inactive phase, it is often best to start from an equilibrium trajectory with a random spin distribution (i.e. infinite temperature) as its initial configuration. If one were to start from an artificially constructed inactive state, the slow dynamics and large escape barrier from any given inactive state results in a long decorrelation time from the initial trajectory (one can think of this as being stuck in a very deep potential well in a thermal problem). This is the same motivation behind techniques such as simulated annealing [26] starting from an effective “high-temperature” state before relaxing to fairly sample from the entirety of a distribution. This is particularly true for systems with many deep, nearly degenerate free energy (or $g(x)$ in our

analogous situation) minima where crossing from one minima to another at low temperature is a difficult endeavour. To illustrate this we implement a scheme in the East model with $dr > 1$ (signifying changes to more than one set of random numbers at a time) in Figure 4.7. The real time taken for the initial relaxation to the inactive phase, measured by the excitation density falling below a value of $\rho_{crit} = 0.012$ for an $N = 30, T = 0.5, K = 7500$ system, is plotted for a range of (integer) values of dr , along with the corresponding acceptance rates. For ease of implementation we modify sequential sets of random numbers. This then corresponds to effectively a small shooting move at a randomly selected point in the trajectory. With a value of $dr = 10$ we find there is more than a 66% reduction in the time taken to relax the East model trajectories to the inactive phase.

Combining the methods above sets the ground work for current ongoing work on a self-tuning x -ensemble TPS scheme. As we have seen $A_{\%}$ is the critical quantity for the efficiency of Metropolis-Hastings algorithms, and the x -ensemble allows for very strong influence on acceptance rates compared to the s -ensemble. It is possible to modify the protocol used in response to the current measured acceptance ratio as one proceeds with TPS. Each time we modify the protocol an interface is created in the Markov chain across which detailed balance is broken, which is an unfortunate consequence. However, if only a small change is made to the protocol, there is little impact on the sampled distribution, and the Markov chain relaxes almost instantaneously back to the “equilibrium”² of the target distribution. As long as changes to the protocol are small, and made infrequently, we can do so without biasing the sampling. This is a precise analogue of the method of simulated annealing, which also uses small, infrequent breakages of detailed balance

²Again one needs to be careful remembering we are dealing with a Markov chain of trajectories, and a protocol moving through the space of trajectories, rather than configurations - this “equilibrium” is not a thermal equilibrium, but rather with the regards to the stationary distribution of trajectories defined by x .

(with respect to configurations in that case) to optimize sampling of thermal phase spaces.

This self-tuning algorithm would then proceed as follows:

- i) Define an “ideal” target acceptance ratio, $A_{\%}^{(\text{tar})}$ (typically around 25% for a many body system)
- ii) Set an initial protocol, typically with large changes to the system initially (e.g. $dr \gg 1$)
- iii) Perform a set (large) number of TPS moves and measure the current acceptance ratio $A_{\%}^{(\text{curr})}$
- iv) Compare $A_{\%}^{(\text{tar})}$ and $A_{\%}^{(\text{curr})}$, and adjust the protocol accordingly i.e. if $A_{\%}^{(\text{curr})} > A_{\%}^{(\text{tar})}$ increase dr by a small amount, else decrease dr by a small amount.

One must naturally enforce an upper and lower bound on dr to prevent failure of the algorithm on edge cases. Such an approach would greatly reduce the labour intensity of a TPS study, as one does not need to worry about the nuances of the most efficient TPS implementation, the algorithm would self-tune towards optimum sampling. This is particularly useful for studying a broad range of x , and particularly, crossing dynamical phase boundaries, where the optimum protocol in each phase may be drastically different. It will also potentially help when the dynamical phase space is “rugged,” where the system can struggle to find the global minimum. One can start from a large dr , $x = 0$ protocol and slowly tune x towards the desired target, correspondingly adjusting dr to facilitate the location of the global minimum. This is exactly the situation simulated annealing was developed to address in thermal systems. We have implemented such a scheme, again for a simple test case of the $T = 0$ quantum two-level system in Fig. 4.8.

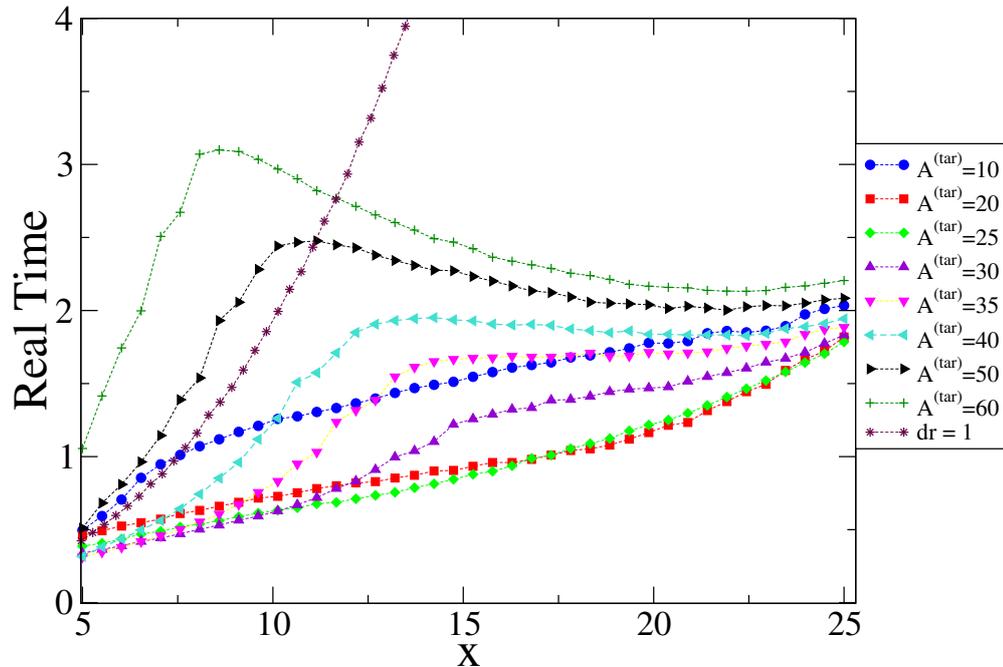


Figure 4.8: Implementation of a self-tuning x -ensemble TPS algorithm in a $T = 0$ quantum two-level system. The real time taken to converge within 0.5% of the analytical rates across a range of values for a range of target acceptance rates. Also included is the original method with $dr = 1$ (maroon asterisks). For large x all schemes converge towards the same time taken as they consistently hit the lower dr bound of 0.03. The plateaus, and subsequent improvements in time taken, for the high acceptance rate targets (40 – 60%, teal triangles, black triangles and green crosses) are caused by encountering this lower dr bound much sooner, and the inability to maintain unnecessarily high acceptance rates.

The specific implementation measures the local acceptance rate over 10^4 TPS moves and slowly adjusts the value of dr accordingly. dr has upper bound of 1 (complete replacement) and a lower bound of 0.03. The 'ideal' acceptance rate appears to be somewhere in the region of 20–25%, suggesting a surprising amount of complexity in Metropolis sampling even a seemingly simple dynamical system (we again reiterate here that ideal acceptance rates have been shown to be 50% for sampling a one-dimensional Gaussian distribution, falling to 23% for an infinite-dimensional Gaussian [93]). We are in the process of refining this approach and applying it to a variety of systems to fully establish its usefulness.

Beyond this there are two further points we wish to discuss. Firstly, it is also possible to incorporate the shooting and shifting methods of the s -ensemble in cases where time-reversed dynamics are available, but such moves are fundamentally less limited in the x -ensemble. One can imagine exotic schemes where, for example, a set of random numbers is cut from the trajectory, a small segment of the trajectory (but not the entirety of the remainder) shifted into its place, and a new random number inserted at the end of the shifted segment. While such a scheme may or may not be useful, it serves to illustrate the flexibility offered by the x -ensemble. At a worst case scenario one can always reduce to an s -ensemble scheme, so we argue well implemented x -ensemble TPS scheme should never fail to perform at least as efficiently as a well implemented s -ensemble scheme. The second point is that, in many places in this work, we have exploited a fixed initial configuration \mathcal{C}_i of a trajectory to aid in sampling of a desired phenomenon. If the trajectory is long enough to (effectively) realise the LD limit of the system, the initial configuration is irrelevant to the statistics of the sampled distribution. However, when not using time-reversed dynamics, one could consider the initial configuration to simply be the first set of random numbers in a trajectory, treated like any other, with attempted changes being made by modifying this set of random numbers. These changes could be large (replacing the initial configuration

completely) or small (flipping a single spin in the initial configuration for example) without breaking detailed balance. This can potentially be extremely useful when sampling from a bistable distribution, where the typical configurations throughout a typical trajectory from each peak are drastically different. This is again, not a possibility that is practical with any s -ensemble scheme. As such, we again emphasise the degree of flexibility offered by the x -ensemble is of great importance to the efficiency with which TPS studies can be performed.

4.6 Outlook

The results of this chapter illustrate the usefulness of the x -ensemble in a numerical setting. It is perhaps natural that an ensemble of fixed event numbers would be well suited to computer simulation, with predefined memory requirements and simple methods of trajectory alteration. This has allowed the application of the x -ensemble in situations where the s -ensemble struggles - notably where the dynamics of systems do not obey detailed balance. Even in scenarios where the s -ensemble has been successfully utilised, the x -ensemble demonstrates improved performance. This performance is further enhanced by the greater flexibility in manipulation of trajectories allowed by the x -ensemble, allowing more precise control over the acceptance rates of TPS schemes. We are hopeful that these developments will allow for the exploration of the dynamical properties of systems that have so far been beyond the reach of traditional methods.

5. DYNAMICAL AND STATIC TRANSITIONS IN PLAQUETTE MODELS OF GLASSES

In this chapter, we investigate plaquette spin models of glasses [35], for which both overlap-fluctuations and dynamical activity-fluctuations can be analysed, by a combination of analytical and computational methods. We concentrate on two models, whose relaxation behaviour is similar to that of the facilitated East model [27, 28, 96] – their relaxation times increase faster than an Arrhenius law at low temperatures, but the equilibrium relaxation time is finite at all positive temperatures, diverging only as $T \rightarrow 0$. We present evidence that these models support *both* dynamic and thermodynamic phase transitions. In the thermodynamic case we consider a coupling between two *annealed* replicas, and transitions *occur only for non-zero (positive) values of the coupling*.

We argue that these results provide a connection between the (apparently quite different) ‘thermodynamic’ and ‘dynamic’ theories of the glass transition. This connection is built on the idea of metastability, which is intrinsically connected to glassy behaviour. The formation of a metastable state in a finite dimensional system requires that small perturbations in that state do not grow: the system

prefers to relax back into the metastable state. This stability to small perturbations may be described in terms of an interfacial cost that acts to penalise local perturbations. Different theories ascribe different origins to these interfacial costs, which might be either static or dynamic, depending on the system of interest, and the kinds of fluctuation being considered. However, the existence of these interfacial costs seems quite generic, and may be useful for rationalising different kinds of phase transition in these systems.

The main results of this chapter are as follows. We analyse the triangular plaquette model (TPM) in two spatial dimensions [33,34], and a three-dimensional variant of this model, which we refer to as the square pyramid model (SPyM). In section 5.2, we show numerical evidence that the TPM in a magnetic field supports a phase transition in the $2d$ -Ising universality class. There is a previously known method of mapping two (annealed) coupled replicas of these systems to a single system in a magnetic field [36,97,98], which is shown in Appendix A. It then follows that the coupled replicas of these systems also support a similar phase transition. In section 5.3, we show that the TPM also supports dynamical “space-time” phase transitions, similar to those in [12,13,99]. In section 5.4, we introduce the SPyM, and show evidence that it supports phase transitions in the coupled replica setting, and dynamical space-time phase transitions. Finally in section 5.5, we discuss the relationships between the thermodynamic and dynamical phase transitions that we have found, and we consider the consequences of these results for theories of the glass transition.

5.1 Plaquette models, and coupled replicas

5.1.1 Models

We consider plaquette spin models defined in terms of classical Ising spins on regular lattices, with energy functions of the form

$$E_J(\sigma) \equiv -\frac{J}{2} \sum_{\mu} \sigma_{i_{\mu}} \sigma_{j_{\mu}} \cdots \sigma_{k_{\mu}}, \quad (5.1)$$

where $\sigma_i = \pm 1$ with i indicating a lattice site ($i = 1, \dots, N$), and where the interactions are in terms of products of spins $\sigma_{i_{\mu}} \sigma_{j_{\mu}} \cdots \sigma_{k_{\mu}}$ around the *plaquettes* μ of the lattice. See [33–35] for a more general overview of the relevant properties of these systems. On a square lattice, one labels each square plaquette with an index μ , and $\{\sigma_{i_{\mu}}, \sigma_{j_{\mu}}, \dots, \sigma_{k_{\mu}}\}$ is the set of four spins on the vertices of plaquette μ . This construction is easily generalised to higher dimensions: for a cubic lattice and cubic “plaquettes”, each term in the energy would involve eight spins. This motivates us to define plaquette variables $\tau_{\mu} = \sigma_{i_{\mu}} \sigma_{j_{\mu}} \cdots \sigma_{k_{\mu}}$.

An interesting model in this class is the TPM [35], where the lattice is triangular and the interactions are between triplets of spins in the corners of upward pointing triangles,

$$E_J(\sigma) \equiv -\frac{J}{2} \sum_{\mu=\Delta} \sigma_{i_{\mu}} \sigma_{j_{\mu}} \sigma_{k_{\mu}}, \quad (\text{TPM}) \quad (5.2)$$

The geometrical setting is shown in Fig. 5.1(a). Our analysis rests on a correspondence between configurations of the spin variables σ_i and the plaquette variables τ_{μ} . If we first consider rhombus-shaped systems whose linear size is an integer power of 2, with periodic boundaries, then there is a one-to-one mapping between spin configurations and plaquette configurations. (It is clear that every spin configuration corresponds to a unique plaquette configuration, but the existence of a spin configuration corresponding to *every* plaquette configuration is

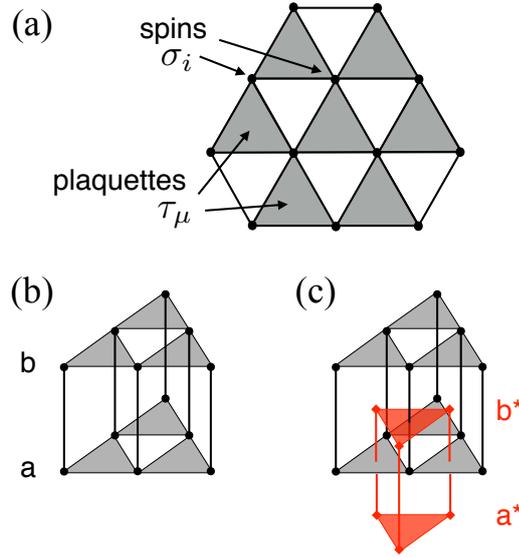


Figure 5.1: (a) Illustration of spins and plaquettes in the TPM. The spins σ_i are located on the vertices of the lattice. The plaquette variables τ_μ are located on the upward pointing triangles (shaded). Each plaquette is associated with three spins and the variable τ_μ is given by the product of these spins. (b),(c) Geometrical illustration of the duality relation (A.9) in the TPM. Panel (b) shows two TPM systems, a and b , with coupling as in (5.3). Panel (c) shows the location of the sites of the dual problem, again two coupled TPMs a^* and b^* . The plaquettes in the dual system bisect the coupling interactions in the direct system, and vice versa.

less trivial [28, 33–35].) For systems of different sizes or with different boundary conditions, the correspondence is not perfectly one-to-one, but these deviations turn out to be irrelevant in the thermodynamic limit. In section 5.4 below, we will also discuss the SPyM, a three-dimensional model with the same one-to-one correspondence, on the body-centred cubic (bcc) lattice.

In cases where the one-to-one mapping holds exactly, the fully polarised state $\sigma_i = 1 \forall i$ is the unique ground state of (5.1). In terms of the plaquette variables, the ground state is $\tau_\mu = 1 \forall \mu$, and the elementary excitation is a “de-

fect”, $\tau_\mu = -1$. Two-body spin correlations vanish in these models [35], although higher-order spin correlations are finite and allow access to a growing length scale at low temperatures [100]. Also, since there is a one-to-one mapping between spins and plaquettes, the thermodynamic properties of these models are those of non-interacting binary plaquette variables [28, 33–35], or a free gas of ‘defective’ plaquettes (with $\tau_\mu = -1$)¹.

However, while the thermodynamic properties of plaquette models are trivial, their (single spin-flip) dynamics is not. This effect arises because flipping a single spin σ_i changes the states of all of the plaquettes in which it participates. The plaquette dynamics is therefore “kinetically constrained” [28, 34, 35] possibly leading to complex glassy dynamics at low temperatures. This is what occurs for example in the TPM whose dynamical properties are similar to those of the East facilitated model [28, 34, 35], displaying “parabolic” super-Arrhenius relaxation, dynamic heterogeneity, and other characteristic features of the glass transition [7].

5.1.2 Coupled replicas

To probe thermodynamic overlap fluctuations, we consider two coupled replicas of a plaquette model [11, 36, 102–104]. The energy function of the combined system is

$$E_{J,\varepsilon}(\sigma^a, \sigma^b) \equiv E_J(\sigma^a) + E_J(\sigma^b) - \varepsilon \sum_i \sigma_i^a \sigma_i^b, \quad (5.3)$$

where σ^a and σ^b are the spin configurations in the replicas a and b . The overlap,

$$Q(\sigma^a, \sigma^b) \equiv \sum_i \sigma_i^a \sigma_i^b, \quad (5.4)$$

¹If the plaquette interactions do not allow for a spin-defect duality, as for example with spins in a cubic lattice with interactions on the square faces (rather than on the cubes), then the static properties may be non-trivial. See for example [101]

measures how similar the two copies are, and the strength of their coupling is given by its conjugate field ε . The coupling (5.3) is denoted *annealed* since both replicas are allowed to fluctuate on an equal footing. The case of *quenched* coupling, in contrast, involves one of the replicas being frozen in an equilibrium configuration. Here we will only consider the case of annealed coupling which is easier to treat both analytically and numerically. Hence the partition function for these two coupled replicas is

$$Z_2(J, \varepsilon) = \sum_{\sigma^a, \sigma^b} e^{-\beta E_{J, \varepsilon}(\sigma^a, \sigma^b)}. \quad (5.5)$$

where the sum runs over the configurations σ^a, σ^b : that is, over all $\sigma_i^a = \pm 1$ and all $\sigma_i^b = \pm 1$. Here and in the following, we sometimes set $\beta = 1$ where there is no ambiguity [for example, the left hand side of (5.5) should strictly be $Z_2(\beta J, \beta \varepsilon)$ but we suppress the dependence on β , for simplicity].

5.1.3 Dualities and phase transitions

The mapping from two coupled replicas to a single system in a field has useful consequences, since we may exploit existing results for plaquette models in magnetic fields. The duality of this model, Eq. (A.6) (see also [98]), implies a duality relation for the free energy $F_1 = -\ln Z_1$,

$$F_1(h, J) + \frac{N}{2} \ln \sinh(2\beta h) = F_1(\tilde{h}, \tilde{J}) + \frac{N}{2} \ln \sinh(2\beta \tilde{h}) \quad (5.6)$$

where \tilde{h} and \tilde{J} are given in Eq. (A.7). Phase transitions appear as singularities in the free energy density $f_1 = \lim_{N \rightarrow \infty} F_1/N$.

From (5.6), if F_1 is singular at (h, J) , it is also singular at (\tilde{h}, \tilde{J}) . This places constraints on the possible phase behaviour of the system. In particular, if the system supports only a single phase transition, it must occur for parameters such that $(h, J) = (\tilde{h}, \tilde{J})$. This condition defines a line in the (h, J) -plane, which is

given by

$$\beta J = -\ln \tanh \beta h. \quad (5.7)$$

On this line one has also $\sinh \beta J \sinh 2\beta h = 1$. Phase transitions that occur on such lines were investigated in [97, 98]: the plaquette models considered there support a single critical point that occurs at some point (J_c, h_c) on this line, with first-order phase coexistence occurring on the part of the line with $J > J_c$.

We note that (5.6) resembles the Kramers-Wannier duality of the Ising model, which allows the position of the critical point to be identified exactly in that model. Here the situation is different because the transition takes place at finite h , in contrast to the Ising transition which is known to take place at zero field, by symmetry. For this reason, the duality of the TPM does not fully determine the position of the critical point, but restricts it to the line (5.7) within the (h, J) plane.

From the mapping in Appendix A, the phase transitions of the TPM in a field correspond to phase transitions in the coupled replica system: the first-order transition line separates a state with low overlap (small ε) from one with high overlap (large ε). For the coupled replicas, the self-dual line is

$$\sinh(\beta J) \sinh(\beta \varepsilon) = 1, \quad (5.8)$$

This situation, where the self-dual line for the coupled-replica system contains a first-order transition region and a critical point, was proposed for the TPM in Ref. [36]. We present numerical evidence for this situation in section 5.2 below.

5.1.4 Other consequences of dualities and symmetries

In this section, we explore some further consequences of the results derived thus far. First, we note that the relation (A.2) means that for a coupled replica system

at parameters (J, ε) , the probability of a particular configuration of the overlap variables q is the same as the probability of finding the configuration $\sigma = q$ for a single system in a field, with parameters $(J', h = \varepsilon)$. From a numerical perspective, the single system in a field is much simpler to simulate, and the result (A.2) means that such a simulation provides direct access to all observables based on the overlap variables. (This result is much stronger than a mapping at the level of free energies.)

Second, for a geometrical interpretation of the duality relation (A.9), we refer to Figs. 5.1(b) and 5.1(c). The original coupled system can be thought of as a lattice consisting of two parallel layers, a and b . The duality relation (A.9) may be interpreted as a mapping between two different two-layer systems, where the plaquette energy scale in one model determines the interlayer coupling in the other, and vice versa. Fig. 5.1(c) illustrates this situation, in which the interlayer ‘bonds’ in the original system intersect the intralayer plaquettes in the dual system, and *vice versa*. This geometrical way of seeing the duality easily generalises to other lattices and plaquette interactions.

Third, the duality relation for a plaquette model in a field can be used to analyse the behaviour of its free energy in the vicinity of a (presumed) critical point. We assume that a critical point exists somewhere on the self-dual line, and that this critical point is in the Ising universality class, as is found generically [97] (see also below). The free energy is singular at the critical point (h_c, J_c) , and the form of this singularity is universal. Given (5.6), it is convenient to define

$$F_{\text{symm}}(h, J) = F_1(h, J) + \frac{N}{2} \ln \sinh(2\beta h), \quad (5.9)$$

The singular behaviour of F_{symm} is the same as that of F_1 since the added term is regular. The reason for introducing F_{symm} is that the duality relation (5.6) now reads simply

$$F_{\text{symm}}(h, J) = F_{\text{symm}}(\tilde{h}, \tilde{J}). \quad (5.10)$$

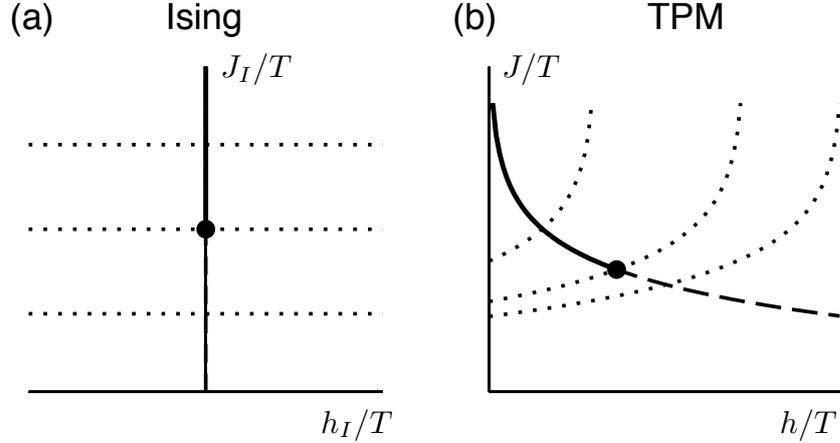


Figure 5.2: Illustration of the relation between critical behaviour of the Ising model and the TPM. (a) Ising phase diagram. the $h_I = 0$ axis is a symmetry line, there is a critical point indicated by a circle, with a first-order (phase coexistence) line for large J_I , indicated by a solid line. Selected lines of constant J_I are indicated by dotted lines. (b) The corresponding situation for the TPM in a field. On the solid/dashed line (5.7), the system has a discrete (Z_2) symmetry: a critical point and phase coexistence both occur on this line, as indicated. The dotted lines are obtained from (5.13) for three different values of h_0 , and correspond to the lines of constant J_I in panel (a). Near the critical point, they indicate the direction of the most relevant renormalisation group flow.

To investigate the universal behaviour of F_{symm} , we introduce fields $\mathcal{H} = \mathcal{H}(h, J)$ and $\mathcal{J} = \mathcal{J}(h, J)$. These fields are defined to be equal to zero at the critical point, and will correspond to the directions in parameter space that are relevant under the renormalisation group. The universal character of phase transitions means that these fields can be chosen such that

$$F_{\text{symm}}(h, J) - F_{\text{symm}}(h_c, J_c) \simeq a_0[F_I(\mathcal{H}, \mathcal{J}) - F_I(0, 0)] \quad (5.11)$$

where a_0 is a constant and $F_I(\mathcal{H}, \mathcal{J})$ is the free energy of an Ising model in a field $h_I = \mathcal{H}$ and with coupling $J_I = J_{c,I} + \mathcal{J}$, in which $J_{c,I}$ is the critical coupling of

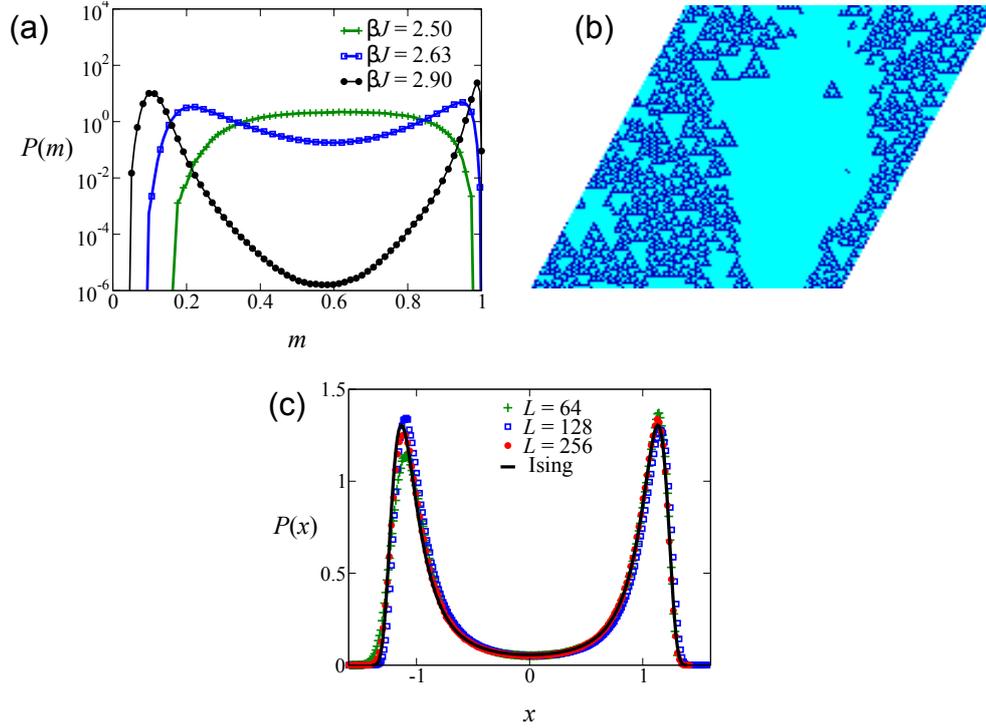


Figure 5.3: Simulations of the TPM in a field. (a) Distribution of the magnetisation at various values of J for state points on the self-dual line (5.8), at system size $L = 128$. The bimodal distribution $P(m)$ indicates a first-order transition, which disappears on reducing J . From (A.2), the same distributions would be obtained when considering the overlap between two coupled TPMs, at appropriate state points. (b) Representative configuration at phase coexistence ($\beta J = 2.9$ and $L = 128$) showing interfaces between regions of small and large magnetisation (corresponding to regions of small and large overlap in the two-replica problem). (c) At our estimated critical point, ($J_c = 2.634$, $h_c = 0.072$) and for various system sizes, we show distributions of the variable x that is obtained by rescaling the order parameter \mathcal{M} to zero mean and unit variance. The full line is the corresponding result for the $2d$ Ising model at criticality, indicating that the critical point of the TPM in a field (and therefore of the two coupled TPMs) is in the $2d$ Ising universality class.

the Ising model. The approximate equality in (5.11) accounts for non-universal contributions which are negligible for sufficiently large systems and close enough to the critical point.

The symmetry of the Ising model under inversion of the field h_I means that $F_I(\mathcal{H}, \mathcal{J}) = F_I(-\mathcal{H}, \mathcal{J})$. The corresponding symmetry relation for the TPM is (5.10), and the operation of this symmetry should invert the magnetic field \mathcal{H} but preserve the coupling \mathcal{J} in the equivalent Ising model. Hence the functions \mathcal{H} and \mathcal{J} should satisfy

$$\mathcal{H}(\tilde{h}, \tilde{\mathcal{J}}) = -\mathcal{H}(h, J), \quad \mathcal{J}(\tilde{h}, \tilde{\mathcal{J}}) = \mathcal{J}(h, J) \quad (5.12)$$

A geometrical interpretation of (5.11) is shown in Fig. 5.2. The self-dual line (5.7) corresponds to the symmetry line $\mathcal{H} = 0$, which corresponds to $h_I = 0$ for the Ising model. We now derive TPM analogs of the dotted lines $J_I = \text{const.}$ in Fig. 5.2(a). That is, we seek curves $J(h)$ such that $\mathcal{J}(h, J(h))$ is constant.

In fact, since (5.11) applies only close to the critical point, it does not fully fix the dependence of \mathcal{H} and \mathcal{J} on (h, J) . However, a consistent choice for the family of curves with constant \mathcal{J} is

$$\beta J_{h_0}(h) = -\ln \tanh(2\beta h_0 - \beta h). \quad (5.13)$$

This family of curves is illustrated by the dotted lines in Fig. 5.2b. Each curve is associated with a parameter h_0 , which is the value of the field h at which it crosses the self-dual line. The dual of any point on the curve $J_{h_0}(h)$ is easily verified to be $(\tilde{\mathcal{J}}, \tilde{h}) = (J_{h_0}(2h_0 - h), 2h_0 - h)$, which also satisfies (5.13), consistent with these two points having the same value of \mathcal{J} (recall (5.12)). It also follows from (5.6) that F_{symm} is a symmetric function of $h - h_0$, when evaluated on these lines, just as $F_I(\mathcal{H}, \mathcal{J})$ is a symmetric function of \mathcal{H} , when evaluated at fixed \mathcal{J} .

These curves are useful because the function $F_I(\mathcal{H}, 0)$ is the cumulant generating function for the magnetisation of the Ising model at criticality. One obtains cumulants by taking derivatives with respect to \mathcal{H} . Differentiating (5.11) in this way, evaluation of the right hand side requires $(\partial h/\partial \mathcal{H})_{\mathcal{J}}$ and $(\partial J/\partial \mathcal{H})_{\mathcal{J}}$. In fact it is sufficient to have the ratio of these derivatives, $(\partial J/\partial h)_{\mathcal{J}}$, but since (5.13) defines lines of constant \mathcal{J} , this is simply $J'_{h_0}(h) = -2/\sinh(4\beta h_0 - 2\beta h)$. Evaluating this derivative on the self-dual line yields $J'_{h_0}(h_0) = 2 \sinh \beta J_0$ where $J_0 = -\ln \tanh h_0$ is the value of the coupling at that point (recall $\sinh 2\beta h_0 \sinh \beta J_0 = 1$ for points on the self-dual line). Hence we have from (5.11),

$$\left(\frac{\partial}{\partial \mathcal{H}}\right)^m F_I(0, 0) \propto \left(\frac{\partial}{\partial h} + 2 \sinh \beta J \frac{\partial}{\partial J}\right)^m F_{\text{symm}}(h_c, J_c) \quad (5.14)$$

The left hand side is proportional to the m th cumulant of the Ising magnetisation. All cumulants with odd m are zero by symmetry. An explicit calculation using the symmetry (5.10) shows that the right hand side is also zero for odd m , as required. (Note also that one expects corrections to (5.14) due to non-universal contributions, as in (5.11)).

To obtain a more direct interpretation of the right hand side of (5.14), recall that $(-\partial_h)^m F_1$ is the m th cumulant of the magnetisation $M = \sum_i \sigma_i$ in the TPM, and $(\partial_J)^m [F_1 + (NJ/2)]$ is the m th cumulant of the number of defects, $N_d = \frac{1}{2} \sum_{\mu} (1 - \tau_{\mu})$. We therefore define an order parameter for the TPM

$$\begin{aligned} \mathcal{M} &= -T \left(\frac{\partial}{\partial h} + 2 \sinh \beta J \frac{\partial}{\partial J}\right) [F_1 + (NJ/2)] \\ &= M - 2N_d \sinh J, \end{aligned} \quad (5.15)$$

Using (5.9) and (5.14), and working at the critical point of the TPM, one sees that cumulants \mathcal{M} may be written as a sum of terms, one of which is proportional to an Ising model cumulant $(\partial_{\mathcal{H}})^m F_I(0, 0)$, while others are non-universal corrections. Close to the critical point, the even cumulants of \mathcal{M} show singular behavior while the corrections remain regular. Hence, for systems large enough that the singular

terms dominate the regular ones, we expect the critical distribution of \mathcal{M} in the TPM at criticality to match the critical distribution of the Ising magnetisation. This prediction will be verified numerically in the following section.

5.2 Numerical results for the TPM in a field

We performed numerical simulations of the TPM in a field, to analyse its phase behaviour. Working always on the self-dual line (5.8) we use continuous time Monte Carlo simulations [105, 106] to sample a reweighted Boltzmann distribution $P(\sigma) \propto b(M(\sigma))e^{-\beta E_J(\sigma) + \beta h M(\sigma)}$, where $b(M)$ is a bias function and $M(\sigma) = \sum_i \sigma_i$ is the magnetisation. We measure the resulting distribution $P_b(M)$ of the magnetisation, but we choose the function $b(M)$ so that this sampled distribution does not include any deep minima (free energy barriers) [107]. The ‘true’ distribution $P(M)$ associated with the unbiased model is then easily obtained as $P(M) \propto P_b(M)/b(M)$.

The bias potential $b(M)$ is chosen such that $b(M) \approx 1/P(M)$ within a range of M close to its mean. Outside this range we take $b(M)$ independent of M (and $b(M)$ is continuous at the edge of the range). This means that the sampled distribution $P_b(M)$ is approximately constant over a range close to its mean. In practical terms, we start at a relatively high temperature T for which sampling is easy and we collect N_s representative configurations σ^r with $r = 1, 2, \dots, N_s$. We typically take N_s in the range $10^3 - 10^4$: these samples are not fully independent from each other but the sampling runs are long enough that the configurations fully decorrelate within each run. For each sample, we store both the magnetisation $M_r = M(\sigma^r)$ and the number of defects $N_{d,r} = N_d(\sigma^r)$ as prescribed by (5.2). This provides an estimate for a suitable bias potential for further simulations at

this temperature:

$$b(M) \propto \frac{1}{\sum_r \delta_{M,M_r}} \quad (5.16)$$

where $\delta_{M,M'}$ is the Kronecker delta, so $1/b(M)$ is the *empirical distribution* of M . Clearly (5.16) can be used only in the range of M for which one has good sampling: for large systems it may also be useful to smooth $b(M)$ by forming a histogram with a suitably chosen bin width.

We now reduce the temperature to $T - \Delta T$, which corresponds to an increase in J/T of $\Delta_J = J/T - J/(T - \Delta T)$. We also reduce the parameter ε/T by Δ_ε so that the system remains on the self-dual line (5.7). It is easily verified that given sufficient sampling,

$$\frac{1}{b_\Delta(M)} \propto \frac{1}{b(M)} \sum_r \delta_{M,M_r} e^{-\Delta_J N_{d,r} - \Delta_\varepsilon M_r} \quad (5.17)$$

converges to the (unbiased) distribution $P(M)$ at the new state point. We choose a value of ΔT that is small enough for this estimate to be reasonably accurate (essentially this requires that the exponential weights in (5.17), do not result in concentration of the probability onto too small a fraction of the samples). Then $b_\Delta(M)$ is used as a bias potential for a simulation at this new state point. Repeating this procedure allows the temperature to be further reduced. The advantage of the method is that the histogram of M being sampled is approximately flat at each stage. In contrast to unbiased simulations ($b(M) = 1$ for all M), this method is useful when $P(M)$ has two peaks separated by a deep minimum, in which case unbiased calculations tend to sample only from one peak or the other, and rarely make transitions between them. However, the flat histogram method facilitates these transitions: see [107] for a review of these kinds of method. We note in passing that this scheme could also be implemented using \mathcal{M} as the order parameter: we would expect similar performance in this case.

Figure 5.3(a) shows numerical evidence that for large J (and small h), the self-

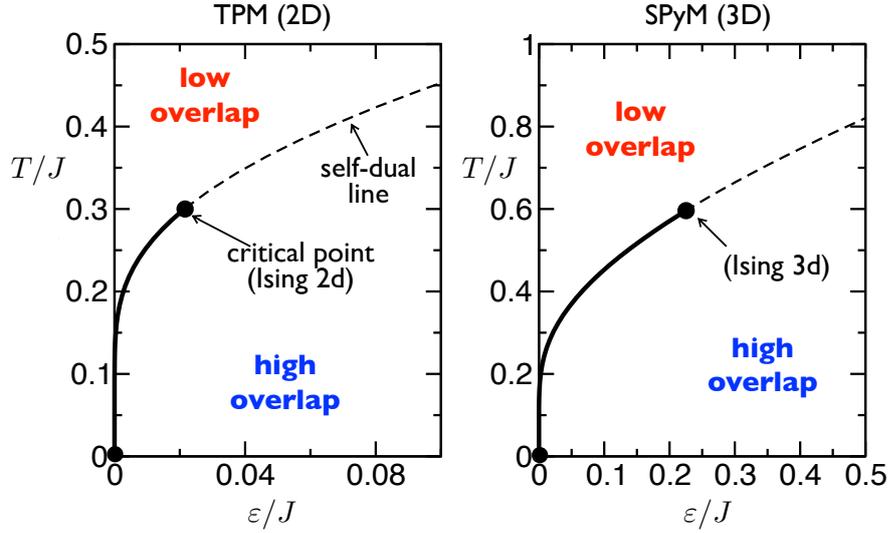


Figure 5.4: Phase diagrams of coupled plaquette models, the two-dimensional TPM (left) and the three-dimensional SPyM (right). The full line corresponds to a line of first-order transitions between a thermodynamic phase of small overlap and one of large overlap between the replicas. This curve is on the self-dual line (5.8) (dashed line). The first-order transition line ends at a critical point that is in the $2d$ Ising universality class for the TPM and the $3d$ Ising universality class for the SPyM.

dual line is associated with a first-order phase transition: see also Ref. [98]. The figure shows the distribution $P(m)$ of the magnetisation density $m = M/N$, for three state points on the self-dual line. At large J the distribution is bimodal, characteristic of first-order coexistence. A typical configuration at these conditions is shown in Fig.5.3(b), showing coexistence of low- and high- m regions, separated by sharp interfaces, as expected for a first-order transition. Based on smaller systems, a previous study [98] speculated that phase separation would not occur for the TPM in a field, but our results show that this does indeed occur large enough systems are considered. Given the mapping (A.2), Fig. 5.3(b) is also a representative configuration of the overlap between two coupled replicas, for suitable (J, ε) .

As J is decreased (or equivalently, temperature and field are increased) along the self-dual line, the bimodality in $P(M)$ becomes less pronounced and eventually disappears. Fig. 5.3(a) indicates that the first-order line terminates at a critical point (J_c, h_c) with $J_c \approx 2.6$. To identify the universality class of this phase transition, we performed a finite-size scaling analysis, using the order parameter \mathcal{M} defined in (5.15). Note that if all spins are up, one has $M = N$ and $N_d = 0$, giving $\mathcal{M} = N$. On the other hand, in a state with $h = 0$ then $M = 0$ and $N_d = 1/(1 + e^J)$, at low temperatures this gives $\mathcal{M} \approx -N$. In general, one expects a crossover between these two limits as h is increased from 0, with the crossover occurring near $\mathcal{M} = 0$.

As discussed in Sec. 5.1.4, one expects the distribution of \mathcal{M} at the critical point to be close to the distribution of the magnetisation in a critical Ising model. This provides a sensitive method for identifying the critical point. In order to match the *shape* of the distribution, it is convenient to subtract the mean of \mathcal{M} and rescale so that the distribution has unit variance. We accomplish this by defining $x = (\mathcal{M} - \langle \mathcal{M} \rangle) / \sqrt{\langle (\mathcal{M} - \langle \mathcal{M} \rangle)^2 \rangle}$. Working at the critical point, we then expect $P(x)$ to be independent of the system size, and that $P(x)$ for the TPM should also match with a similarly rescaled distribution of the magnetisation of the $2d$ Ising model. This allows the critical point to be estimated as the temperature for which $P(x)$ is independent of L and matches with its known Ising form.

We find these conditions to be satisfied for a coupling $J/T = 2.634$ (with h/T chosen to be on the dual line). The resulting scaling collapse is shown in Fig. 5.3(c), which includes data for the TPM at three system sizes, and data for a 2d Ising model at criticality [108]. The data collapse is not perfect but correspondence with the universal (Ising) form is increasingly good as the system size increases, consistent with the singular terms in the free energy becoming dominant at the system size increases (recall Sec. 5.1.4). These data confirm the Ising universality

class of this phase transition. Note that this collapse implies that cumulant ratios such as $\langle \delta \mathcal{M}^4 \rangle / \langle \delta \mathcal{M}^2 \rangle^2$ must take the appropriate universal values, independent of system size. Hence the data in Fig. 5.3c is sufficient to ensure that cumulant ratio crossings (see for example [97]) also take place at this estimated critical temperature. (Note also that while the critical distribution of the order parameter is bimodal, the transition is *second order* [108]: the separation of the peaks vanishes as the system size $L \rightarrow \infty$ and the depth of the trough between them remains constant, in contrast to first-order transitions for which the peaks spacing remains constant and $P(M) \rightarrow 0$ in the trough.)

We also calculated the ratio of the susceptibilities $\chi = L^{-d} \langle \delta \mathcal{M}^2 \rangle$ at criticality, for the system sizes $L = 128$ and $L = 256$. We find $\chi(L = 256) / \chi(L = 128) = 3.36$. Theory predicts that the susceptibility should scale as $\chi \propto L^{\gamma/\nu}$ where $(\gamma, \nu) = (7/4, 1)$ are the susceptibility and correlation length exponents for $2d$ -Ising universality. This yields a prediction for the ratio of our susceptibilities of $256^{7/4} / 128^{7/4} = 2^{7/4} \approx 3.364$, which is consistent with our results.

Bringing together these results, we arrive at the phase diagram shown in Fig. 5.4. We re-introduce the temperature $T = \beta^{-1}$ as an explicit parameter and plot the phase diagrams as a function of T/J and ε/J since this is conventional representation in supercooled (glassy) liquids. The form of this phase diagram was proposed in [36]. However, the results here now provide both the correct location and universality class of the critical point as compared to that work, whose arguments were based on the incorrect assumption that the TPM in a field was equivalent to the generalised Baxter Wu model [109]. Our finding that the critical point is in the Ising class is interesting, since this is the expected result from other general arguments [102], and what is observed in simulation of coupled liquids [103, 104].

A key feature of the plaquette models is that the first-order transition line does

not intersect the $\varepsilon = 0$ axis except at $T = 0$ [36]. However, the duality line at low temperatures is $\beta\varepsilon \simeq 2e^{-\beta J}$, so the field $\varepsilon = \varepsilon^*$ at which the first-order phase transition takes place is very small. Indeed, simple extrapolation of the first-order line from high-temperature data might lead one to propose a phase transition at $\varepsilon = 0$ and some $T = T_K > 0$, as predicted in mean-field theories [11]. In this case the scenario shown in Fig. 5.4 might be hard to distinguish from the mean-field picture. However, direct simulations near this proposed phase transition (if possible) would demonstrate that there are no diverging fluctuations as near the first-order transition line, in contrast to the mean-field scenario where length and time scales diverge at T_K .

5.3 Active-inactive dynamical transitions in the TPM

The TPM falls into a category of glassy models that are thermodynamically simple but where glassy behaviour arises because of non-trivial dynamical pathways to the equilibrium state at low temperature [27, 28, 34, 110]. In fact, the dynamics of the TPM [34] is closely related to that of a two-dimensional East model [27, 28, 96]. Many kinetically constrained models, including the East model (see section 4.3.3), display dynamical phase transitions – phase transitions in the space of trajectories – between a phase with a high dynamical activity, K , and one with low dynamical activity [12].

For these lattice models, the models evolve in time by flips of spin variables (for plaquette models, these are the σ_i variables). Each time a spin changes its state, this changes the configuration of the system: the activity K is again defined as the total number of configuration changes in a trajectory [16, 47], with a time-intensive

activity,

$$k(s) \equiv \langle K \rangle_s / \tau. \quad (5.18)$$

Like the East model, the dynamical relaxation of the TPM is hierarchical, due to energy barriers to relaxation that are logarithmic in the linear size of relaxing regions [28]. This in turn leads to a “parabolic” [111] super-Arrhenius law for the typical relaxation time in equilibrium as a function of temperature. Given the similar dynamical properties of the TPM and the East model dynamics, a natural question is whether the TPM also displays active-inactive dynamical phase transitions.

In order to answer this question numerically we again make use of transition path sampling (TPS) to efficiently sample trajectories in the s -ensemble. These systems are computationally expensive, so the efficiency provided by the x -ensemble is of great benefit, however for clarity of comparison to prior results we convert back to the s -ensemble, as in Chapter 4.

Figure 5.5 shows the results of the s -ensemble analysis of the TPM. It shows the average activity for a system of size $N = 8 \times 8$ as a function of s , at temperature $T = 0.5$ (note this is the TPM in the absence of field). As the length of the trajectories is increased the change in $\langle K \rangle_s$ becomes more pronounced, as seen in the corresponding susceptibilities. This is indicative of a first-order transition at some $s_c \gtrsim 0$. Similar size scaling is observed by changing system size, as shown in the insets. The dependence of this transition on the temperature is discussed in Sec. 5.4.3 below, together with similar results for the (three-dimensional) SPyM.

We note at this point that these transitions are inherently dynamical in nature. For example, one might also consider a plaquette model that involves in time by dynamical rules in which single plaquette variables may change their state independently. In this case the dynamics is that of a (grand canonical) free gas of

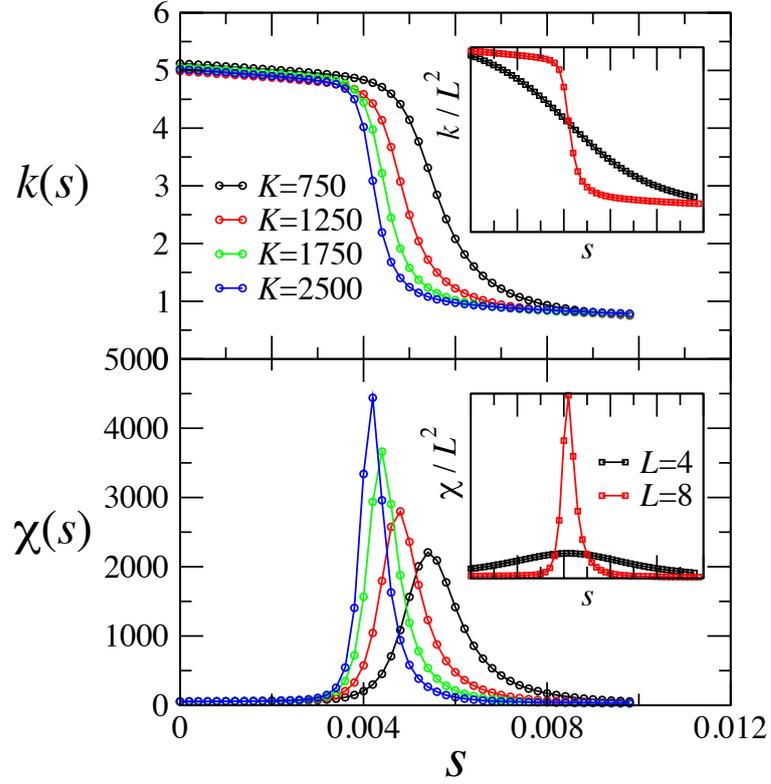


Figure 5.5: Average activity $k(s)$, and the associated susceptibility $\chi(s)$ in the TPM at $T = 0.5$. The main panels show data for system size $L = 8$. These results were obtained by the x -ensemble method (see text), using trajectories with fixed numbers of events K , as shown. On increasing the trajectory length, the crossover from active to inactive behaviour becomes increasingly sharp and the susceptibility peak increases. The behaviour for smaller systems ($L = 4$) is shown in the insets, with both quantities normalised by the system size. In the absence of a phase transition, one expects both $k(s)/L^d$ and $\chi(s)/L^d$ to be independent of L , so the sharper crossover at $L = 8$ is again consistent with an underlying phase transition.

defective plaquettes and the statistics of K are simply those of a Poisson process (the same situation is observed when considering kinetically constrained models from which the constraint has been removed [12]). Similarly, the transition in the coupled replica system of the previous section is also destroyed if one measures the

overlap at the defect level (that is, one replaces Q in (5.4) by $Q_d = \sum_{\mu} \tau_{\mu} \tau'_{\mu}$). In summary, observation of the phase transitions discussed in this chapter requires a suitable choice of order parameter (Q not Q_d), and the dynamical phase transitions also depend on the dynamical rules by which the systems evolve.

5.4 Overlap and activity transitions in a three-dimensional plaquette model

In order to explore whether the static and dynamical transitions found above for the TPM are present in dimensions other than two it is of interest to generalise the TPM to higher dimensions. One of the reasons is that if one wishes to consider plaquette models to study “quenched” coupled replicas [11, 102, 103] or “random pinning” [112–114] the distinction between two and three dimensions may be very significant, due to the inability of two-dimensional systems in random fields to support first-order static transitions [115]. Here we introduce a three-dimensional model that is similar to the TPM.

5.4.1 Model

The model we consider is defined on a three-dimensional body-centred cubic (bcc) lattice. The “plaquettes” are upward-pointing square-based pyramids, each containing five spins. Considering the standard bcc unit cell, one such pyramid is formed by the spin at the centre of the cube together with the four spins at the corners of the lower face: see Fig. 5.6. We call this model the square pyramid

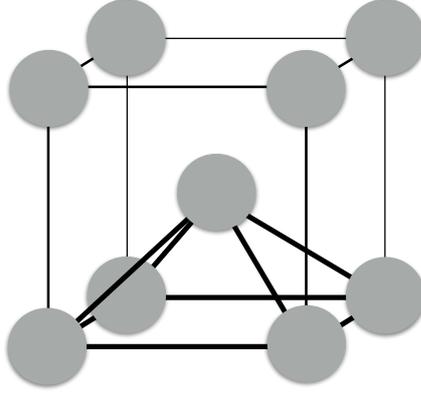


Figure 5.6: (a) The SPyM consists of spins (grey circles) on the sites of a BCC lattice, which interact in quintuplets at the vertices of upward pointing square pyramids. One such pyramid is indicated; the central spin also participates in four other upward pointing pyramids whose apexes are the four spins on the upper face of the cube.

plaquette model, or SPyM. Its energy function reads,

$$E_J(\sigma) \equiv -\frac{J}{2} \sum_{\mu} \sigma_{i_{\mu}} \sigma_{j_{\mu}} \sigma_{k_{\mu}} \sigma_{l_{\mu}} \sigma_{m_{\mu}} \quad (\text{SPyM}) \quad (5.19)$$

where μ runs over all the pyramidal plaquettes on the lattice, and the location of the five interacting spins $\sigma_{i_{\mu}} \cdots \sigma_{m_{\mu}}$ is shown in Fig. 5.6. This is “model 1” of [97].

Just like the TPM, the SPyM has a one-to-one correspondence between spin and plaquette configurations. An alternative model [35, 97] may be defined on a face-centred cubic lattice, in which the plaquettes are tetrahedral pyramids (“model 2” of [97]). However, in this case each interaction involves four spins so the system has a global spin-flip symmetry, and the spin-plaquette correspondence is not exact in finite (periodic) systems. However, these deviations from the one-to-one correspondence are irrelevant in the thermodynamic limit.

Returning to the SPyM, we explicitly demonstrate the one-to-one correspondence between spins and plaquettes, by a general method that applies also to the TPM.

The total number of configurations of the spin variables is 2^N , and this is also equal to the total number of configurations of the plaquette variables, by construction of the model. For any spin configuration, the configuration of the plaquette variables is uniquely specified since the τ_μ are defined in terms of the spins σ_i . However, it might be that some plaquette configurations can be achieved by more than one spin configuration, in which case the mapping is not one-to-one (this happens for example in the square plaquette model [100] with periodic boundary conditions). In that case the equal numbers of spin and plaquette configurations means that there must exist plaquette configurations that cannot be realised by any realisation of the spin variables. To rule this out and establish the one-to-one mapping, we now show how a spin configuration may be constructed for any given plaquette configuration.

We choose as basis vectors for the lattice $\vec{a}_1 = (1, 0, 0)$, $\vec{a}_2 = (0, 1, 0)$ and $\vec{a}_3 = (-1, -1, \sqrt{2})/2$. We focus on systems whose sites are at $l\vec{a}_1 + m\vec{a}_2 - n\vec{a}_3$ with $l, m, n \in \{0, 1, 2, \dots, L-1\}$, with periodic boundaries (so for example sites with $n = L-1$ are neighbours of those with $n = 0$). We indicate the location of the μ -th pyramid by the position of the spin at the apex. The plaquette variable τ_μ for $\mu = (i, j, k)$ is then

$$\begin{aligned} \tau_{(i,j,k)} = & \sigma_{(i,j,k)}\sigma_{(i,j,k-1)}\sigma_{(i-1,j,k-1)}\sigma_{(i,j-1,k-1)} \\ & \times \sigma_{(i-1,j-1,k-1)}. \end{aligned} \quad (5.20)$$

Following the same reasoning as in [35] we can invert this relation in terms of a ‘‘Pascal pyramid’’: the idea is to demonstrate that introducing a single defect into the system corresponds to flipping a particular set of spins.

Starting from the ground state, we demonstrate the procedure by introducing a single defect at the origin: this affects those spins in upper layers which lie on the sites of an inverted Pascal pyramid (or fractal pyramid). Assuming that the

central spin in Fig. 5.6 is at the origin, we flip that spin, which introduces a defect in the pyramid below it. In order to avoid any other defects, we also flip the four spins on the top face of the cube shown in Fig. 5.6, which ensures that there are no defects in any of the pyramids pointing upward from the origin. Iterating this procedure for all other layers, the final spin configuration is

$$\sigma_{(i,j,k)} = 1 - 2 \left[\binom{k}{i} \binom{k}{j} \pmod{2} \right], \quad (5.21)$$

where $\binom{n}{r} = \frac{n!}{r!(n-r)!}$ are combinatorial numbers, and $0 \leq i, j \leq k$ (all other spins $\sigma_{(i,j,k)} = 1$). Given periodic boundary conditions, this procedure determines all spins in the system: on setting the final layer of spins, it may be that defects in the final layer are unavoidable. However, for systems whose linear size L is a power of 2, it is easily shown that this procedure produces a final state with exactly one defect.

Now observe that for any spin configuration, flipping the set of spins for which $\sigma_{(i,j,k)} = -1$ in (5.21) inverts the state of the plaquette variable just below the origin, leaving all other plaquette variables constant. Similarly, to flip the state of any other plaquette, one applies a spatial translation to the same set of spins, and flips all the spins within this translated set. Hence, by repeatedly applying this procedure, one can generate a spin configuration that corresponds to any given configuration of the plaquette variables. This establishes the one-to-one correspondence between spin and plaquette configurations, by the argument given above.

This correspondence between spin and defect configurations means that the thermodynamics of the SPyM is that of a free binary gas of plaquettes. Furthermore, the relaxational dynamics is similar to that of a (three-dimensional) East model. Figure 5.7 shows the decay of the energy at low temperatures starting from a $T = \infty$ configuration. We see the characteristic hierarchical decay of both the

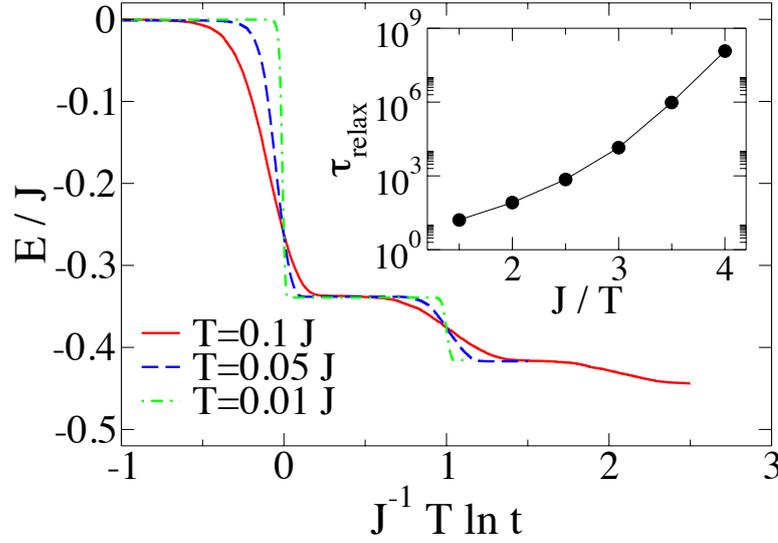


Figure 5.7: Relaxation of the energy of the SPyM at low temperature starting from a random configuration (the system size is $L = 16$). The curve shows the characteristic plateaus indicative of hierarchical relaxation, as in the East model and the TPM. Inset: average relaxation time as a function of inverse temperature, showing super-Arrhenius behaviour.

East model and the TPM: the energy decays in steps with characteristic time scales $\tau_n = e^{n\beta J}$ with $n = 0, 1, 2, \dots$. These steps become apparent when plotting as a function of the rescaled time variable $(T/J) \ln t$ [116]. The inset to Fig. 5.7 shows that the equilibrium relaxation time of the SPyM is super-Arrhenius, as in the East model and the TPM.

5.4.2 Phase transition in (annealed) coupled replicas

The SPyM possesses the exact duality described in Appendix A. In particular, the properties of the SPyM in a field were studied in Ref. [97] (“model 1” of that paper), where it was found that on the self-dual line there is a first-order transition between phases of small and large magnetisation terminating at a critical point in

the 3D Ising class. From those results we can directly infer the phase diagram of the two coupled SPyMs via the mapping of Appendix A. The result is shown in Fig. 5.4. This phase diagram is similar to that of the TPM, except that the range of phase coexistence is larger and the critical point occurs at higher temperature.

5.4.3 Evidence for a dynamical (space-time) phase transition

As well as the phase transition for coupled replicas in the SPyM, we also present evidence for a space-time phase transition, similar to that shown for the TPM in Fig. 5.5. The results for the SPyM are shown in Fig. 5.8, for temperature $T/J = 0.65$ and linear system size $L = 4$. There is good evidence for a sharp transition at $s = s^* > 0$, as found in the TPM.

In Fig. 5.8(b,c), we show how the crossover in activity varies with the temperature T/J , for both the TPM and the SPyM. Simple estimates [117, 118] indicate that if the inactive state is metastable and relaxes to equilibrium via some kind of nucleation process with rate γ_{nuc} per unit volume, then $s^* \approx \gamma_{\text{nuc}}/\delta k$, where δk is the activity difference (per unit space-time) between the active and inactive states [117]. We attribute the existence of the transition in this model to a stable inactive state with almost no defects. We expect that the rate for relaxation back to equilibrium is a strongly decreasing function of temperature, which is consistent with the increasing s^* as T/J increases. (The activity difference Δk between active and inactive states increases with T/J but this dependence is much weaker than that of the relaxation rate.)

Comparing Figs. 5.4 and 5.8, a natural question is whether the crossover lines identified in Fig. 5.8(b) and 5.8(c) are indeed first-order phase transition lines, and

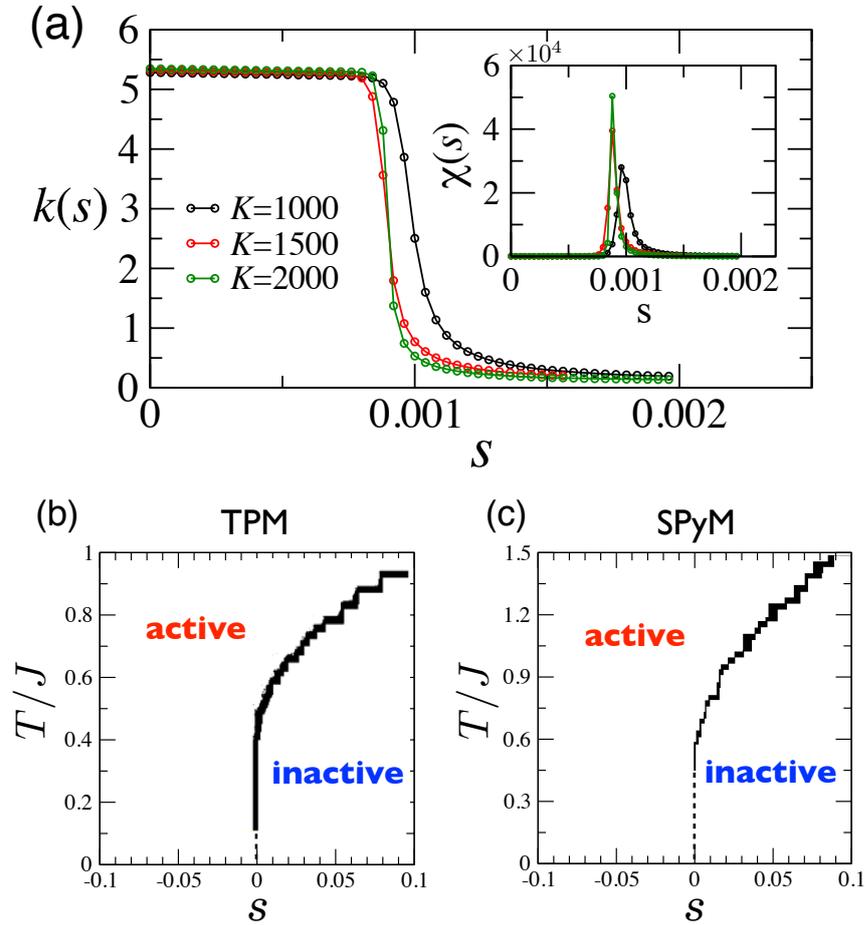


Figure 5.8: (a) Activity as a function of s in the SPyM for system size $L = 4$, various trajectory lengths, and $T/J = 0.65$. The inset shows the corresponding susceptibility. (b) and (c) s -ensemble phase diagrams for the TPM and SPyM. The full curves are an estimate of the transition point from the simulations. The dashed lines are extrapolations in the low temperature regime inaccessible to numerics.

whether the true phase behaviour involves a first-order line for low temperature that terminates at a critical point. Our numerical results are not sufficient to answer this question, due to the significant computational expense of sampling the s -ensemble. However, we expect on general grounds that the situation should be similar to the softened Fredrickson-Andersen model considered in [99], in which case the first-order line would indeed end at a finite-temperature critical point. This scenario is consistent with the results presented here.

5.5 Discussion

5.5.1 Connection of phase transitions to long-lived metastable states

The phase behaviour shown in Figs. 5.4 and 5.8 reveals striking similarities between thermodynamic transitions (for coupled replicas) and dynamic transitions (based on dynamical activity). The two transitions are distinct and we do not believe that they are related by any exact mapping. Nevertheless, we argue in the following that these transitions are connected to the existence of long-lived metastable states, which are intrinsically linked to the glassy behaviour in these systems.

Consider first transitions for annealed replicas. If we work at the phase coexistence point, but within the high-overlap phase, the system occupies low-energy states. For $\varepsilon = 0$, these states would be metastable (with finite lifetimes). This metastability means that when localised low-overlap regions are generated by thermal fluctuations, these regions tend to shrink, just as small fluctuations tend to shrink within classical nucleation theory. Escape from the metastable state

requires a collective process that operates on some finite length scale. As ε is increased from zero, this length scale increases, as does the associated free energy barrier: both diverge at the coexistence point where the high-overlap phase becomes stable.

The situation for dynamical phase transitions is similar, except that one should think of trajectories of the system as $(d+1)$ -dimensional objects that exist in space-time. If one works at the dynamical phase coexistence point (some $s = s^* > 0$) then inactive trajectories dominate the s -ensemble. During these trajectories, the system remains localised in low-energy metastable states, with small thermal fluctuations of the activity, associated with space-time “bubbles” [50]. If the trajectory length τ is less than the time required for escape from the metastable state, similar inactive trajectories can be generated with unbiased ($s = 0$) dynamics, by taking initial conditions from low-energy metastable states. Transformation of such a trajectory into a *typical* equilibrium trajectory involves the introduction of an active space-time bubble, which subsequently grows to macroscopic size. The connection with metastability arises because if one introduces a small active bubble within an inactive trajectory, one expects to incur a cost in probability (if this were not the case, the state would not be metastable since it would readily relax back to equilibrium). As in the case of overlap fluctuations, the critical bubble size and the probability barrier increase as s is increased from zero, diverging at the coexistence point.

We argue that this analogy between phase coexistence phenomena induced by s - and ε -fields provides a qualitative explanation of the similarity between Figs. 5.4 and 5.8, in that both are linked to the existence of metastable states that can be observed in unbiased ($s = 0 = \varepsilon$) systems. The relevant metastable states have low energy: both the inactive state of Fig. 5.8 and the high-overlap state of Fig. 5.4 have much lower defect concentrations than the equilibrium average value

$c \approx e^{-\beta J}$. For large s , the system minimises its propensity for dynamical activity by removing defects, so that $N_d/N \approx 0$; for large coupling ε it is easy to show that $N_d/N \approx e^{-2\beta J}$ since the system has an effective temperature $1/(2\beta)$ [119]. (Considering the partition function (5.5), for large ε we have $\sigma^a \approx \sigma^b$ in which case the Boltzmann factor in (5.5) reduces to $e^{-2\beta E_J(\sigma^a) + \varepsilon N}$, and the distribution over configurations σ^a is equilibrium-like at temperature $1/(2\beta)$. The general case considered in [119] involves m coupled replicas in the case where the coupling is strong enough to localise all replicas in the same metastable state, but weak enough that it has a negligible effect on intra-state fluctuations.)

Evidence for phase coexistence induced by s and ε -fields have both been presented in atomistic models [13, 103]. By contrast, in kinetically constrained models (KCMs), dynamical phase coexistence can be induced by the s -field [12] but there is no such transition as a function of ε . Metastable states can be identified in KCMs [118] – here a *state* is defined [120] as a region of configuration space for which the equilibration time within the state is much smaller than the time to escape from it. (As argued in [120], this is a robust definition in finite-dimensional systems at $\varepsilon = 0$, for which mean-field constructions break down.) Given this definition of metastability, phase coexistence at positive s may be expected, as argued above. However, this metastability does not lead to phase coexistence for any ε : in fact the free energy for two coupled KCMs can be obtained exactly and the statistics of the overlap are trivial in this case. Clearly the static construction based on coupled replicas does not reveal the metastability in this model: the reason is that the kinetic constraints generate metastable behavior by reducing the ability of the system to evolve from high-overlap to low-overlap states. (To see this, note that if one removes the kinetic constraints, the metastable states disappear but the results for coupled replicas do not change.)

A similar situation occurs when considering plaquette models with different over-

lap measures and different rules for their dynamical evolution. As discussed in Sec. 5.3, changing the dynamical rules of the plaquette models so that plaquette (defect) variables flip independently destroys the dynamical phase transition. (Viewing the plaquette models in a defect representation, this is equivalent to removing a kinetic constraint [28].) Also, changing the definition of overlap from Q to $Q_d = \sum_{\mu} \tau_{\mu} \tau'_{\mu}$ means that the phase transition for the coupled replicas is lost. (This also happens in plaquette models on random graphs for which exact mean-field calculations are possible [121].)

The general point here is that the natural dynamical definition of metastability [120] coincides with the presence of phase transitions for coupled replicas only if the natural dynamics of the system leads to unconstrained local changes in the overlap. So if the dynamics has (unconstrained) single spin flips then the natural overlap parameter is the spin overlap and one observes both static and dynamic transitions in plaquette models. If the dynamics has independent changes to defect variables and one uses the defect overlap then both static and dynamical transitions are lost, so the two constructions are still consistent with each other. If one uses spin flip dynamics and an overlap based on defects then the dynamical transition survives but the coupled replicas are insensitive to it. (This is similar to the KCM case: the dynamical rules for defect variables in plaquette models are constrained but the spin-overlap is not sensitive to this constraint). Finally if one uses independent-defect dynamics but an overlap based on spins then the dynamical transition is lost but the static one survives: the dynamical and static constructions probe different aspects of the system in this case and give different results.

We end this section by noting that the inactive and high-overlap states in Figs. 5.4 and 5.8 are structurally distinct from equilibrium states at temperature T , and also from the active and low-overlap states. For example, as noted above, con-

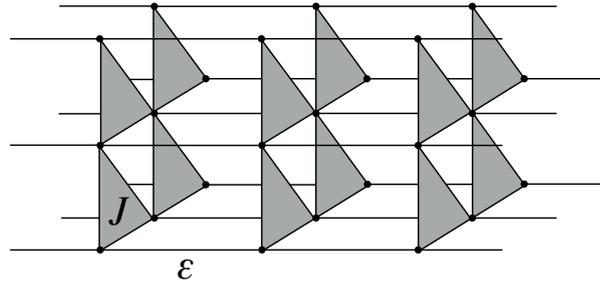


Figure 5.9: Three dimensional stack of coupled two-dimensional TPMs.

figurations deep inside the high-overlap phase have energies and spin correlation functions representative of equilibrium at temperature $T/2$ while the low overlap phase is close to an equilibrium state at temperature T . The states observed at finite ε are not exactly the equilibrium states at temperature T or $T/2$ that are found for $\varepsilon = 0$ or $\varepsilon \rightarrow \infty$, but their properties are qualitatively similar, and there is a significant jump in the energy $E_J(\sigma)$ on crossing the first-order transition line. This is quite different from ensembles with a quenched coupling between replicas [122, 123], where the structures of high- and low-overlap states should be statistically (almost) indistinguishable, with (at most) a small discontinuity in $E_J(\sigma)$ at the transition (see also [112]).

5.5.2 Connection between multiple coupled replicas and biased-activity ensembles

The connection between Figs. 5.4 and 5.8 can be further motivated through a generalisation of the coupled two-replica system discussed above. Given a plaquette model in dimension d , consider the associated $d + 1$ system composed of many replicas of the d -dimensional system arranged parallel to each other along

the extra dimension, see Fig. 5.9. Such system of n coupled replicas has an energy

$$E_n(\sigma^1, \sigma^2, \dots) \equiv E_J(\sigma^1) + E_J(\sigma^2) + \dots - \varepsilon \sum_i (\sigma_i^1 \sigma_i^2 + \sigma_i^2 \sigma_i^3 + \dots). \quad (5.22)$$

Using the methods of Refs. [36, 98, 109] it is easy to prove that the partition sum of the $(d + 1)$ -dimensional problem also has an exact duality:

$$Z_n(J, \varepsilon) = (\sinh \beta J \sinh \beta \varepsilon)^{Nn} Z_n(J^*, \varepsilon^*), \quad (5.23)$$

where we have assumed periodic boundary conditions in the transverse direction, and (J, ε) and (J^*, ε^*) are related again by (A.9). Similar results were found in [124, 125] for other classes of plaquette models. Given this duality we expect the phenomenology of this many-replica system to be similar to that of two replicas, except that any phase transitions should be in the $(d + 1)$ -dimensional Ising universality class (assuming that both longitudinal and transverse dimensions are taken to infinity in the thermodynamic limit).

The partition sum Z_n has a natural transfer matrix representation in the transverse direction, $Z_n = \text{Tr}(\mathbb{T}^n)$, with

$$\mathbb{T} = \cosh^N(\beta J/2) e^{N\beta \varepsilon} \bigotimes_{\mu} \left[1 + \tanh(\beta J/2) \sigma_{i_{\mu}}^z \sigma_{j_{\mu}}^z \cdots \sigma_{k_{\mu}}^z \right] \bigotimes_i (1 + e^{-2\beta \varepsilon} \sigma_i^x), \quad (5.24)$$

where $\sigma^{x,z}$ are Pauli matrices. This in turn can be related to the generator of (imaginary time) quantum evolution in the usual manner [124, 125] when the transverse coupling is large and the longitudinal one small [$e^{-2\beta \varepsilon}, \tanh(\beta J/2) \ll 1$], so that $Z_n \propto \exp(-tH)$, with

$$H \equiv -h \sum_i \sigma_i^x - g \sum_{\mu} \sigma_{i_{\mu}}^z \sigma_{j_{\mu}}^z \cdots \sigma_{k_{\mu}}^z, \quad (5.25)$$

where

$$\delta t h = e^{-2\beta \varepsilon}, \quad \delta t g = \beta J/2, \quad t = n \delta t. \quad (5.26)$$

The Hamiltonian (5.25) generates dynamics in the transverse direction. While it is not derived from a stochastic operator it has the basic features of the generator [12, 16] for the dynamical ensemble defined in (2.39): an off-diagonal part (the σ^x terms) that perform configuration changes and a diagonal part (the σ^z) plaquette terms associated to the escape rate. The parameter s in the s -ensemble operator controls the relative strength of the diagonal and off-diagonal terms [12, 16], in analogy with the balance between h and g in (5.25). Furthermore, the duality (5.23) implies a duality $h \leftrightarrow g$ in (5.25), with the possibility of a dynamical transition at that self-dual point $g = h$. This connection between a static transition in the $d + 1$ -dimensional problem (5.22) (itself closely connected to the static transition in the two-replica plaquette system) and a dynamical transition in the d -dimensional system (5.25) provide another rationalisation of the similarities between Figs. 5.4 and 5.8.

5.5.3 Outlook

Plaquette spin models have several features that make them attractive for studies of the glass transition. As we have shown here, exact results can be derived, which guide numerical studies of phase behaviour and many-body correlations. The models are also computationally much less demanding than atomistic models of supercooled liquids, so that (for example) finite size scaling over a large range of system sizes can be performed, to analyse phase transitions. The equilibrium relaxation of the models follows a dynamical facilitation scenario, in which point defects play a central role. However, there are strong many-body correlations, and the statistics of overlap fluctuations are rich and complex, as anticipated in the theory of Franz and Parisi [11]. In this sense, the models provide a bridge between different theories. Indeed, as argued in [100], one might describe plaquette models by a modified form of RFOT, but with two important caveats: (i) the analogue

of the Kauzmann transition occurs at zero temperature in these models (ii) the interfacial cost associated with growing droplets of a new state within a typical equilibrium state scales logarithmically in the droplet size (not as a power law as anticipated by RFOT).

Looking forward, we hope that further work on plaquette models (particularly in $d = 3$) will show to what extent mean-field [11] and RFOT ideas can be modified to apply in this setting. We can imagine that the apparently different physical pictures envisaged by thermodynamic and dynamical theories of the glass transition [7] might *both* be applicable in these models. In that case, it is not clear whether some new results would be required to discriminate between the theories, or whether they might in fact offer complementary descriptions of the same phenomena.

6. META-WORK AND AN ANALOGOUS JARZYNSKI RELATION

6.1 Jarzynski relation in trajectory space

Given that the thermodynamics of trajectories is a generalisation of equilibrium thermodynamics to dynamical ensembles, it is natural to expect there will be an analogous extension of the fundamental non-equilibrium relations encoded by the fluctuation theorems to trajectory ensembles. In this chapter we study an analogous Jarzynski equality in the thermodynamics of trajectories by identifying an analogous quantity to work. This allows for the computation of the large deviation function $g(x)$ of the x -ensemble. We study two systems, the quantum two-level system of Fig. 3.2 to provide a simple illustration of the approach, and the micromaser of section 4.3.2, where the rich dynamical phase space allows us to investigate the behaviour of the Jarzynski relation as one crosses first-order discontinuities. This is a situation which has thus far received comparatively little attention (see for example Refs. [126, 127] for a numerical and Ref. [128] for a mean-field study in the case of the Ising model).

6.1.1 Meta-dynamics: Dynamics in the space of trajectories

In the thermodynamics of trajectories it is of interest to be able to determine the function $g(x)$ over a range of values x . Instead of performing many “equilibrated” simulation runs at fixed x , we aim to extract the function $g(x)$ while changing x . To this end we require the notion of a meta-dynamics and a meta-time, which for convenience we take as integer, enumerating the sequence of generated trajectories $\vec{\mathbf{X}} \equiv (\mathbf{X}_0, \dots, \mathbf{X}_N)$. The meta-dynamics that generates these trajectories is required to obey detailed balance (see Eq. 4.1) with respect to the distribution $P_x(\mathbf{X})$ defined by Equation 3.7. The natural candidate for this meta-dynamics is transition path sampling (see Chap. 4). For the purpose of efficiency we stick to the x -ensemble scheme.

6.1.2 Meta-work and the Jarzynski relation

Equation (3.7) has the form of an equilibrium Boltzmann distribution, where $E_x(\mathbf{X}) = \mathbf{x}\tau(\mathbf{X})$ can be identified as the analogue of an “energy”. Suppose that we change x along the sequence $\vec{\mathbf{X}}$: We start with a value x_0 for the biasing field and generate the initial trajectory \mathbf{X}_0 . We then change the value of x_0 to x_1 and generate the next trajectory \mathbf{X}_1 of the sequence with a single TPS step and so on. The change of the “energy” along the whole sequence is

$$\Delta E \equiv E_{x_N}(\mathbf{X}_N) - E_{x_0}(\mathbf{X}_0) = W + Q, \quad (6.1)$$

which can be split into two sums

$$Q \equiv \sum_{i=0}^{N-1} [E_{x_{i+1}}(\mathbf{X}_{i+1}) - E_{x_{i+1}}(\mathbf{X}_i)], \quad W \equiv \sum_{i=0}^{N-1} [E_{x_{i+1}}(\mathbf{X}_i) - E_{x_i}(\mathbf{X}_i)]. \quad (6.2)$$

These sums are identified as “heat” Q and “work” W , respectively. In particular, the meta-work

$$W = \sum_{i=0}^{N-1} (x_{i+1} - x_i) \tau(\mathbf{X}_i) \quad (6.3)$$

sums the incremental changes of the “energy” due to a change of the field x for the same trajectory.

We can now prove the Jarzynski relation following standard arguments by combining the form of the path probability Eq. (3.7) with Eq. (4.1). Consider the average

$$\langle e^{-W} \rangle = \int \mathcal{D}\mathbf{X}_0 \cdots \mathcal{D}\mathbf{X}_N P_{x_0}(\mathbf{X}_0) P_{x_1}(\mathbf{X}_0 \rightarrow \mathbf{X}_1) \cdots P_{x_N}(\mathbf{X}_{N-1} \rightarrow \mathbf{X}_N) e^{-W} \quad (6.4)$$

The first integral reads

$$\begin{aligned} & \frac{1}{Z_K(x_0)} \int \mathcal{D}\mathbf{X}_0 P_0(\mathbf{X}_0) P_{x_1}(\mathbf{X}_0 \rightarrow \mathbf{X}_1) e^{-x_1 \tau(\mathbf{X}_0)} \\ &= \frac{Z_K(x_1)}{Z_K(x_0)} \int \mathcal{D}\mathbf{X}_0 P_{x_1}(\mathbf{X}_0) P_{x_1}(\mathbf{X}_0 \rightarrow \mathbf{X}_1) = \frac{Z_K(x_1)}{Z_K(x_0)} P_{x_1}(\mathbf{X}_1). \end{aligned} \quad (6.5)$$

Unraveling all terms thus leads to

$$\langle e^{-W} \rangle = \frac{Z_K(x_N)}{Z_K(x_0)}, \quad (6.6)$$

which is the analogous Jarzynski relation for the meta-work in canonical ensembles of trajectories.

It is important to remember that there are two layers to the dynamics in this case, and that the trajectory analogue of the Jarzynski applies at the level above the physical dynamics occurring in a trajectory. This can be confusing as the conventional, thermodynamic Jarzynski relation uses the concept of a trajectory linking two configurations of a system. Here however, the trajectories play the roles of (meta-) configurations, with a meta-dynamics - the dynamic process of moving through the trajectory space with TPS - playing the role of the trajectories.

For clarity we refer to quantities of interest in the trajectory ensembles as meta-quantities of their thermodynamic analogues (meta-work meta-free energy etc.).

6.1.3 Computing the meta-free energy $g(x)$

From Eq. (6.6), we can extract the change of the trajectory (or meta-) free energy

$$\Delta g \equiv g(x_N) - g(x_0) = \lim_{K \rightarrow \infty} \frac{1}{K} \ln \langle e^{-W} \rangle \quad (6.7)$$

from the meta-work. Because in practice finite K must be used, it should be noted we are in essence calculating a finite- K meta-free energy $g(x, K)$. However by using a sufficiently large value of K this value approaches its large-deviation limit, $g(x)$ while still allowing the computation to be done in finite time. Using this result, the meta-free energy $g(x)$ of the x -ensemble can be calculated from simulation in the following way. A trajectory with fixed number of events K is created and equilibrated to the desired starting value x_0 using the x -ensemble TPS algorithm. The system then moves along the “forward” path up to the desired maximum value x_N in a series of steps. For simplicity, we consider a linear protocol $x_i = x_0 + i(x_N - x_0)/N$ although other protocols might be more suitable. Each step corresponds to a single change to the trajectory whether the proposed change is accepted under the Metropolis criterion or not.

This process is repeated M times until a good distribution of meta-work for both the forward and the reverse process (going from x_N to x_0) is built up. The meta-free energy difference between x_N and x_0 can then be computed with an iterative Bennett’s Acceptance Ratio (BAR) method [129, 130],

$$\Delta g^{(k+1)} = -\ln \frac{\sum_{j=1}^M \left[1 + e^{W_{\uparrow,j} - \Delta g^{(k)}} \right]^{-1}}{\sum_{j=1}^M \left[e^{W_{\downarrow,j}} + e^{-\Delta g^{(k)}} \right]^{-1}}, \quad (6.8)$$

where the sum over j denotes the sum over the work values for each repetition of forward (\uparrow) and reverse (\downarrow) process. The work values are random numbers with probability distributions $P_{\uparrow}(W)$ and $P_{\downarrow}(W)$, respectively.

As is the case in thermodynamic problems, there need be some overlap in the work distributions for the forward and reverse processes, but the rate at which these processes occur need not be slow enough to ensure equilibrium at all points (resulting in completely overlapping work distributions). Strictly speaking, the large-deviation function $g(x)$ is defined in the limit of $K \rightarrow \infty$. In practice, for the numerical estimation of $g(x)$, the length of individual trajectories as defined by the number of events K is not critical to the result, provided the meta free-energies are scaled per event. Furthermore, while short trajectories of low K necessarily require less computation time, they also necessarily have much larger fluctuations in work distributions, requiring more repetitions to build a reasonable distribution numerically, meaning there is some trade off in efficiency. Note however that a positive aspect of these fluctuations is that the broadening of work distributions can lead to an increase in their overlap. These considerations indicate that the optimal trajectory length, and number of steps to calculate the effective meta free-energies as efficiently as possible, are highly system dependent.

6.2 Application to open quantum systems

For the purpose of demonstrating the validity of the analogous Jarzynski equality (6.6), we consider again open quantum systems whose dynamics are described by Lindblad master equations as in Eq. 2.31 Throughout, \hbar is set to unity and we simulate the systems are using continuous-time Monte Carlo algorithms [30] (see Chapter 4).

6.2.1 Two-Level System

We consider a laser-driven two-level system, which exchanges photons with a radiation bath as in section 4.3.1. Again the counted events are photon emissions and absorptions under the action of the Lindblad operators L_1, L_2

We consider first the zero-temperature case ($\lambda = 0$), for which there is only one jump - described by action under L_1 (photon emission). The large deviation function in this case (from Eq. 3.43) reads

$$g(x) = -3 \ln \left(1 + \frac{x}{2} \right) \quad (6.9)$$

Fig. 6.1(a) provides a numerical test of the Jarzynski relation (6.6) for trajectories with $K = 20$ events. Due to the simplicity of the system, with only one possible outcome for each event, this is sufficient for a good agreement with the large- K limit¹ of the meta-free energy - $g(x)$ (see Fig. 6.1). We have sampled $M = 5000$ trajectories for the forward and backward protocol, where trajectories started from an initial $x_0 = 0$ (equilibrium) state to a final state ranging between $x = -1$ and $x = 1.5$, with $N = 1000$ TPS step moves for each direction. As criterion to stop the BAR iterations, we chose the threshold 10^{-5} for the fractional change of the estimated $g(x)$ between iterations. For this system convergence is reached very fast taking typically 2-3 iterations, and there is a good agreement between the results obtained from the Jarzynski relation and the exact results.

We now consider the finite temperature case with parameters $\gamma = 6\Omega$, $\lambda = 2\Omega$. Here action under both L_1 and L_2 occurs, and so there are two jump possibilities. Fig. 6.1(b) provides a numerical test of the Jarzynski relation in this case. Analyt-

¹Actually, given that there is only one possible event, drawn from one waiting time distribution, then provided we start in the ground state of the system, $K = 1$ is enough to realise the large- K limit of the system.

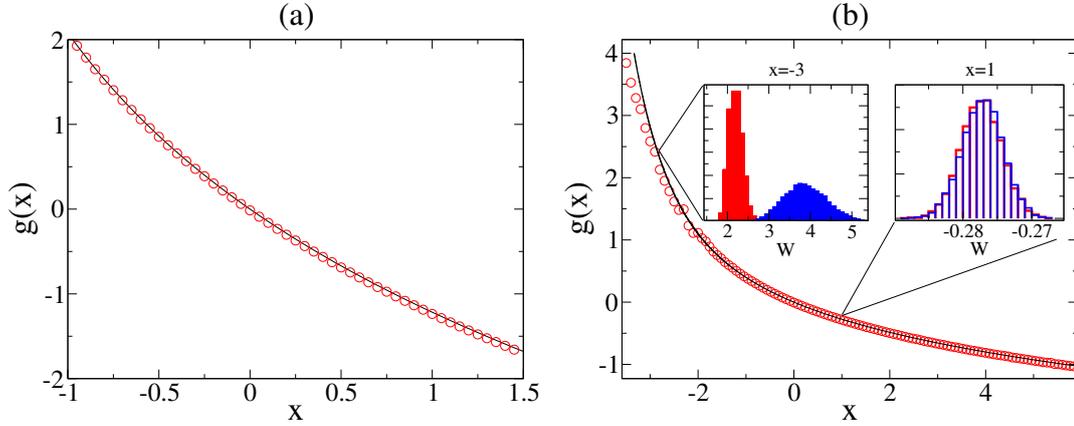


Figure 6.1: 2-level system. (a) Comparison of the meta-free energy $g(x)$ obtained numerically via the trajectory Jarzynski relation (symbols) to the exact analytical result (6.9) (solid line) for a range of x , in the zero temperature case, with $\gamma = 4\Omega$. (b) Same as in (a), but now for the finite temperature case, with $\gamma = 6\Omega$ and $\lambda = 2\Omega$. The statistical error is smaller than the symbol sizes. Insets to (b): Sampled histograms for the meta-work distribution $P_{\uparrow}(W)$ for the forward (red) and $P_{\downarrow}(-W)$ for the backward process (blue), at the two final values of x shown.

ical results are again obtained from the largest eigenvalue of the deformed master operator corresponding to the s -ensemble, and inverted to give the x -ensemble meta-free energy $g(x)$. The exact expression is available but cumbersome and rather unilluminating to be given explicitly. Note that the true $g(x)$ diverges close to $x \simeq -3.5$ [cf. with the zero temperature case, Eq. (6.9), where the limiting value is $x = -2$]. Again, $M = 5000$ iterations were used for trajectories of $K = 20$ events but with now $N = 5 \times 10^5$ TPS step moves for each iteration. There are now two possible outcomes for each event, corresponding to photon emission and absorption. However $K = 20$ still converges to the large- K limit satisfactorily. While there is a good agreement between the results obtained from the Jarzyn-

ski relation and the exact results for a broad range of x_N , we have extended the plotted range of x values to demonstrate that the numerical estimate for $g(x)$ starts to divert from its analytical prediction as we approach the divergence. For $x < 0$ the “pressure” is negative, selecting rare trajectories with large trajectory length τ . Our numerical procedure breaks down because it takes an increasing amount of time to equilibrate the system at the final x for the backward iterations. For the forward-backward protocol, N has to be sufficiently large to generate work distributions that sufficiently overlap in order for Eq. (6.8) to work. This is demonstrated in the inset of Fig. 6.1(b). This is a general feature of the Jarzynski relation. Although in principle it holds for any driving speed and any protocol, application to data requires either to sample extreme work values sufficiently or to generate distributions from forward and backward protocols that overlap.

6.2.2 Micromaser

We again consider the micromaser (see section 4.3.2), which provides a useful test of a pseudo-many-body system, as well as a system with many first-order phase transitions in the x -ensemble. A detailed account of the model can be found in Ref. [94]. The events being counted are the actions under any of the four Lindblad terms.

Despite being a system with a single degree of freedom, the micromaser has a rich dynamical behaviour due to the combination of an infinite dimensional Hilbert space and the non-linear jump operators L_1 and L_2 . In particular, it displays a number of distinct dynamical phases and transitions between them [95, 131, 132]. (Strictly speaking, these are sharp crossovers which only become singular in the limit of $r \rightarrow \infty$; see [74, 131].) As a result it provides a useful testing ground for the trajectory Jarzynski relation.

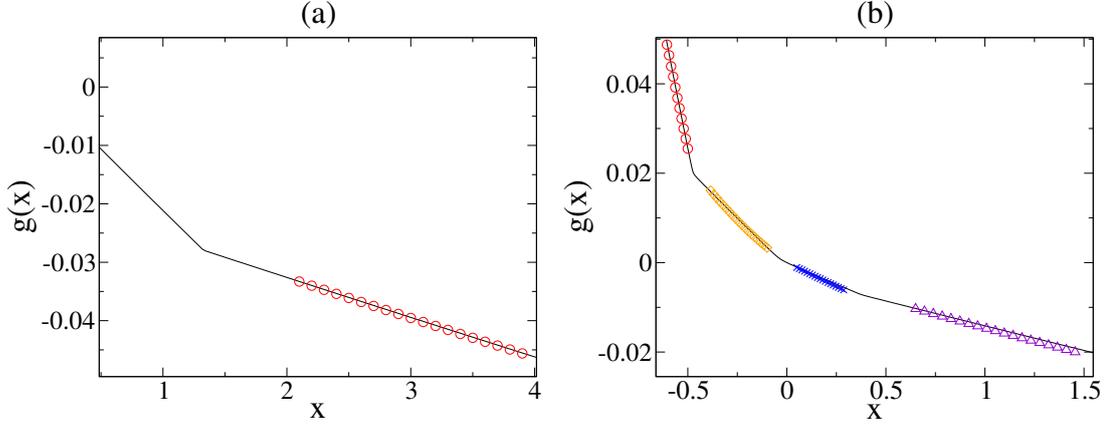


Figure 6.2: Micromaser. (a) Comparison of the meta-free energy $g(x)$ obtained numerically via the trajectory Jarzynski relation (symbols) to results obtained by direct diagonalisation of the master operator (solid line) for $\alpha = 1.2\pi$, where the system is initially equilibrated to $x = 2$. (b) Same as in (a), but now for $\alpha = 4\pi$. Different simulations, equilibrated to different initial values of x are denoted by different symbols.

We first attempt to compute meta-free energy differences within a single phase. Fig. 6.2(a) provides a numerical test of the Jarzynski relation for a pump parameter of $\alpha = 1.2\pi$. The trajectories are initially equilibrated to a non-equilibrium dynamical phase with $x_0 = 2$, and the Jarzynski protocol run for trajectories of $K = 1000$ jumps, with $M = 5000$ iterations and $N = 60000$ TPS step moves per iteration. The computed meta-free energy differences are compared to results obtained from direct diagonalisation of the master operator, as in [95], and a good agreement is found between the two methods. Provided the existence of phases, and the boundaries between them, is known, a complete picture of meta free energy differences can be constructed even when there are multiple dynamical phases. For example, with the pump parameter taking a value of $\alpha = 4\pi$, four distinct phases occur, see Fig. 6.2(b), and $g(x)$ can be computed within phases by

initially equilibrating the trajectories to a value of x within the required phase. Again trajectories of $K = 1000$ jumps were used, with $M = 5000$ iterations and $N = 60000$ TPS step moves per iteration.

6.2.3 Driving across a first-order phase transition

We finally examine the behavior of the Jarzynski relation using a protocol $x_0 \rightarrow x_N$ that crosses a phase boundary x^* at a finite speed. In the quasi-stationary limit of $N \rightarrow \infty$, we obtain from the definition Eq. (3.3) the well-known expression

$$\ln \frac{Z_K(x_N)}{Z_K(x_0)} = \int_{x_0}^{x_N} dx \frac{\partial \ln Z_K(x)}{\partial x} = - \int_{x_0}^{x_N} dx \langle \tau \rangle_x \quad (6.10)$$

for thermodynamic integration, where the subscript emphasizes that the average is calculated from equilibrated trajectories at fixed x . Eq. (6.10) is known to fail in the presence of a discontinuous phase transition, not because the equation is wrong but because of the way a simulation is carried out in practice. Typically, one will apply a small change $x_i \rightarrow x_{i+1}$, let the system relax, and then record data to calculate the average. Crossing x^* , the system will not immediately adapt to the new state but follow the metastable branch due to the cost of nucleating the new stable phase, thus violating the assumption that the calculated mean corresponds to the true equilibrium mean. In the micromaser, sharp crossovers occur at certain values of the biasing field between phases that can be characterised by either their average emission rate, or the closely related expected photon occupation of the cavity [95, 131, 132]. When considering these transitions in the context of the x -ensemble, different phases have significantly different average trajectory lengths for the same fixed number of quantum jumps. Just like in ordinary first-order transitions, pronounced metastability may prevent from estimating meta free-energies accurately with (6.10). This can occur when the transition at x^* is between phases with very different activities. In this case, if trajectories are

prepared in the less active phase (for example starting from $x = 0$ and increasing x), the barrier to nucleate the more active phase when $x > x^*$ can be prohibitive for practical simulation. The nucleation event can be promoted externally, for example by altering the photon occupation of the cavity by temporarily increasing the pump parameter (or similar “parallel tempering”). But without such external interference the timescale for nucleating the new stable phase is often beyond what can be reasonably simulated.

One could hope that the Jarzynski relation, given that it applies to arbitrarily fast non-equilibrium protocols, would provide a way out of this problem since trajectories can be sampled at finite rate for the change in x . In practice, however, even with slow driving speeds it is problematic to compute free energy differences across first-order phase boundaries. Results for the micromaser are shown in Fig. 6.3 (for a pump parameter of $\alpha = 1.2\pi$ and with $\frac{\gamma}{\kappa} = 0.15$ corresponding to a temperature $T = 0.5$). Trajectories with $K = 2000$ jumps were sampled for $M = 5000$ iterations, with $N = 10^6$ TPS step moves for each iteration. For the chosen parameters, the system undergoes a first-order transition at $x^* \simeq 1.34$. The computed free energy difference using the Jarzynski relation gets locked to the phase that is stable for $x < x^*$ but which becomes metastable for $x > x^*$. This is evident by the fact that the computed free energy follows the path of the eigenvalue that dominates for $x < x^*$, but which becomes subdominant at $x > x^*$.

Fig. 6.3(b) shows parameters that are more suitable for a cross-phase meta-free energy computation. The system is set to zero temperature - removing the possibility of the cavity absorbing a photon from the bath and thus simplifying the simulation - and the atom beam rate is reduced - which serves to reduce the difference in the dynamic properties of the two phases [95]. The sampling is also improved by doubling the number of iterations. While there is improvement over the finite temperature case of Fig. 6.3(a), the calculated meta-free energy still

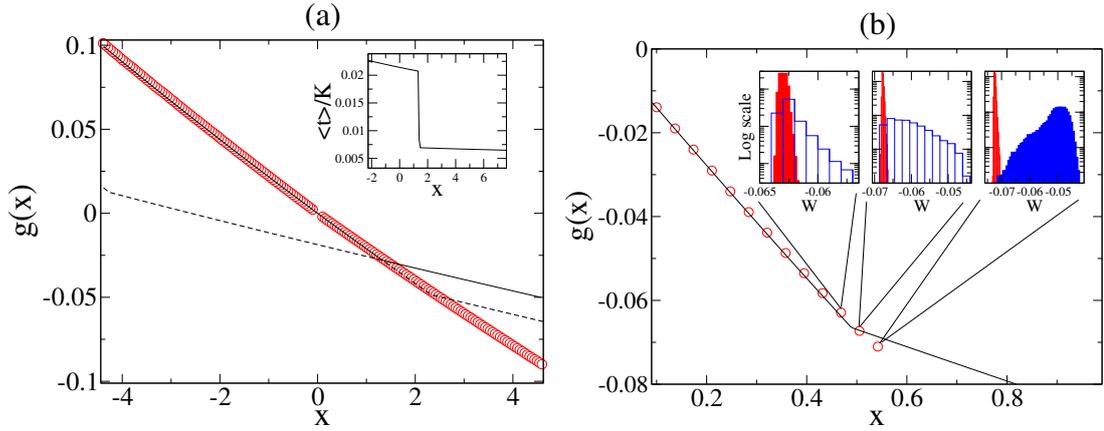


Figure 6.3: Micromaser with cross-phase Jarzynski protocol. (a) Comparison of the numerical meta free energy, $g(x)$, obtained numerically via the Jarzynski relation (symbols) to results obtained by direct diagonalisation (solid line) in a micromaser with pump parameter $\alpha = 1.2\pi$, at a finite temperature ($\gamma/\kappa = 0.15$). The second largest eigenvalue (dashed line) is plotted to illustrate the meta free energy calculation being locked to the metastable branch after the transition. Inset to (a): the expected waiting time per event showing the differing dynamic properties of the two phases. (b) Same as in (a), but now at zero temperature ($\gamma/\kappa = 0$). Insets to (b): sampled meta-work distributions for the forward (red) and backward (blue) process for the three points shown.

deviates from its expected value across the phase boundary. The cause can be understood by looking at the meta-work distributions for the forward and reverse processes, see insets to Fig. 6.3(b). For the conditions shown, the driving is slow enough for the forward and reverse meta-work distributions to overlap immediately before the phase transition. However as the phase boundary is crossed the two become separated. A small residual spike of the reverse distribution lies within the bulk of the forward distribution, corresponding to a small fraction of cases where the reverse process starts in the metastable phase. This occurs precisely because the simulation cannot be done in the “thermodynamic limit” of

$K \rightarrow \infty$ and $r/(\kappa - \gamma) \rightarrow \infty$, i.e. the transition is not strictly a phase transition but a very sharp crossover [131]. Thus when differences in the meta free energy $g(x)$ is computed with the BAR method, it only sees the metastable phase. It is worth noting that these attempts to compute a cross-phase free energy difference took two orders of magnitude more computation than any of the single-phase free energy computations.

6.3 Outlook

The results of this chapter further underpin the thermodynamics approach to dynamics. Not only ensembles of dynamical trajectories can be studied by generalising equilibrium statistical mechanics via large deviation methods, but also non-equilibrium statistical mechanics tools can be generalised and applied to uncover properties of such ensembles. By considering the analogous Jarzynski relation we have shown that the large-deviation function that encodes the properties of one trajectory ensemble can be obtained by considering the statistics of the meta-work performed as the parameter that characterises the ensemble is driven.

A further interesting observation is the following. The general relation between forward and backward processes that underpins most integral fluctuation theorems is a straightforward consequence of probability conservation [58]. Few integral fluctuation relations are “non-trivial” in the sense of conveying actually useful information about the problem studied. This occurs when one can write the stationary distribution in terms of “weights” that encode their functional dependence on the objects that form the ensemble under consideration (usually configurations; trajectories in our case), and a “free-energy”. For ensembles of configurations, these include the Jarzynski relation proper [37, 38] and the Hatano-Sasa relation [55] for

driven stationary states. We note that the class of trajectory ensemble problems we studied here adds to this small group. These are cases where the “normalisation constant” of the stationary probability distribution also has physical meaning, as it is given by the large-deviation function which is the generating function for moments and cumulants of time-integrated observables (and thus play the role of trajectory free-energies).

While the trajectory Jarzynski relation allows for the direct computation of the generating functions for the dynamics of a system, it is unclear if it will be a computationally efficient method of doing so. There are doubts about whether the conventional Jarzynski relation is more efficient at calculating free energy differences than other methods (see for example Refs. [133,134]), and while there are proposed methods that show some potential (for example Ref. [135]), in general it is not clear if Jarzynski-based methods are optimal [136]. This, coupled with the issues in crossing a phase boundary, mean the trajectory Jarzynski relation is unlikely to be the optimal method of calculating meta-free energy differences. Never the less the importance of the Jarzynski relation to non-equilibrium statistical mechanics makes the investigation of its dynamical analogue worthwhile. It also serves to emphasise the analogy between dynamical ensembles and thermodynamics and provide a basis for exploring other useful thermodynamic relations in the context of dynamical ensembles.

7. DYNAMICAL LEE-YANG ZEROS IN THE x -ENSEMBLE

Although ensembles in the thermodynamics of trajectories are formally equivalent to equilibrium statistical mechanics, trajectory phases are in general difficult to probe either experimentally or numerically. The problem is two fold: firstly, these trajectory phase transitions occur in the limit of long times or large K which may be difficult to access in practice. Secondly, and perhaps more importantly, the “counting fields” which drive these trajectory phase transitions are generally not simply related to physically accessible parameters [12, 13, 16, 19, 137–140] and so are hard to tune in either an experiment or simulation [17]. While transition path sampling has been shown to be an effective numerical technique for exploring dynamical phase spaces, it can still be problematic and computationally expensive to implement. One can imagine situations where one does not require the full dynamical phase diagram of a system, but rather the locations of any dynamical phase transitions (for example one could potentially combine the location of phase space singularities with the trajectory Jarzynski relation shown in Chapter 6 to efficiently compute the generating function of a system).

A potential solution to this question is based on a generalization of the Lee-Yang theory [39, 40] of phase transitions to these dynamical systems [41–43]. Using this

approach one may extract the zeros of the moment generating function (MGF) using its short time high-order cumulants [64–66]. In equilibrium statistical mechanics, the zeros of the partition function are complex for a finite system, but move towards real values of the order parameter when the system size is taken to infinity. Analogously in the dynamical case, the zeros of the MGF move towards real values of the fields s, x when trajectory observation time is taken to infinity.¹ Thus the location of trajectory transition points can be extrapolated by tracking the movement of these complex zeros for short, but increasing, time. This approach was applied successfully to the s -ensemble in both classical stochastic systems, which had both 1st order [42] and continuous trajectory phase transitions [43], and open quantum systems [141]. Motivated by the success of this approach to the s -ensemble we extend it to the x -ensemble (i.e. the ensemble of trajectories where the overall time of a trajectory is not fixed, see Chapter 3) and use it to identify trajectory transition points using high-order cumulants of the overall time length of trajectories at small values of K . The efficiency of the x -ensemble in simulating a range of systems, as shown in Chapter 4 potentially makes it the ensemble of interest for identifying the trajectory transition points of a system.

In this chapter we develop an analogous dynamical Lee-Yang theory for the x -ensemble and apply it to an open quantum system, the “blinking” 3-level system, as well as the 1d East model of section 4.3.3. The 3-level system displays dynamical intermittency [18, 30, 142] which was shown recently to be the result of two complex conjugate transition points lying close to the origin in the complex s plane, when analysed within the s -ensemble framework [141]. Contrastingly, the x -ensemble zeros do not converge to any location in the complex x plane indicating intermittency does not manifest as complex conjugate transition points in this

¹Strictly speaking both system size and time should be taken to infinity for this convergence.

trajectory ensemble. An analytic expression for the x -ensemble LD function is extracted via exact diagonalization and is confirmed not to possess any trajectory transition points. The emergence of complex singularities in one ensemble and not the other highlights that although the s - and x -ensembles encode equivalent information the former ensemble couples to a dynamical observable which captures this intermittency more directly.

7.1 Dynamical Lee-Yang zeros of the x -ensemble

Taking inspiration from the ideas of Lee and Yang [39, 40], we focus on the dynamical zeros ($x_j(K)$) of the MGF given in Eq. 3.5 as a function of the number of configuration changes K . Because a trajectory phase transition corresponds to the crossing of dominant eigenvalues of the tilted transition operator, see Chapter 3 and Eqs. 3.10, 3.26, and 3.34, when we are close to a transition point $x = x_c$, we can approximate the MGF by the two eigenvalues of the transfer matrix with the largest real parts of the transition operator, cf. Eq. 3.9

$$Z_K(x) \simeq c_0(x)e^{K\lambda_0(x)} + c_1(x)e^{K\lambda_1(x)}. \quad (7.1)$$

Here the initial conditions are encoded in the coefficients $c_0(x)$ and $c_1(x)$, however these are unimportant in the large K limit. Rearranging Eq. (7.1) we find the zeros of the MGF are given by

$$\lambda_0(x) = \lambda_1(x) + \frac{\log[c_1(x)/c_0(x)] + i2\pi n}{K}, \quad (7.2)$$

where $n \in \mathbb{Z}$. In the large K limit the second term of Eq. (7.2) vanishes, indicating that with increasing K the zeros of the MGF will converge to where $\lambda_0(x) = \lambda_1(x)$, i. e. a trajectory transition point x_c .

Here we note that the difference between dynamical Lee-Yang zero approach of Refs. [42, 43, 141] is that we are using the total number of configuration changes

K to define our ensembles. This quantity is a non-negative integer and thus the motion of the Lee-Yang zeros in the complex x plane with increasing K will not be smooth. Having established a form of dynamical Lee-Yang theory for fluctuating time ensembles, we will relate these Lee-Yang zeros to the high-order cumulants of the observation time. Using the Hadamard factorization theorem [143] we rewrite the MGF in terms of the Lee-Yang zeros:

$$Z_K(x) \propto \prod_j \left[\frac{x_j(K) - x}{x_j(K)} \right], \quad (7.3)$$

where $x_j(K)$ denotes the j th zero of the MGF after K events.

For finite K and real x , the MGF is both real and positive implying that the Lee-Yang zeros appear in complex conjugate pairs. Ignoring any analytic prefactors in Eq. (7.3) the cumulant generating function (CGF) may be approximated as: $G(x, K) \simeq \sum_j \log(x_j(K) - x) - \log(x_j(K))$, hence the Lee-Yang zeros of the MGF appear as logarithmic singularities of the CGF (as expected). According to Darboux's theorem [144,145] these singularities will determine the functional form of the high-order cumulants of the observation time. Differentiating the CGF with respect to x , we find the zeros of the MGF and high-order cumulants are related via [64–66]

$$\langle\langle \tau^m \rangle\rangle \simeq (-1)^{(m-1)} (m-1)! \sum_j \frac{e^{-im \arg[x_j(K)]}}{|x_j(K)|^m}, \quad (7.4)$$

where we have written the complex zeros in polar notation, $x_j(K) = |x_j(K)|e^{-i \arg[x_j(K)]}$, and as before $\langle\langle \cdot \rangle\rangle$ indicates cumulants, eg. $\langle\langle \tau^2 \rangle\rangle = \langle \tau^2 \rangle - \langle \tau \rangle^2$, etc.

Considering high-order cumulants, that is those with large orders m , the sum in Eq. (7.4) is dominated by the pair of Lee-Yang zeros closest to the origin, which we label $x_0(K)$ and $x_0^*(K)$. Thus we make a further approximation that the high-order cumulants are given solely by the contribution from this pair [42,43,64–66,145,146],

$$\langle\langle \tau^m \rangle\rangle \approx (-1)^{(m-1)} (m-1)! \frac{2 \cos[m \arg x_0(t)]}{|x_0(t)|^m}. \quad (7.5)$$

The key observation is that this approximation allows us to determine the position of the leading Lee-Yang zero pair from high-order cumulants of the observation time. Formally this pair is found by solving the matrix equation [42, 43, 66, 147],

$$\begin{bmatrix} 1 & -\frac{\kappa_m^{(+)}}{m} \\ 1 & -\frac{\kappa_{m+1}^{(+)}}{m+1} \end{bmatrix} \cdot \begin{bmatrix} -(x_0 + x_0^*) \\ |x_0|^2 \end{bmatrix} = \begin{bmatrix} (m-1)\kappa_m^{(-)} \\ m\kappa_{m+1}^{(-)} \end{bmatrix}, \quad (7.6)$$

where the entries of the matrices are given by the ratios of the cumulants

$$\kappa_m^\pm(K) \equiv \frac{\langle\langle \tau^{m\pm 1} \rangle\rangle}{\langle\langle \tau^m \rangle\rangle}. \quad (7.7)$$

From Eq. (7.6) we will extract the leading Lee-Yang pair, x_0 and x_0^* , and with increasing K these zeros will converge to the trajectory transition points closest to the origin. We will now apply this approach to a classical stochastic model of a glass former that has a 1st order trajectory transition at the origin and a dissipative open quantum system that displays dynamical intermittency.

7.2 East Model

To exemplify the ideas described in the previous section, we study the trajectory phase transitions in the fluctuating time ensembles of the 1d East Model [27]. From the s -ensemble it is known that the glassy dynamics of this model is a consequence of a dynamical phase transition occurring at $s_c = 0$ between phases of distinct activity. This first order dynamical transition is between an equilibrium liquid phase (or active phase) and an inactive glassy phase, the glassy dynamics then results from the system effectively existing close to phase coexistence between these regimes. In Chapter 4 we show this phase coexistence picture holds when one considers the x -ensemble as well, see Fig. 7.1(a), where the transition between the glassy inactive phase and liquid active phase also occurs at $x_c = 0$. We now

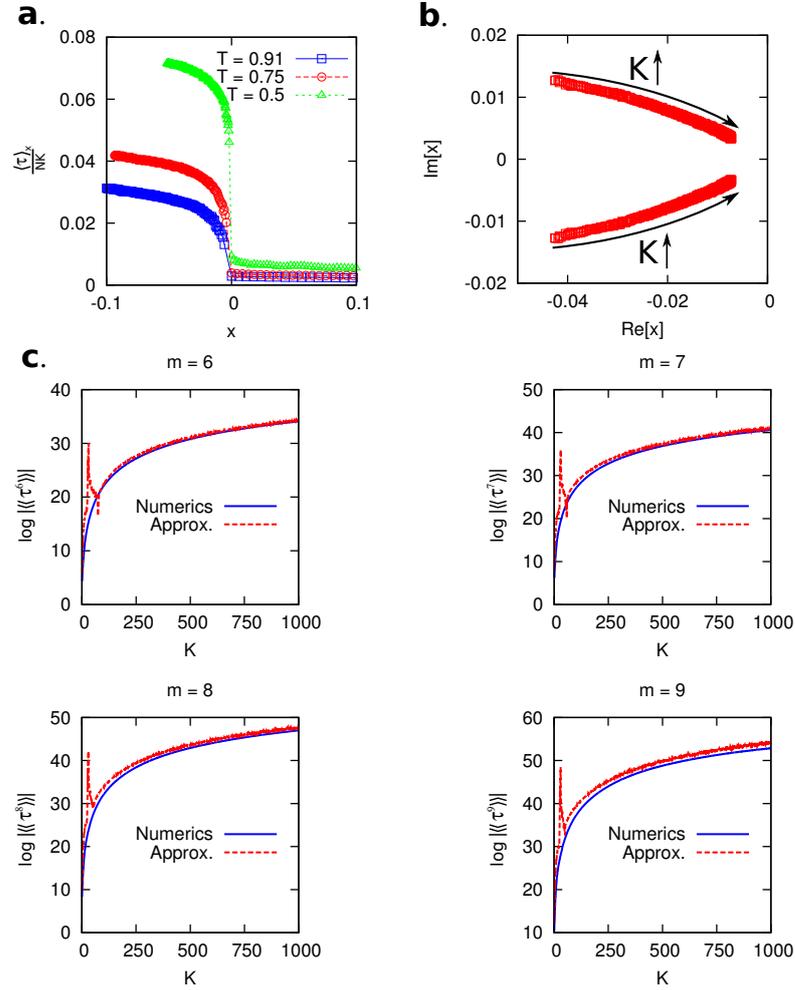


Figure 7.1: East Model results for $N = 60$ spins. (a) The observation time activity as a function of x displays a discontinuity at $x = 0$. This marks a first order transition from an inactive glassy phase, $x < 0$, to an active liquid phase. (b) The leading Lee-Yang zero pair extracted from the high-order cumulants. These converge to points close to the origin with increasing K , this highlights one infer the presence of a trajectory transition point at the origin in the limit of $K \rightarrow \infty$ from these small K cumulants. (c) The approximated cumulants (dashed lines) are found to match their numerically extracted values (full lines) for small K , confirming the validity of the approximation.

show why this may be expected by considering the relationship between the LD functions $\theta(s)$ and $g(x)$. Differentiating Eq. (3.14) with respect to x twice we find

$$g''(x) = \frac{-\theta''(g(x))}{[\theta'(g(x))]^3}. \quad (7.8)$$

First and second order trajectory transition points are marked by diverging $g''(x)$ or $\theta''(s)$ depending on the ensemble of trajectories one is considering. We know that $g(0) = 0$ and $\theta''(0) \rightarrow \infty$ due to probability conservation and the presence of the first order trajectory transition at $s_c = 0$ respectively. Combined with a finite activity at $s = 0$, i. e. $-\theta'(0)$ is finite, one can readily see that $g''(0)$ will diverge and therefore $x_c = 0$ is a trajectory transition point in the fluctuating time trajectory ensemble.

Having understood the trajectory phase behaviour of this model within the x -ensemble we apply the proposed method of Sec. 7.1 to extract the trajectory transition point using the observation time cumulants for fixed jump number K . Figure 7.1(c) shows high-order observation time cumulants as a function of K extracted directly from simulation (full lines). Using the approximation in Eq. (7.5) we extract the leading Lee-Yang zero pair, $x_0(K)$ and $x_0^*(K)$, from these cumulants and plot them as a function of K in Fig. 7.1(b). These zeros are plotted for $N = 60$ and $T = 0.5$ and are found to approach the origin with increasing jump number K . As the trajectory phase transition only occurs in the limit of large N and K these zeros converge to a pair of complex conjugate transition points close to the origin. In the thermodynamic limit these transition points merge and form the transition point $x_c = 0$. However, already for finite N and K one can see evidence of a trajectory phase transition at $x_c = 0$.

To test the validity of the approximation we use the extracted Lee-Yang pair to estimate the cumulants via Eq. (7.5) and compare it with the numerical cumulants.

These estimated cumulants (dashed lines) are plotted in Fig. 7.1(c) and show that over the K range of interest the approximation fits the numerically extracted cumulants well. This confirms that we are indeed extracting the leading Lee-Yang zero pair of the system. However, going to even larger K values the approximation in Eq. (7.5) begins to break down as higher order zeros begin to move towards the origin and contribute to the cumulants. We demonstrate here that one may be able to infer the position the trajectory transition points before this occurs; if one were to use even higher order cumulants this method is also expected to give better results.

7.3 Dissipative Three-level system

We now consider a driven dissipative 3-level quantum system as shown in Fig. 7.2(a). This system is driven by two resonant lasers with Rabi frequencies Ω_1 and Ω_2 respectively which results in a system Hamiltonian:

$$H = \sum_{i=1}^2 \Omega_i (c_i + c_i^\dagger), \quad (7.9)$$

where $c_i = |0\rangle\langle i|$ and $c_i^\dagger = |i\rangle\langle 0|$. In addition to this Hamiltonian the system also possesses a single decay channel $|1\rangle \rightarrow |0\rangle$ with an associated decay rate κ . This manifests as a single jump operator in the systems evolution,

$$L_1 = \sqrt{\kappa} c_1, \quad (7.10)$$

and hence the value of i in Eq. (2.31) is fixed as 1. In this section we apply the x -ensemble approach to the emission waiting times from this system. In particular we are interested in the parameter regime where the transition timescales of the $|0\rangle \rightleftharpoons |1\rangle$ processes are much shorter than the $|0\rangle \rightleftharpoons |2\rangle$ transition timescales, i. e. $\Omega_1, \kappa \gg \Omega_2$. In this parameter regime the photon emission trajectories are

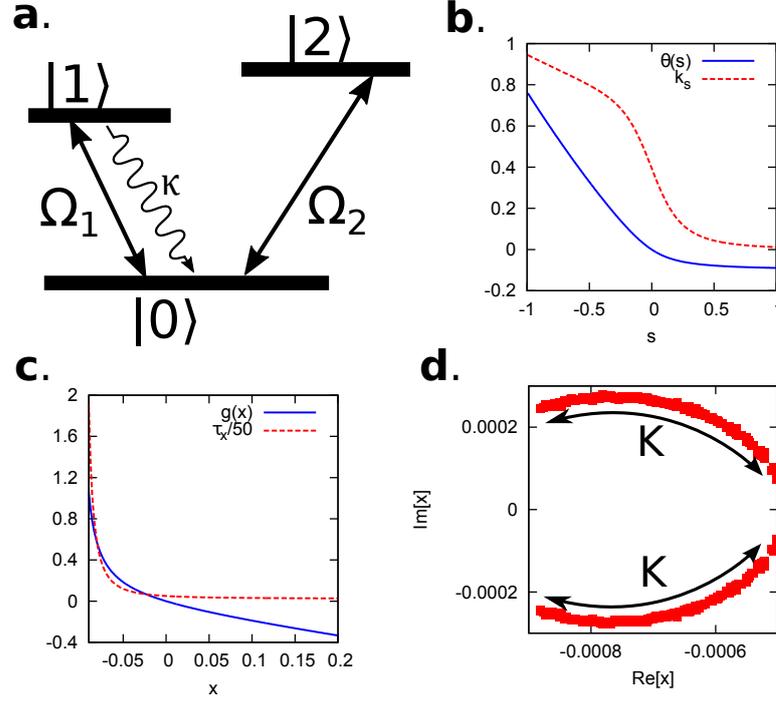


Figure 7.2: We consider $\kappa = 4\Omega_1$ and fix $\Omega_1 = 1$ and $\Omega_2 = 0.15$. (a) Schematic of a dissipative 3-level system. The system is driven by two lasers of frequency Ω_1 and Ω_2 which drive transitions from $|0\rangle \rightarrow |1\rangle$ and $|0\rangle \rightarrow |2\rangle$ respectively. There is a single decay channel $|1\rangle \rightarrow |0\rangle$ with an associated decay rate κ . (b) The large deviation function $\theta(s)$ and associated dynamical activity $k(s)$ as a function of s . A rounded crossover is seen at $s \sim 0$ where the system changes from behaving as an active two-level system comprised of the states $|1\rangle$ and $|0\rangle$, i. e. an active phase, to a photon inactive phase. (c) The large deviation function associated with the fluctuating time trajectory ensemble along with the observation time is plotted as a function of x . Although the system is intermittent this does not manifest as a crossover as in the s -ensemble case. (d) Dynamical Lee-Yang zeros of the fluctuating time ensemble MGF extracted from cumulants of order $m = 6, 7, 8, 9$. These are found to oscillate with increasing K and do not converge to a single point. This is indicative of a lack of a trajectory transition point existing in the complex x plane.

composed of large temporal regions where the system is very photon active and large inactive regions where few photons are emitted [30, 142]. This intermittency manifests as a dynamical crossover in the activity $k(s)$ as a function of s [18]. Due to the small scale of the system the LD function $\theta(s)$ and the corresponding activity may be extracted via exact diagonalization and are shown in Fig. 7.2(b). The rounded crossover marks a transition from a highly active phase for $s \lesssim 0$, where the dynamics is dominated by the transition channel $|1\rangle \rightarrow |0\rangle$, to an inactive phase ($s \gtrsim 0$) where the dark state $|2\rangle$ dominates.

Recently it has been shown that this crossover is the result of a complex conjugate pair of transition points in the complex s plane [141]. Moreover the crossover was shown to sharpen as the transition points moved towards the real s axis, i. e. in the limit $\kappa \rightarrow 0^+$. We now ask two questions: firstly does the dynamical intermittency manifest in the x -ensemble as a crossover and secondly do complex x transition points exist, similar to the s -ensemble case. To answer these questions we begin by diagonalizing the transfer matrix, see Eq. (3.34), to find the dynamical free energy

$$g(x) = \log \left[\frac{4\kappa\Omega_1^2(x(x(2x + \kappa) + 2\Omega_1^2) + 2(x + \kappa)\Omega_2^2)}{(x(x(x + \kappa) + 4\Omega_1^2) + 4(x + \kappa)\Omega_2^2) \dots (4x^3 + 4x^2\kappa + 2\kappa\Omega_1^2 + x(\kappa^2 + 4(\Omega_1^2 + \Omega_2^2)))} \right] \quad (7.11)$$

We note that this function is not well-defined for all values of $x \in \mathbb{R}$. From Eq. (3.14) we see that the domain of $g(x)$, i. e. $[x_l, x_u]$, is simply the range of $\theta(s)$:

$$x_l = \min_{s \in \mathbb{R}} \theta(s), \quad (7.12)$$

$$x_u = \max_{s \in \mathbb{R}} \theta(s).$$

Furthermore it is straightforward to relate the observation time activity to the photon activity,

$$g'(x) = \frac{1}{\theta'(g(x))}. \quad (7.13)$$

From the above Equation it is clear that a crossover in $\theta'(s)$ may not necessarily result in a crossover in $g'(x)$ due to the nontrivial functional form of $g(x)$. We evaluate the LD function $g(x)$ and find no crossover in the observation activity despite the system being intermittent, see Fig. 7.2(c).

However this does not necessarily imply no transition points are present in the complex x plane. To show that the susceptibility $g''(x)$ may still be a smooth function of real x but possess complex trajectory transition points we factorize the MGF in terms of its zeros

$$Z_K(x) \simeq \prod_j \frac{e^{-x_j(K)} - e^{-x}}{e^{-x_j(K)}}, \quad (7.14)$$

where we once again denote the Lee-Yang zeros as $x_j(K)$. From this we approximate the dynamical free energy as

$$g(x) = \lim_{K \rightarrow \infty} \frac{\log Z_K(x)}{K} \quad (7.15)$$

$$\simeq \log\left(\frac{e^{-x_c} - e^{-x}}{e^{-x_c}}\right) + \log\left(\frac{e^{-x_c^*} - e^{-x}}{e^{-x_c^*}}\right), \quad (7.16)$$

where we simply replace the Lee-Yang zeros by the complex transition point pair x_c and x_c^* . Furthermore we've assumed that in the limit of large K the diverging Lee-Yang zero density at these transition points is compensated by the factor of $1/K$. From this simple approximation we find the susceptibility at $x = \text{Re}[x_c]$ is given by

$$g''(x = \text{Re}[x_c]) \simeq \frac{1}{2 \sin^2(\text{Im}[x_c])/2}. \quad (7.17)$$

This equation implies that although there would be a peak in the susceptibility when $\text{Im}[x_c] \ll 1$, corresponding to a crossover, in general this may not be present despite complex transition points being present. Such a crossover becomes a sharp discontinuity when the transition points move onto the real x axis, i. e. $\text{Im}[x_c] \rightarrow 0$. Having concluded that despite the lack of a crossover complex transition points

may still be present we apply our Lee-Yang zero method described in Sec. 7.1. The high-order observation time cumulants are calculated using Quantum Jump Monte Carlo [30] as in Chapter 4. Here we recall that the motion of the Lee-Yang zeros relies only on the finite K cumulants at $x = 0$ and so are potentially experimentally accessible.

The leading Lee-Yang zero pair $x_0(K)$ and $x_0^*(K)$ are extracted using cumulants of order $m = 6, 7, 8, 9$ in a parameter regime where the system is intermittent, e. g. $\Omega_1 = 1, \kappa = 4$ and $\Omega_2 = 0.15$. In Fig. 7.2(d) we show the motion of this pair as a function of K and find that they do not converge to any complex pair of transition points, instead their positions oscillate as a function of K . This lack of convergence indicates that the Lee-Yang approach of Sec. 7.1 is not applicable to the x -ensemble of this model and $g(x)$ does not possess any complex transition points as well as no crossover as a function of $x \in [x_l, x_u]$.

However in Ref. [141] it was shown that in this intermittent regime there exists a pair of complex conjugate s -ensemble transition points close to the origin, s_c and s_c^* . At these transition points the dynamical susceptibility diverges, i. e. $\theta''(s_c^{(*)}) \rightarrow \infty$. This susceptibility is related to the x -ensemble activity and susceptibility via

$$\theta''(s) = -\frac{g''(\theta(s))}{[g'(\theta(s))]^3}, \quad (7.18)$$

one would then naively expect there to exist a pair of x -ensemble transition points to exist at $x_c^{(*)} = \theta(s_c^{(*)})$. In this instance we instead find $g''(\theta(s_c^{(*)}))$ is finite but $g'(\theta(s_c^{(*)})) \rightarrow 0$. Thus a transition point in the s -ensemble may not result in an equivalent transition point in the x -ensemble but a point of zero activity. We note that the observation time activity $g(x)$ is finite for all finite real $x \in [x_l, x_u]$. This means a s -ensemble transition point $s_c \in \mathbb{R}$ results in a real x -ensemble transition point at $x_c = \theta(s_c)$ through Eq. (7.18), as seen in the East model studied in Sec. 7.2.

We have shown that there is not necessarily a one-to-one mapping between complex transition points of the s - and x -ensembles. For completeness we confirm that $g(x)$ possesses no complex x transition points analytically. Differentiating Eq. (7.11) twice with respect to x we identify the points at which the susceptibility $g''(x)$ diverges to be the zeros of a high-order polynomial in x . Thus in our parameter regime of interest ($\Omega_1 = 1, \kappa = 4$ and $\Omega_2 = 0.15$) we identify the transition points

$$x_c = \{-2.26285, -1.9515, -1.64015, -1.22567, \\ -1.17993, -0.868574, -0.663688, -0.110638, -0.096998\}. \quad (7.19)$$

Crucially all of these transition points lie outside the domain of $g(x)$ (with the smallest lying on the domain boundary of $g(x)$) and hence the cumulants and extracted zeros are not influenced by their presence. Thus the lack of convergence of the Lee-Yang zeros confirms the lack of x -ensemble transition points for this model despite intermittency. This highlights that although equivalent information is captured in both the s - and x -ensembles, in the former signatures of the intermittent stationary state are more apparent. This is analogous to equilibrium thermodynamics where coupling to the “correct” order parameter reveals information about the system more readily than other order parameters would.

While there are no peaks in the fixed- K susceptibility, $g''(x)$, indicating a lack of a crossover in the x -ensemble, it is important to note that physical quantities of interest exhibit the same behaviour in both ensembles. Of particular interest, the Mandel Q parameter - a measure of the degree of photon bunching/antibunching - can be defined for both the s -ensemble,

$$Q_s = \frac{\langle K^2 \rangle_s - \langle K \rangle_s^2}{\langle K \rangle_s} - 1 \quad (7.20)$$

and x -ensemble. [142]

$$Q_x = \frac{\langle \tau^2 \rangle_x - 2\langle \tau \rangle_x^2}{\langle \tau \rangle_x^2} \quad (7.21)$$

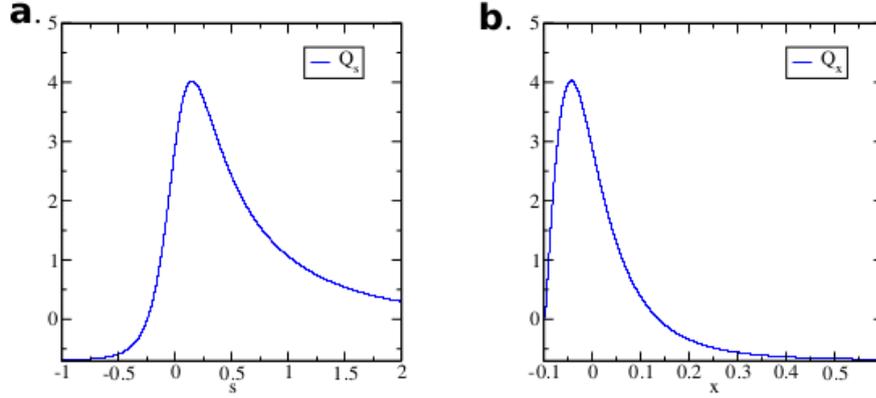


Figure 7.3: We consider $\kappa = 4\Omega_1$ and fix $\Omega_1 = 1$ and $\Omega_2 = 0.15$. Plots of the Mandel Q Parameter for (a) the s -ensemble and (b) the x -ensemble. Despite the information being encoded in different ways, the same physical properties are realised by both ensembles

In both ensembles the Q parameter is peaked close to $s/x = 0$ on the inactive side ($x < 0, s > 0$), with the same limiting behaviour ($Q = 0$ in the low-activity limit, $Q = -2/3$ in the high-activity limit, see Fig. 7.3). This illustrates that while the two ensembles encode information about the system in different ways, they correspond to the same physical dynamics.

7.4 Outlook

While this chapter demonstrates the applicability of the Lee-Yang method of locating trajectory transition points to the x -ensemble, the most notable outcome is where it fails in the quantum three-level system. It serves as an important reminder that while different trajectory ensembles encode the same physical dynamics, the manner in which the information is encoded is critical to what is uncovered. While the s -ensemble reveals a dynamical phase transition in the complex s plane, it is a transition to a state with zero dynamical activity (i.e. $\theta'(s) = 0$). The x -ensemble,

by its very construction, is incapable of encoding trajectories with zero activity and so no trajectory transition points are found. This can be of great consequence to the understanding of dynamical behaviours, as even systems where trajectories in the long-time limit can not physically realise zero activity can have their high order cumulants (and therefore the tails of their distributions) influenced by complex transition points to states of zero activity. This is precisely the case in the quantum three-level system.

Since it is true that the x -ensemble can not encode trajectories with precisely zero activity, the inverse must be true for the s -ensemble. A trajectory with zero activity simply has $K = 0$ in the s -ensemble, while the order parameter diverges in the x -ensemble, $\tau \rightarrow \infty$. Conversely, states with infinite activity have $\tau = 0$ while K diverges. While a dynamical system with infinite activity appears inherently unrealisable, there are two points to note. Firstly the presence of a transition point in the complex x plane can influence the high-order cumulants of the system in an analogous fashion to the three-level system. Secondly this is inherently an issue of an order parameter diverging, and one might choose an ensemble whose order parameter can physically diverge, as in the x -ensemble.

This chapter then serves as a useful illustration of the need for, and benefit of exploring a variety of dynamical ensembles. Thermodynamics has a long history of selecting the ensemble most suited to the problem at hand, and with care one can navigate the shortcomings of any particular ensemble. While the thermodynamics of trajectories has thus far been focused on the s -ensemble, it is potentially beneficial to take inspiration from thermodynamics and explore a variety of alternate ensembles.

8. CONCLUSIONS AND FUTURE WORK

The key message of this thesis is the benefits that can be reaped by exploring alternative dynamical ensembles. While observing a system for a set amount of time and counting events of importance is the most natural method of defining a trajectory ensemble - with intuitive analytical structure - flexibility in mathematical descriptions of physical systems is an important endeavour. The analytical structure of fixed time (s -) ensembles easily extends from master equation descriptions of dynamics, whereas the fluctuating time (x -) ensembles we introduced in Chapter 2 adapt far more naturally to the structure of computer simulations. Thermodynamics has a long history of selecting the most appropriate ensemble to the task at hand, with a wide variety of order parameters being used to define thermal phase spaces, and we have explored similar, potentially beneficial, situations in the thermodynamics of trajectories. Recent studies using more exotic order parameters (for example see [148]) have also uncovered further depth in systems already studied by the s -ensemble, emphasising the benefit of this flexibility. Furthermore, the ensemble correspondence we illustrated allows for the results obtained in the x -ensemble to be converted to the s -ensemble and vice versa, and more recent work in Ref. [149] has demonstrated the full equivalence of these ensembles. This further illustrates the scope of flexibility allowed by alternative

dynamical ensembles.

s -ensemble path sampling schemes have struggled in situations where time-reversal symmetry is broken, as they are dependent on being able to produce time-reversed dynamics. Even when this is not the case they provide limited ability to control acceptance rates - which is of key importance to the efficiency of Metropolis-Hastings algorithms. We explored an alternative x -ensemble TPS scheme which was shown to provide significant computational benefits, including better algorithmic scaling, in a number of stochastic systems, both quantum and classical. Furthermore this scheme shows the potential to develop into a self-tuning algorithm that precisely controls acceptance rates to maximize efficiency. Since TPS often involves some trial and error to find the optimal methods for any given system, a self-tuning algorithm would speed up initial TPS investigations of new systems.

After this exploration of x -ensemble TPS we turned to plaquette models of glasses, where the efficiency provided by x -ensemble scheme allowed the exploration of the dynamical phase space of a three-dimensional glass former. It was found to be strikingly similar to the thermodynamic one, and while we do not believe they are directly related, we argue both are intrinsically related to long-lived metastable states. Since metastability and dynamical heterogeneities are a universal features of glass formers this is a good point to try and resolve competing theories of glassiness. Bringing the argument firmly onto a three-dimensional footing where both dynamical and thermodynamic arguments can be applied is an important step. In general, the thermodynamic arguments (such as RFOT) are compelling in the mean-field limit (i.e. infinite dimensions) where they are exact. In contrast, kinetically constrained dynamics is supposed to describe the dynamics in low (two and three) dimensions where dynamical fluctuations are prevalent. Our three-dimensional version of a plaquette model (the SPyM model) provides a good test bed for the study of thermodynamical behaviour of coupled systems, in the spirit

of RFOT, together with the study of dynamical behaviour associated to facilitated dynamics. The models studied in Chapter 5 are therefore particularly useful to the discussion of the nature of the glass transition. Utilising a highly optimised TPS scheme to locate the dynamical critical point in the SPyM (if there is one as we assume) will be an interesting endeavour, as sampling the distribution close to a critical point is difficult. We are hopeful the self-tuning x -ensemble algorithm described in Chapter 4 will provide the capability to accomplish this.

We then looked to expand upon the thermodynamics of trajectories as a whole with a trajectory analogue of the Jarzynski equality. This allowed for the computation of dynamical free energies from the “meta-dynamics” through the space of trajectories conducted by TPS. While this is unlikely to be the most efficient means of calculating these LD functions, it is a useful demonstration of an important relation in non-equilibrium statistical physics. We expect other important relations from statistical physics will be similarly developed for the thermodynamics of trajectories. We further comment on the possibility that a combination of numerical techniques, perhaps the histogram reweighting methods used in Chapter 5, will enable a cross-phase meta-free energy calculation using the Jarzynski protocol, as this is a problem that has received comparatively little attention. Failing that, the Lee-Yang zeros method studied in Chapter 7 can be used to locate trajectory phase transition points, and single-phase free energy calculations used instead to capture the full LD functions. This method of Lee and Yang also provided a useful reminder of the importance of careful selection of ensembles. Studying a “blinking” quantum three-level system revealed no trajectory phase transitions in the x -ensemble, where prior studies in the s -ensemble had found complex s transition points. While these two ensembles are equivalent in the “thermodynamic limit” (with some small technical nuances, see Ref. [149]) of $K \rightarrow \infty$, $\tau \rightarrow \infty$, the transition point previously found in the s -ensemble is a transition to a dynamical state of zero activity. As such, information about this state can not be encoded in the

x -ensemble. We argue that, if anything, this further stresses the need for a variety of dynamical ensembles constructed from different order parameters. One can then analyse if an order parameter is likely to capture any behaviours of interest in a system under study.

Finally we close with a general comment on the thermodynamics of trajectories as a whole. That the mathematical formalism of a centuries old theory can be so usefully adapted to a fundamentally new setting is a testament to the power of statistical mechanics. We look forward to further development of these ideas, with the hope that non-equilibrium physics can be placed on equally solid mathematical footing as its equilibrium counterpart.

A. MAPPING COUPLED TPMS TO SINGLE SYSTEM IN FIELD

In [36], a mapping was derived between the free energies of the coupled system (5.5) and a single plaquette model in a magnetic field. Here, we present a mapping between (sets of) configurations of these systems, which extends that analysis, as well as recovering the same mapping between free energies.

We introduce overlap variables $q_i = \sigma_i^a \sigma_i^b$ on each site: our aim is to calculate the statistical weight of a particular configuration of these variables. This weight is

$$W_2(q|J, \varepsilon) = \sum_{\sigma^a, \sigma^b} e^{-\beta E_{J, \varepsilon}(\sigma^a, \sigma^b)} \prod_i \delta(q_i - \sigma_i^a \sigma_i^b). \quad (\text{A.1})$$

We now perform the sum over the σ variables. If we sum over σ^b first we obtain,

$$W_2(q|J, \varepsilon) = \sum_{\sigma^a} \exp \left[\frac{\beta J}{2} \sum_{\mu} \sigma_{i_{\mu}}^a \sigma_{j_{\mu}}^a \cdots \sigma_{k_{\mu}}^a \right. \\ \left. \times (1 + q_{i_{\mu}} q_{j_{\mu}} \cdots q_{k_{\mu}}) + \beta \varepsilon \sum_i q_i \right].$$

For the summation over σ^a we replace $\sigma_{i_{\mu}}^a \sigma_{j_{\mu}}^a \cdots \sigma_{k_{\mu}}^a$ by τ_{μ}^a . Then we use the characteristic feature of the model, that plaquette and spin configurations are in a one-to-one correspondence, so we replace the sum over the σ_i^a with a sum over

the τ_μ^a .

$$W_2(q|J, \varepsilon) = \sum_{\tau^a} \exp \left[\frac{\beta J}{2} \sum_{\mu} \tau_\mu^a (1 + q_{i_\mu} q_{j_\mu} \cdots q_{k_\mu}) + \beta \varepsilon \sum_i q_i \right],$$

Performing the sum, we arrive at

$$W_2(q|J, \varepsilon) = (4 \cosh \beta J)^{N/2} \cdot e^{-\beta E_{J'}(q) + \beta \varepsilon \sum_i q_i}. \quad (\text{A.2})$$

with

$$\beta J' = \log \cosh(\beta J) \quad (\text{A.3})$$

We recognise the exponential term in (A.2) as the statistical weight of a configuration $\sigma = q$ for a single plaquette model with energy scale J' , in a magnetic field $h = \varepsilon$.

To explore the consequences of this property for the free energy, we observe $Z_2(J, \varepsilon) = \sum_q W_2(q|J, \varepsilon)$, so that

$$Z_2(J, \varepsilon) = (4 \cosh \beta J)^{N/2} \cdot Z_1(J', \varepsilon), \quad (\text{A.4})$$

where

$$Z_1(J, h) = \sum_{\sigma} e^{-\beta E_J(\sigma) + \beta h \sum_i \sigma_i}. \quad (\text{A.5})$$

is the partition function of a single plaquette model in a field h . In addition, this latter system is known to have an exact duality [97, 98]

$$Z_1(J, h) = (\sinh \beta J \sinh 2\beta h)^{N/2} Z_1(\tilde{J}, \tilde{h}), \quad (\text{A.6})$$

where

$$e^{-\beta \tilde{J}} = \tanh(\beta h), \quad e^{-2\beta \tilde{h}} = \tanh(\beta J/2). \quad (\text{A.7})$$

From Eqs. (A.4)-(A.7) the duality of the coupled plaquette system follows:

$$Z_2(J, \varepsilon) = (\sinh \beta J \sinh \beta \varepsilon)^N Z_2(J^*, \varepsilon^*), \quad (\text{A.8})$$

with

$$e^{-\beta\varepsilon^*} = \tanh(\beta J/2), \quad e^{-\beta J^*} = \tanh(\beta\varepsilon/2). \quad (\text{A.9})$$

This duality is precisely the one obtained in [36] for the two coupled replicas of the TPM. (Note that if $\tanh y = e^{-2x}$ then $\tanh x = e^{-2y}$, which follows from the definition of the tanh function, and facilitates inversion of these duality transforms.)

Bibliography

- [1] L. Peliti. *Statistical mechanics in a nutshell*. Princeton University Press, 2011.
- [2] D. Chandler. *Introduction to Modern Statistical Mechanics*. Oxford University Press, Oxford, 1987.
- [3] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler. Transition path sampling: throwing ropes over rough mountain passes, in the dark. *Annu. Rev. Phys. Chem.*, 53:291, 2002.
- [4] M. D. Ediger, C. A. Angell, and S. R. Nagel. Supercooled liquids and glasses. *J. Phys. Chem.*, 100:13200, 1996.
- [5] A. Cavagna. Supercooled liquids for pedestrians. *Phys. Rep.*, 476(51), 2009.
- [6] L. Berthier and G. Biroli. Theoretical perspective on the glass transition and amorphous materials. *Rev. Mod. Phys.*, 83:587, Jun 2011.
- [7] G. Biroli and J. P. Garrahan. Perspective: The glass transition. *J. Chem. Phys.*, 138:12A301, 2013.
- [8] V. Lubchenko and P. G. Wolynes. Theory of structural glasses and supercooled liquids. *Annu. Rev. Phys. Chem.*, 58:235, 2007.

-
- [9] G. Brambilla, D. El Masri, M. Pierno, L. Berthier, L Cipelletti, G. Petekidis, and A. B. Schofield. Probing the equilibrium dynamics of colloidal hard spheres above the mode-coupling glass transition. *Phys. Rev. Lett.*, 102:085703, 2009.
- [10] L. Berthier, G. Biroli, J. Bouchaud, L Cipelletti, and W. van Saarloos. *Dynamical heterogeneities in glasses, colloids and granular media*. Oxford University Press, Oxford, 2011.
- [11] S. Franz and G. Parisi. Phase diagram of coupled glass systems: a mean-field study. *Phys. Rev. Lett.*, 79:2486, 1997.
- [12] 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. *Phys. Rev. Lett.*, 98:195702, May 2007.
- [13] L. O. Hedges, R. L. Jack, J. P. Garrahan, and D. Chandler. Dynamic order-disorder in atomistic models of structural glass formers. *Science*, 323(5919):1309, 2009.
- [14] D. Ruelle. *Thermodynamic formalism*. Cambridge University Press, 2004.
- [15] M. Merolle, J. P. Garrahan, and D. Chandler. Space-time thermodynamics of the glass transition. *Proc. Natl. Acad. Sci. USA*, 102(31):10837, 2005.
- [16] V. Lecomte, C. Appert-Rolland, and F. van Wijland. Thermodynamic formalism for systems with markov dynamics. *J. Stat. Phys.*, 127(1):51, Apr 2007.
- [17] J. P. Garrahan, R. L Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland. First-order dynamical phase transition in models of glasses: an approach based on ensembles of histories. *J. Phys. A: Math. Theor.*, 42(7):075007, 2009.

-
- [18] J. P. Garrahan and I. Lesanovsky. Thermodynamics of quantum jump trajectories. *Phys. Rev. Lett.*, 104(16):160601, April 2010.
- [19] E. Pitard, V. Lecomte, and F. van Wijland. Dynamic transition in an atomic glass former: A molecular-dynamics evidence. *Europhys. Lett.*, 96(5):56002, 2011.
- [20] C. Giardinà, J. Kurchan, V. Lecomte, and J. Tailleur. Simulating rare events in dynamical processes. *J. Stat. Phys.*, 145:787, 2011.
- [21] M. Gorissen, J. Hooyberghs, and C. Vanderzande. Density-matrix renormalization-group study of current and activity fluctuations near nonequilibrium phase transitions. *Phys. Rev. E*, 79:020101, Feb 2009.
- [22] R. L. Jack and P. Sollich. Large deviations and ensembles of trajectories in stochastic models. *Prog. Theor. Phys. Supp.*, 184:304–317, 2010.
- [23] J. P. Eckmann and D. Ruelle. Ergodic theory of chaos and strange attractors. *Rev. Mod. Phys.*, 57:617, Jul 1985.
- [24] A. Dembo and O. Zeitouni. *Large Deviation Techniques and Applications*. Springer, second edition, 1998.
- [25] H. Touchette. The large deviation approach to statistical mechanics. *Phys. Rep.*, 478(1-3):1, July 2009.
- [26] D. P. Landau and K. Binder. *A guide to Monte Carlo simulation in statistical physics*. Cambridge University Press, 2009.
- [27] J. Jckle and S. Eisinger. A hierarchically constrained kinetic ising model. *Zeitschrift fr Physik B Condensed Matter*, 84:115, 1991.
- [28] F. Ritort and P. Sollich. Gglass dynamics of kinetically constrained models. *Adv. Phys.*, 52:219, 2003.

-
- [29] C. W. Gardiner. *Handbook of stochastic methods*. Springer, 1986.
- [30] M. B. Plenio and P. L. Knight. The quantum jump approach to dissipative dynamics in quantum optics. *Rev. Mod. Phys.*, 70:101, 1998.
- [31] C. W. Gardiner and P. Zoller. *Quantum Noise*. Springer, 2004.
- [32] G. E. Crooks and D. Chandler. Efficient transition path sampling for nonequilibrium stochastic dynamics. *Phys. Rev. E*, 64:026109, 2001.
- [33] M. E. J. Newman and C. Moore. Glassy dynamics and aging in an exactly solvable spin model. *Phys. Rev. E*, 60:5068, 1999.
- [34] J. P. Garrahan and M. E. J. Newman. Glassiness and constrained dynamics of a short-range nondisordered spin model. *Phys. Rev. E*, 62:7670, 2000.
- [35] J. P. Garrahan. Glassiness through the emergence of effective dynamical constraints in interacting systems. *J. Phys. Condens. Matter*, 14:1571, 2002.
- [36] J. P. Garrahan. Transition in coupled replicas may not imply a finite-temperature ideal glass transition in glass-forming systems. *Phys. Rev. E*, 89:030301, 2014.
- [37] C. Jarzynski. Nonequilibrium equality for free energy differences. *Phys. Rev. Lett.*, 78:2690, 1997.
- [38] C. Jarzynski. Equilibrium free-energy differences from nonequilibrium measurements: A master-equation approach. *Phys. Rev. E*, 56:5018, 1997.
- [39] T. D. Lee and C. N. Yang. Statistical theory of equations of state and phase transitions. ii. lattice gas and ising model. *Phys. Rev.*, 87:410, Aug 1952.
- [40] C. N. Yang and T. D. Lee. Statistical theory of equations of state and phase transitions. i. theory of condensation. *Phys. Rev.*, 87:404, Aug 1952.

- [41] R. A. Blythe and M. R. Evans. Lee-yang zeros and phase transitions in nonequilibrium steady states. *Phys. Rev. Lett.*, 89:080601, 2002.
- [42] Christian Flindt and Juan P. Garrahan. Trajectory phase transitions, lee-yang zeros, and high-order cumulants in full counting statistics. *Phys. Rev. Lett.*, 110:050601, 2013.
- [43] J. M. Hickey, C. Flindt, and J. P. Garrahan. Trajectory phase transitions and dynamical lee-yang zeros of the glauber-ising chain. *Phys. Rev. E*, 88:012119, 2013.
- [44] G. Lindblad. On the generators of quantum dynamical semigroups. *Comm. Math. Phys.*, 48:119, 1976.
- [45] C. W. Gardiner and P. Zoller. *Quantum Noise*. Springer, 2004.
- [46] Belavkin. A stochastic posterior schroedinger equation for counting non-demolition measurement. *Lett. Math. Phys.*, 20:85, 1990.
- [47] M. Baiesi, C. Maes, and B. Wynants. Fluctuations and response of nonequilibrium states. *Phys. Rev. Lett.*, 103:010602, Jul 2009.
- [48] J. L Lebowitz and H. Spohn. A gallavotti-cohen-type symmetry in the large deviation functional for stochastic dynamics. *J. Stat. Phys.*, 95:333, 1999.
- [49] P. G. Debenedetti and F. H Stillinger. Supercooled liquids and the glass transition. *Nature*, 410:259, 2001.
- [50] D. Chandler and J. P. Garrahan. Dynamics on the way to forming glass: bubbles in space-time. *Annu. Rev. Phys. Chem.*, 61(1):191, 2010.
- [51] J. Zhao, S. L. Simon, and G. B. McKenna. Using 20-million-year-old amber to test the super-arrhenius behaviour of glass-forming systems. *Nat. Commun.*, 4:1783, 2013.

-
- [52] G. Gallavotti and E. G. D. Cohen. Dynamical ensembles in nonequilibrium statistical mechanics. *Phys. Rev. Lett.*, 74:2694, 1995.
- [53] J. Kurchan. Fluctuation theorem for stochastic dynamics. *J. Phys. A: Math. Gen.*, 31:3719, 1998.
- [54] G. E. Crooks. Entropy production fluctuation theorem and nonequilibrium work relation for free energy differences. *Phys. Rev. E*, 60:2721, 1999.
- [55] T. Hatano and S. Sasa. Steady-state thermodynamics of langevin systems. *Phys. Rev. Lett.*, 86:3463, 2001.
- [56] C. Bustamante, J. Liphardt, and F. Ritort. the nonequilibrium thermodynamics of small systems. *Physics Today*, 43, 2005.
- [57] E. M. Sevick, R. Prabhakar, S. R. Williams, and D. J. Searles. Fluctuation theorems. *Annu. Rev. Phys. Chem.*, 59:603, 2008.
- [58] U. Seifert. Stochastic thermodynamics, fluctuation theorems, and molecular machines. *Rep. Prog. Phys.*, 75:126001, 2012.
- [59] D. J. Evans, E. G. D. Cohen, and G. P. Morris. Probability of second law violations in shearing steady states. *Phys. Rev. Lett.*, 71(15):2401, 1993.
- [60] D. J. Evans and D. J. Searles. Equilibrium microstates which generate 2nd law violating steady-states. *Phys. Rev. E.*, 50(2):1645, 1994.
- [61] D. J. Evans and D. J. Searles. The fluctuation theorem. *Adv. Phys.*, 51(7):1529, 2002.
- [62] S. Vaikuntanathan and C. Jarzynski. Dissipation and lag in irreversible processes. *Europhys. Lett.*, 87:60005, 2009.
- [63] G. E. Crooks. Nonequilibrium measurements of free energy differences for microscopically reversible markovian systems. *J. Stat. Phys.*, 90:1481, 1998.

-
- [64] C. Flindt, C. Fricke, F. Hohls, T. Novotný, K. Netocny, T. Brandes, and R. J. Haug. Universal oscillations in counting statistics. *Proc. Natl. Acad. Sci. USA*, 106(25):10116, June 2009.
- [65] Christian Flindt, T. Novotný, A. Braggio, and A.-P. Jauho. Counting statistics of transport through coulomb blockade nanostructures: High-order cumulants and non-markovian effects. *Phys. Rev. B*, 82:155407, Oct 2010.
- [66] D. Kambly, C. Flindt, and M. Büttiker. Factorial cumulants reveal interactions in counting statistics. *Phys. Rev. B*, 83:075432, Feb 2011.
- [67] R. M. L. Evans. Rules for transition rates in nonequilibrium steady states. *Phys. Rev. Lett.*, 92:150601, 2004.
- [68] C. Maes and K. Netocny. Canonical structure of dynamical fluctuations in mesoscopic nonequilibrium steady states. *Europhys. Lett.*, 82:30003, 2008.
- [69] J. Kurchan. Six out of equilibrium lectures. *arXiv:0901.1271*, 2009.
- [70] T. Nemoto and S. Sasa. Thermodynamic formula for the cumulant generating function of time-averaged current. *Phys. Rev. E*, 84:061113, 2011.
- [71] A. Budini. Large deviations of ergodic counting processes: A statistical mechanics approach. *Phys. Rev. E*, 84(1):011141, July 2011.
- [72] Thomas Speck and David Chandler. Constrained dynamics of localized excitations causes a non-equilibrium phase transition in an atomistic model of glass formers. *J. Chem. Phys.*, 136(18):184509, 2012.
- [73] T. Speck, A. Malins, and C. P. Royall. First-order phase transition in a model glass former: coupling of local structure and dynamics. *Phys. Rev. Lett.*, 109:195703, 2012.

-
- [74] C. Catana, M. van Horsen, and M. Guta. Asymptotic inference in system identification of the atom maser. *Phil. Trans. Royal Soc. A*, 370:5308, 2012.
- [75] T. Bodineau and C. Toninelli. Activity phase transition for constrained dynamics. *Commun. Math. Phys.*, 311:357, 2012.
- [76] J. K. Weber, R. L. Jack, and V. S. Pande. Emergence of glass-like behavior in markov state models of protein folding dynamics. *J. Am. Chem. Soc.*, 135:5501, 2013.
- [77] C. P. Espigares, P. L. Garrido, and P. I. Hurtado. Dynamical phase transition for current statistics in a simple driven diffusive system. *Phys. Rev. E*, 87:032115, 2013.
- [78] A. S. J. S. Mey, P. L. Geissler, and J. P. Garrahan. Rare-event trajectory ensemble analysis reveal metastable dynamical phases in lattice proteins. *Phys. Rev. E*, 89:032109, 2014.
- [79] S. Vaikuntanathan, T. R. Gingrich, and P. L. Geissler. Dynamic phase transitions in simple driven kinetic networks. *Phys. Rev. E*, 89:062108, 2014.
- [80] J. K. Weber, D. Shukla, and V. S. Pande. Heat dissipation guide activation in signaling proteins. *Proc. Natl. Acad. Sci. USA*, 112:10377, 2015.
- [81] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England. Dissipation bounds all steady-state current fluctuations. *Phys. Rev. Lett.*, 116:120601, 2016.
- [82] R. L. Jack, I.R. Thompson, and P. Sollich. Hyperuniformity and phase separation in biased ensembles of trajectories for diffusive systems. *Phys. Rev. Lett.*, 114:060601, 2015.
- [83] M. Ueda and S. Sasa. Replica symmetry breaking in trajectories of a driven brownian particle. *Phys. Rev. Lett.*, 115:080605, 2015.

-
- [84] C. De Bacco, A. Guggiola, R. Khn, and P. Paga. Rare events statistics of random walk on networks: localization and other dynamical phase transitions. *arXiv:1506.08436*, 2015.
- [85] J. Szavits-Nossan and M. R. Evans. Inequivalence of nonequilibrium path ensembles: the example of stochastic bridges. *J. Stat. Mech.*, page P12008, 2015.
- [86] G. Verley. Nonequilibrium thermodynamic potentials for continuous-time markov chains. *Phys. Rev. E*, 93:012111, 2016.
- [87] M. V. S. Bonana and C. Jarzynski. Conditional reversibility in nonequilibrium stochastic systems. *Phys. Rev. E*, 93:022101, 2016.
- [88] T. Nemoto, F. Bouchet, R. L Jack, and V. Lecomte. Population-dynamics method with a multicanonical feedback control. *Phys. Rev. E*, 93:062123, 2016.
- [89] T. Speck. Thermodynamic formalism and linear response theory for nonequilibrium steady states. *arXiv:1601.03540*, 2016.
- [90] R. L Jack and R. M. L. Evans. Absence of dissipation in trajectory ensembles biased by currents. *arXiv:1602.03815*, 2016.
- [91] R Chetrite and H. Touchette. Nonequilibrium microcanonical and canonical ensembles and their equivalence. *Phys. Rev. Lett.*, 111:120601, 2013.
- [92] D. Frenkel and B. Smit. *Understanding molecular simulation: from algorithms to applications*. Elsevier, 2002.
- [93] G. O. Roberts, A. Gelman, and W. R. Gilks. Weak convergence and optimal scaling of random walk metropolis algorithms. *Ann. Appl. Probab.*, 7(1):110, 1997.

-
- [94] B. G. Englert. Elements of micromaser physics. arXiv:quant-ph/0203052.
- [95] J. P. Garrahan, A. D. Armour, and I. Lesanovsky. Quantum trajectory phase transitions in the micromaser. *Phys. Rev. E*, 84(2):021115, August 2011.
- [96] J. P. Garrahan and D. Chandler. Coarse-grained microscopic model of glass formers. *Proc. Natl. Acad. Sci. USA*, 100:9710, 2003.
- [97] J. R. Heringa, H. W. J. Blte, and A. Hoogland. Phase transitions in self-dual ising models with multispin interactions and a field. *Phys. Rev. Lett.*, 63:1546, 1989.
- [98] S. Sasa. Thermodynamic transition associated with irregularly ordered ground states in lattice gas model. *J. Phys. A*, 43:465002, 2010.
- [99] Y. S. Elmatad, R. L. Jack, D. Chandler, and J. P. Garrahan. Finite-temperature critical point of a glass transition. *Proc. Natl. Acad. Sci. USA*, 107(29):12793, 2010.
- [100] R. L. Jack and J. P. Garrahan. Caging and mosaic length scales in plaquette spin models of glasses. *J. Chem. Phys.*, 123:164508, 2005.
- [101] D. A. Johnston. Goniherdic (and fuki-nuke) order. *J. Phys. A*, 45:405001, 2012.
- [102] S. Franz and G. Parisi. Universality classes of critical points in constrained glasses. *J. Stat. Mech.*, page P11012, 2013.
- [103] L. Berthier. Overlap fluctuations in glass-forming liquids. *Phys. Rev. E*, 88:022313, 2013.
- [104] G. Parisi and B. Seoanne. Liquid-glass transition in equilibrium. *Phys. Rev. E*, 89:022309, 2014.

-
- [105] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz. A new algorithm for monte carlo simulation of ising spin systems. *J. Comp. Phys.*, 17:10, 1975.
- [106] M. E. J. Newman and G. T. Barkema. *Monte Carlo methods in statistical physics*. Oxford University Press, Oxford, 1999.
- [107] A. D. Bruce and N. B. Wilding. Computational strategies for mapping equilibrium phase diagrams. *Adv. Chem. Phys.*, 127:1, 2003.
- [108] D. Nicolaides and A. D. Bruce. Universal configurational structure in two-dimensional scalar models. *J. Phys. A*, 21:233, 2988.
- [109] Y. Deng, W. Guo, J. R. Heringa, H. W. Blöte, and B. Nienhuis. Phase transitions in self-dual generalizations of the baxter-wu model. *Nucl. Phys. B*, 827:406, 2010.
- [110] G. H. Fredrickson and H. C. Andersen. Kinetic ising-model of the glass-transition. *Phys. Rev. Lett.*, 53(13):1244, 1984.
- [111] Y. S. Elmatad, D. Chandler, and J. P. Garrahan. Corresponding states of structural glass formers. *J. Phys. Chem.*, 113:5563, 2009.
- [112] C. Cammarota and G. Biroli. Ideal glass transitions by random pinning. *Proc. Natl. Acad. Sci. USA*, 109:8850, 2012.
- [113] R. L. Jack and L. Berthier. Random pinning in glassy spin models with plaquette interactions. *Phys. Rev. E*, 85:021120, 2012.
- [114] W. Kob and L. Berthier. Probing a liquid to glass transition in equilibrium. *Phys. Rev. Lett.*, 110:245702, 2013.
- [115] M. Aizenman and J. Wehr. Rounding of first-order phase transitions in systems with quenched disorder. *Phys. Rev. Lett.*, 62:2503, 1989.

-
- [116] P. Sollich and M. R. Evans. Glass time-scale divergence and anomalous coarsening in a kinetically constrained spin chain. *Phys. Rev. Lett.*, 83:3238, 1999.
- [117] R. L Jack, L. O. Hedges, J. P. Garrahan, and D. Chandler. Preparation and relaxation of very stable glassy states of a simulated liquid. *Phys. Rev. Lett.*, 107:275702, 2011.
- [118] R. L Jack. Counting metastable states in a kinetically constrained model using patch repetition analysis. *Phys. Rev. E*, 88:062113, 2013.
- [119] R. Monasson. Structural glass transition and the entropy of metastable states. *Phys. Rev. Lett.*, 75:2847, 1995.
- [120] J. Kurchan and D. Levine. Order in glassy systems. *J. Phys. A*, 44:035001, 2011.
- [121] L. Foini, F. Krzakala, and F. Zamponi. On the relation between kinetically constrained models of glass dynamics and the random first-order transition theory. *J. Stat. Mech.*, page P06013, 2012.
- [122] C. Cammarota, A. Cavagna, I. Giardina, G. Gradenigo, T. S. Grigera, G. Parisi, and P. Verrocchio. Phase-separation perspective on dynamic heterogeneities in glass-forming liquids. *Phys. Rev. Lett.*, 105:055703, 2010.
- [123] L. Berthier and R. L Jack. Evidence for a disordered critical point in a glass-forming liquid. *Phys. Rev. Lett.*, 114:205701, 2015.
- [124] C. Xu and J. E. Moore. Strong-weak coupling self-dual in the two-dimensional quantum phase transition of $p + ip$ superconducting arrays. *Phys. Rev. Lett.*, 93:047003, 2004.

-
- [125] C. Xu and J. E. Moore. Reduction of effective dimensionality in lattice modes of superconducting arrays and frustrated magnets. *Nucl. Phys. B*, 716:487, 2005.
- [126] C. Chatelain and D. Karevski. Probability distribution of the work in the two-dimensional ising model. *J. Stat. Mech.*, page P06005, 2006.
- [127] A. K. Hartmann. High-precision work distributions for extreme nonequilibrium processes in large systems. *Phys. Rev. E*, 89:052103, 2014.
- [128] A. Imparato and L. Peliti. Work distribution and path iintegral in general mean-field systems. *Europhys. Lett.*, 2005.
- [129] C. H. Bennett. Efficient estimation of free energy differences from monte carlo data. *J. Comp. Phys.*, 22:245, 1976.
- [130] M. R. Shirts, E. Blair, G. Hooker, and V. S. Pande. Equilibrium free energies from nonequilibrium measurements using maximum-likelihood methods. *Phys. Rev. Lett.*, 91:140601, 2003.
- [131] M. van Horssen and M. Guta. Large deviations, central limit and dynamical phase transitions in the atom maser. *arXiv:1206.4956*, 2013.
- [132] I. Lesanovsky, M. van Horssen, M. Guta, and J. P. Garrahan. Characterization of dynamical phase transitions in quantum jump trajectories beyond the properties of the stationary state. *Phys. Rev. Lett.*, 110:150401, 2013.
- [133] T. Z. Mordasini and J. A. McCammon. Calculations of relative hydration free energies: a comparative study using thermodynamic integration and an extrapolation method based on a single reference state. *J. Phys. Chem. B.*, 104:360, 2000.
- [134] G. Hummer. Fast-growth thermodynamic integration: Error and efficiency analysis. *J. Chem. Phys.*, 114:7330, 2001.

-
- [135] F. M. Ytreberg and D. M. Zuckerman. Single-ensemble nonequilibrium path-sampling estimates of free energy differences. *J. Chem. Phys.*, 120:10876, 2004.
- [136] C. Dellago and G. Hummer. Computing equilibrium free energies using non-equilibrium molecular dynamics. *Entropy*, 16:41, 2014.
- [137] Ivan P. Levkivskyi and Eugene V. Sukhorukov. Noise-induced phase transition in the electronic mach-zehnder interferometer. *Phys. Rev. Lett.*, 103:036801, Jul 2009.
- [138] D. A. Ivanov and A. G. Abanov. Phase transitions in full counting statistics for periodic pumping. *Europhys. Lett.*, 92(3):37008, November 2010.
- [139] Dmitri A. Ivanov and Alexander G. Abanov. Characterizing correlations with full counting statistics: Classical ising and quantum xy spin chains. *Phys. Rev. E*, 87:022114, Feb 2013.
- [140] Y. Utsumi, O. Entin-Wohlman, A. Ueda, and A. Aharony. Full-counting statistics for molecular junctions: Fluctuation theorem and singularities. *Phys. Rev. B*, 87:115407, Mar 2013.
- [141] J. M. Hickey, C. Flindt, and J. P. Garrahan. Intermittency and dynamical lee-yang zeros of open quantum systems. *Phys. Rev. E*, 90:062128, 2014.
- [142] E. Barkai, Y. Jung, and R. Silbey. Theory of single-molecule spectroscopy: beyond the ensemble average. *Annu. Rev. Phys. Chem.*, 55:457, 2004.
- [143] S. G. Krantz. *Handbook of complex variables*. Birkhuser, 1999.
- [144] R. B. Dingle. *Asymptotic Expansions: Their Derivation and Interpretation*. Academic Press, London, 1973.

-
- [145] M. V. Berry. Universal oscillations of high derivatives. *Proc. R. Soc. A*, 461(2058):1735, 2005.
- [146] R. S. Bhalerao, N. Borghini, and J. Y. Ollitrault. Analysis of anisotropic flow with lee–yang zeroes. *Nucl. Phys. A*, 727(3):373, 2003.
- [147] J. Zamastil and F. Vinette. Determination of singularities of a function from its perturbation expansion. *J. Phys. A: Math. Gen.*, 38:4009, 2005.
- [148] J. M. Hickey, S. Genway, I. Lesanovsky, and J. P. Garrahan. Thermodynamics of quadrature trajectories in open quantum systems. *Phys. Rev. A*, 86:063824, Dec 2012.
- [149] J Kiukas, M. Guta, I. Lesanovsky, and J. P. Garrahan. Equivalence of matrix product ensembles of trajectories in open quantum systems. *Phys. Rev. E*, 92:012132, 2015.