

TRAJECTORY ENSEMBLE METHODS FOR UNDERSTANDING
COMPLEX STOCHASTIC SYSTEMS

BY

ANTONIA SILKE JULIA SUSANNA MEY, BSC.

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

July 2013

ABSTRACT

This thesis investigates the equilibrium and dynamic properties of stochastic systems of varying complexity. The dynamic properties of lattice models – the 1-d Ising model and a 3-d protein model – and equilibrium properties of continuous models – particles in various potentials – are presented. Dynamics are studied according to a large deviation formalism, by looking at non-equilibrium ensembles of trajectories, classified according to a dynamical order parameter. The phase structure of the ensembles of trajectories is deduced from the properties of large-deviation functions, representing dynamical free-energies.

The 1-d Ising model is studied with Glauber dynamics uncovering the dynamical second-order transition at critical values of the counting field 's', confirming the analytical predictions by Jack and Solich. Next, the dynamics in an external magnetic field are studied, allowing the construction of a dynamic phase diagram in the space of temperature, s-field and magnetic field. The dynamic phase diagram is reminiscent of that of the 2-d Ising model. In contrast, Kawasaki dynamics give rise to a dynamical phase structure similar to the one observed in kinetically constrained models.

The dynamics of a lattice protein model, represented by a self avoiding walk with three different Hamiltonians, are studied. For the uniform Gō Hamiltonian all dynamics occurs between non-native and native trajectories, whereas for heterogeneous Hamiltonians and Full interaction Hamiltonians a first-order dynamical transition to sets of trapping trajectories is observed in the s-ensemble. The model is studied exhaustively for a particular sequence, constructing a qualitative phase diagram, from which a more general dynamic behaviour is extrapolated.

Lastly, an estimator for equilibrium expectations, represented by a transition matrix in an extended space between temperatures and a set of discrete states obtained through the discretisation of a continuous space, is proposed. It is then demonstrated that this estimator outperforms conventional multi-temperature ensemble estimates by up to three orders of magnitude, by considering three models of increasing complexity: diffusive particles in a double-well potential, a multidimensional folding potential and a molecular dynamics simulations of alanine dipeptide.

L'erreur, c'est de croire que le bonheur et le but.

— Francois Lelord (*Le voyage d'Hector*)

ACKNOWLEDGMENTS

I would like to thank everyone who has supported me throughout my time as a PhD student. First of all my parents, who made it possible for me to study in the first place and helped me get to where I am now. I would also like to thank Ralf Banisch for tolerating my many mood swings in the last few months.

All my friends in Nottingham who ensured that I had a great time living there and help me develop as a person. Special thanks to my brilliant house mate Christina Voss and our tea and video evenings after a long day of work. Frank Hellmann with whom I enjoyed many creative hours making theatre come to life. In general I would like to thank everyone who has been involved in the theatre with me, allowing me to have such a great time, doing some creative projects next to all the science. In particular I would like to thank Anya Skatova who has become a good friend and Ollie Braeunling and Celine Siragna. My other passion being dancing, I would also like to thank the LBSS team for a great time in Nottingham and lots of fun dance competitions.

Dan Rayneau-Kirkhope for the many coffees. Joe Shearring for all the entertaining online messaging even though we were only sitting a door apart from each other. Nick McMullin for always listening to my problems and Lama Hamadeh for being such a good friend. I would also like to thank Mark Patton for many fun evenings and trips to Ikea, we started at the same time and finished at the same time. Physics would have not been the same without the tea room, so I would like to thank everyone making tea breaks and other social activities a lot of fun (Pete, Suzie, Jenn, Jara, Daphne, Hector, Douglas, Lester, Kunal, Kuldeep, Ryan ..). I also want to thank Michael Cann for many hours of help with C++ and general moral support.

As I spend a significant amount of time using not only my work station but also the HPC facilities in Nottingham, I would like to thank the University for providing such facilities and in particular Collin Bannister for always replying to my request and helping me install all the software needed for my computational work.

On my long journey to completion of this thesis I met some wonderful people who have been great friends as well as colleagues, thank you Yael Elmatad and Anna Schneider. Special thanks to John Chodera for always having a patient ear, a lot of interesting science discussions or cocktail-related ones and generally giving great advice. I would also

like to thank Frank Noé and everyone in the CMB group for being so patient with me while I was finishing and not doing any exciting new science!

Last but not least, I would like to thank my supervisor, Juan P. Garrahan, for giving me advice and helping me make connections in Berkeley.

CONTENTS

List of Figures	vii
List of Tables	vii
Listings	vii
List of Acronyms	vii
List of Symbols	vii
1 INTRODUCTION	1
2 THEORETICAL BACKGROUND	8
2.1 Equilibrium Stochastic Processes	8
2.1.1 Introduction: Coin Tossing	8
2.1.2 Canonical Ensemble	10
2.1.3 Example: Single Particle in a Harmonic Potential	12
2.2 Markov State Models	13
2.2.1 Formal Definition of a Markov Process	14
2.2.2 Transfer Operator Approach	15
2.2.3 Discrete Approximation of the Transfer Operator – the Transition Matrix	16
2.2.4 Transition Matrix Estimation	17
2.2.5 Large Systems with Continuous State Spaces	19
2.2.6 Example: Single Particle in a Double Well	21
2.2.7 Markov Model Estimation	22
2.3 Phase Transitions and their Classification	22
2.3.1 The Order Parameter	23
2.3.2 Landau Theory of Second-Order Transitions	23
2.4 Non-equilibrium Properties	25
2.4.1 Large Deviation Theory	25
2.4.2 The Dynamic Ensemble	29
2.4.3 Applications	31
3 DYNAMIC PHASES OF THE 1-d ISING MODEL	32
3.1 The Ising Model	32
3.2 Equilibrium and Non-equilibrium Properties	33
3.2.1 Monte Carlo Simulations	34
3.2.2 Universality	37
3.2.3 The 2-d Ising Universality Class	38
3.2.4 Finite Size Scaling – Measuring Exponents	39
3.2.5 Non-Equilibrium Properties – Spinodal Decomposition	41
3.3 Dynamic Properties of the 1-d Ising Model	44
3.3.1 Theoretical Predictions	45
3.3.2 Simulation Setup	47
3.3.3 Simulation Results of the Dynamic Properties	50

3.3.4	Simulation Results of the Dynamic Properties with External Magnetic Field	55
3.3.5	Results from a Hysteresis Study	56
3.3.6	Results for the Dynamic Phase Behaviour Using Kawasaki Dynamics	58
3.4	Conclusion and Outlook	62
4	LATTICE PROTEIN MODELS	64
4.1	Introduction to Proteins	64
4.1.1	Protein Structure	65
4.1.2	States or Phases of a Protein	68
4.1.3	The Search for the Native State: the Levinthal Paradox	69
4.1.4	Experimental Techniques	69
4.1.5	Computational Techniques	70
4.2	The Lattice Protein Model	71
4.3	Equilibrium Properties of the 2-d Gō Model	72
4.3.1	Hamiltonian	72
4.3.2	MC Simulations	73
4.3.3	Equilibrium Simulation Results	75
4.4	The 3-d Lattice Model	76
4.4.1	Hamiltonian and Native State Definition	76
4.4.2	Equilibrium Simulation Results	77
4.5	The Lattice Protein Model in the s-ensemble	80
4.5.1	MC Simulations Setup	80
4.5.2	Gō Results	81
4.5.3	HeGō and Full Interaction Potential Results	85
4.5.4	Summary	95
4.6	Conclusion	95
5	EQUILIBRIUM ESTIMATIONS	97
5.1	Molecular Dynamics: Sampling the Equilibrium	97
5.2	Advanced Sampling Techniques	100
5.2.1	Parallel Tempering	101
5.2.2	Simulated Tempering	102
5.3	Multiple State Equilibrium Estimation	103
5.3.1	Defining the Extended Ensemble	103
5.3.2	The Choice of Scaling Variable	105
5.3.3	Estimating \tilde{T} from Simulation Data	107
5.4	Results	107
5.4.1	Double Well Potential	108
5.4.2	Folding Potential	110
5.4.3	Alanine Dipeptide	113
5.5	Discussion and Conclusion	116
6	SUMMARY AND OUTLOOK	118
A	APPENDIX: EXPERIMENTAL TECHNIQUES	122
B	APPENDIX: AMINO ACIDS AND ADDITIONAL RESULTS	124
B.1	Amino Acid Letter Code	124
B.2	Additional Results	124

BIBLIOGRAPHY 127

LIST OF FIGURES

Figure 2.1	Harmonic potential and equilibrium probabilities	13
Figure 2.2	Markov model example at lagtime $\tau = 10$: double well	21
Figure 2.3	First- and second-order phase transition	23
Figure 2.4	Landau free energy	24
Figure 2.5	Large deviation rate function	26
Figure 3.1	1-d Ising spin chain	33
Figure 3.2	Correlation times	40
Figure 3.3	Domain growth	42
Figure 3.4	Average domain size	44
Figure 3.5	1-d Ising dynamic phase diagram	46
Figure 3.6	Sampling the s-ensemble	49
Figure 3.7	Susceptibilities of energies for Glauber dynamics	51
Figure 3.8	Susceptibilities of activity for Glauber dynamics	53
Figure 3.9	Susceptibilities of the integrated magnetisation for Glauber dynamics	55
Figure 3.10	Glauber dynamics – first-order field driven transition	56
Figure 3.11	Hysteresis loops	58
Figure 3.12	Kawasaki dynamics	59
Figure 3.13	Scaling behaviour from Monte Carlo simulations with Kawasaki dynamics in the s-ensemble	61
Figure 3.14	Phase diagram from Kawasaki dynamics	61
Figure 4.1	Amino acid – chemical structure	65
Figure 4.2	Ramachandran diagram	68
Figure 4.3	Exemplary native state of 2-d lattice protein	73
Figure 4.4	Moveset for 2-d lattice protein	74
Figure 4.5	Equilibrium simulation properties of the 2-d Gō model	75
Figure 4.6	Exemplary native state of the 8mer amino acid chain in 3-d	76
Figure 4.7	Native structure of 48mer and exemplary sequence	78
Figure 4.8	Nativeness and mean first passage time behaviour for all interaction models in 3-d.	79
Figure 4.9	Gō model: Equilibrium activity result	82
Figure 4.10	Gō model: joint probability of κ and integrated nativeness	82
Figure 4.11	Gō model: native activity with respect to s	83

Figure 4.12	Gō model: phase diagram	84
Figure 4.13	HeGō and Full potential: Equilibrium dynamics	85
Figure 4.14	HeGō: simulation results with s-ensemble biasing for 3 different temperatures	86
Figure 4.15	HeGo: Joint probability density plot of native activity and integrated nativeness with first passage time distributions and representative trapping structure	88
Figure 4.16	s-ensemble Trajectory root mean square distance (RMSD) analysis	90
Figure 4.17	Folding trajectory RMSD analysis	91
Figure 4.18	HeGō: $\langle \kappa_g \rangle_s$ biasing at T_{fold}	92
Figure 4.19	Full potential in the s-ensemble	93
Figure 4.20	Dynamic phase diagram of the HeGō and Full model	94
Figure 5.1	Double well potential and Langevin dynamics	100
Figure 5.2	Replica swapping in parallel tempering	101
Figure 5.3	Parallel tempering example	102
Figure 5.4	Double well potential and simulated tempering results	108
Figure 5.5	Estimator results from parallel tempering and random swapping data	109
Figure 5.6	Distribution overlap of potential energies of a double well	111
Figure 5.7	Toy folding potential and idealised folding funnel	111
Figure 5.8	Estimator results of parallel tempering and simulated tempering data of folding model	112
Figure 5.9	Molecular structure and free energy surface of alanine dipeptide	113
Figure 5.10	Relative error convergence of chosen 4 discrete states	115
Figure 5.11	Convergence of stationary probabilities from estimator and histogram counts	116
Figure B.1	HeGōs-ensemble biasing for sequence 2	125

LIST OF TABLES

Table 2.1	Number of times a state transition is observed in the light switching example from a 1000s trajectory.	18
Table 5.1	Angle minima Θ_0 for von Mises potential	114

LISTINGS

Listing 3.1	Pseudocode: Metropolis algorithm	35
Listing 4.1	Pseudocode: Monte Carlo (MC) algorithm for lattice protein simulation	74

ACRONYMS

MFPT	mean first passage time
TPS	transition path sampling
MSM	Markov State model
CAO	contact appearance order
RMSD	root mean square distance
RMS	root mean square
NMR	nuclear magnetic resonance
FRET	Förster resonance energy transfer
MC	Monte Carlo
MD	molecular dynamics
REMD	replica exchange molecular dynamics (Monte Carlo dynamics)
PT	parallel tempering
WHAM	weighted histogram analysis method
LDT	large deviation theory
PCCA	Perron-Cluster Cluster Analysis
ST	simulated tempering
RS	random swapping
TRAM	transition matrix analysis method
MGF	moment generating function

LIST OF SYMBOLS

$(\mathbf{X}_t)_{t \in I}$	trajectory of any stochastic process, page 9
B_{ij}	interaction energy matrix with entries of -1 a.u. between amino acid residues i and j which are native contacts, page 72
Δ_{ij}	matrix with entries of unity if residue i and j are nearest neighbours for a given lattice protein conformation, page 72
$\tilde{\mathbf{T}}$	transition matrix in temperatures and states $\mathbb{R}^{KM \times KM}$, page 104
\ddot{x}	acceleration, page 98
\dot{x}	velocity, page 98
$\hat{\mathbf{T}}(\tau)$	transition matrix estimate, page 18
\mathbf{T}	transition matrix in $\mathbb{R}^{K \times K}$, where K is the number of discrete states, page 17
$\mathbb{E}(x_i - \mu_n)^2, \sigma$	variance or second central moment, page 9
$\mathbb{E}(x_n), \langle x_n \rangle, \mu_n$	expected value or first central moment, page 9
\mathbf{C}	count matrix, page 17
\mathbf{D}	diagonal matrix holding temperature jumps from a lower to a higher temperature in the extended ensemble, page 104
$\mathbf{F}(\mathbf{x})$	force acting on a particle, page 19
\mathbf{I}	diagonal matrix converting a rate matrix to the transition matrix of the extended ensemble, page 106
\mathbf{K}	generator of a reversible stochastic process, page 5
\mathbf{S}	diagonal matrix ensuring stochasticity of the extended ensemble matrix, page 105
\mathbf{U}	diagonal matrix holding temperature jumps from a lower to a higher temperature in the extended ensemble, page 104

$\mathbf{X}_{t_{\text{obs}}}$	trajectory of a stochastic process observed for a length of time t_{obs} , page 5
\mathbf{k}	wave vector, page 43
\mathcal{B}	set of all possible events (Borel σ -algebra), page 8
$\mathcal{F}(\beta)$	free energy of a given system, page 11
$\mathcal{H}(\mathbf{x})$	Hamiltonian of a given system, page 14
\mathcal{H}_N	N particle Hamiltonian, page 27
S	entropy of a given system, page 11
$\mathcal{T}(\tau)$	transfer operator of a stochastic process at a given lagtime τ , page 15
\mathbf{x}, \mathbf{y}	vectors in state space Ω , page 14
\mathbf{x}_0	initial position of the system in Ω at time $t = 0$, page 5
Λ_A	characteristic function for being in a set A , page 5
Ω	state space – set of all possible configurations the system can take, page 5
Υ	scaling variable in the extended ensemble proportional to a ratio of partition functions at different temperatures, page 106
Ξ	scaling variable associated to down jumps in the extended ensemble, page 106
α	scaling exponent associated to the specific heat, page 38
β	inverse temperature: $\beta = \frac{1}{k_B T}$, page 10
$\boldsymbol{\eta}(t)$	Gaussian random noise, page 19
$\boldsymbol{\pi}$	stationary probability vector of a discrete process, page 17
χ	susceptibility, page 38
χ_e	susceptibility of the integrated energy, page 52
χ_κ	susceptibility of the normalised activity κ , page 53
ϵ	relative error of the stationary estimate of a given state, page 110
ϵ_{std}	standard error from the mean, page 50
γ	scaling exponent associated to the susceptibility, page 38

κ	normalised activity, as an intensive variable, page 30
κ_g	general activity: Incremental count of all accepted Monte Carlo moves over the whole trajectory, page 80
λ_i	the i^{th} dominant eigenvalue of the transfer operator \mathcal{T} , page 16
μ_1	first moment of the circularly averaged structure factor, page 43
μ_0	observed mean of a given realisation, page 9
ν	scaling exponent associated to the correlation length, page 38
ω_i	microstate of a larger system, page 10
ϕ, ψ	dihedral angles in a peptide, page 67
$\pi(x)$	equilibrium probability of the system depending on position x , page 12
$\psi(s)$	dynamical free energy, page 30
ψ_i	the i^{th} dominant eigenfunction of the transfer operator \mathcal{T} , page 16
ρ	order parameter, page 23
σ_i	spin in the Ising model at the i^{th} lattice position, page 33
τ	lag time, page 5
τ_{corr}	correlation time, page 38
$\tilde{\chi}$	scaling function of the susceptibility, page 40
ν	scaling exponent associated to the domain size, page 41
ξ	correlation length, page 38
$D(t)$	spin domain size depending on time, page 41
$D_c(t)$	spin domain size approximated by the two point correlation function, page 42
$D_s(t)$	spin domain size approximated by the structure factor, page 43
E	time integrated energy, page 51
E_g	native energy, or ground state energy of the lattice protein., page 79
E_t	total energy of the system, page 11

$G_{ij}(\mathbf{r})$	spatial correlation function dependent on the distance vector \mathbf{r} , page 36
$I(r)$	large deviation rate function, page 26
I_0	zeroth order Bessel function, page 114
J	coupling constant, page 33
K	activity observable for a trajectory \mathbf{X}_t , page 29
L	length of the lattice in any orthogonal lattice direction, page 34
L_p	number of amino acids, of a lattice protein, page 72
M	time integrated magnetisation, page 51
N_c	number of native contacts of a lattice protein, page 72
N_s	number of spins on the lattice, page 34
$P(A_n \in B)$	probability of a random variable A_n depending on n random variables of taking a value in a set B . The corresponding rate function is I_B ., page 26
$P(E_g)$	probability of being in the native state, page 73
P_s	probability of observing a given trajectory in the s-ensemble, page 30
$P_t(K)$	probability of observing a trajectories of length $t = t_{\text{obs}}$ with activity K , page 30
P_{PT}	parallel tempering acceptance criterion, page 101
P_{ST}	simulated tempering acceptance criterion, page 103
Q	time integrated nativeness, page 81
$R(t)$	auto correlation function in time, page 36
$S(\mathbf{k}, t)$	structure factor, page 43
T	temperature of the system, page 1
T_{Ref}	reference temperature of the Gaussian random noise, page 97
T_c	critical temperature in a phase transition, page 23
T_{fold}	temperature at which 50% of the native contacts are formed/broken, page 76
$U(\Theta)$	van Mises potential, page 114
$U(x)$	exemplary 1-d potential, page 12

V	volume, page 60
$V(\langle E_N \rangle \in du)$	volume of the microstate ω_i that falls into the energy interval du , page 28
$Z(\beta)$	thermodynamic partition function serving as a probability normalisation constant, depending on the inverse temperature β (sometimes written as Z_m), page 10
$Z_t(s)$	dynamic partition function, dependent on biasing field s , page 30
s_c	critical biasing parameter s , page 46
$\langle E_N \rangle$	average energy per particle in an N particle system, page 27
$\hat{S}(k_n, t)$	circularly averaged structure factor of circular bin n , page 43
$\hat{T} = \frac{T-T_c}{T_c}$	dimensionless relative temperature, page 38
K	number of discrete states of a Markov model, making up the MSM state space, page 17
M	number of temperatures of the extended ensemble, page 104
t, h	the event of observing tails or heads respectively, page 8
\tilde{m}	instantaneous magnetisation, page 36
t_{obs}	length of time a process is observed, page 5
ζ	friction constant, page 19
c	specific heat, page 38
$f(\beta)$	free energy divided by the inverse temperature to ensure it is a convex function of β , page 28
$g(k)$	cumulant generating function of the k^{th} cumulant, page 27
g_m	weight factor used in simulated tempering simulations, page 103
h	external magnetic field, page 33
k_B	Boltzmann constant, page 1
m	mass of a particle, page 1
n	number of times any random experiment is repeated, page 8

n_c	number of native contacts present of a given lattice configuration, page 75
n_{ex}	number of times replicas are attempted to be exchanged, page 101
$p(\mathbf{x}, \mathbf{y}; dt) d\mathbf{y}$	transition probability between \mathbf{x} and \mathbf{y} within time dt , page 14
$p_t(\mathbf{x})$	probability of being in \mathbf{x} at time t , page 15
q	nativeness, i.e. the ratio of current native contacts over the total native contacts, page 75
s	biasing or counting field conjugate to the activity K , page 30
t_i	implied timescale at a given lagtime τ obtained from the i^{th} dominant eigenvalue, page 16
u_i	potential energy of the microstate ω_i , page 10
v	speed of a particle, page 1
z	dynamic scaling exponent, page 38
S	set comprising all discrete states that make up the state space, page 19

INTRODUCTION

Statistical mechanics is an area of study concerned with probabilistic systems with a large number of degrees of freedom. This means it aims to describe the properties of systems, where no deterministic solutions are available due to either an intrinsic randomness or the sheer complexity of the system. The variety of systems studied in the framework of statistical mechanics is very broad, such as particle or spins systems, animal populations and stock markets [7]. Especially in more recent years, these well known concepts have seen more application to non-classical areas such as biological, chemical, nano-mechanical and social systems [8].

This area of research originated in the study of ideal gases and was later driven by simple physical problems such as the behaviour of magnets, or the Brownian motion of particles and other systems that obey Maxwell-Boltzmann statistics (see equation (1.1)).

With the popularity of quantum mechanics gaining momentum, the statistical behaviour of fermions and bosons at low temperatures was studied and Bose-Einstein statistics were introduced around 1925 [9]. In contrast to fermionic particles, i.e. particles that are governed by the Pauli exclusion principle, Enrico Fermi and Paul Dirac developed the theory known as Fermi-Dirac statistics in 1926 [10, 11]. This is, for example, applicable to electrons. Interestingly, this behaviour is only observed at low enough temperatures or very dense systems; for high temperatures and low densities Boltzmann statistic describes particle behaviour well. In this thesis, we will concentrate on systems which can generally be described by Boltzmann statistics.

As systems with a large number of degrees of freedom are considered, the notion of microscopic and macroscopic properties is important. In the presence of microscopic fluctuations it is natural to consider macroscopic observables in order to obtain a probabilistic description of the system. One can describe average quantities of gas particles in a box based on their temperature. It may not be possible to distinguish individual particles and say exactly what velocity they have at a given point in time, but it is possible to deduce the average speed of all particles. In fact, the speed $v = |\mathbf{v}|$ follows the Maxwell-Boltzmann distribution:

$$f(v) = \sqrt{\frac{2}{\pi}} \left(\frac{m}{k_B T}\right)^3 v^2 \exp\left(\frac{-mv^2}{2k_B T}\right). \quad (1.1)$$

More generally, the equilibrium behaviour can be expressed in terms of the **partition function**, which will be discussed in Chapter 2. Therefore, in order to obtain an equilibrium description of a system of in-

terest, the knowledge of the partition function is essential. For example, when considering a system comprised of ideal gas particles the partition function can be written down exactly, but in more complex systems this is not necessarily possible. An alternative approach uses computational sampling techniques in order to estimate the equilibrium behaviour of a system. The next section will give an overview of commonly used computational techniques in order to achieve equilibrium sampling.

COMPUTATIONAL APPROACHES

With the advent of computers, algorithms have been developed that allow sampling of systems with large **state spaces** in order to obtain information regarding their equilibrium behaviour. These algorithms rely on **ergodic theory**, which means they make use of the fact that the time average of the system is the same as the ensemble average in the limit of infinite observation time and therefore sampling in time can be used in order to obtain estimates of the average equilibrium behaviour [12]. In particular, as the last 60 years have seen tremendous improvements in computational modelling due to the fast advancement of computational power (processor speed and data storage possibilities), which lead to algorithmic improvements and could be applied to increasingly complex systems. Generally, one can distinguish between two types of algorithms: **Monte Carlo (MC) simulations** and **molecular dynamics (MD) simulations**. Both aim to sample the equilibrium properties of the system of interest by exploring the state space, i.e. the space of all possible configurations the system of interest can take.

MC simulations are based on accepting newly proposed configurations of a system according to a probability function and were initially used to approximate high dimensional integrals. In this thesis **MC** simulations will be applied to stochastic systems on a lattice and will be discussed in detail in chapter 3 and chapter 4. The interested reader is also referred to a very exhaustive book by Landau and Binder [13] on the topic of **MC** simulations.

The other type of algorithm, **MD** simulations, is often applied to molecular systems, such as simulations of liquids, hard spheres or biological molecules to name a few. In this approach Newton's equations of motion are integrated to evolve the system in time. Mathematically, this means a time discretisation of the underlying differential equations is used in order to model it on a computer. These equations are taken from either Lagrangian and Hamiltonian mechanics or a stochastic counterpart such as Langevin or Brownian dynamics. Each particle of the system has a momentum and velocity. Generally, in order to initialise the system at a given temperature, velocities are randomly drawn from the Maxwell-Boltzmann distribution (see equation (1.1)).

At each time step the discretised differential equation will give information about the new velocity of each particle, based on the underlying **potential landscape**. In fact, the derivative of the potential is the force exerted on the particle at that point in **phase-space** – the space that carries the information of position and velocity of the particle. Depending on the system, model potentials can be anything from Lennard Jones type interactions, to Coulomb potentials, to model intermolecular interactions, to bond, angle and torsion potentials in order to model intramolecular interactions. This, in turn, leads to a vast area of research, that of parametrisation of interaction force fields based on experiments and quantum mechanical calculations.

In order to make valid statements about the equilibrium behaviour of the system, simulations look at the **average behaviour** of observables and their **fluctuations**. For the example of an Ising model **MC** simulation (discussed in Chapter 3), this could be the average magnetisation; in a **MD** simulation the average position of a particle/molecule in a potential may be of interest.

RARE EVENTS SAMPLING

Often the systems of interest do not have a straight forward equilibrium behaviour. For example, the well studied Ising model (used for studying ferromagnetic behaviour) exhibits a phase transition between a paramagnet and a ferromagnet. Near the critical point of the phase transition the convergence of the **MC** dynamics is very slow. Further, molecular systems are often plagued by metastability, which leads to inefficient sampling of the whole state space and thus the convergence to the global equilibrium is very slow or even infeasible within an appropriate simulation time. This sampling problem is even more pronounced at lower temperatures, as even less kinetic energy is available to the system. Therefore it becomes less likely to overcome a potential barrier with insufficient kinetic energy in the system. In the following, two commonly used methods that allow the extraction of the equilibrium behaviour while distorting the dynamics of the system will be briefly introduced; these are: Replica Exchange dynamics and Metadynamics.

Multi Ensemble Algorithms: Replica Exchange

In the following, a temperature ensemble is considered but the methodology has been employed to various different ensembles. At low temperatures a metastable system does not have sufficient kinetic energy in order to overcome potential barriers in the potential energy landscape. This is manifested by not accepting **MC** steps, or the oversampling of a metastable conformation in **MD** simulations. Metastable barriers can, however, often be overcome at higher simulation tempera-

tures (they vanish at infinite temperature). This property has been exploited by a family of ensemble algorithms, evolving different replicas of the system at the same time. Each replica samples the system at a different thermodynamic state, i.e. with different temperatures. These replicas are exchanged periodically according to a criterion ensuring that each replica samples from the global equilibrium distribution.

This was first proposed in 1986 by Swendsen and Wang with an application to a spin glass system [14] and has since seen various different flavours. In its most general form this algorithm is referred to as **replica exchange molecular dynamics (Monte Carlo dynamics) (REMD)** and can be explored in more detail, for example in a review by Earl [15]. The replica approach has found vast application in both **MC** simulations as well as **MD** simulations. These ensemble algorithms may generate correct equilibrium distributions e.g. different temperatures, but at the expense of generating unphysical dynamics. Generally, replica exchange algorithms are a very useful class of algorithms for sampling equilibrium properties by reducing correlation times for sampling independent configurations in comparison to using a single replica. **REMD** simulations will be discussed in greater detail in conjunction with the results of a novel estimator presented in chapter 5 .

Metadynamics

A different class of algorithms that allow fast and efficient exploration of the equilibrium properties of a complex system is metadynamics. For an extensive discussion of metadynamics see the review by Lao et al. [16].

Metadynamics is an algorithm that biases the system away from phase space regions that have already been sampled, by adding Gaussian functions to the potential function in these regions. From the sum of all the added Gaussian functions the free energy of the system and other interesting properties can be inferred. This type of algorithm is suitable for both **MC** and **MD** dynamics and has been successfully employed to many different systems [17, 18]. However, one drawback again is that it suffers from generating unphysical dynamics, but gives accurate information about the ensemble expectations. It can thus only be used for the estimation of ensemble averages (equilibrium properties). Its advantage over the ensemble algorithms is that only a single replica is required, thus computational cost is smaller than that of ensemble algorithms.

ENSEMBLES OF TRAJECTORIES BASED METHODS

Until now, we have discussed methods that are used in order to study the thermodynamic and thus stationary properties of systems. However, questions such as what is the rate of going from one metastable

state to a different metastable state are of great interest. This would for example allow the distinction of folding pathways in the protein folding problem. Generally, the dynamics of a stochastic system are given by the Fokker-Planck equation, i.e. the time evolution of the probability densities of the system from some initial point in time to some later point in time. Mathematically, this can be expressed by a **master equation** of the form:

$$\frac{dP}{dt} = \mathbf{K}P. \quad (1.2)$$

Here, P is a probability vector for all possible states at time t and \mathbf{K} is a rate matrix or **generator** containing transition rates between states. In most cases, just as it is impossible to write down the partition function, it is impossible to write down a rate matrix, or generator of the stochastic process exactly. Instead one has to think of methods that will generate an ensemble of reactive trajectories (by reactive, trajectories starting in an initial set A and reacting to a final set B , are meant). An example of such a method is transition path sampling (**TPS**) and was proposed by Geissler *et al.* in the early 2000's. [19, 20].

Ensembles of Reactive Trajectories

TPS is a method used to facilitate the generation of reactive trajectories in systems, where one tries to get from one thermodynamic state to a different one, for example in protein folding from an unfolded state to a folded state. The basic idea is to generate a Markov Chain, or a random walk on the space of dynamical trajectories using **MC** sampling. Newly proposed trajectories are only accepted if they are reactive, i.e. if they make it from A to B . The **TPS** algorithm has been referred to as "throwing ropes over high mountain passes in the dark" [19]. This can be explained in the sense that the throwing ropes is an attempt to get from one metastable state to a different one. This is particularly difficult in high dimensional state spaces, thus the "throwing" occurs in the dark, as no perfect route is known. In more mathematical language **TPS** can be explained as follows. A trajectory consists of a sequence of states observed with lag time τ , where τ is the interval of the time discretisation: $\mathbf{X}_{t_{\text{obs}}} = (\mathbf{x}_0, \mathbf{x}_\tau, \mathbf{x}_{2\tau}, \dots, \mathbf{x}_{t_{\text{obs}}})$. The probability of observing a given trajectory depends on the path ensemble. Thus if one defines substates $A, B \in \Omega$ and $A \cap B = \emptyset$ in a state space Ω , the probability of a reactive trajectory going from A to B is given by:

$$P_{AB}(\mathbf{X}_{t_{\text{obs}}}) = \Lambda_A(\mathbf{x}_0)P(\mathbf{X}_{t_{\text{obs}}})\Lambda_B(\mathbf{x}_{t_{\text{obs}}})Z_{AB}^{-1}. \quad (1.3)$$

Here, $\Lambda_S(\mathbf{x})$ (S is either set A or B) is a characteristic function which takes on the value of unity if $\mathbf{x} \in S$ and is zero otherwise. Z_{AB} is a normalisation constant that ensures $P_{AB}(\mathbf{X}_{t_{\text{obs}}})$ integrates to unity. All trajectories generated in this way obey the true dynamics of the

system. Generating trajectories in this way, the ensemble of reactive trajectories is explored.

This scheme will find application in chapter 3 and 4 in an adapted form in order to sample rare events in dynamic ensembles.

Ensembles of Non-Equilibrium Trajectories

Until now, we have given a detailed introduction to sampling equilibrium and dynamic properties of a system. As the thesis is mainly concerned with looking at non-equilibrium fluctuations in ensembles of trajectories, a short introduction will be given in the following. In the study of non-equilibrium dynamics, the evolution of the system depends not only on the current state, but also on time. To be more precise, such methods look at the time dependent fluctuations of ensembles of trajectories, i.e. a dynamic ensemble given for example by an observable capturing the so called **activity** of the system. In order to study such fluctuations, much of the methodology used for studying equilibrium properties can be reused. Trajectories can be described by such an activity observable, which will have certain expectations based on the choice of dynamics and set observational time of the system. Analogous to the thermodynamic partition function, a dynamic partition function over path averages, or trajectories can be defined. This will give information about the likelihood of observing a trajectory with a given dynamic observable. Therefore, understanding the physical conditions for rare trajectories is essential. In this case a rare trajectory does not mean a trajectory with a rare event such as a transition from A to B as discussed in the context of TPS, but trajectories whose activity is unlikely, i.e. a trajectory from the tails of an activity distribution. This concept will be introduced in much greater detail in chapter 2 and 3, but essentially in order to effectively sample rare trajectories a trick much akin to umbrella sampling or biasing the trajectories is needed. This is referred to as the *s*-ensemble.

Reweighting dynamic trajectories according to the *s*-ensemble has been very successful in describing glass forming system, or highly frustrated quantum systems [2, 4, 21, 22] amongst others and gives insight in the influence of non-equilibrium fluctuations on the equilibrium properties of the system. For example, in the case of glasses these fluctuations give rise to the slow relaxation times of glasses in the liquid state. In this thesis this formalism will be applied to the 1-d Ising chain in chapter 3 and a lattice model of a protein in chapter 4.

THESIS OVERVIEW

In the following, a brief overview of the subsequent chapters is presented:

Chapter 2 gives an introduction to the mathematical and physical concepts of statistical mechanics, with particular emphasis on the tools needed in the later chapters. The canonical ensemble is introduced, the Markov model theory and construction is outlined and Landau theory of phase transitions is briefly discussed. Lastly, the dynamic s-ensemble formalism and large deviation theory are explored.

In chapter 3 the Ising model and MC algorithms are introduced in detail. The thermodynamic properties of the 2-d Ising model are outlined, with particular interest in the scaling behaviour and universality class. Then the results of a novel and extensive computational study of the 1-d Ising model are presented, where a dynamic phase transition with the same universality class of the 2-d Ising model is observed. This is done using two different move-set MC algorithms: Glauber-Metropolis dynamics and Kawasaki dynamics.

The subsequent chapter 4 is concerned with the dynamic properties of a lattice protein model, a computational study we contrived and conducted using a variety of different interaction potentials. This is investigated using again the dynamic s-ensemble formalism: equilibrium trajectories are reweighted uncovering dynamic complexity in trajectory space. Trapping states are identified and it can be shown that these are able to influence the equilibrium dynamics significantly, by prolonging mean first passage times to the native state of the lattice protein.

Chapter 5 moves away from the MC dynamics, and will instead consider the results obtained from MD simulations of toy potentials as well as alanine dipeptide, in order to introduce a novel estimator (transition matrix analysis method (TRAM)) we proposed for obtaining the stationary behaviour of highly metastable systems. Here the use of Markov models is vital.

In chapter 6 the thesis is concluded with a discussion of all the results obtained and an outlook for potential future work is given.

THEORETICAL BACKGROUND – MARKOV MODELS,
LARGE DEVIATIONS AND THE s -ENSEMBLE

This chapter will give an introduction to the theoretical tools used for the analysis of equilibrium and dynamic properties in a variety of stochastic systems. In particular, it will introduce the concept of Markov models, the theory of Landau phase transitions and the ideas behind the large deviation formalism. Large deviations are then used in order to introduce the concepts of the s -ensemble, which allows to probe non-equilibrium fluctuations in ensembles of trajectories.

2.1 EQUILIBRIUM STOCHASTIC PROCESSES

This section will set the stage for the discussion of stochastic processes by defining a few commonly used terms. A stochastic process can be seen as a collection of random variables. This collection of random variables can be probed for its equilibrium as well as dynamic properties. In general, when we discuss random variables in the following we assume them to be independent identically distributed (iid).

2.1.1 Introduction: Coin Tossing

Thinking of one of the simplest stochastic processes, tossing coins comes to mind. We can define the "Laplace coin", whose probability of either heads (h) or tails (t) occurring is $\frac{1}{2}$. If we now consider a series of n coin flips, there are many different possible outcomes; 2^n to be precise. Thus, for two flips the set of possible outcomes already contains four entries:

$$\Omega = \{\{h, h\}, \{h, t\}, \{t, h\}, \{t, t\}\}.$$

Each of the possible outcomes is also termed a **microstate**.

In statistics, there are two important laws which govern a stochastic process such as coin tossing. The first one is the **law of large numbers** which states that, as the number of attempts (n) takes the limit $n \rightarrow \infty$, the sample average will almost surely converge to the true value [23]. Let us assume that x_n is a sequence of random coin tosses of length n , then it holds that

$$\lim_{n \rightarrow \infty} \langle x_n \rangle = \frac{1}{n} (x_1 + x_2 + \dots + x_n) = \mu. \quad (2.1)$$

Here μ is the true value or expected value and x_i is the i^{th} random coin toss. Therefore, an interesting choice of **observable** may be the probability of how many heads are found in a single sample, i.e.

$P(h_n) = \frac{1}{n} \sum_i x_i = \frac{1}{2}$ as n tends to infinity. Furthermore, the number of heads observed after n tosses, i.e. h_n can be seen as a **macrostate** of the system, whose probability of occurring is just governed by combinatorics:

$$P(h_n) = \frac{n!}{h_n!(1-h_n)! 2^n}. \quad (2.2)$$

Now, it is possible to consider a second important law, the **central limit theorem**, which states that random samples of length n with a finite mean and variance will converge in law to a normal distribution as $n \rightarrow \infty$. In this way the true value of the probability for each outcome can be estimated [23]. This can be approached in the following way: Rather than letting the number of samples go to infinity, instead the probability distribution of possible outcomes for an experiment of a given length n is observed. The **first central moment** or **sample average** of the probability distribution gives information about the most likely outcome often denoted by $\mathbb{E}(x_n)$, $\langle x_n \rangle$ or μ_n . The **variance** or **second central moment** gives information on the accuracy of the estimate given by $\mathbb{E}(x_i - \mu_n)^2$ or simply denoted by σ . The distribution is constructed from many repetitions or **realisations** of length n . For the case of the coin flipping, the probability density obtained from many realisations will be Gaussian, meaning it is of the form:

$$P(\mu_O) = \frac{1}{\sigma\sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{\mu_O - \mu}{\sigma} \right)^2 \right], \quad (2.3)$$

where σ is the variance and μ the mean of the distribution and μ_O the observed mean of a single realisation. The central limit theorem now tells us that if the number of realisations is large enough, the mean of the distributions of observed means ($P(\mu_O)$) will approach the true mean of the stochastic process.

In a more mathematical description, such a discrete stochastic process can be defined on a probability space (Ω, \mathcal{B}, P) . The probability space consists of Ω being the set of all possible outcomes, \mathcal{B} is the set of events and has to form a Borel σ -algebra over Ω , in the case of the coin flips a power set $\mathcal{B} = 2^\Omega$. The last part of the triple is the probability measure P such that $P : \mathcal{B} \rightarrow [0, 1]$. The probability of the whole sample space is $P(\Omega) = 1$. Thus, a stochastic process is a collection of random variables such that: $(X_t)_{t \in I}$ on Ω .

A physical interpretation of this is a **trajectory** in time t . If the random variables and the time increment are discrete, then such a set of events is known as a discrete Markov jump process, i.e. $I = \mathbb{N}$. In a situation of continuous time and a continuous state space then the process is referred to as a continuous time Markov process and $I = \mathbb{R}_+$. For more details on Markov processes see [24].

Let $\omega_i \in \Omega$ be a microstate. Thus, having defined microstates and macrostates, it is possible to apply the most fundamental postulate in statistical mechanics. In equilibrium the probability of being in a particular microstate is given by $P(\omega_i) = 1/|\Omega|$, assuming that Ω is the

set of all possible states. This means each microstate is equally likely. This is known as the microcanonical ensemble [25]. Of particular interest might be measurable observables, i.e. macrostates of the system, and in particular how likely a certain outcome is to be observed – such as a given number of heads h_n . This depends on how many ways there are to obtain the average divided by all possible outcomes (microstates). In the following section, we will explore how one can move from the microcanonical description of a system, where it is assumed that each state is equally likely, to a canonical description – where each states likelihood is dictated by its underlying potential.

2.1.2 Canonical Ensemble

The **canonical measure** or canonical/NVT ensemble of constant particle number, phase space volume and temperature will be introduced in this section. By using the maximum entropy assumption, the probability of being in microstate ω_i is defined as:

$$P_i = \frac{1}{Z} \exp(-\beta u_i), \quad (2.4)$$

where $Z(\beta)$ is called the partition function of the system, serving as a normalisation constant and u_i the energy of state ω_i and β the inverse temperature (see formal definition below). Z does always depend on β , but this will not be explicitly stated in the following. The derivation is achieved through the use of entropy maximisation, where we can define the entropy as $S = \ln \Omega$. Entropy is also defined later on by equation (2.12) Furthermore, an energy and particle constraint can be posed on the system leading to the use of Lagrange multipliers, where the energy constraint (see equation (2.6)) and a number of particle constraint (see equation (2.5)) are assumed such that:

$$\alpha \sum_i dn_i = 0 \quad (2.5)$$

and

$$\beta \sum_i u_i dn_i = 0, \quad (2.6)$$

with α and β being the Lagrange multipliers. Furthermore, the multiplier β associated to the energy constraint is proportional to the inverse temperature of the system, or to be more precise:

$$\beta = \frac{1}{k_B T}. \quad (2.7)$$

Generally, the Boltzmann constant k_B will be set to unity in this thesis; where this is not the case this will be explicitly stated. For a more detailed sketch of the derivation, refer to [25]. The normalisation constant Z is of particular interest, as it relates a lot of equilibrium prop-

erties of the system. It is given by the following sum over state space:

$$Z = \sum_i \exp(-\beta u_i). \quad (2.8)$$

For systems with degenerate energy states a factor called the density of states is introduced, but will not be further discussed here. The expression $\exp(-\beta u_i)$ is also known as the **Boltzmann factor**, representing the unnormalised probability of a given state. This allows the determination of some thermodynamic/equilibrium properties with respect to the partition function.

For example, quantities of interest may be the ensemble average of the total energy given by:

$$\langle E_t \rangle = -\frac{\partial \ln Z}{\partial \beta}, \quad (2.9)$$

and the fluctuations in the total energy:

$$\langle (E_t - \langle E_t \rangle)^2 \rangle = \frac{\partial^2 \ln Z}{\partial \beta^2}. \quad (2.10)$$

Another important observable is the free energy of the system, which is defined as:

$$\mathcal{F} = -k_B T \ln Z, \quad (2.11)$$

which is closely related to the entropy

$$\mathcal{S} = -\frac{\partial \mathcal{F}}{\partial T}. \quad (2.12)$$

Thus, it can be concluded that the knowledge of the partition function gives useful information about the equilibrium properties of the system. So far only systems with a discrete number of states (the coin tossing experiment) were considered. It is evident, however, that one can easily think of situations where Ω is not discrete but continuous. An example of such a system would be that of N identical gas particles, with $\Omega = \mathbb{R}^{6N}$. In that case the sum in equation (2.8), turns into an integral over all of state space, i.e.:

$$Z = \int_{\Omega} d\omega \exp(-\beta U(\omega)). \quad (2.13)$$

This integral is in most cases not easily solvable and thus finding ways of estimating this integral are important.

In the following we will consider the example of a diffusive particle in a 1-d potential, for which an exact solution of the integral of the partition function is available, in order to illustrate some the properties of the partition function.

2.1.3 Example: Single Particle in a Harmonic Potential

In the following a 1-d harmonic potential with a diffusive particle will be considered, in order to demonstrate the usefulness of the knowledge of the partition function for a continuous state space $\Omega = \mathbb{R}$. The example chosen here is exactly solvable. Let $U(x) = (x - 2)^2$ be the potential, depending on the – here unitless – position x on state space and let the corresponding partition function Z depending on the inverse temperature β be given by:

$$Z(\beta) = \int_{-\infty}^{+\infty} dx \exp(-\beta(x - 2)^2) = \sqrt{\frac{\pi}{\beta}}, \quad (2.14)$$

which is a standard Gauss integral in this case. The expectation of the energy is given by:

$$\langle U \rangle = -\frac{\partial \ln Z}{\partial \beta} = \frac{\partial \ln \sqrt{\pi/\beta}}{\partial \beta} = -\frac{1}{2} k_B T. \quad (2.15)$$

The equilibrium probability, $\pi(x)$ of each point in state space Ω is given by:

$$\pi(x) = \frac{\exp(-\beta U(x))}{Z(\beta)}. \quad (2.16)$$

This probability $\pi(x)$ is always normalised to unity by the partition function such that:

$$\int_{-\infty}^{+\infty} dx \pi(x) = 1. \quad (2.17)$$

Figure 2.1(a) shows the harmonic potential $U(x)$. Figure 2.1(b) shows the equilibrium probability for 3 different values of the inverse temperature β in reduced units, evaluated according to equation (2.16). Alternatively, figure 2.1(b) can also be obtained by "sampling" the potential. This would require observing the position of the particle in the potential as it performs a diffusive motion (e.g. Brownian motion), from this, the average position (equilibrium probability) of the particle can be extracted. The potential itself carries all necessary information for the equilibrium behaviour of the diffusive particle in this harmonic potential. Probing equilibrium properties by means of sampling is only necessary if equation (2.16) cannot be evaluated exactly.

Evaluating equation (2.16) in real world systems – one could think of a protein molecule, or a carbon nano tube – often poses a problem, as the dimensionality is too high and the state space is too complex, such that the solution can be written down exactly. This problem is also known as the **curse of dimensionality**. Hence, ways of estimating the equilibrium behaviour need to be found. This is often achieved by means of sampling and further reducing the dimensionality of the system of interest. Generally the dimensionality reduction is achieved by projecting the high dimensional continuous state space – which

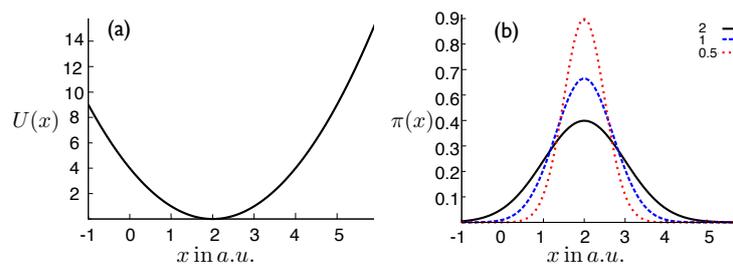


Figure 2.1: (a) Potential $U(x)$. (b) The normalised equilibrium distribution of the harmonic potential (a) for three different temperatures $1/\beta = \{0.5, 1, 2\}$ in reduced units.

is unknown – onto a single or even multiple **reaction coordinates**, along which a histogram can be constructed. Using this approach a lot of information is evidently lost, therefore instead one can think of clustering the data into a high dimensional histogram and in this way captures more of the properties of the original system. For evaluating stationary and dynamic properties of a system discretised in such a way, the theory of a Markov State model (**MSM**) is very useful and will be discussed in the following section.

2.2 MARKOV STATE MODELS

Markov state models (MSMs) are stochastic models which obey the Markov property, which means that the system does not carry any memory of a past state or future state in a time evolution process. An example of such a process would be a random walker, or the diffusive particle in the 1-d well discussed in section 2.1.3. **MSMs** use the notion of the Markov property in order to obtain a quantitative description of the stochastic process. In fact from the Markov property it follows that the system is fully described by the generator of the stochastic process.

The construction of these models will allow the reconstruction of equilibrium and kinetic behaviour of the particular stochastic process of interest, therefore it is a very useful tool. Generally, the idea behind the construction of an **MSM** is to approximate the generator \mathbf{K} of the stochastic process. The generator – from the Master Equation of the process – is an operator that contains the transition probabilities between states of the system. In most systems of interest, e.g. molecular systems, the curse of dimensionality does not allow for the generator to be written down exactly, or be approximated by means of finite element methods. Instead, it can be approximated by data obtained from simulations which are obtained evolving the system according to the underlying stochastic dynamics. These then form the basis of the **MSM**. In the following we will first introduce the **transfer operator** and its relation to the generator and discuss some of its very useful properties.

Then we will move towards a discretisation of the state space and introduce the discrete equivalent of the transfer operator, the **transition matrix**. Next it will be discussed how one can arrive at an estimate of a transition matrix from an actual observed trajectory. This will be discussed using a discrete example only allowing two possible states. Lastly, it will be shown how such a model can be applied to continuous state spaces and an example will be given of a model where the state space is given by $\Omega = \mathbb{R}$. It will then be alluded to how this can be extended to much larger state spaces, e.g depending on the number of atoms (N) contained in a system, such that $\Omega = \mathbb{R}^{3N}$.

2.2.1 Formal Definition of a Markov Process

We assume there exists a time continuous stochastic process $(\mathbf{X}_t)_{t \geq 0}$ on a probability space (Ω, \mathcal{B}, P) . If the Markov property for \mathbf{X}_t holds, meaning that the history of the system does not influence the future of the system the transition probability between states in Ω fully describe the dynamics. Let us denote two different states in Ω in a continuous setting as \mathbf{x} and \mathbf{y} . In the following transition probabilities between these points:

$$p(\mathbf{x}, \mathbf{y}; dt) d\mathbf{y} = P[\mathbf{x}_{t+dt} \in \mathbf{y} + d\mathbf{y} | \mathbf{x}(t) = \mathbf{x}] \quad \forall \{\mathbf{x}, \mathbf{y}\} \in \Omega \quad . \quad (2.18)$$

Equation (2.18) gives the probability of the system transitioning from initial state \mathbf{x} at time t to state \mathbf{y} after a time interval dt . Furthermore, $\mathbf{y} + d\mathbf{y}$ could be defined as a whole region (subset of Ω) out of which, or into which probability can flow. It is also useful to assume that the Markov process is **ergodic**. This means all points \mathbf{x} on Ω are dynamically connected and as $t \rightarrow \infty$ every \mathbf{x} will be visited infinitely often [12]. The probability density for each element \mathbf{x} in the infinite limit is the stationary probability of the system given by equation (2.16), where $U(\mathbf{x})$ can in such a case be described by the Hamiltonian $\mathcal{H}(\mathbf{x})$. Throughout the thesis $U(X)$ and $\mathcal{H}(x)$ are often used interchangeably. From ergodicity, the uniqueness of the stationary distribution follows. The Hamiltonian can, for example, be describing the energy terms of a protein molecule. The last assumption made is that the dynamics are reversible, which means that the **detailed balance** condition must hold:

$$\pi(\mathbf{x})p(\mathbf{x}, \mathbf{y}; dt) = \pi(\mathbf{y})p(\mathbf{y}, \mathbf{x}; dt) \quad . \quad (2.19)$$

Thus, detailed balance means that the probability of going from one state to another and the reverse are the same, weighted by the equilibrium probability $\pi(\mathbf{x})$. Again, the notion of a trajectory can be used and is defined as a sequence of states \mathbf{x} from time $t = 0$ to time $t = t_{\text{obs}}$ and denoted by \mathbf{X}_t . For a more rigorous definition and discussion of Markov processes, see [24].

2.2.2 Transfer Operator Approach

In the following, we will first discuss how the dynamics on the full/continuous state space behaves, before considering its discretised analogy. As with any stochastic process, predicting the behaviour of a single realisation is impossible, due to the randomness of the process. However, it is possible to say something about the averaged properties of the ensemble. Therefore, it is of interest to investigate the time evolution of a probability density, resulting from an ensemble of different realisations. The Master equation (1.2) allows to propagate a probability density in time. This is achieved by the generator \mathbf{K} of the stochastic process. For convenience, we will define the transfer operator \mathcal{T} in this section, which is related to the generator in the following way:

$$\mathcal{T}(\tau) = \exp(\tau\mathbf{K}) , \quad (2.20)$$

where τ is a lagtime, at which the process is observed. The timestep of a numerical integrator dt can, for example, represent such a lagtime. The transfer operator propagates a weighted probability density f_t , related to the probability distribution $p_t(\mathbf{x})$ by the equilibrium probability $\pi(\mathbf{x})$:

$$p_t(\mathbf{x}) = f_t(\mathbf{x})\pi(\mathbf{x}) , \quad (2.21)$$

with $p_t(\mathbf{x})$ being the probability of being in \mathbf{x} at time t , in the following way:

$$f_{t+\tau}(\mathbf{y}) = \mathcal{T}(\tau) \circ f_t(\mathbf{y}) = \frac{1}{\pi(\mathbf{y})} \int_{\mathbf{x} \in \Omega} p(\mathbf{x}, \mathbf{y}; \tau) \pi(\mathbf{x}) f_t(\mathbf{x}) . \quad (2.22)$$

The transfer operator is a non-negative operator and therefore fulfils the requirements of the Perron-Frobenius theorem. This allows the deduction of a few useful well known properties with respect to Markov models [26]:

1. The powers of $\mathcal{T}(\tau)$, i.e. $\mathcal{T}(\tau)^n$, obey the Chapman-Kolmogorov equation, such that

$$\mathcal{T}(\tau)^n = \mathcal{T}(n\tau) . \quad (2.23)$$

2. The continuous eigenvalue spectrum has some interesting properties if the transfer operator models a reversible process. All eigenvalues are real-valued and between $[0, 1]$, i.e the largest eigenvalue will always be $\lambda_1 = 1$ and all subsequent eigenvalues will be real valued and decreasing as $\lambda_i \geq \lambda_{i+1} \geq 0$ in all cases considered here. The set of m dominant eigenvalue and eigenfunction pairs can be written as:

$$\mathcal{T}(\tau) \circ \psi_i(\mathbf{x}) = \lambda_i \psi_i(\mathbf{x}) . \quad (2.24)$$

These dominant eigenvalues and eigenfunction are of particular interest to the system, and are again dependent on the lagtime τ . This leads to the fact that the spectrum holds information about metastable behaviour of the system.

3. Metastability of the system can be observed through a gap, i.e. $\lambda_m \gg \lambda_{m+1}$ in the eigenvalue spectrum, such that a distinctive separation of fast and slow processes is possible. All slow processes (the m dominant eigenfunction and eigenvalue pairs, see equation (2.24)) will have eigenvalues close to the dominant value of unity and all fast processes can be found in a ball around 0.
4. The eigenvalues corresponding to the slow processes can be associated with an implied timescale (t_i) of the underlying dynamics of the system through the following equation:

$$t_i = -\frac{\tau}{\ln \lambda_i} . \quad (2.25)$$

This timescale is related to the mean first passage time between the slow regions of the system (e.g. between two metastable basins of a double well potential as seen in figure 2.2.)

Of course a lot more can be said about the transfer operator, therefore for a more detailed discussion the reader is referred to [26].

Unfortunately the transfer operator suffers the same curse of dimensionality as the generator \mathbf{K} , in the sense that it can only be approximated for systems of low dimensionality, where an operator approximation by means of finite elements computation is possible [27].

2.2.3 Discrete Approximation of the Transfer Operator – the Transition Matrix

Until now Markov processes in continuous time and on a continuous state space were considered. Most of the time, however, trajectories on a set of discrete states are used. This leads to an approximation of the transfer operator by the so called transition matrix, holding transition probabilities between discrete states. For the estimation of this transition matrix one can generally distinguish between two different cases: The first is an estimate from a trajectory obtained from a discrete Markov jump process, where the observed dynamics has predefined discrete states. The second is an estimate obtained from a trajectory generated on a continuous state space. In the latter case the state space requires discretisation first, in order to turn the state continuous trajectory into a discrete trajectory.

A discrete trajectory is a sequence of integer number, each of which identifies a unique state in state space (i.e a naturally discrete state) or a region on the continuous space. Let us assume that we have $\{1, \dots, k\} \in K$ such states. This means that an exact number of K

states can be identified and the transfer operator can be approximated by a transition matrix $\mathbf{T}(\tau) \in \mathbb{R}^{K \times K}$. The transition matrix is a non-negative row stochastic matrix and is required to be irreducible, for all cases discussed in this thesis. For matrices of this kind the Perron-Frobenius theorem holds and properties ((1)-(4)) from section 2.2.2 are also valid for the transition matrix. Again, \mathbf{T} will also depend on the chosen lagtime τ of the system, and will be implicitly assumed, but will be neglected again in the notation, i.e. \mathbf{T} means it has a certain lagtime τ . In the following one property will be highlighted as it will play a more prominent role in chapter 5. The eigenvector corresponding to the trivial eigenvalue $\lambda_1 = 1$ represents the stationary probability of the system, and thus the following holds:

$$\boldsymbol{\pi} \mathbf{T}^\top = \mathbf{T}^\top \boldsymbol{\pi} \quad (2.26)$$

The vector $\boldsymbol{\pi}$ has K elements and holds the equilibrium probability of each state as its entries.

2.2.4 Transition Matrix Estimation

The question is how to obtain the transition matrix from the observed trajectories of a simulation or experiment. The transition matrix is given by the discrete transition probabilities between each configuration visited in the simulation at a given lagtime τ . How these can be estimated will be discussed in the following sections.

2.2.4.1 Count Matrix – Stochastic Switch Example

Let us assume we have a system with some unknown underlying properties, but we can observe a light switching on and off in a random manner. Thus we can define a state space consisting of two states: $\Omega = \{\text{on}, \text{off}\}$. A measurement is taken every second and after n seconds a trajectory (\mathbf{X}_t) which may look like this is observed:

$$\mathbf{X}_t = \{t_1 = \text{on}, t_2 = \text{on}, t_3 = \text{on}, t_4 = \text{off}, t_5 = \text{off}, \dots, t_n = \text{on}\} \text{ .}$$

Now, the underlying stochastic process which we do not know anything about, dictates the switching of the light, but with the recorded observation we have counts of how many times each state, as well as switches between states are observed. An example of this is found in table 2.1. In this way we can define a count matrix $\mathbf{C}(\tau)$. Let us assume that we have taken 1000 measurements, i.e. a trajectory that is 1000s long and we construct $\mathbf{C}(\tau)$ at a "native" lagtime, i.e. $\tau = 1\text{s}$ in this case. It is very straight forward defining a count matrix from table 2.1:

$$\mathbf{C}(\tau = 1) = \begin{pmatrix} 300 & 100 \\ 100 & 500 \end{pmatrix} \text{ .} \quad (2.27)$$

State Transition	Observed number of times
on \rightarrow off	100
off \rightarrow off	500
off \rightarrow on	100
on \rightarrow on	300

Table 2.1: Number of times a state transition is observed in the light switching example from a 1000s trajectory.

Taking the limit of an infinitely long trajectory, the counts out of and into states at the lagtime τ will be exactly the correct transition probabilities, if the underlying dynamics is Markovian. As it is not possible to generate infinitely long trajectory, conditional probabilities are used in order to estimate a transition matrix from the observed counts, given a certain lagtime.

2.2.4.2 Maximum Likelihood Estimation

The aim is to estimate a most likely transition matrix based on the observed trajectory at lagtime τ . The conditional probability for a transition matrix, from the observed counts is given by:

$$p(\mathbf{T}|\mathbf{C}) \propto \prod_{i,j} T_{ij}^{C_{ij}} . \quad (2.28)$$

The maximum likelihood estimator of the transition matrix is trivial and is recovered by row normalising the count matrix $\mathbf{C}(\tau)$. The estimate for the transition matrix $\hat{\mathbf{T}}$ is given by the following entries (not the hat is not always explicitly written. All transition matrices from simulations are assumed to be estimates):

$$\hat{T}_{ij} = \frac{C_{ij}}{\sum_j C_{ij}} , \quad (2.29)$$

which is just

$$\hat{\mathbf{T}} = \arg \max p(\mathbf{T}|\mathbf{C}) , \quad (2.30)$$

assuming that transition matrices are uniformly distributed. Looking back at the light switch example a maximum likelihood transition matrix just becomes:

$$\hat{\mathbf{T}}(\tau = 1) = \begin{pmatrix} 0.75 & 0.25 \\ 0.17 & 0.83 \end{pmatrix} , \quad (2.31)$$

resulting in an equilibrium probability of $\boldsymbol{\pi} = \{0.4, 0.6\}$, i.e 40% of the time the light is on and 60% of the time the light is off. This works well for a 2×2 system, but for a system with more than two states,

the problem arises that with this method of estimation, \mathbf{T} is not necessarily reversible, meaning that the entries do not necessarily fulfil detailed balance. The detailed balance condition for the entries of the transition matrix estimate mean that $\pi_i T_{ij} = \pi_j T_{ji}$ is true, while π_i is the stationary probability of being in state i . Generally reversibility is desired, as it ensures only real valued eigenvectors. A method for the reversible estimate of a maximally likely transition matrix based on the observed counts has been proposed by Prinz *et al.* [28] and involves an iterative pairwise update of the matrix entries, weighted by the stationary probability distribution for each updated state.

2.2.5 Large Systems with Continuous State Spaces

Until now only the small two states switch example, which was discrete by construction, was considered. However, Markov models can also be very useful for classifying equilibrium and dynamic behaviour of continuous stochastic systems. Thus as an example Markov models can be built from MD simulations. Here the task arises to discretise the dynamics in time but also reduce the size of the state space to a manageable set of representable states, for example, in the case of a large protein system. The dynamics (Langevin dynamics are considered here) need to be discretised so that the stochastic differential equation can be integrated computationally. The Langevin equation is given by:

$$m \frac{d^2 \mathbf{x}}{dt^2} = -\mathbf{F}(\mathbf{x}) - \zeta m \frac{d\mathbf{x}}{dt} + \boldsymbol{\eta}(t) , \quad (2.32)$$

where m is the mass of the particles, \mathbf{x} the particle positions, $\mathbf{F}(\mathbf{x})$ the force acting on each particle, ζ a friction coefficient and $\boldsymbol{\eta}(t)$ a random white noise. Equation (2.32) can, for example, be integrated by a leap frog algorithm, discussed in more detail in chapter 5. The size of the time step is dictated intrinsically by the system and for molecular simulations generally in the range of femtoseconds [29]. In this way the continuous Markov process now becomes a time discrete Markov chain. With this a molecular system can be evolved in time and at every time step a snapshot of the state space \mathbf{x} can be taken. This makes up a trajectory \mathbf{X}_t . Now in order to propagate a set of path ensembles from time $t = 0$ to some later time, the knowledge of the transfer operator is essential, but is generally not known exactly. Instead, it can be approximated by a transition matrix, but now a discretisation of the continuous state space will be necessary, i.e. many points in continuous state space are considered as one discrete state. The overall space discretisation is ideally done in such a way that the reduced state space becomes more manageable but also retains the important dynamic and static properties of the system one wishes to study. Generally a crisp

discretisation is chosen, meaning that each point in the whole state space Ω is sorted into a clustered state $S = \{S_1, \dots, S_k\}$, such that:

$$\bigcup_{i=1}^K S_i = \Omega \quad \text{and} \quad S_i \cap S_j = \emptyset \quad \forall i \neq j . \quad (2.33)$$

In this case K is the total number of discrete states. Discretisation introduces three main problems.

- The choice of discretisation:
Finding a good discretisation which resolves all relevant slow dynamics of the system can be difficult. There are a number of different clustering algorithms available, such as k-center clustering or k-means clustering [30] and many more. K-means is a standard clustering algorithm and clusters a d-dimensional vector into K sets, and minimises the within cluster sum of squares according to a Voronoi partitioning. Mathematically speaking the following expression is minimised:

$$\arg \min_S \sum_{i=1}^K \sum_{\mathbf{x}_j \in S_i} \|\mathbf{x}_j - \mu_i\|^2 , \quad (2.34)$$

where the set S contains K sets, which make up the discretisation of the real d-dimensional vector \mathbf{x} and μ_i is the mean of the i^{th} set. These sets can be further coarse grained into metastable sets. This can, for example, be achieved by the Perron-Cluster Cluster Analysis (PCCA) [31]. One can think of many other methods of clustering data, but these do not have immediate relevance for this thesis. Many research efforts are being made in identifying ideal clustering for Markov model analyses.

- Loss of Markovianity:
Through the clustering, the trajectory can be discretised according to the different clusters and can be described by a natural number assigning the current state of the trajectory to its corresponding cluster (see the two state switch example). Through the projection of the high dimensional state space onto the K discrete clusters, the Markov property of the stochastic system is lost. In order to overcome this problem, generally the systems studied are considered at "sufficiently large" lagtimes, for which the system becomes Markovian again. Establishing such lagtimes is not straight forward.
- Projection error:
From the projection of the whole state space Ω onto the discrete set of states, a projection error has to be accounted for and estimated. However, it can be shown that for a fine discretisation the projection error is sufficiently small. For more details on how to estimate projection errors, refer to [28].

Once a discrete trajectory is obtained, it can be used to construct a count matrix $\mathbf{C}(\tau)$ from which the transition matrix can be estimated, as discussed previously in section 2.2.4, and in this way approximating the transfer operator $\mathcal{T}(\tau)$.

2.2.6 Example: Single Particle in a Double Well

For the purpose of illustration, an example of a double well will be considered. The double well is illustrated in figure 2.2(a), where the state space, here in the 1-d case, is discretised into 50 states (given by the points along the curve). The dotted vertical lines represents every 10th state boundary). The results shown in figure 2.2 are obtained from simulations at $1/\beta = 10k_B T$ (k_B is set to unity, as well as all other constants, thus working in reduced units), where a single particle is diffusing according to Langevin dynamics in the potential. The trajectory of 1×10^5 data points in continuous space was discretised into 50 states on which a Markov model transition matrix was estimated. The first 4 eigenvalues including the eigenvalue gap are depicted in figure 2.2(b). The eigenvector corresponding to the eigenvalue of $\lambda_1 = 1$ is that of the equilibrium distribution and seen in figure 2.2(c) – denoted in this case by ψ_i . The second largest eigenvalue corresponds to the exchange between the two basins, whereas all other dynamics are diffusive properties within the two basins. Figure 2.2(d) shows the second eigenvector corresponding to λ_2 . The sign change observed in ψ_2 shows that the slow dynamics occurs from the left basin to the right basin and the rate at which this occurs is captured in the eigenvalue λ_2 , as given by equation (2.25). The transition matrix, eigenvalues and eigenvectors were estimated with the Markov model software EMMA [32].

This means that with the aid of Markov models, even for complex systems it is relatively straight forward to define an equilibrium probability for the occupation of different discrete states of the system. Furthermore, if the system exhibits a two state or even multiple state dynamics, it is possible to estimate a lower bound of the timescales for these which can then even be compared to experimental results [33].

2.2.7 Markov Model Estimation

In order to facilitate the estimation of Markov models and transition matrices, a software package was put together by some members of the Computational Molecular Biology research group at the FU-Berlin and myself. My contribution to this software project was mainly to provide a Java implement of existing algorithms, fitting in the overall design framework of the software and help with the design and testing of the tutorial material. The software EMMA is available for download at www.simtk.org/emma/ and can be used to build and validate Markov models from simulation or experimental data. All relevant in-

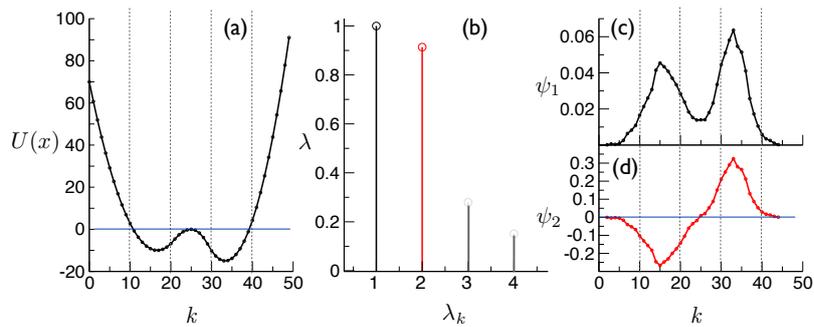


Figure 2.2: (a) The double well potential with discretisation, indicated by the points along the curve. Vertical lines are state boundaries at every 10th state. (b) Shows the first 4 eigenvalues of the 50×50 transition matrix at $\tau = 10$. (c) Shows the right eigenvectors corresponding to the first eigenvalue, i.e the stationary distribution and (d) shows the second eigenvector. Blue lines indicate $U(x) = 0$ in (a) and $\psi_2 = 0$ in (d).

formation in conjunction with the software is published in the Journal of Computation and Theory [32]. The software is command line driven with a series of command line tools, suitable for Windows and UNIX systems, that allow the construction of a Markov model. The following steps are usually carried out.

- Clustering and assignment of the data to the clusters, either by a provided clustering algorithm or by manually selected cluster centres.
- Probing the connectivity of the discrete sets and if necessary restricting the estimation of the transition matrix to the largest set of connected data.
- Construction of a transition matrix and its analysis, which provides eigenvalues and the corresponding eigenvectors as well as the matrix itself and the stationary distribution.
- Estimation of implied timescales or transition paths, if required.

There are some further features available not discussed here, as they have no further relevance to this thesis.

2.3 PHASE TRANSITIONS AND THEIR CLASSIFICATION

Moving away from Markov models, the next section will be concerned with further exploring the phase space of stochastic systems, in particular considering phase transitions. Many physical systems exhibit phase transitions. The freezing transition of water is an obvious example that comes to mind. It is generally not trivial to study phase transitions analytically other than using mean field approaches. Ehrenfest

introduced a classification of phase transitions, based on the derivatives of the free energy of the system [34]: a discontinuity in the first derivative is a first-order transition, a discontinuity in the second derivative is a second-order transition, and so on. Only the first- and second-order (continuous) phase transition will be discussed here.

2.3.1 The Order Parameter

In order to observe a phase transition a parameter that captures the phase behaviour needs to be defined. This parameter is referred to as the **order parameter** of the system. It is a parameter that changes abruptly if the system undergoes a first-order phase transition, or continually if it is a second-order transition. Figure 2.3 shows a sketch of the order parameter ρ with respect to temperature for a first- (a) and second-order (b) transition. An example of a first-order transition would be the transition between liquid and vapour where the density serves as the order parameter and abruptly changes as the volume or pressure of the system is varied. This transition ends in a tricritical point of a second-order transition, at a certain combination of parameters of temperature and pressure of the system, where all phases are present at the same time – i.e. liquid, vapour and solid. At the critical point a clear distinction of the liquid phase from the gas phase is not possible. However, a commonly discussed example of second-order phase transitions are magnets, where at a low enough temperature spontaneous symmetry breaking occurs and a transition from a disorder paramagnet to an order ferromagnet occurs. This is, for example, discussed in greater detail in [13, 35]. In the case of the magnet a

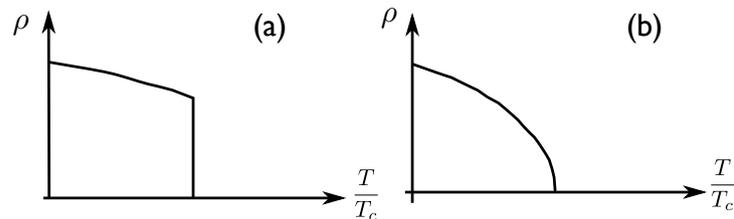


Figure 2.3: (a) First-order phase transition, with a discontinuity in the order parameter ρ . (b) Second order or continuous phase transition, with an analytic function for the order parameter ρ . The critical temperature is denoted by T_c .

spontaneous magnetisation is observed below the critical temperature, and thus the magnetisation serves as an order parameter. It should be noted however, that it is not always straight forward to define an order parameter for a given system, as it requires physical insight into the nature of the system.

2.3.2 Landau Theory of Second-Order Transitions

In the late 1950's, Lev Landau constructed a theory of second-order phase transitions which relies on the fact that the free energy of the system can be described entirely by the order parameter of the system. Any microscopic knowledge of the system is disregarded [36]. From the free energy, the order of the phase transition is fixed and the function describes all phases of the system. Also, in equilibrium the derivative of the free energy with respect to the order parameter is zero, $\frac{\partial \mathcal{F}}{\partial \rho} = 0$.

Expanding the free energy using a Taylor expansion with respect to the order parameter ρ gives information about the phase behaviour close to the critical temperature.

$$\mathcal{F}(T, \rho) = \mathcal{F}_0 + \rho + a\rho^2 + b\rho^3 + c\rho^4 + \dots, \quad (2.35)$$

where all coefficients ($a - c$) are smooth functions of temperature. The order parameter ρ will be vanishing for values above a critical temperature and therefore a minimum of \mathcal{F} at $\rho = 0$ must exist. When, considering a magnetic system, due its symmetry $\mathcal{F}(\rho) = \mathcal{F}(-\rho)$, only even powers of the order parameters need to be considered and a free energy for the second-order phase transition can be evaluated to be:

$$\mathcal{F}(\rho, T) = \mathcal{F}_0(T) + a(T - T_c)\rho^2 + c\rho^4, \quad (2.36)$$

where c is a constant that is larger than 0. Figure 2.4 shows the Landau free energy with respect to the order parameter for the constant $a > 0$, $a = 0$ and $a < 0$. In the case where the cubic term in equation (2.35)

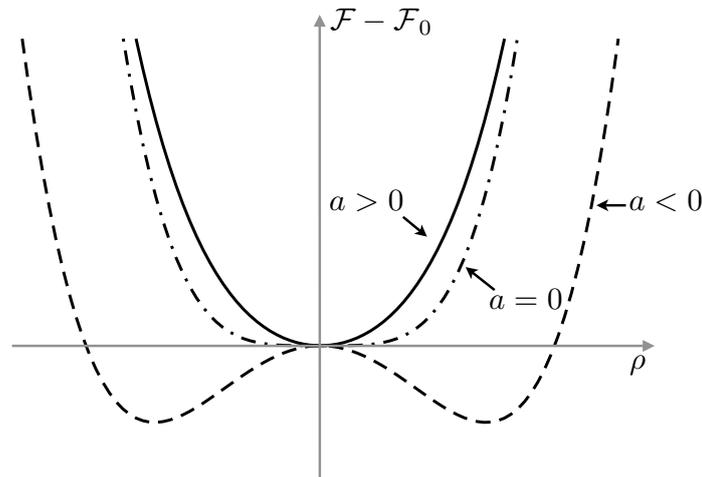


Figure 2.4: The free energy as an expansion in the powers of the order parameter for a system with symmetry. The behaviour for subcritical (continuous line), critical (broken-dotted line) and supercritical (broken line) temperatures is shown.

cannot be neglected, as in the case of a liquid gas transition, the Landau approximation does not give clear information about the nature of the transition. Now both a second-order as well as first-order transition can be present. For a first-order transition, the order parameter is discontinuous at the critical temperature, thus the expansion in ρ does not capture the phase behaviour any more and a different approach needs to be taken to evaluate the free energies. Such approaches are, for example, that of Renormalisation group theory. This is discussed in [37].

2.4 LARGE DEVIATION THEORY AND THE S-ENSEMBLE

In the following section, the s -ensemble formalism will be introduced which allows the investigation of dynamic phases in many different systems. This is done by means of biasing path averages according to their activity or other time integrated observables. First, some aspects of large deviation theory (LDT) will be discussed, as the s -ensemble formalism relies on this theory in order to define a dynamic free energy. Once the LDT is established, the dynamic ensemble is defined further by drawing analogies to the well known canonical ensemble and the section will be concluded with an overview of applications for this methodology.

2.4.1 Large Deviation Theory

In this subsection the notion of LDT will be introduced. Loosely speaking, as the name suggests, it looks at the large deviations of a probability distribution, i.e. its tails. In terms of computational sampling these represent the rare events which occur in a system. These can be the escape rates from a metastable basin. A trivial example would be the double well potential, see figure 2.2(a) of section 2.2.6. The formulation of LDT has been mainly attributed to Donsker and Varadhan [38, 39, 40] in the 1960's and 70's, but the principle has been previously applied, for example, by Cramér [41]. In fact, Boltzman used this technique in 1877, when he calculated equilibrium macrostates in 2-d turbulences [42]. For a physicist, studying statistical mechanics, the LDT gives information on how to deal with the exponential decay of probabilities of large deviations in stochastic processes.

2.4.1.1 Coin Tossing Revisited

In order to explain this concept further the coin tossing experiment from section 2.1 is revisited. Let us assume that μ_{O} is the sample average after n tosses. The question that poses itself now is, what is the dominant behaviour of observing a sample mean $r = \mu_{\text{O}}$ as n grows large. The probability of observing any sample mean is just the proba-

bility of observing a macrostate, as given by equation (2.2) multiplied by the number of coin tosses, thus:

$$P(\mu_O = r) = \frac{n!}{(rn)![(1-r)n]!} \frac{1}{2^n} , \quad (2.37)$$

where the number of observed heads is $h = rn$. Using Stirling's approximation $\ln n! \approx n \ln n - n$ results in the probability $P(\mu_O = r)$ being approximated by an exponential with a rate function

$$I(r) = \ln 2 + r \ln r + (1-r) \ln(1-r)$$

for its dominant behaviour, that is:

$$P(\mu_O = r) \approx \exp(-nI(r)) , \quad (2.38)$$

provided that n is large. The rate function $I(r)$ is what is sought after in LDT and constitutes its main result. For the case of the coin tossing experiment, the rate function gives information about how likely it is to observed averages away from the value of $P(\mu_O) = \frac{1}{2}$. As n gets large, the likelihood of observing averages larger or smaller than $\frac{1}{2}$ falls exponentially. In figure 2.5 the rate function is plotted and increases as r moves away from the mean, thus according to the definition of the probability of observing events far away from the mean from equation (2.38), this probability exponentially decreases if r moves away from 0.5.

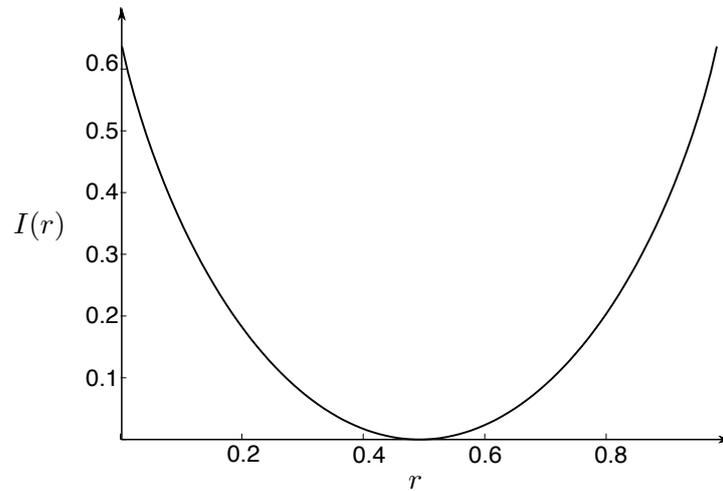


Figure 2.5: Rate function $I(r)$

2.4.1.2 A Formal Definition

Generally, the aim is to establish whether a LDT for a random number/process exists and what the corresponding rate function is. A large deviation principle exists if the probability P_n can be approximated

by an exponential, i.e. $P_n \approx \exp(-nI)$. Let A_n be a random variable that depends on a sum of n random variables itself. According to [43], a large deviation principle is satisfied, if A_n has the probability $P(A_n \in B)$ of taking a value in a set B , which has a dominant exponential behaviour, such that the limit

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \ln P(A_n \in B) = I_B \quad (2.39)$$

exists and I_B is the corresponding rate function. See Touchette's review [43] for more details on estimating upper bounds if the limit does not exist, or how to deal with continuous random variables. Evaluating the rate function is not always as straight forward as in the example of the coin tossing where the rate function appears directly out of the application of the Sterling approximation to the exact state probabilities. A different approach can be taken which will always work for strictly convex and differentiable scaled cumulant generating functions $g(k)$. Note, that the cumulant generating function is the logarithm of the moment generating function (MGF) and can be defined by the following limit:

$$g(k) = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \langle \exp(nkA_n) \rangle , \quad (2.40)$$

where A_n is a real random variable and $k \in \mathbb{R}$. The average can also be written in integral form such that:

$$\langle \exp(nkA_n) \rangle = \int_{\mathbb{R}} \exp(nkb) P(A_n \in db) . \quad (2.41)$$

The shorthand notation $db = [b, b + db]$, meaning that $P(A_n \in db)$ is the probability of A_n being in the interval $[b, b + db]$, a more lenient approach that using an exact value b . If for the random variable A_n a large deviation principle exists, i.e

$$P(A_n \in db) \approx \exp(-nI(b)) , \quad (2.42)$$

then the Gärtner-Ellis theorem states that the rate function $I(b)$ is given by a Legendre-Fenchel transform [43], provided the scaled cumulant generating function fulfils the requirement of being differentiable and convex. Thus $I(b)$ is given by

$$I(b) = \sup_{k \in \mathbb{R}} \{kb - g(k)\} . \quad (2.43)$$

Here, \sup is the supremum. For more details on the Gärtner-Ellis theorem and the Legendre-Fenchel transform see the discussion as presented in the review by Touchette [43].

2.4.1.3 Large Deviations: Free Energy and Entropy

The large deviation principle has many useful features and many different applications. In the following, the relation between the rate

function and the free energy and entropy will be sketched out in the case of looking at the large deviations of the mean energy of a system of interest. The notation was mainly borrowed from Touchette [43], and for brevity not all the steps are presented. In general the review is highly recommended for further details. With the following set up it is possible to relate the entropy density and free energy to each other through the Legendre transform. A system with N particles and ω_i microstates a set or state space of $\Omega_N \subseteq \mathbb{R}^N$ containing all microstates. The interaction energy between particles is defined through a Hamiltonian $\mathcal{H}_N(\omega_i)$ with an average energy per particle $\langle E_N \rangle = \mathcal{H}_N(\omega_i)/N$. Now we are interested in the large deviations of the probability of observing an average energy $\langle E_N \rangle$ in an infinitesimal energy interval $[u, u + du] = du$, i.e.

$$P(\langle E_N \rangle \in du) = \int_{[\omega_i \in \Omega_N: \langle E_N(\omega_i) \rangle \in du]} P(d\omega_i) . \quad (2.44)$$

This then defines a volume of microstate ω_i comprising the energy interval du , such that:

$$V(\langle E_N \rangle \in du) = \int_{[\omega_i \in \Omega_N: \langle E_N(\omega_i) \rangle \in du]} d\omega_i . \quad (2.45)$$

Therefore, it is possible to define a rate function, which captures the behaviour of observing average particle energies outside of the energy interval du , by:

$$I(u) = \lim_{N \rightarrow \infty} -\frac{1}{N} \ln P(\langle E_N \rangle \in du) \quad (2.46)$$

which gives:

$$I(u) = \ln |\Omega| - \lim_{N \rightarrow \infty} \frac{1}{N} \ln V(\langle E_N \rangle \in du) . \quad (2.47)$$

As $\ln |\Omega|$ is just a number, the microcanonical entropy can be defined as the negative of the rate function:

$$\mathcal{S}(u) = \lim_{n \rightarrow \infty} \frac{1}{N} \ln P(\langle E_N \rangle \in du) \quad (2.48)$$

and

$$I(u) = -\mathcal{S}(u) . \quad (2.49)$$

From thermodynamics it is known that the entropy and free energy are related by a Legendre transformation. It is possible to write down an N particle partition function, very similar to the single particle partition function given by equation (2.14). Thus it is given by:

$$Z_N(\beta) = \int_{\Omega_N} \exp(-\langle E_N(\omega_i) \rangle) d\omega_i \quad (2.50)$$

$$= \int_{\Omega_N} \exp(-\beta \mathcal{H}_N(\omega_i)) d\omega_i . \quad (2.51)$$

With this it is straight forward to write down the free energy again:

$$\mathcal{F}(\beta) = \lim_{N \rightarrow \infty} \frac{1}{N} \ln Z_N(\beta) \quad (2.52)$$

This can be reformulated by changing measure from $d\omega_i$ to $P(d\omega_i)$, see [43] for more details:

$$\begin{aligned} f(\beta) &= \mathcal{F}(\beta)/\beta \\ &= \lim_{N \rightarrow \infty} -\frac{1}{N} \int_{\Omega_N} \exp(-\langle E_N(\omega_i) \rangle) P(d\omega_i) \\ &= \lim_{N \rightarrow \infty} -\frac{1}{N} \ln \langle \exp(-N\beta(-\langle E_N(\omega_i) \rangle)) \rangle . \end{aligned} \quad (2.53)$$

Furthermore, it can be shown that the cumulant generating function is the negative of the convex free energy $f(\beta)$, which will be omitted here. However, this fact implies that the convex free energy, and entropy density are related by the Legendre-Fenchel transform in the following way:

$$f(\beta) = \inf_{\mathbf{u}} \{\beta \mathbf{u} - \mathcal{S}(\mathbf{u})\} \quad (2.54)$$

and the Gärtner-Ellis Theorem results in:

$$\mathcal{S}(\mathbf{u}) = \inf_{\beta} \{\beta \mathbf{u} - f(\beta)\} , \quad (2.55)$$

as is shown on page 31 of the review by Touchette [43]. In the following these ideas will find application in the context of a dynamic ensemble.

2.4.2 The Dynamic Ensemble

In the following, we will introduce the s -ensemble. For an extensive and mathematically rigorous discussion, the reader is referred to [3]. In section 2.1 the canonical ensemble and the associated partition function as well as free energy were introduced (see equation (2.35)). The free energy is an intensive variable, whose singularities carry information about phase transitions in the system as outlined in section 2.3, provided the thermodynamic limit is taken. In order to investigate the dynamic properties we are interested in defining a dynamic free energy. Analogously to the canonical ensemble, the singularities of the dynamical free energy would then carry information about phase transitions in the dynamic ensemble. Let us first of all define what we mean by dynamic ensemble. Instead of the previously defined microstates and associated macrostate, i.e. time independent instantaneous observables of a system, we want to introduce time dependent observables. We will do this in the language of ensembles of trajectories, as the ideas presented here will find application in computational studies. For an equilibrated system at a given temperature, the path ensemble $P[\mathbf{X}_t]$ of trajectories \mathbf{X}_t can be sampled straightforwardly by generating dynamical trajectories starting from an equilibrated initial state.

Now, rather than looking at the typically observed trajectories of a system we are interested in the rare fluctuations within the trajectories. Therefore, it is useful to define a dynamic observable which describes a system of a given observation time $t = t_{\text{obs}}$. Such an observable can be denoted by the activity K , where each trajectory can be described by a single K , resulting in a distribution of activities $P_t(K)$ in an equilibrium setting. The choice of the dynamic activity is system dependent. The systems considered in this thesis all present obvious choices for the activity. When we say we look at the rare fluctuations of a system, we mean trajectories whose average activity is far away from the mean of the distribution $P_t(K)$, i.e. we are interested in a large deviation property of observing trajectories of a given length whose activity is unlikely. In order to make such observation more likely in simulations or more general to obtain a quantitative understanding of the dynamic ensemble we introduce a modified ensemble of trajectories which is biased according to the activity:

$$P_s[\mathbf{X}_t] \equiv \frac{P[\mathbf{X}_t]e^{-s K[\mathbf{X}_t]}}{Z_t(s)} . \quad (2.56)$$

The parameter s is a biasing “counting” field conjugate to the activity $K[\mathbf{X}_t]$ [2]. The exponential factor in equation (2.56) biases the probability of trajectories towards those which are less (more) active when $s > 0$ ($s < 0$) compared to the unbiased ensemble. The normalisation factor

$$Z_t(s) \equiv \sum_{\mathbf{X}_t} P[\mathbf{X}_t]e^{-s K[\mathbf{X}_t]} \quad (2.57)$$

is the moment generating function (MGF) for K , that is,

$$\langle K^n \rangle = (-)^n \partial_s^n Z_t(s)|_{s=0} \quad (2.58)$$

and can be thought of as a dynamical partition function associated to the ensemble of trajectories biased with s .

In analogy with an equilibrium statistical mechanics problem, the MGF $Z_t(s)$ is the object of interest. Now depending on some arbitrary observation time t . At long times ($t \rightarrow \infty$) the MGF acquires a large-deviation form [2, 43],

$$Z_t(s) \sim e^{t\psi(s)} . \quad (2.59)$$

The large deviation function $\psi(s)$ can be thought of as a dynamical free energy, which through a Legendre transform determines the probability $P_t(K)$ of observing an activity K over time t , at long times. Just like the free-energy in an equilibrium problem, the analytic structure of $\psi(s)$ as a function of s tells us about *dynamical* phases and possible phase transitions (or crossovers in the case of systems of finite extent) between them. The biasing of the ensemble of trajectories according to the counting field s is what we define as the s -ensemble.

Often we will consider scaled activities, allowing for different system sizes and different observation times to be compared to each other by defining the following intensive activity:

$$\langle \kappa \rangle_s \equiv (Nt)^{-1} \langle K \rangle_s = (Nt)^{-1} \sum_{X_t} P_s[X_t] K[X_t] . \quad (2.60)$$

The normalised activity presented here as κ will be associated to two different observables in chapter 3 and chapter 4. In chapter 3 it will represent the number of accepted MC moves for the Ising model and in chapter 4 it will represent the number of native contact being broken and formed in a lattice protein model. They will be discussed in more detail again in the appropriate chapters.

2.4.3 Applications

Investigating the large deviations of trajectories was originally done in the context of kinetically constrained models, in an attempt to get insight into the dynamical behaviour of glassy systems [1]. It was proposed that the increasing relaxation time in glass forming materials can be described using the s -ensemble formalism, by attributing it to a discontinuous dynamic phase transitions in trajectory space, i.e. a non-equilibrium dynamic phase transition. Here, the glassy features of the system get attributed to a coexistence of dynamically mobile and dynamically arrested phases. This type of analysis has been carried out for various systems ranging from simple binary spin models (e.g. Frederick Anderson models) [3, 2], spin glass models such as the random orthogonal model [44], or more complex models such as Lennard Jones mixtures [4]. All these models can be attributed with an interesting space time phase behaviour, resulting in coexistence regions of the different dynamic phases.

In the following two chapters this formalism is applied to the dynamics of the 1-d Ising model and those of a lattice protein model. These models are of interest, as for the first time a magnetic model (Ising model) is investigated. Furthermore the lattice protein is the model of highest complexity yet, investigated under this formalism.

This chapter aims to give a brief introduction to the Ising model and commonly studied equilibrium and non-equilibrium thermodynamic properties associated to it. It will then move towards a dynamical description of the system by employing the formalism of the s -ensemble as introduced in chapter 2. This means the trajectory space of the 1-d Ising model is studied, looking at the rare dynamic fluctuations of the trajectories. The investigation was carried out producing trajectories from two different types of MC movesets, which are Glauber and Kawasaki dynamics. The critical behaviour of the system as presented in [45] is studied computationally and predicted results were verified. In addition to the theoretical results available, the systems will also be considered in the presence of an external magnetic field. The work presented here was done in collaboration with Ernesto Loscar and Juan Garrahan. The work was published in the journal of statistical mechanics [22] and was adapted to this thesis where appropriate.

3.1 THE ISING MODEL

The Ising model is a well studied model of a simple magnet in theoretical/statistical physics [13, 46, 47, 48]. In this model spins are placed on a lattice of a given dimensionality, which point either up or down (± 1), representing magnetic dipoles. The alignment, or anti-alignment of the spins will give the entire lattice its magnetic property, as in an actual magnet. This leads to two extreme cases, one where all spins are either all up or down, representing the lowest energetic states possible – the ferromagnetic state of the system–, the other where spins are randomly pointing up or down and no bulk magnetisation is present – the paramagnetic state of the system. These extremes will occur at $T = 0$ and $T = \infty$. The interesting behaviour is how this change between order and disorder occurs, especially depending upon the dimensionality of the model. This is also known as its phase behaviour.

The model is named after Ernst Ising, who discussed the 1-d case in his PhD thesis in 1928 [49]. Most famous is probably the Ising model in two dimensions, as it displays a second-order phase transition, with the magnetisation as the order parameter. This will be discussed in greater detail later. In this case an exact analytical solution of the phase behaviour is available, originally proposed by Onsager [50]. Since the model was first introduced, it has seen many variations but the underlying idea is always the same, i.e. the arrangements of spins on a lattice of choice. Generally, only short range interaction energies be-

tween neighbouring spins are considered. The Hamiltonian is given by the following simple expression:

$$\mathcal{H} = -J \sum_{i,j} \sigma_i \sigma_j - h \sum \sigma_i , \quad (3.1)$$

where $J > 0$ is a coupling constant which is assumed to be 1 for the rest of this chapter. The sum runs over all neighbouring spins σ_i , taking a value of 1 or -1 respectively. For the 1-d case this is just the nearest neighbours along the chain, on a 2-d lattice this would be the neighbours in the two directions of the orthonormal basis, i.e \mathbf{i} and \mathbf{j} . Depending on the lattice type a set of nearest neighbours can always be defined. The second sum is depending on h which is an external magnetic field that couples to each spin on the lattice. An illustration of the 1-d lattice with up and down pointing spins can be found in figure 3.1.

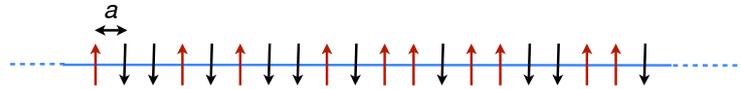


Figure 3.1: Spin chain with spins σ_i represented by arrows pointing up or down indicating positive or negative spin values. The spacing of the spins is equidistant. The lattice width a is indicated.

In the 1-d model no phase transition occurs, as demonstrated by Ising. The reason for this is that the free energy can be described by an analytic function for all temperatures other than zero. Another explanatory approach to this observation is the Mermin-Wagner theorem, which gives information about second-order transitions in magnetic systems between paramagnetic-ferromagnetic states above absolute zero [51, 52]. The theorem states that such a transition cannot be present in **isotropic** systems – i.e. structurally homogeneous systems – whose dimensionality is $d < 3$, where the 1-d Ising is an example of this.

3.2 EQUILIBRIUM AND NON-EQUILIBRIUM PROPERTIES OF THE 2-d ISING MODEL

The 2-d Ising model however, exhibits an interesting phase behaviour, as the Mermin-Wagner theorem does not apply in this case due to the special symmetry of the system. A second-order phase transition between an ordered ferromagnetic and a disordered paramagnetic phase can be observed. The mean magnetisation per spin sharply drops to zero, at a critical temperature T_c , as the specific heat per spin c diverges at that temperature. The specific heat and the phase behaviour will be discussed in more detail in section 3.2.3. From Onsager's an-

alytic solution [50] an expression for the critical temperature can be derived in units of the Boltzmann constant and the coupling constant.

$$T_c = \frac{2}{\ln(1 + \sqrt{2})} \approx 2.269 k_B J . \quad (3.2)$$

For any higher dimension, there are no analytical solutions available other than mean field approaches and one has to resort to computational methods in order to adequately describe the system. Various different MC algorithms can be used to characterise the system. The simplest approach is by means of the Metropolis algorithm.

3.2.1 Monte Carlo Simulations

Before looking at different sampling algorithms it is useful to define a couple of variables which will describe the model, dependent on the dimensionality of the system. The length L describes the length of a lattice side, thus for a 1-d chain $L = N_s$, where N_s is the number of spins in the system. In fact the number of spins scales with the dimension of the system such that: $N_s = L^d$, for any lattice type considered in this thesis. There are a variety of very fast algorithms sampling the behaviour of the Ising model most efficiently. Rather than giving a literature overview of the different simulation algorithms available, this section serves as an introduction to the two most straight forward single spin algorithms available: The Metropolis algorithm and the Kawasaki algorithm. The first one is used for modelling a ferromagnet and the latter for properties of lattice gases. In the model of the lattice gas the number of up and down pointing spins is fixed as only exchanges of spins are allowed as valid moves, and thus there is no change in overall magnetisation of the system during the period of observation. For the ferromagnetic model the magnetisation is not conserved, as spins are randomly attempted to be flipped. Each of the two algorithms sample equilibrium properties of the system and in this way approximate the partition function. In the subsequent section only the 2-d model on a square lattice will be considered, but this can in principle be extended to n -dimensions and different types of lattices.

3.2.1.1 Metropolis Dynamics and Non-Conserved Magnetisation

As for any statistical system, the knowledge of the partition function Z is vital to fully describe the equilibrium behaviour of the system, as previously discussed in chapter 2. In the case of the Ising model the state space is given by $\Omega = 2^{N_s}$, whose elements representing specific possible lattice configurations $I = \{\sigma_i\}_{i \in N_s}$. It is evident that even for modest system sizes the exact enumeration of the partition function soon becomes non-trivial. Thus a Monte Carlo (MC) algorithm can serve as an approximation to the partition function.

Listing 3.1: Pseudocode: Metropolis algorithm

```

Step1: Compute the instantaneous energy of the system.
Step2: Choose a lattice site at random, with a pseudo random
        number generator.
Step3: Flip the spin, if the proposed new configuration and
        its associated energy is accepted from the Metropolis
        criterion
Step 4: Repeat 1-3 until the maximum number of iterations is
        reached.

```

Originally, MC methods were developed in order to solve complex integrals of high dimensional functions and only in the 1940's started to be applied to statistical sampling problems [53]. The most famous algorithm was introduced by Metropolis in 1953 in order to simulate hard sphere gases [54]. This can easily be applied to the Ising model where a spin is chosen at random and is proposed to be flipped to its opposite orientation. The flip will occur, if the acceptance criterion is fulfilled. Whether this acceptance criterion is fulfilled depends on the change in total energy $\Delta\mathcal{H} = \mathcal{H}_{\text{new}} - \mathcal{H}_{\text{old}}$ of the system which is dependent upon the neighbouring spin configurations, known from the Hamiltonian.

$$P_{\text{accept}} = \begin{cases} \exp(-\beta\Delta\mathcal{H}) & \text{if } \Delta\mathcal{H} \geq 0 \\ 1 & \text{otherwise} \end{cases} \quad (3.3)$$

The Hamiltonian \mathcal{H} gives the instantaneous energy of the system. It is evaluated according to equation (3.1) and $\beta = \frac{1}{k_B T}$ representing the inverse temperature of the system, as previously defined in chapter 2. The algorithm obeys detailed balance, which means that the probability going from configuration I to J and J to I is equally likely when weighted according to the equilibrium probability of each state

$$p_{IJ}\pi_I = p_{JI}\pi_J \quad \forall I, J \in \mathcal{S} \quad (3.4)$$

(Here I, J indicated one certain configuration of the entire lattice.) In practice, the algorithm consists of the steps listed in listing 3.1.

In a unit of time, N_s spin flips will be attempted which is referred to as a MC lattice sweep. Initially some simulation data will be discarded as the system needs time to equilibrate to the simulation temperature. After the initial equilibration time, measurements of equilibrium fluctuations of the system at the temperature of interest can be conducted. From a large set of equilibrium simulations, observables of the system for the given temperature can be deduced without the knowledge of the partition function. The equilibrium expectation of any observable of the system, for example the magnetisation, can be accurately

predicted from long simulation runs. Any average of an observable is given by:

$$\langle O \rangle = \frac{\sum_{i=0}^n O \exp(-\beta \mathcal{H})}{\sum \exp(-\beta \mathcal{H})} = \frac{1}{t_{\text{obs}}} \sum_{i=0}^{t_{\text{obs}}} O_i , \quad (3.5)$$

where the sum runs over n MC lattice sweeps (the number of lattice sweeps n is equivalent to the observation time t_{obs}), also representing the unit of time. Equation (3.5) is valid, provided detailed balance is enforced and the system is ergodic. After each MC lattice sweep any observable such as instantaneous energy and magnetisation is calculated and can be recorded. The instantaneous magnetisation is given by:

$$\tilde{m} = \sum_i^{N_s} \sigma_i . \quad (3.6)$$

How does this information relate to the phase behaviour of the Ising model? Furthermore, if a phase transition only occurs in an infinite system, which can not be emulated on a computer, how can this be captured? The latter problem can be resolved to some extent. In order to remove spatial boundaries in a finite lattice of a given size, periodic boundary conditions can be introduced. The resulting boundary effects can be neglected far away from the critical temperature even at modest system sizes. For the Ising model this can be as small as a 32×32 system. However, when approaching T_c the spatial correlation will eventually be of the size of the system. Secondly, the choice of order parameter is important. An obvious choice for the Ising model is the magnetisation. The magnetisation can be used to define an auto-correlation function, which in turn allows to classify the correlation times of the system after a given time t :

$$R(t) = \int dt' [\tilde{m}(t')\tilde{m}(t'+t) - \langle \tilde{m} \rangle^2] , \quad (3.7)$$

with $R(t)$ being the auto correlation function at time t and $\langle \tilde{m} \rangle$ denotes the average magnetisation of the system and gives an insight into the timescales of the system. Another useful measure is the two point correlation function which gives insight into the spatial correlations at equilibrium:

$$G_{ij}(\mathbf{r}) = \langle \sigma_i \sigma_j \rangle - \langle \sigma_i \rangle \langle \sigma_j \rangle , \quad (3.8)$$

where \mathbf{r} is the distance between the two spins at different lattice sites. From $G_{ij}(\mathbf{r})$ one can extract the correlation length, which diverges in the limit of $L \rightarrow \infty$ at T_c . This nicely illustrates the spontaneous phase separation and hints at the phase transition, due to the divergence of the cluster size, as the critical temperature is approached.

3.2.1.2 Kawasaki Algorithm: Conserving the Magnetisation

As mentioned, a second common algorithm conserves the magnetisation of the system which also allows the measurement of equilibrium properties through equations (3.5) and (3.8). In order to conserve the magnetisation of the system, a different approach other than flipping spins needs to be taken, as with every spin flip the overall magnetisation of the system changes. The Kawasaki algorithm was introduced in the mid 1960's exactly to do that [55]. Kawasaki proposed to exchange neighbouring spins as a basic MC move, thus preserving the magnetisation. This results in a diffusive movement of spins around the lattice. The acceptance criterion for exchanging spins is still given by the Metropolis acceptance criterion from equation (3.3). This algorithm also fulfils detailed balance and the change in energy can easily be calculated locally, such that summing the overall energy of the lattice is not necessary. The Hamiltonian is also unchanged given by equation (3.1). When picking pairs of spins which should be exchanged, also aligned spin pairs need to be counted as a MC move, otherwise detailed balance would be violated.

Finding domain walls, i.e. anti-aligned neighbouring spins, can be a quite slow process. In particular, if the simulation temperature is below the critical temperature of the 2-d Ising model and thus large domains of aligned spins have formed. There are a variety much more efficient algorithms, circumventing the diffusive motion of the particles on the lattice. A common approach in order to improve this are continuous time algorithms, see [56].

3.2.2 Universality

Universality is the reason why second-order, or critical phase transitions are such an interesting area of study, but did not properly emerge until the 1960's with the development of the Renormalisation group [46]. Universality refers to the scaling behaviour of the system near the critical point of a phase transition. For the purpose of investigating universality a bit further, we can define a set of observables describing the system, which are the correlation length ξ , the specific heat c and the susceptibility χ . The susceptibility of the magnetisation gives information about the magnetic fluctuations of the system and can be used as a scaling variable. The susceptibility is given by:

$$\chi_m = \frac{\beta}{N_s} (\langle \tilde{m}^2 \rangle - \langle \tilde{m} \rangle^2) . \quad (3.9)$$

The specific heat per spin, gives information about the fluctuation of the energy of the system and can also be used as a scaling property of the system. The specific heat is given by:

$$c = \frac{1}{k_b T^2 N_s} (\langle \mathcal{H}^2 \rangle - \langle \mathcal{H} \rangle^2) . \quad (3.10)$$

The correlation length ξ simply holds the information over what distance \mathbf{r} the spins are correlated with each other. The correlation length diverging as the critical temperature of the phase transition is approached. As χ and c implicitly depend on the correlation of the spins, these too serve as scaling variables thus a set of different scaling functions for these variables can be deduced. The correlation length scales as:

$$\xi \sim |\hat{T}|^{-\nu} , \quad (3.11)$$

with $\hat{T} = \frac{T-T_c}{T_c}$ being a dimensionless quantity giving information on how close to the critical temperature the system is and ν is defined as a scaling exponent. This exponent reflects what universality means, as it is independent of the lattice type or coupling constant J and is valid for any 2-d Ising model. In the same way two more universal critical exponents can be defined:

$$c \sim |\hat{T}|^{-\alpha} \quad (3.12)$$

and

$$\chi \sim |\hat{T}|^{-\gamma} , \quad (3.13)$$

where χ is the susceptibility of the system, c the specific heat and α and γ in the literature commonly used symbols for the associated exponents. Essentially any phase transition will fall into a universality class and can be identified further according to its set of scaling exponents. A fourth exponent can also be defined which is known as the dynamic exponent and describes the scaling behaviour of the divergence of the correlation time τ_{corr} . This depends on the correlation length of the system such that:

$$\tau_{\text{corr}} \sim \xi^z , \quad (3.14)$$

with z being this dynamic exponent. It can be used as a measure for the efficiency of the simulation algorithm. For example, if z is large, then the simulation efficiency is low, as the correlation time becomes large around T_c and critical slowing down is making the simulation slow near the transition region. If z is small critical slowing is much more negligible and thus the simulation algorithm used is more suitable for measuring critical points.

In the following the universality class of the 2-d Ising model will be briefly discussed followed, by an overview of the commonly used finite size scaling methods to classify phase transitions.

3.2.3 The 2-d Ising Universality Class

The 2-d Ising universality class is presented here, as it will find application in section 3.3.2.1. In fact the *dynamic* phase behaviour of the

1-d Ising model, which will be explored in later sections falls into the 2-d Ising universality class. This is not so surprising as the time acts as a second dimension to the spatial dimension. However, this will be discussed in the appropriate section. With Onsager's solution to the 2-d Ising model, a set of scaling exponents for the phase transition can be worked out. His ansatz involved a complicated transfer operator approach. From this the critical temperature could be determined exactly, as given by equation (3.2). The heat capacity per spin can be evaluated which diverges as T approaches T_c , given by:

$$\frac{c}{k_B} = \frac{1}{\pi} \left(\frac{2J}{k_B T_c} \right)^2 \left[-\ln \left(1 - \frac{T}{T_c} \right) + \ln \left(\frac{k_B T_c}{2J} \right) - \left(+ \frac{\pi}{4} \right) \right] . \quad (3.15)$$

For the three thermodynamic scaling exponents as introduced in section 3.2.3 the following values can be deduced:

$$\alpha = 0 , \quad (3.16)$$

$$\gamma = \frac{7}{4} \quad (3.17)$$

$$\text{and } \nu = 1 . \quad (3.18)$$

These scaling exponents are used to classify the phase behaviour of the 2-d Ising model. There are a couple more scaling exponents which are all related to each other, but will not be discussed here. For more details on the scaling behaviour in the 2-d Ising model [46].

3.2.4 Finite Size Scaling – Measuring Exponents

As was discussed in section 3.2.3 the exact solution of the 2-d Ising model is known and thus its scaling exponents. It is still of interest to compare how well a simulation performs, which is achieved by estimating these. The exponents are a good indicator of the performance quality of an algorithm near the critical point.

3.2.4.1 Measuring Exponents

With the critical temperature for the 2-d Ising model known, it is relatively easy to estimate the dynamic exponent, given by equation (3.14). The simplest approach is conducted by repeating simulations at the known critical temperature, i.e. $T_c = 2.69$ for different lattice sizes. Due to the divergence of the correlation length, the correlation time will eventually reach the size of the lattice for temperature values close enough to the critical temperature, thus the scaling behaviour becomes:

$$\tau_{\text{corr}} \sim L^z . \quad (3.19)$$

In the book by Newman and Barkema [47], a simple MC simulation was run in order to show this scaling behaviour. The figure found in this textbook is shown in a slightly adapted version in figure 3.2. From

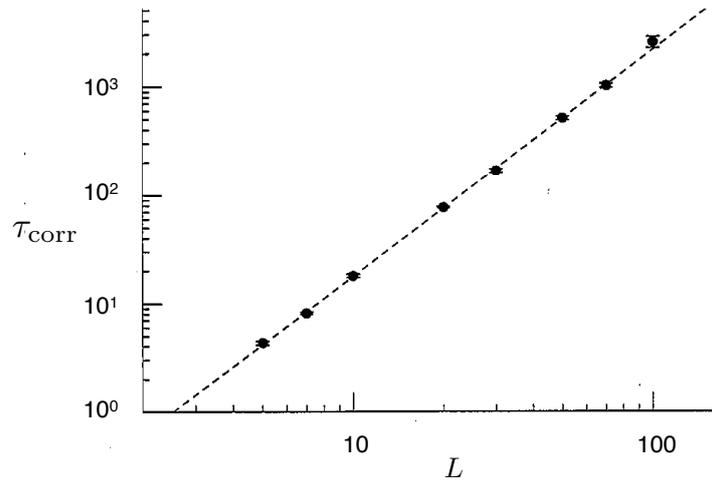


Figure 3.2: Dependence of the correlation time versus the lattice size on a log-log scale.

the fit in figure 3.2 z is estimated at $z = 2.09$. It is also mentioned that an even higher estimate can be found in [57], suggesting that the Metropolis algorithm does not scale well near the critical temperature of the system.

Of course, often the situation is very different and the critical behaviour of the system is not known. Therefore a different approach to finite size scaling will have to be used instead. For this purpose the magnetic susceptibility is a useful observable and its scaling behaviour is given by equation (3.13) which depends on the critical temperature. This can be rewritten such that it depends on the correlation length of the system

$$\chi \sim \xi^{\gamma/\nu} . \quad (3.20)$$

Evidently, ξ is also an unknown quantity but the diverging property of ξ near the critical point can be used to define a scaling function, which can be measured. The scaling function is given by:

$$\chi = \xi^{\gamma/\nu} \chi_0\left(\frac{L}{\xi}\right) . \quad (3.21)$$

The dimensionless function χ_0 has the following properties:

$$\chi_0(x) = \text{constant}, \quad \forall x \gg 1$$

and

$$\chi_0(x) \sim x^{\gamma/\nu}, \quad \text{as } x \rightarrow 0 .$$

Hence, the scaling behaviour for the magnetic susceptibility, making use of equation (3.11) is given by:

$$\chi_{\text{fin}} = L^{\gamma/\nu} \tilde{\chi}\left(L^{1/\nu} \hat{T}\right) . \quad (3.22)$$

Now the magnetic susceptibility is measured for each system size at a set of temperatures \hat{T} , which are close to the presumed critical temperature. This can be gauged from an initial temperature scan of the system. As the name suggests, the scaling function should look the same for each of the evaluated curves, if the exponents are correctly estimated. This can be done, by calculating $\tilde{\chi}(L^{1/\nu}\hat{T})$ for the different system sizes and then varying the exponents as well as the critical temperature until all curves collapse on top of each other. This kind of scaling function can be calculated for the specific heat and magnetisation and thus will yield good estimates for the other scaling exponents. For more details on this see for example [13]. Generally, a Renormalisation group approach allows for an accurate calculation of these scaling exponents and depending on the situation maybe a more desirable approach.

3.2.5 Non-Equilibrium Properties – Spinodal Decomposition

Everything discussed so far assumed the Ising system to be in equilibrium. Therefore in the following section non-equilibrium behaviour will be briefly discussed. We will follow closely ideas as presented in [47] and represent well established results. However, all simulations were setup and carried out by the author in order to illustrate the non-equilibrium properties which can be observed in spin systems, they also served as a point of validation for the implementations of the algorithms. A non-equilibrium simulation setup can be achieved in the following way: the system is initially at a high or infinite temperature, where all spins are randomly arranged and the system is in its paramagnetic state. Then the temperature is quenched to a temperature below the critical temperature. This quench will drive the system out of equilibrium. The relaxation of the system will result in the growth of up and down pointing spin domains. This process is also known as spinodal decomposition. It also serves as a simple model of de-mixing of colloids. The process of spinodal decomposition can also be classified through a universality class, where the universality class depends on the algorithm used for the relaxation process. Hence, this is often referred to as a dynamic universality class. Considering the Metropolis MC algorithm for the Ising model. The domain growth is well described by a power law behaviour, such that:

$$D \sim t^\nu , \quad (3.23)$$

where t is the time (in this particular case the number of MC lattice sweeps) and $D(t)$ is the size of a spin domain at a given time t . The exponent ν depends on the coarsening property of the simulation algorithm. Figure 3.3(a)-(d) shows the progressing spinodal decomposition of an Ising model at $0.5T_c$ of lattice size $L = 500$, after $t = 1$ (a), $t = 10$ (b), $t = 100$ (c), and $t = 1000$ (d). Each time step is a MC sweep of the entire lattice. The growth of the domains can be observed and is

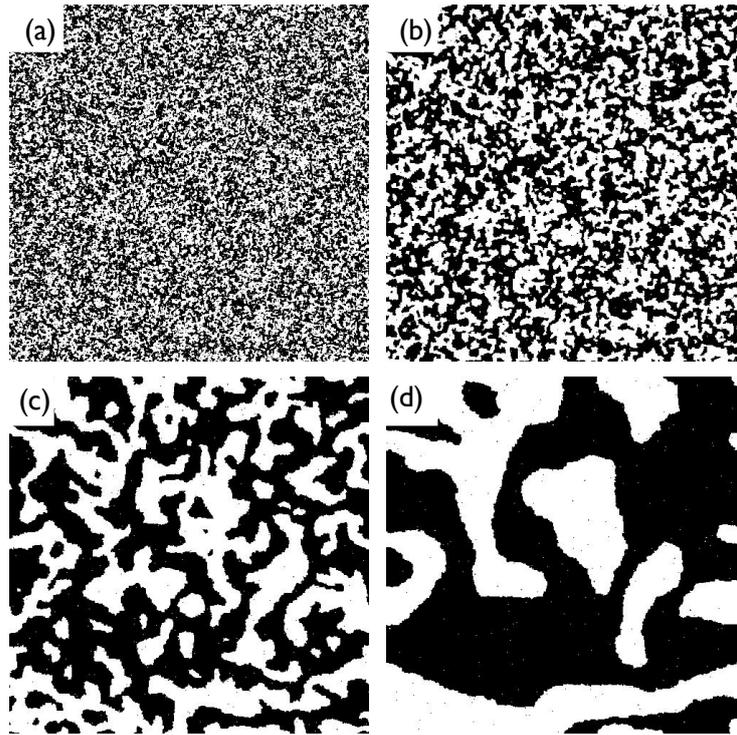


Figure 3.3: 500×500 Ising model after 1 (a), 10 (b), 100 (c), 1000 (d) MC lattice sweeps at $0.5T_c$, where white and black are domains of spins with values 1 and -1 respectively.

energetically favourable, as the total area of domain walls decreases reducing the overall energy of the system. This is due to the fact that the energetic cost of a domain wall depends on the surface tension. The surface tension is reduced by minimising the number of interfaces.

Quantitatively, the size of the domain $D(t)$ can be measured in two ways, through the two point correlation function, defined by equation (3.8) or the two dimensional structure factor, defined by equation (3.24). The distance $D_c(t)$ over which the correlation function falls to a value of $1/e = 0.368$ gives information about the domain size, in terms of the correlation function.

In crystallography the structure factor is used to estimate regularity in crystals, here we utilise these properties for the estimation of domain sizes. The structure factor will exhibit a peak corresponding to a change in spin orientation which can be used to define the domain size $D_s(t)$ with respect to the structure factor. The structure factor is given by:

$$S(\mathbf{k}, t) = \left\langle \left| \frac{1}{N} \sum_{\mathbf{r}_i} s(\mathbf{r}_i, t) \exp(i\mathbf{k} \cdot \mathbf{r}_i) \right|^2 \right\rangle. \quad (3.24)$$

Here, the angular brackets indicate an average over independent simulation runs. Furthermore, \mathbf{k} is defined as:

$$\mathbf{k} = \frac{2\pi}{L}(m\mathbf{i} + n\mathbf{j}) \quad m, n = 1, 2, \dots, L .$$

A circular average of the two dimensional structure factor will result in a smooth estimate, where the average is obtained by defining a spherical shell as a bin [58]. In this way the structure factor can be discretised and averaged. The shell is defined to be $n - \frac{1}{2} \leq (L/2\pi) |\mathbf{k}| < n + \frac{1}{2}$ for the circularly averaged structure factor of the form:

$$\hat{S}(k_n, t) = \sum_{\mathbf{k}} S(\mathbf{k}, t) / \sum_{\mathbf{k}} 1, \quad k_n = 2\pi n/L \quad (3.25)$$

and n runs over half the lattice, i.e. $n = 1, 2, \dots, L/2$ and k_n is the averaged absolute value of the wave vector falling into the given bin. In the end, the first moment of the circularly averaged structure factor, gives information about the location of the peak and thus an estimate of the domain size. The first moment μ_1 has the following relation to the domain size.

$$D_s(t) = \frac{2\pi}{\mu_1} \quad (3.26)$$

and is given by:

$$\mu_1(t) = \sum_{k_n=0}^{k_c} k_n \hat{S}(k_n, t) / \sum_{k_n=0}^{k_c} \hat{S}(k_n, t) . \quad (3.27)$$

The cut-off value is given by $k_c = \frac{L}{2}$. In figure 3.4 the scaling behaviour of the domain size growth over time from the two point correlation function and the circularly averaged structure factor are shown. Both methods result in a straight line on a log-log plot with an exponent of $\nu \approx 0.5$. The multiplicative factor can be neglected in the scaling behaviour and just indicates that both methods for the domain size estimation are not equivalent but exhibit identical scaling behaviour.

This establishes a scaling behaviour for the domain size growth from non-equilibrium simulations for the example of the Metropolis MC algorithm. The existence of such magnetic domains is a well known and well studied property of real magnets, where magnetic domains are formed if a ferromagnet is cooled below the Curie temperature. For more information and a detailed introduction to more realistic magnets than the simplistic Ising model, the reader is referred to a textbook by Cullity [59].

With this, the introduction to the commonly used tools in computational studies of the critical behaviour of phase transitions in the Ising model is concluded. In the next section the previously established s-ensemble methods will be applied to the 1-d Ising chain. With the aid of the introduced finite size scaling method within this section 3.2.4, the dynamic phase space behaviour will be investigated.

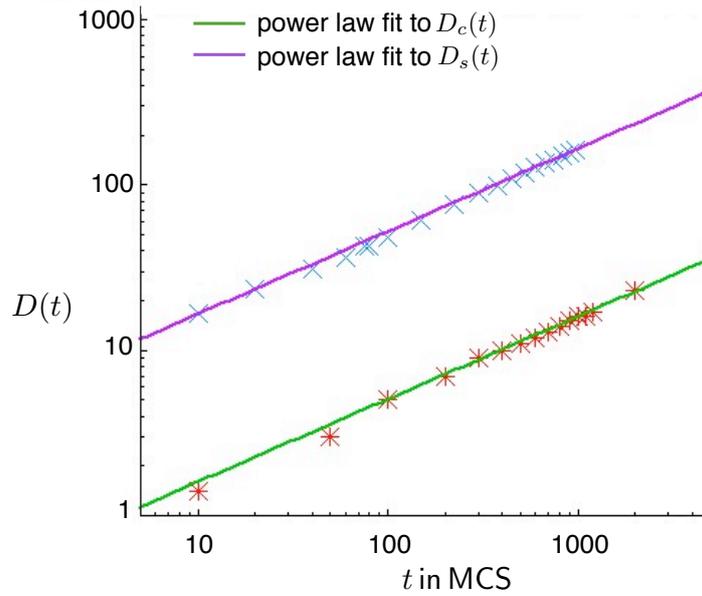


Figure 3.4: Log-log plot of $D_s(t)$ (crosses) and $D_c(t)$ (stars) with respect to t in units of MC lattice sweeps. Averages were taken over 5 independent runs and errors are within the size of the symbol. The power law fit results in a scaling exponent of $\nu \approx 0.5$.

3.3 DYNAMIC PROPERTIES OF THE 1-d ISING MODEL IN THE S-ENSEMBLE

In various model systems of glasses, such as kinetically constrained systems, it has been shown that looking at the ensemble of trajectories, interesting dynamic phase behaviour can be uncovered [5, 2, 4]. Therefore, it is of interest to apply these ideas to magnetic models as well, where the 1-d Ising model is a very simplistic example. For this purpose we are interested in looking at ensembles of trajectories generated from MC simulations and in particular looking at atypical trajectories, i.e. rare fluctuations in the distribution of typical trajectories. This is achieved by biasing equilibrium trajectories according to the previously introduced formalism of the s-ensemble. Furthermore, this computational investigation is supported by a set of theoretical results giving good predictions of the expected dynamic phase behaviour [45].

All results presented in the following sections are novel and the author contributed in the development of the simulation algorithm and analysis of the results. The exhaustive and novel computational study is published in the journal of statistical mechanics [22].

3.3.1 Theoretical Predictions

The computational study was motivated and facilitated through a recent publication by Jack et al [45] which looks into the dynamic properties of the 1-d Ising model using an analytic toolset. The authors predict a dynamic phase transition in an energy biased ensemble and reveal a set of critical values of the biasing parameter at which point the dynamic phase transition occurs. In [45] this is achieved by the introduction of a modified path ensemble with a biasing parameter g – known as the g -ensemble. This path ensemble makes use of an energy constraint, as opposed to an activity constraint found in the previously introduced s -ensemble. The g -ensemble introduces a dynamic free energy $\phi(g)$ in a similar fashion as discussed in chapter 2. This free energy can also be obtained through the large deviation formalism for the dynamic free energy with no external magnetic field. The second derivative of this dynamical free energy diverges, resulting in a line of critical points of the biasing parameter g . Furthermore, it is shown that the scaling behaviour of the observed transition falls into the 2-d Ising universality class. The analytical solutions, as presented in [45], were derived by mapping the 1-d spin chain onto a corresponding quantum problem [60]. An equivalence between the activity constrained ensemble (with biased field s) and the energy constrained ensemble (with biased field g) can be established, since in the Ising problem of interest, the biasing of the energy is akin to biasing in terms of the escape rate of the corresponding continuous-time Markov chain, see e.g. [3]. With this equivalence known, it is straight forward to probe the dynamic phase behaviour by means of MC simulations. The known critical behaviour from the g -ensemble can be directly applied to the s -ensemble, which in turn is readily accessible through MC simulations.

3.3.1.1 The Phase Diagram

The purpose of this section is to give an overview of the theoretically predicted dynamic phase diagram. Let us consider three variables $\{s, T, h\}$ spanning a parameter space, which contain all relevant information of the phase space of the dynamic system. Looking at the evolution of the system in time, a dynamic phase diagram can be constructed on this parameter space, where s is the biasing field, T the temperature of the system and h an external magnetic field. From the equivalence of the g -ensemble and s -ensemble an exact solution for the dependence of s and T can be deduced. This is:

$$s(T) = -\ln[\tanh(2J/k_B T)] \quad , \quad (3.28)$$

as derived in: [45]. This gives the line of critical points as seen in figure 3.5(a) which represents a schematic of the phase diagram of the dynamic 1-d Ising model.

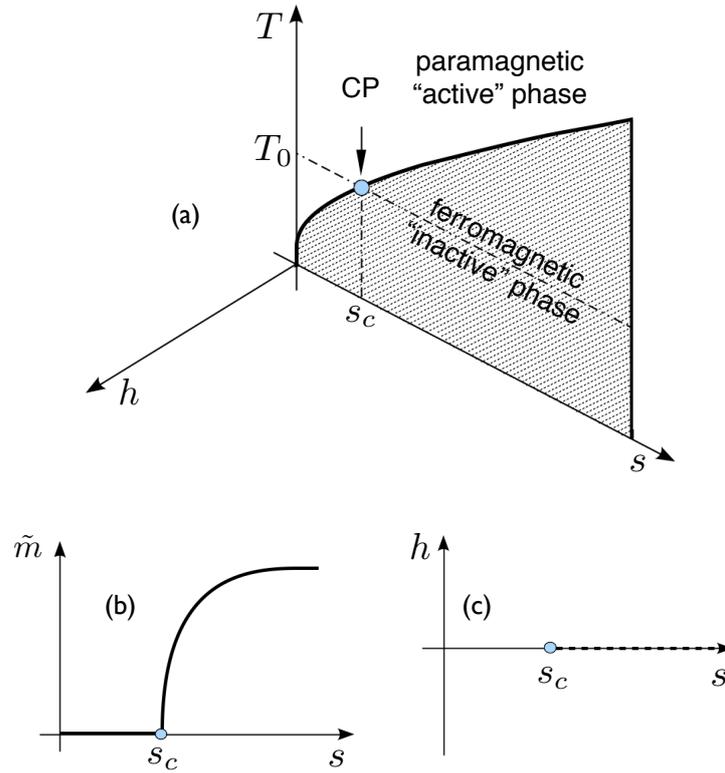


Figure 3.5: Schematic representation of the dynamic phase diagram of the 1-d Ising model in the s -ensemble with magnetic field h and the application of Glauber dynamics. (a) The phase space is defined as the parameter space $\{T, s, h\}$. The continuous line in the plane $h = 0$ represents the critical points given by the theoretical solution of equation (3.28) separating paramagnetic (and highly active) and ferromagnetic (and low-activity) dynamical phases. The dashed line corresponds to the state where $T = T_0 = \text{const}$ and $h = 0$ and contains a critical point (CP) marked by a circle. The dotted area below the critical line represents a surface of first-order transitions, when varying the external magnetic field h . (b) $\tilde{m}(s|T = T_0)$ gives the spontaneous magnetisation, where the symmetry is broken at $s = s_c$. (c) The plane (h, s) for $T = T_0$. The critical point $s = s_c$ divides the line in two phases, a paramagnetic phase for $s < s_c$ and a ferromagnetic phase $s > s_c$. For $s > s_c$, the dashed line indicates points of first-order transitions.

The parameter space was chosen in such a way that it allows a clear distinction of different dynamic phases of this system. In figure 3.5(a), the line of critical points that lie in the plane of $h = 0$, as predicted by equation (3.28), divides this plane into two different phases of activity. As a separator of these phases acts a critical point at $s = s_c$. For values of $s < s_c$ the system is in a paramagnetic and active phase, whereas for $s > s_c$ the system is in a ferromagnetic and less active phase.

The dashed line in figure 3.5(a) is an isotherm at constant temperature $T = T_0$, including the critical point (CP) for a critical value of

$s = s_c$, only considering the plane where $h = 0$. The temperature T_0 in this picture is similar to the critical temperature T_c for the 2-d model previously discussed. As indicated, this line contains a critical point corresponding to $s_c = s(T_0)$ given by equation (3.28).

Figure 3.5(b) shows the magnetisation of the Ising chain as a function of the biasing field s . Here the critical point $s = s_c$ separates the ferromagnetic phase for $s > s_c$ with non-zero spontaneous magnetisation $\bar{m} \neq 0$ and, the paramagnetic phase for $s < s_c$ with $\bar{m} = 0$. Tuning the biasing field s , once the critical point is reached, the symmetry of the system is spontaneously broken. The inclusion of the magnetic field in the isothermal behaviour gives rise to a third dimension and therefore the plane $T = T_0$ can be considered as shown in the 3.5(c). Here for $s < s_c$ the magnetic field acts on a paramagnet, while for $s > s_c$ the magnetic field acts on a ferromagnet 3.5(c).

Figure 3.5(c) is analogous to that of a 2-d Ising model assuming that $s \rightarrow \beta$. In fact, by taking $s = s_c$ and using the temperature T (instead of s) as the control parameter, the critical behaviour is in complete analogy with the 2-d Ising model, and thus it is not surprising that the model belongs to the same universality class as the 2-d Ising model. This will be verified subsequently. In the presence of an external magnetic field h the system is expected to exhibit a first-order phase transition along the dashed line, as shown in figure 3.5(c), where $s > s_c$. In the three dimensional parameter space this results in a surface of first-order transitions, which is indicated in figure 3.5(a), by the dotted area below the critical line. The goal is to computationally verify the phase space behaviour as laid out in this section, based on the theoretical work by [45]. For this purpose a series of MC simulations are conducted in order to study the second-order critical behaviour as well as the first-order behaviour in the presence of an external magnetic field h .

3.3.2 Simulation Setup

In this section, we will present the simulation setup used for the numerical study. Motivations for certain choices of system size will be given, where appropriate.

3.3.2.1 Glauber Dynamics

As the theoretical work was based on Glauber dynamics, the same dynamical approach is used in the computational study. Until now only the Metropolis acceptance criterion given by equation (3.3) was discussed. The Glauber acceptance criterion is not much different from

the Metropolis criterion and can be employed instead of it [61]. It is given by:

$$P_{\text{accept}} = \frac{1}{1 + e^{\beta\Delta\mathcal{H}}} , \quad (3.29)$$

where β is the inverse temperature $\beta = 1/k_B T$ and the Boltzmann constant k_B is assumed to be 1. This acceptance function results in a smoother crossover as opposed to the Metropolis criterion, as energetically favourable moves can be rejected. The MC simulation is carried out in the same way as described in algorithm 3.1 substituting the Glauber acceptance criterion.

3.3.2.2 Sampling the s -ensemble

Trajectories, or paths of a certain 'temporal' length t_{obs} can be generated by employing the single spin flip dynamics of algorithm 3.1 using the Glauber acceptance criterion. A single timestep Δt corresponds to a lattice sweep, i.e. N_s attempted spin flips. From these, path average properties can be deduced, such as the average magnetisation, as mentioned previously. In order to bias the trajectories according to the s -ensemble, a dynamic measure of the system needs to be defined. To this end, we define a dynamic activity K as the incremental count of accepted spin flips. This is then normalised to $\kappa = \frac{K}{N_s t_{\text{obs}}}$ in terms of the observation time and lattice size. When generating many independent realisations of trajectories of a given observation time t_{obs} , each trajectory is associated with a value of an activity κ . Typical trajectories will have a certain distribution of activities centred around a mean. However, here we are not interested in the typical trajectories, but rare trajectories which have a larger than average, or lower than average activity. In order to bias the path averages evaluated according to their average activity towards low or high activity, the activity of an old and newly generated trajectory are compared and the new trajectory is accepted according to the Metropolis criterion in terms of the s -ensemble:

$$P_{\text{accept}} = e^{-s\Delta\kappa} . \quad (3.30)$$

Unlike in the case of the inverse temperature, s is unitless as κ is just a normalised incremental count. The old trajectory stays the same until a newly proposed trajectory is accepted, much like the spin configuration stays the same if a spin flip is rejected. As this is driving the system towards the tails of the activity distribution, generating trajectories which will be accepted is not easy. In order to facilitate this, a modified transition path sampling (TPS) method is proposed, as used in [4]. In the introduction 1 this method was briefly outlined. The basic idea, as illustrated by figure 3.6, is to choose uniformly a point along a given trajectory – the shooting point – and choose with a 0.5 probability a shooting direction – forward or backward – from

which the dynamics is run until a new trajectory reaches the desired length of t_{obs} .

Therefore, if we already have a trajectory which exhibits some inactivity (which we are trying to bias towards) a new trajectory will also have low activity, as only part of it was regenerated. This makes it more likely for accepting the new trajectory in the s -ensemble according to equation (3.30). This allows a more effective way of generating

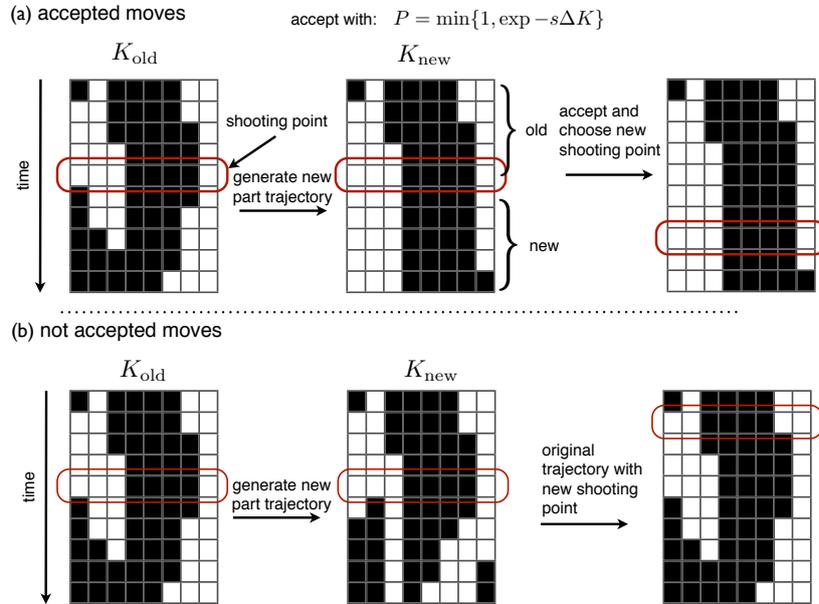


Figure 3.6: The schematic illustrates the TPS schedule employed for generating new trajectories, which are similar to the old trajectory. A shooting point is randomly chosen along the original trajectory (red frame). Up to that point all of the original trajectory is kept and dynamics are restarted at the shooting point until the desired trajectory length is obtained. The direction of the newly generated trajectory is arbitrarily chosen. (a) Outlines the schedule for an accepted trajectory and (b) for a rejected trajectory

new trajectories for the modified path ensemble, than regenerating completely new trajectories each time. Furthermore, a description of the system in this path space now extends to the time dimension. This sets the ground for a computational exploration and quantitative study of the dynamic phase space of the 1-d Ising model as proposed in section 3.3.1.1.

3.3.2.3 Choice of Parameters and Approach

In order to study how the system scales with the observation time t_{obs} and the system size $L = N_s$ some initial parameter choices have to be made. For this purpose, before the simulations for the ensemble of trajectories study was setup, some initial trial simulations using Glauber dynamics were carried out in order to investigate the effect of the finite system size. In an ideal world one can use an infinitely large

system ($L \rightarrow \infty$, i.e. the thermodynamic limit), but as this is not possible on a computer, instead a 'sufficiently' long spin chain with periodic boundary conditions is considered. To this end, the 1-d Ising chain was simulated for various different chain lengths, from 30 to 500 spin sites. From these short trial runs a system size of $N_s = 64$ was deemed sufficiently large enough to mimic a thermodynamic limit. Therefore the system size was fixed, allowing only time as a scaling variable.

In order to justify the observation time t_{obs} as a scaling variable similar to that of the system size, an associated quantum problem needs to be considered. The dynamic evolution of the system can be described by a time evolution operator, which is the equivalent to a quantum Hamiltonian. From this point of view the length (t_{obs}) of the trajectories is equivalent to the inverse of the temperature $T^{(q)}$ of the quantum problem presented in [45]. It has previously been shown that there is a mapping between the classical 2-d Ising model and a quantum chain [60]. The mapping demonstrates that the inverse of $T^{(q)}$ is equivalent to the size in one spatial direction in the 2-d Ising model (that is $1/T^{(q)} \leftrightarrow L$). Therefore it can be correctly assumed that t_{obs} is the equivalent of the linear size L in the classical 2-d Ising model (i.e. $t_{\text{obs}} \leftrightarrow L$). The scaling relations for the lattice size as discussed in section 3.2.4 can simply be replaced by the observational time t_{obs} .

In the light of simulating a stochastic system, generally a very large number of realisations is used in order to estimate averages. Most of the time, unless otherwise stated, data points are generated from $N_{\text{realisation}} = 5 \times 10^6$. As the standard error of the mean is given by:

$$\epsilon_{\text{std}} = \frac{\sigma}{\sqrt{N_{\text{realisation}}}} \quad , \quad (3.31)$$

it is often the case that the ϵ_{std} will fall into the size of the data points shown, if errors are not explicitly stated.

3.3.3 Simulation Results of the Dynamic Properties

In the following section the results of the simulations with respect to the phase diagram shown in figure 3.5 will be discussed.

3.3.3.1 Dynamic Continuous Phase Transition

As was established in the discussion of section 3.3.1.1 describing figure 3.5, there are two ways in which the dynamical critical point can be reached: by setting either s or T to a constant value. Trajectories are generated as described in section 3.3.2.1 using a fixed temperature T and zero external magnetic field. From each accepted trajectory of a given value of the biasing parameter s a set of time extensive variables is obtained, which classify the dynamic properties of the system according to a given observational time t_{obs} . These are the

activity κ representing the total number of spin flips observed in a given trajectory (i.e. the number of configurational changes), the integrated energy $E = \int_0^{t_{\text{obs}}} dt' \mathcal{H}(t')$, and the integrated magnetisation $M = \int_0^{t_{\text{obs}}} dt' \tilde{m}(t')$, where $\tilde{m}(t) = \sum_{i=1}^{N_s} \sigma_i(t)$.

Drawing from the experience of the study of the thermodynamic behaviour of the 2-d Ising model, finite size scaling is a valid approach in order to classify the universality class of the dynamic phase transition. As the scaling behaviour of dynamic properties is the observable of interest, the temporal size is the scaling variable of choice, as was motivated in section 3.3.2.3.

In order to determine the critical temperature of a continuous phase transition the peaks of the susceptibilities of an observable – here integrated energy or activity – can be used. According to the previous discussion χ is expected to scale according to [13]:

$$\chi(t_{\text{obs}}) = t_{\text{obs}}^{\alpha/\nu} f\left(\frac{(T - T_c)}{T_c} t_{\text{obs}}^{1/\nu}\right), \quad (3.32)$$

where α and ν are the previously defined exponents from the scaling relations defined in equations (3.13). Furthermore, the location of the peaks define an effective size-dependent transition temperature which is expected to vary as:

$$T_c(t_{\text{obs}}) = T_c(\infty) + A t_{\text{obs}}^{-1/\nu}. \quad (3.33)$$

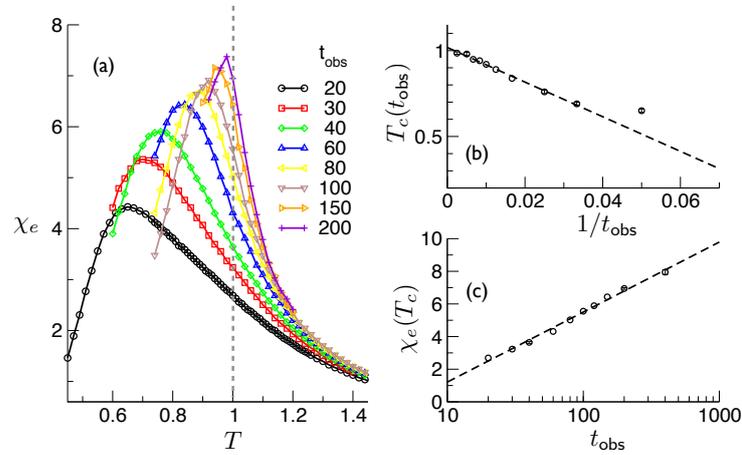


Figure 3.7: (a) Susceptibility of the energy as a function of T at constant $s = 0.036\ 635\ 37$. Vertical dashed line corresponds to $T_c = 1.00$ given by equation (3.28). The standard error falls into each data point. (b) Effective critical T , taken as the value of T at the maximum of each of the curves of (a), with respect to $1/t_{\text{obs}}$. The extrapolated value, for $t_{\text{obs}} \rightarrow \infty$, is $T_c = 1.02 \pm 0.02$, from a linear fit. (c) Susceptibility χ_e at $T_c = 1.00$ versus t_{obs} . The divergence is logarithmic, as in the case of the specific heat in the 2-d Ising model.

The scaling given by the equation (3.33) is expected in terms of T or in terms of s , respectively. Thus, both approaches were used for probing the scaling behaviour.

Figure 3.7 shows the simulation results for a fixed value of s and varying the temperature T . The critical value for $s = s_c$ was chosen in such a way that it would correspond to an equivalent temperature of $T \equiv 1$. From equation (3.28) the relationship between s and T is given such that $s_c = 0.03663537$. By defining a generalised susceptibility of the energy χ_e in such a way that

$$\chi_e = \frac{\langle E^2 \rangle_s - \langle E \rangle_s^2}{T^2 N t_{\text{obs}}}, \quad (3.34)$$

the scaling behaviour can be observed.

All data is shown from averages over 5×10^6 independent trajectories and $N_S = 64$. In figure 3.7(a) the susceptibility of the energy is plotted with respect to temperature, for increasing observational times t_{obs} . Observed peaks sharpen as t_{obs} increases. This leads to an extrapolation of an 'effective critical temperature' $T_c(t_{\text{obs}})$ according to t_{obs} . The peak of χ_e was determined with a quadratic least square fit from data points around the peak. From the functional value obtained from the fit, the maximum was extracted. The linear regression fit and the maximum were calculated with the plotting program Grace, using the standard error of the fit as the error for the extrapolated value. The obtained effective critical temperatures, from the peak values were plotted in figure 3.7(b) with respect to $1/t_{\text{obs}}$. The dotted line is a linear fit to peaks, ignoring the data point from the shortest observation time and allows an extrapolation to the critical temperature T_c for infinite observation times. The expected critical temperature is $T_c = 1$. The linear interpolation gives a value of $T_c = 1.02 \pm 0.02$, therefore the MC simulation agrees with the predicted value within errors of the fit. One should bear in mind that we have neglected the error resulting from the least square fit to the peaks. This also supports the fact that the scaling behaviour belongs to the 2-d Ising universality class with a critical exponent $\nu = 1$. Figure 3.7(c) shows the scaling behaviour of the susceptibility versus t_{obs} on a log-linear scale. As the dynamical free energy $\phi(g)$ must have the same critical properties as the thermodynamic free energy $\mathcal{F}(T)$ of the 2-d Ising model, its second derivative must be analogous to the behaviour of the specific heat. It was shown in [45] that the second derivative of $\phi(g)$ has a logarithmic divergence when the critical point is approximated [45]. Figure 3.7(c) shows that a logarithmic behaviour is observed, which means that equation (3.32) holds with an exponent of $\alpha = 0$.

The next step is to consider a case where an isothermal of $T = 1$ is used and instead the biasing field s is varied. This will also approach the critical behaviour. Figure 3.8(a) shows the behaviour of the susceptibility of the activity χ_k . This is in principle analogous to the be-

haviour of the energy fluctuations and thus can be defined in the following way:

$$\chi_{\kappa} = \frac{\langle K^2 \rangle_s - \langle K \rangle_s^2}{N t_{\text{obs}}} . \quad (3.35)$$

Figure 3.8(a) shows the increase of the peaks of the susceptibility as the observation time of the system increases. The critical biasing field s_c is indicated by the grey dashed line. Using the same approach as before, an effective critical s can be deduced from the peaks, using again a quadratic approximation of the peak. The position of the peaks and therefore the effective critical value of s is plotted with respect to $1/t_{\text{obs}}$ in 3.8(b). The extrapolation results in a critical value for $s_c = 0.036 \pm 0.01$, again using a linear regression fit with the Grace plotting tool, ignoring the data point coming from the shortest observation time. The scaling exponent is as before consistent with the 2-d Ising universality class $\nu = 1$. Also figure 3.8(c) shows the susceptibility χ_{κ} value for $s = 0.03663537$ in a log-linear scale. This logarithmic divergence of the susceptibility of the activity means that the scaling behaviour for exponent $\alpha = 0$. This indicates a complete equivalence between the inverse temperature β and the biasing field s . The finite time effects of the fluctuations of the activity or energy respectively mean that both E and κ can capture the dynamic phase behaviour.

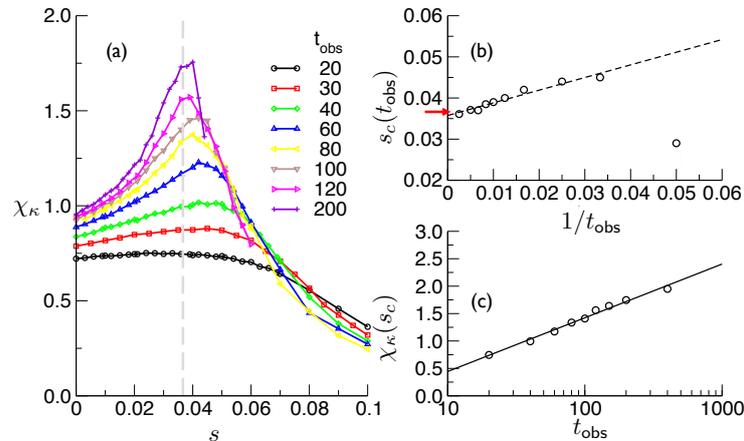


Figure 3.8: (a) Susceptibility of the activity as a function of s at constant temperature $T = 1$. (b) Convergence of the effective critical s . Values for s are taken from the peaks of (a). The extrapolated value is $s_c = 0.036 \pm 0.001$. The arrow indicates the theoretical value from equation (3.28) (c) Susceptibility at $s = s_c = 0.0366$ constant versus t_{obs} . The critical divergence of the χ_{κ} is logarithmic.

3.3.3.2 Magnetic Properties

So far the Ising chain was only considered in a 'neutral' environment. The obvious next step is to apply an external magnetic field h to the

system. This magnetic field couples to the system, as given by equation (3.1). Therefore the magnetic field is applied to the 1-d chain, when the system is already close to the critical point. The magnetic susceptibility of the 2-d Ising system can be expressed by the two point correlation function:

$$\begin{aligned}\chi_{\tilde{m}} &= \frac{\langle \tilde{m}^2 \rangle - \langle \tilde{m} \rangle^2}{kTN_s} \\ &= \frac{1}{kTN_s} \sum_{i,j=1}^{N_s} G_{i,j} ,\end{aligned}\quad (3.36)$$

where $G_{i,j}$ is given by equation (3.8). Analogously, a correlation function in space and time representing the space-time magnetic susceptibility, can be defined as:

$$\begin{aligned}\chi_M &= \frac{\langle M^2 \rangle_s - \langle M \rangle_s^2}{k_B TN_s t_{\text{obs}}} \\ &= \frac{1}{k_B TN_s t_{\text{obs}}} \sum_{i,j=1}^{N_s} \int_0^{t_{\text{obs}}} \int_0^{t_{\text{obs}}} dt'' dt' G_{i,j}(t', t'') .\end{aligned}\quad (3.37)$$

This leads to the definition of the instantaneous susceptibility in the s-ensemble:

$$\chi_{\tilde{m}}^{(s)} = \frac{\langle \tilde{m}^2 \rangle_s - \langle \tilde{m} \rangle_s^2}{k_B TN_s} .\quad (3.38)$$

Incorporating the magnetic information in the MC simulation is trivial and readily implemented. Again averages were taken over 5×10^6 realisations and a lattice size of $N = 64$ and this time a fixed $t_{\text{obs}} = 200$. The results are depicted in figure 3.9.

For external magnetic fields of $0.004 \lesssim h$, the behaviour is very clear, meaning that the thermodynamic behaviour of the system is observed for different lattice sizes and observational times with no influence from finite size effects. However, if the external field is chosen to be small, finite size effects arising from finite size N_s and finite observational time t_{obs} are non-negligible. As a result the system is no longer well behaved enough to be studied with the amount of effort possible within the scope of this work. Figure 3.9(a) shows the static susceptibility $\chi_{\tilde{m}}^{(s)}$ with a power law fit representing the data well in the region of $0.004 \lesssim h \lesssim 0.2$ for the external magnetic field, with a fitted exponent of 0.90 ± 0.05 . For values of $h \geq 0.2$ a deviation from the power law is expected as the system moves away from the critical behaviour. The observed exponent is consistent with what is expected from the 2-d Ising model as:

$$\chi_{\tilde{m}}^{(s)} \propto h^{-(1-1/\delta)}\quad (3.39)$$

and $\delta = 15$ therefore the exponent is close to 0.933.

In contrast figure 3.9(b) looks at the susceptibility depending on the temporal aspect of the system. Qualitatively the behaviour is similar

to that of the static susceptibility, but a power law fit gives a very different exponent of $\Delta = 1.50 \pm 0.05$. A dynamical critical behaviour can be defined as:

$$\chi_M \propto h^{-\Delta} . \quad (3.40)$$

The exponent Δ can be written as: $\Delta = (\gamma + \nu z)(\beta\delta)^{-1}$. This allows an estimate of the dynamical exponent $z \simeq 1.06 \pm 0.09$, which is within the error of the expected value of $z = 1$. It should also be noted that this numerical experiment was repeated for the following values of observation time and system size: $N_s = 64, 100$ and $t_{\text{obs}} = 160, 200, 300$, obtaining the same results as presented in figure 3.9 within error bars. It is interesting to note that the static criticality of the 2-d Ising model reappears in the dynamic 1-d case, as was shown by means of finite time scaling of the energy and activity.

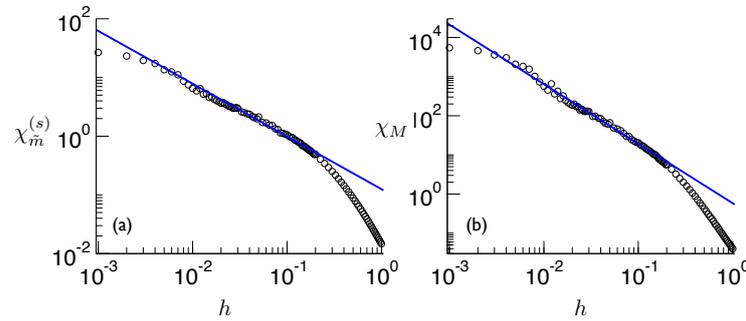


Figure 3.9: (a) Susceptibility of the instantaneous magnetisation corresponding to spatial fluctuations. The continuous line is a fit of a power law that gives an exponent 0.90 ± 0.05 , in good agreement with the exponent of the magnetic susceptibility of the 2-d Ising model $1 - 1/\delta \approx 0.933$. (b) Susceptibility of the integrated magnetisation corresponding to spatio-temporal fluctuations. The continuous line is a fit of a power law that gives an exponent of 1.50 ± 0.05

3.3.4 Simulation Results of the Dynamic Properties with External Magnetic Field

In the phase diagram it was shown that a first-order phase transition is expected under isothermal conditions, if s is greater than the critical value $s_c(T_0)$ (see figure 3.5(c)). The application of an external magnetic field will induce a spontaneously broken symmetry. Unfortunately the mapping of the master equation onto the quantum problem in the presence of an external magnetic field is non-trivial. However, a computational study of the first-order phase transition is carried out instead. Let $\{T, s, h\}$ be the parameter space in which the isothermal plane of $T = 1$ is considered. In the presence of no external magnetic field, i.e the plane $h = 0$ two magnetic phases coexist: a paramagnetic disordered phase for $s < s_c$ and an ordered, ferromagnetic phase for $s > s_c$. This is illustrated in the schematic of the phase diagram 3.5(c).

The MC setup is as before, but a lattice with $N_s = 200$ spins was used and averages were taken over 1×10^6 realisations. It was found that a larger system was necessary to decrease finite size effects. Due to the increased system size the simulation became more costly, therefore fewer number of realisations were considered. Results are shown, considering three different cases of the biasing field s , in figure 3.10. The external field h was varied for setups of $s < s_c$, $s = s_c$ and $s > s_c$, shown in figure 3.10(a), (b) and (c) respectively.

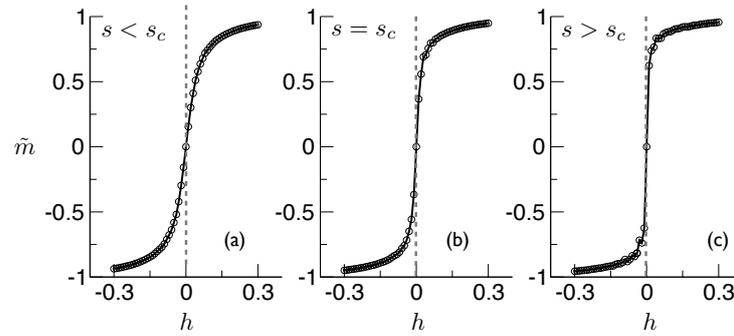


Figure 3.10: Stationary results for the magnetisation with an external magnetic field. The magnetisation with respect to external magnetic field is plotted, varied over a range of $h = -0.3$ to $h = 0.3$ with a constant temperature ensemble at $T = 1$, with a critical value for $s_c = 0.03663537$. Depicted are three different cases of s . In (a) $s = 0.02$ is thus subcritical, in (b) s corresponds the critical value and in (c) s is supercritical as $s = 0.05$. Errors are within the size of the symbols.

In figure 3.10(a) a smooth curve $\tilde{m}(h)$, in the vicinity of $h = 0$, for a low biasing field $s = 0.02$ is observed, representing the paramagnetic phase. In contrast figure 3.10(c) depicts a sharp transition at $h = 0$ well into the ferromagnetic phase, as the choice of biasing parameter s is supercritical with a value of $s = 0.05$. In figure 3.10(b) the exact critical value is used for the biasing field, showing critical behaviour of the system. The critical value was taken to be: $s_c = 0.03663537$. The field driven first order transition ends in the critical point which is precisely the critical value of the biasing parameter s . This point is reached at $s = s_c$ ($T = 1, h = 0$) and thus gives the limits for the ferromagnetic phase ($s > s_c$).

3.3.5 Results from a Hysteresis Study

Hysteresis is an effect observed in many different parts of nature, for example in ferromagnetic material. These ferromagnetic materials will make use of the hysteresis property in very commonly used objects, such as credit cards, magnetic tapes and hard disks, to name a few. Hysteresis basically describes the relationship between the induced mag-

netic flux from an external magnetic field and the resulting magnetisation in the system. For a magnetic material this has a microscopic effect. The dipoles align within the presence of a magnetic field. Once the magnetic field is removed the alignment persists partially and results in a magnetised system. The system will not be able to equilibrate on its own accord back to the non-magnetic state. Instead this can only be achieved by increasing the temperature of the system. The interesting aspect is the relationship between the induced flux from the external magnetic field and the resulting magnetisation which is obeying a non-linear relationship. There are theoretical models available which allow the study of hysteresis [62].

The non-linearity of the relationship gives rise to hysteretic loops, which are constructed through measurements of the magnetisation, while slowly changing the external magnetic field. Eventually a point of magnetic saturation will be reached, where all spins are aligned in one direction. When from the point of saturation the external magnetic field is reduced to 0 the system will retain some residual magnetisation. With regard to this study presented here, rather than varying an external magnetic field, the biasing parameter s can be used to tune the area of the hysteretic loops, as it is the source of the induced magnetisation of the system in the first place.

A typical set of hysteresis loops is seen in figure 3.11, for a ferromagnetic system. The external magnetic field h is varied and the resulting magnetisation \tilde{m} is plotted. Here the area of the loop also depends on the biasing field s . In order to obtain the results of figure 3.11 a set of simulations was carried out in the following way: Simulations were started from a representative trajectory of the stationary distribution at a given value of external field h , which is then perturbed by $h \pm \Delta h$ (with $\Delta h = 0.01$). In order to reach a stationary state for the new value of h , n_{relax} steps in the TPS algorithm are taken for equilibration. Then measurements of the observables are taken for n_{obs} steps. Figure 3.11 shows the resulting loops of the magnetisation versus the magnetic field. Again the lattice size was chosen to be $N_s = 200$ and the number of relaxation steps was $n_{\text{relax}} = 1 \times 10^2$. Averages were taken over 1×10^5 realisations. It was observed that the number of relaxation steps taken for each value of h before conducting the measurement, was sufficient for the system to equilibrate to the new value of the external magnetic field.

The top panel of figure 3.11(a)-(c) depicts the hysteresis loop with increasing supercritical s with a trajectory length of $t_{\text{obs}} = 60$, that is $s = 0.05$ for (a), $s = 0.07$ for (b), and $s = 0.08$ for (c). The bottom panel of figure 3.11(d)-(f), shows an identical plot but now with an increased trajectory length of $t_{\text{obs}} = 80$. It can be observed, that the area of the loops increases, as s is increased. This is due to the fact that the metastability of the system increases and thus the relaxation repre-

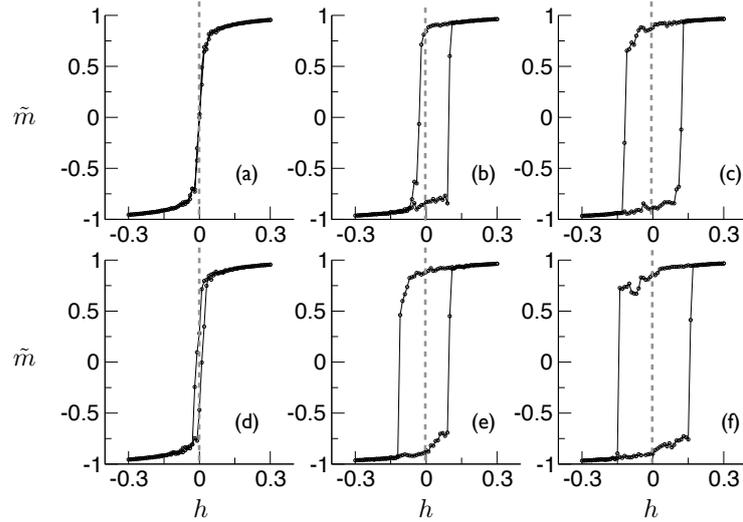


Figure 3.11: magnetisation versus the magnetic field in the form of hysteresis loops at constant temperature $T = 1$ around the first-order transition at $h = 0$. The biasing parameter s is fixed to 3 supercritical values: for (a) and (d) $s = 0.05$, for (b) and (e) $s = 0.07$, and for (c) and (f) $s = 0.08$. The observational time is increased from the top panels of $t_{\text{obs}} = 60$ to $t_{\text{obs}} = 80$ in the bottom panels. Errors are omitted for clarity of the loops. In all plots $h = 0$ is indicated by grey dashed line.

sented by the area of the loop decreases. This supports the previously made argument that a ferromagnetic phase is observed and indicates the validity of the dashed line in figure 3.5(c). Furthermore, the surface (dotted area) below the continuous line is therefore made out of first-order dynamic phase transition points, as depicted in figure 3.5(a).

3.3.6 Results for the Dynamic Phase Behaviour Using Kawasaki Dynamics

Up to this point only single spin flip dynamics, or non-conserved order parameter (magnetisation) dynamics have been considered [63, 45]. Glauber dynamics were chosen out of convenience of the available theoretical predictions, though Metropolis dynamics result in a qualitatively equivalent behaviour. An obvious question to ask is, how the dynamic phase behaviour is affected by the underlying dynamics. So far only single spin flip dynamics were considered, therefore we now turn to the previously introduced Kawasaki dynamics in order to consider spin exchange dynamics instead 3.2.1.2. The Kawasaki dynamics is employed in conjunction with the s -ensemble biasing using the adapted TPS schedule, as described in section 3.3.2.2, carrying out a similar set of simulations as discussed for the Glauber dynamics. The results of this set of simulations will be discussed in this section.

Just to serve as a reminder, neighbouring spins are exchanged according to the Metropolis acceptance criterion, which is given by:

$$P_{\text{accept}} = \min(1, e^{-\Delta\mathcal{H}/k_B T}). \quad (3.41)$$

where $\Delta\mathcal{H}$ is the change in the energy due to the exchange.

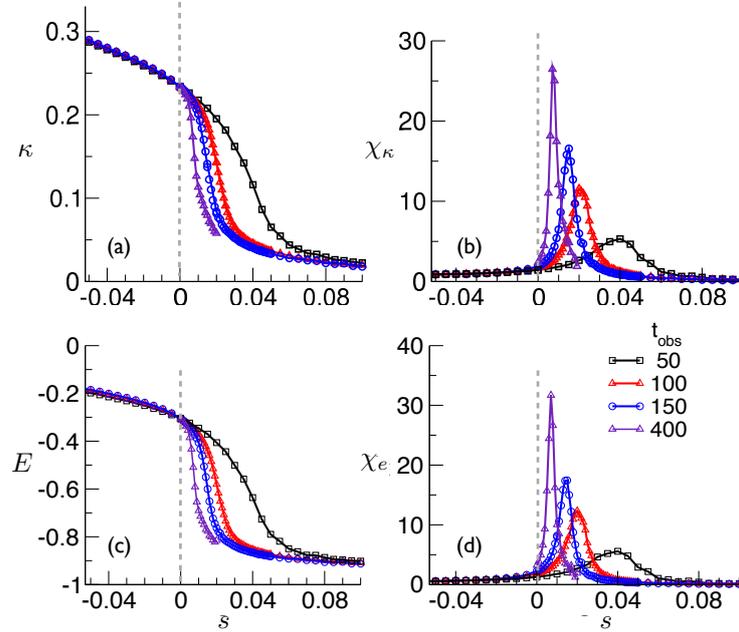


Figure 3.12: Results of the MC simulations using the s -ensemble with Kawasaki dynamics at constant temperature ($T = 3$). (a) Finite size scaling study with t_{obs} as the scaling variable for the activity κ with respect to s and (c) the energy E with respect to s . In (b) and (d) are the respective susceptibilities as a function of s . Errors are within the size of the symbols.

In figure 3.12 the results of the Monte Carlo simulations using the s -ensemble with Kawasaki dynamics for the Ising chain are shown. The chosen parameter set for the simulations was: The temperature was set to $T = 3$, the lattice size $N_s = 64$ and the total magnetisation of the system was zero $\tilde{m} = 0$. As mentioned before, Kawasaki dynamics can be very slow if the simulation temperature is low. Therefore, a simulation temperature that was sufficiently effective at accepting moves was chosen for illustration here. However, other temperatures were looked at and will be commented on later on in this section. Figure 3.12(a) and figure 3.12(b) show the activity and the susceptibility of the activity, given by equation (3.35), respectively versus the biasing field s for different t_{obs} . At the same time figure 3.12(c) shows the energy with respect to s and figure 3.12(d) the susceptibility of the energy with respect to s . Once again finite size scaling was used in order to investigate the behaviour of the dynamic phases. The observational time serves as a scaling parameter and is slowly increased, much in

the same way as presented in section 3.3.2.1. As t_{obs} increases, the observed susceptibilities of activity and energy sharpen in their peaks. For both quantities this behaviour is qualitatively the same, which can be reasoned in the following way. The energy is proportional to the number of domain walls which is the number of sites where anti-aligned spins are neighbouring each other. In order for a change in energy to occur, two neighbouring spins which do not point in the same direction need to exchange, thus activity and energy are closely linked. However spins can exchange within a domain, which will in fact affect the activity but not the overall energy of the system.

From both energy and activity it can be assumed that the observed transition is a first-order transition. The scaling behaviour in a first-order transition very much depends on the size and dimensionality of the system, this is discussed in more detail in [13]. If the peaks χ_k^{max} are chosen to be an effective critical point, depending on the size of the system, the expected scaling behaviour can be written as:

$$s_c(t_{\text{obs}}) = s_c(\infty) + AV^{-1} , \quad (3.42)$$

and

$$\chi_k^{\text{max}} \propto V . \quad (3.43)$$

The constant A is of no further relevance. The volume V is in this case not only dependent on the spatial coordinates but also the observational time. Especially, when taking the spatial dimension as constant, as was done in this case, the volume of the system is directly proportional to the observational time. The volume $V \sim t_{\text{obs}}$, which can be substituted into equations (3.42) and (3.43).

Figure 3.13 illustrates that the finite size analysis holds for the activity of the system. In Figure 3.13(a) the location s_c of the maximum χ_k is plotted for two Ising chains of length $N_s = 64$ and $N_s = 100$ against $\frac{1}{t_{\text{obs}}}$. The effective critical value s_c decreases with increasing t_{obs} and with increasing N_s . Using equation (3.42) an extrapolated value for the limit of $t_{\text{obs}} \rightarrow \infty$ can be calculated. For $N_s = 64$ the value $s_c = 0.002 \pm 0.001$, and $s_c = 0.001 \pm 0.001$, for $N_s = 100$ are obtained. Errors are again those from the goodness of the fitted linear regression. The obtained values are also consistent with the energy analysis.

Figure 3.13(b) shows the height of the peak in the susceptibility of the activity versus the observational time. For the smaller system size the power law given by the equation (3.43) is not observed at all. For the larger sizes there is an intermediate region for which the power law behaviour holds. This shows that for small t_{obs} there are boundary effects in the time of the trajectories, and for longer t_{obs} there are boundary effects in the space of the spin chain.

From the analysis, a phase diagram for the dynamic phase behaviour can now be constructed. There is a clear separation between active and inactive states through a first-order dynamic phase transition. From

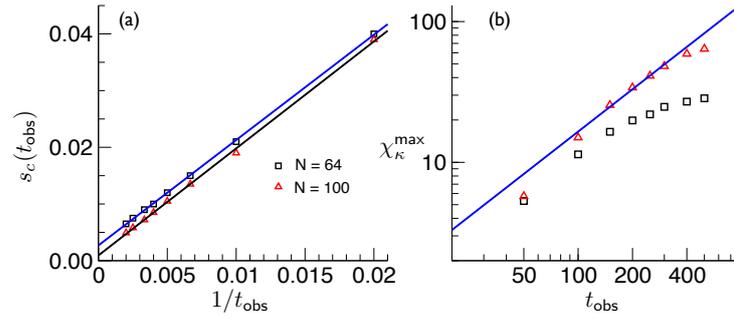


Figure 3.13: Black squares are results from $N_s = 64$, obtained from figure 3.12 (with some additional observation times omitted for clarity in the figure 3.12) and red triangles for $N_s = 100$ (not shown in figure 3.12) (a) Estimation by means of a least square fit of the critical value s_c from figure 3.12(b) with respect to the inverse observation time. The straight lines are fits used for the extrapolations of s_c to the limit $t_{\text{obs}} \rightarrow \infty$ give the following values: $s_c = 0.002 \pm 0.001$ ($s_c = 0.001 \pm 0.001$) for $N_s = 64$ ($N_s = 100$). (b) Height of the peaks of χ_κ from figure 3.12 on a log-log scale with respect to t_{obs} . The straight line is the expected power law behaviour $\propto t_{\text{obs}}$ from the finite time/size theory for first order transitions.

the extrapolation it can be seen that for negative s an active phase is observed and for positive s an inactive phase is observed, with the critical value for $s_c = 0$. This means under non-biased simulation conditions one is at a coexistence of active and inactive dynamical phases, which is a similar observation as seen in idealised kinetically constrained models [2, 3]. A schematic representation of the phase behaviour in the Kawasaki dynamics is shown in figure 3.14.

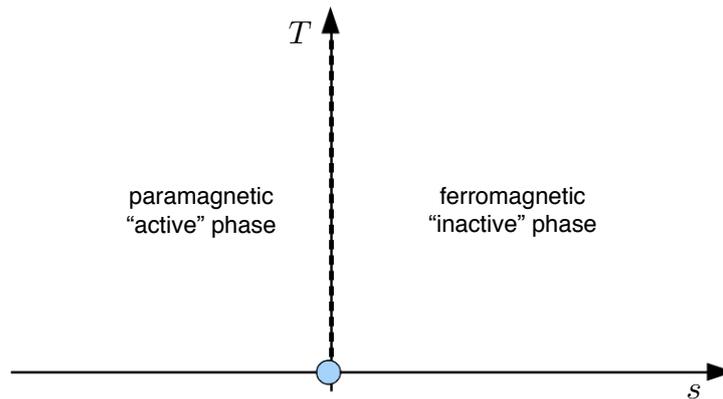


Figure 3.14: Phase diagram of the Ising chain using Kawaksaki dynamics. The dashed line at $s = 0$ represents points of a first-order phase transition between the indicated phases.

The simulation can never truly reconstruct this because of much more prevailing finite size effects in this conserved order parameter model. This further explains the similarity in behaviour of the energy and activity in the inactive phase, as figures 3.12(a) and 3.12(c) are $\sim 1/N_s$. From the way the model is constructed, there will always be two domain walls present from the periodic boundary conditions, which in both cases separate two equal spin domains (one domain, where all spins are pointing up and one where all spins are pointing down). This results in a minimum energy $E - E_0 = \frac{2}{N_s}$, where E_0 is the idealised ground state energy. The two separate domains thus represent a ferromagnet with no activity in the thermodynamic limit. In order to extrapolate this behaviour in a more general way, various different temperatures were considered, which all assert the ferromagnetic/paramagnetic phase transition at a critical value of $s_c = 0$. At low temperatures this transition becomes increasingly difficult to access, as the difference between the phases will be harder to resolve. In the limit $T \rightarrow 0$ the discontinuity disappears and the transition ends in the point $s = 0, T = 0$. Once again, this is analogous to what occurs in kinetically constrained models [2, 3].

3.4 CONCLUSION AND OUTLOOK

Both thermodynamic and dynamic properties of the Ising model in 1-d and 2-d were discussed in detail in this chapter, with particular interest in the previously unobserved dynamic phase behaviour of the 1-d Ising model. In particular, it was observed that in the 1-d Ising model a dynamic phase transition is present for dynamics using single spin flip algorithms with either Glauber or Metropolis acceptance criteria, as well as for diffusive dynamics (Kawasaki dynamics). This allows a direct comparison of the phase diagram of the Glauber dynamics (see figure 3.5) and the Kawasaki dynamics (see figure 3.14). Their differences can be highlighted as follows: a second-order phase transition is observed in the dynamic phase diagram of the Glauber dynamics in the plane of $\{T, s\}$, with a line of critical points given by equations (3.28), where no comparable transition is observed in the Kawasaki dynamics. Instead, a first-order transition at $s = 0$ is found, meaning that in the unbiased ensemble, the system is at a critical state. The similarity found however, is that the phases observed (separated by a first and second-order transition respectively and governed by the activity as the order parameter) are an active paramagnetic and an inactive ferromagnetic phase. Furthermore, looking at the behaviour with an external magnetic field in the case of the Glauber dynamics there is a surface of first-order lines between two symmetric ferromagnetic low activity phases. The transition between them is driven by the magnetic field. The surface ends in a critical line which belongs to the 2-d Ising universality class.

In the future, it will be of interest to extend the study of the dynamic behaviour with s -ensemble biasing further to the 2-d spatial system, where no known analytical solution for the s -ensemble phase behaviour is available. Or even further extend this problem to a more general class of magnetic models, where the spin configurations are allowed to take other values than -1 and 1 , for example. One might wonder why the rare trajectory fluctuations, only reached by an artificial bias are of any interest at all. In fact, the knowledge of the dynamic information of the system is very useful, as their dynamic phase behaviour can play an important role in the equilibrium dynamics of certain types of systems. The existence of rare trajectories away from the equilibrium exhibiting a dynamic phase transition, means that these rare trajectories can occasionally be observed in equilibrium. It is believed that these fluctuations are in fact the reason for the slow relaxation times observed in glasses [4, 1, 44].

The novel work presented in this chapter was carried out by Antonia S.J.S. Mey with advice given by Juan P. Garrahan and Phillip Geissler. A publications regarding this work can be found on the preprint server arXiv [64]. Furthermore, all simulation code developed by the author for the lattice model will be made available on the author's website and can be found under the following link:
<http://www.antoniamey.co.uk>.

This chapter introduces work regarding simulations of a lattice protein and is organised in the following way: A broader introduction to proteins and how these are investigated computationally and experimentally is given. Then an overview of the past work carried out on the lattice model studied here, is presented.

In this thesis the lattice model serves as an approximation to real proteins in order to study their non-equilibrium dynamic behaviour, which would be too computationally costly for an all atom systems. The simplified model allows the study of the dynamic phase space in the s -ensemble in great detail and results of a first-order like dynamic transition in trajectory space between arrested glass-like and active trajectories, for models with heterogeneous interaction potentials.

4.1 INTRODUCTION TO PROTEINS

In recent years, physicists have become more interested in biological systems. Many known physical concepts have since been applied to biological systems. One area of interest in particular, is the concept of protein dynamics and protein folding. This has been studied for many years, but is still poorly understood with regard to even some of the most basic concepts.

In 1952 Sanger *et al.* [65] managed to successfully decoded the first amino acid sequence. He looked at the sequence of insulin, a hormone which plays an important role in the regulation of the carbohydrate metabolism. This was the first step towards understanding the molecular structure of proteins. He showed that the structure of a protein is based on a linear poly peptide chain, which is made of individual amino acids. This sequence, also referred to as the primary structure of a protein, is unique for each protein. Pauling managed to correctly predict the secondary structure of proteins stabilised through their hydrogen bonds [66]. In the 1970's Anfinsen managed to show that a denatured protein can recover its biological active form – also known as

its **native state** – after a denaturing event has taken place – by means of temperature or pH [67]. This suggests that the primary structure alone contains all the necessary information for the protein to find its native 3-d structure.

What complicates the problem of understanding this is the vast amount of different proteins found in nature. Proteins have a broad range of different biological functions and are structurally very different from each other. The protein data base counts 84645 different protein structures on September 18th 2012, most of which have been established through X-ray crystallography and NMR studies [68]. Each of these structures has a unique native state, which is believed to be governed by the primary structure alone. The real problem lies in being able to deduce from any given sequence what its native structure might be. The solution to this problem is still unknown. Various experiments have been carried out **in vivo** and **in vitro**, which conclude that most proteins fold 'spontaneously' into their native state. Only very few are aided by helper molecules called chaperones.

Vast efforts are being made in understanding the underlying concepts in more detail, both experimentally and through computational modelling. The work presented here does not aim to solve the protein folding problem, but looks at the likelihood of observing trajectories away from the folding pathway using computational models. This aims at clarifying how rare dynamic fluctuations can influence the equilibrium behaviour of proteins.

4.1.1 Protein Structure

There are 20 amino acids commonly found in eukaryotes [69] which allows for a large diversity in protein sequences. In particular, taking into account that on average there are between 100-500 monomers in a protein.

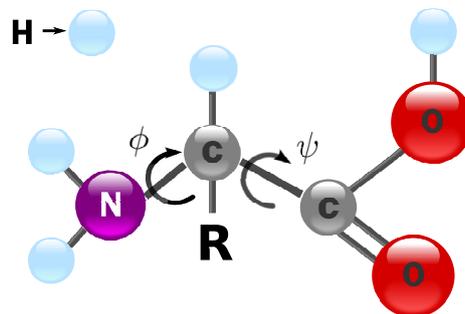


Figure 4.1: Ball and stick diagram of a generic amino acid residue. The R group is the variable group, which can be replaced by one of the 20 possible residues.

Figure 4.1 shows the chemical structures of these amino acids. Each amino acid has a basic structure in common which will later form the **backbone** of a protein. This consists of a central chiral carbon atom C_{α} , which is covalently bonded to a carboxy group (COOH), an amine group (NH₂) and a hydrogen atom. For the C_{α} to form its preferred tetrahedral conformation there is one more ligand space available. This ligand allows to distinguish between the different amino acids and is referred to as the R group – or **residue**, as seen in figure 4.1. The R group depending on the ligand can have different chemical characteristics, which can be grouped in the following way, according to each amino acids name:

- **Hydrophobic residues:** Alanine, Glycine, Valine, Leucine, Isoleucine, Phenylalanine, Thyrosine, Tryptophane, Methionine, Proline
- **Charged residues:** Aspartic Acid, Glutamic Acid, Arginine, Histidine, Lysine
- **Polar residues:** Serine, Threonine, Asparagine, Glutamine, Cysteine

Glycine is a special case, with the side chain consisting of a single hydrogen atom, which makes it non-chiral. For all other residues the C_{α} serves as the chiral centre, due to its 4 different ligands. This means for each amino acid there are two possible enantiomers, a right (R-) and a left (L-) one [69].

In order to describe the 3-d structure of a protein, 4 different 'structure levels' have been introduced.

- The **primary structure** is just the sequence in which the amino acids appear along the protein chain.
- The **secondary structure** is characterising a set of 3-d recurring structural patterns. There are two main types which are observed in protein conformations. These are the α -helix and the β -sheet, as well as various turns and twists.
- The **tertiary structure** describes how the secondary structure elements fit in with each other to form the three-dimensional protein, i.e. whether a protein is for example spherical or elongated.
- The **quaternary structure** can only be found in large proteins which consist of more than one chain. It describes the way individual tertiary domains interact with each other.

4.1.1.1 Secondary Structure

Most proteins consist to 60% or more out of α -helices and β -sheets, which means most of the amino acids are involved in hydrogen bonding in order to stabilise these structures. Secondary structure elements

are believed to be one of the first ones to be formed in a folding process, as they are stabilising the 3-d structure [70]. The existence of these structures was predicted by Pauling *et al.* [66]. He motivated this based on an energetic argument, saying that a large number of hydrogen bonds is favourable and therefore structures that optimise this should be favoured. This is exactly achieved in the secondary structure elements of a protein.

α -Helix

The α -helix is a helical structure with 3.6 residues per turn. Here by residue we mean a generic amino acid. It is stabilised through hydrogen bonding between the carboxy group oxygen and the amine group hydrogen, not involved in the formation of the amide bond (peptide bond). The L-enantiomers of the amino acids will result in a right turned helix.

β -Sheet

The β -sheets are a quasi-two-dimensional structure. This structure is realised through the alignment of the peptide chain in a parallel or anti-parallel fashion and in this way making up a flat sheet and is again stabilised by hydrogen bonding between residues. For more chemical details and other possible secondary structure elements the reader is referred to the book by Berg *et al.* on Biochemistry [69].

These two structures are very distinct, and therefore it can be seen whether an α -helix or a β -sheet is present purely based on a projection onto the two rotational angles or torsional angles ϕ and ψ , as seen in figure 4.1. A projection of proteins onto the dihedral angle space was introduced by Ramachandran in his Ramachandran diagram in the 1960's [71]. In mathematical terms, the Ramachandran diagram is just a function f in terms of the angles ϕ and ψ such that $f : [-\pi, \pi] \times [-\pi, \pi] \rightarrow \mathbb{R}_+$. In a slight adaptation of Ramachandran's original figure 4.2, it can be seen that only certain regions in the ϕ - ψ space are physically populated with a high probability, creating a clear distinction between the α -helix and the β -sheet. Apart from the α -helix and β -sheet torsional combinations, also the left handed α -helix indicated by $L\alpha$ and a different helical structure, the 3_{10} -helix, can be identified in the diagram, but have no further relevance in this thesis. The construction of such a diagram will find more application in chapter 5.

4.1.1.2 Tertiary Structure

The tertiary structure is referred to as the arrangement of secondary structure elements with each other. They are commonly stabilised by salt bridges or sulphide bridges and give a single domain protein its native configuration. The mechanism the protein undergoes in order to

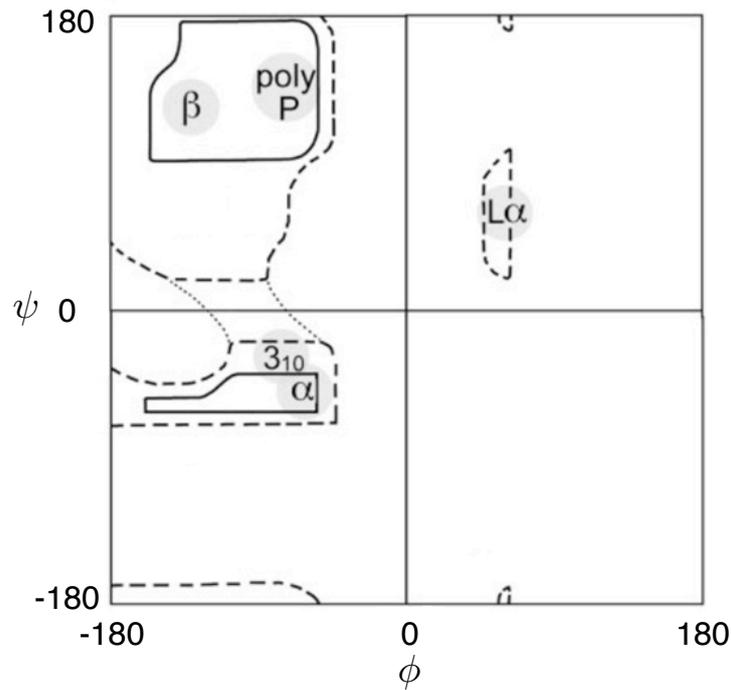


Figure 4.2: Slightly adapted Ramachandran digram as presented in the paper by Ramachandran *et. al.* [71], in the space of the dihedral amino acid angles ϕ and ψ . Angle combinations in the top left corner are most likely to belong to a β -sheet, as indicated, whereas α -helices tend to have angles of around $-50^\circ, -50^\circ$. Other structures are discussed in the text.

reach this structure is still not fully understood due to a gap between possible observation time and the available resolution in experiment. Computationally this is also plagued by the necessity of long simulations and thus not easily attackable. Insufficient experimental resolution, and insufficient computational power are common problems when investigating rearrangements of of protein and can sometimes be referred to as the timescale gap.

4.1.2 States or Phases of a Protein

Proteins, depending on their chemical environment, can have different spatial configurations. These configurations can be distinguished as states or phases, where through the change of a regulating order parameter a different state can be observed.

- The denatured protein

An agent such as temperature or pH can influence the configuration of a protein drastically. When the temperature is changed, the energetics are altered, which can result in the breaking of the hydrogen bonds which stabilise the α -helices and β -sheets. As these structures form the foundation on which a protein structure lies, the conformation is lost and a so called denatured, or

unfolded state is reached. Normally a protein can return to its original configuration even after being denatured.

- The native protein

This represents the biologically active state of the protein. Usually this will be a very compact and low energy state for the protein. Most proteins have a unique native structure which allows for example in the case of enzymes an efficient binding site for ligand molecules.

4.1.3 *The Search for the Native State: the Levinthal Paradox*

The peptide bond has double bond characteristics due to the sharing of π -electrons between the oxygen, carbon and nitrogen atoms. This has an advantage as it restricts conformational space in a way that rotations along that bond become impossible [69]. Therefore, the peptide bond is planar and thus gives rise to the restricted ϕ/ψ space discussed in section 4.1.1.1. Even with these restrictions, there is still a puzzling paradox observed, known as the Levinthal paradox. Assume a peptide with 100 residues, where every residue has only two possible states A and B there would still be 2^{100} possible conformations the peptide could take. Furthermore, assume that the time it takes for the interchange of conformations is in the order of 10^{-13} s. It would take about 10^9 years for the peptide to cover the entire conformational space. However, it has experimentally been shown that proteins are able to fold in a few micro to milliseconds [72]. Hence it can be concluded that the protein folding is not a random search through conformational space for the native conformation, but somehow it must follow energetically favourable pathways [73].

4.1.4 *Experimental Techniques*

Computational power needed for modelling such complex systems only became widely available in the last 20 years. This meant that experiments had a head start and only recently experiments and computational models are being compared to each other to gain more insights. Some of the commonly used experimental techniques are methods, such as X-ray diffraction, nuclear magnetic resonance (NMR), single molecule pulling, Förster resonance energy transfer (FRET) and dStorm [74, 75]. For more details on these see the appendix section on them A. The main challenge faced in experiments is that either it is possible to look at low resolution ensemble averages or high resolution static images. Techniques such as single molecule pulling experiments [76, 77] and dStorm microscopy are working towards overcoming these problems.

4.1.5 *Computational Techniques*

In order to study protein folding and dynamics computationally, many different models have been developed. Nowadays with more computational power available, all atom simulations become much more feasible. However, even simple lattice models can vastly improve an understanding of some of the most fundamental concepts of protein folding. The following will briefly comment on commonly used computational models in order to study these kind of systems stochastically.

4.1.5.1 *Lattice Models*

Lattice models can be thought of as beads on a chain performing a self avoiding random walk on a predefined lattice. The beads represent the amino acids and the links between them the backbone structure of the protein. Generally, one can choose different types of ergodic movesets and different energetic models in order to perform MC simulations on the lattice [78, 79, 80, 81, 82].

4.1.5.2 *Off-lattice Models*

A very widely used class of model are atomistic or coarse grained off lattice models using continuous molecular dynamics.

Coarse grained models

Coarse grained models have the advantage of not having to simulate every single atom in the system and thus allowing for much faster simulations than in the case of atomistic simulations. The chosen degree of coarseness can vary a lot, from where amino acids are represented by single beads, as was the case of the lattice simulations, or in less coarse approaches, beads represent chemical groups which have a certain characteristic [83]. A widely used coarse grained approach is called Martini, which is a simulation package grouping 3 to 4 atoms together and thus tries to be as close to atomistic simulations as possible, but at the same time allowing an integration timestep that is a 10 fold larger than that of atomistic simulations [84]. This is a very useful approach for very large systems, such as protein and membranes.

Atomistic models

For atomistic models, there is a number of simulation packages available, such as GROMACS or NAMD to name a couple [85, 86]. Here, proteins are presented in an atomistic fashion, solvated in water or a different solvent. For large systems it is also possible to simulate the solvent implicitly in order to speed up simulations.

In order to attribute for the interaction energies between different atoms various force fields have been developed. A forcefield is a set

of classical potentials for various interactions (bonded, non-bonded, dihedral... etc.). The parametrisation of such forcefields is non-trivial and uses results from quantum mechanical calculations on the molecular scale, hard encoded into the classical potentials which are then verified according to experimental data. Forcefield development is a very active field and new improvements to existing forcefield parameter are frequently published, as often the desired agreement between simulations and experimental results is still not ideal. Choosing the right forcefield for a simulation can be a difficult task on its own. Once the choice has been made, generally the system is evolved in either an NVT or NPT ensemble with thermostats (e.g Berendsen or Nosé Hoover [87]) and barostat (e.g. Berendsen [88]).

The main problem in molecular simulations that arises is that the required timestep is small (order of 2fs), yet interesting molecular events occur on a millisecond timescale – also known as the timescale problem. Thus, it is often difficult to obtain simulations that are long enough for observing a number of infrequent events, such as the folding of a peptide structure. Therefore the development of Markov state models is very important as it allows a much more efficient estimation of long time scales of the system without having to simulate them, as briefly outlined in chapter 2.

As simulating large systems on a standard workstation is still non-trivial, mass parallelisation projects have been formed such as folding@home [89]. Or the development of super computers such as Anton in the D.E. Shaw research facility [90].

For the purpose of the investigation presented here, and further justified due to the lack of available computational power, the choice of model system was that of a lattice protein, as data generation is fast and abundant and readily done on a normal workstation or small cluster. The following sections will now turn to the discussion of the 2-d and 3-d lattice model in greater detail and will present the results obtained from these models.

4.2 THE LATTICE PROTEIN MODEL

Generally, when studying **lattice proteins**, the protein is interpreted as a **self avoiding random walk** on a lattice, with each vertex being occupied by an amino acid and each edge connecting the amino acids in a way the backbone would do in a real protein. The links along the formed chain are unbreakable, representing the strong peptide bonds. Amongst the lattice models one can find different variants such as the Hydrophobic Polar (HP) model by Ken Dill or an even older lattice model, the Gō model [79, 91]. Here variants of the Gō model on a quadratic and cubic lattice will be considered.

The Gō model, as proposed by Abe and Gō in the early 1980's, relies on the predefinition of a **native state**, in order to impose a Hamiltonian

carrying the information of the interaction potential [79]. This follows from the idea that the native state is an energetically favourable state and will drive the folding process. The model, by construction, allows the study of the dynamics and equilibrium properties within the restraints of the model but is not concerned with the prediction of the native state from a given amino acid sequence, which presents a wide research area on its own [92]. However, it is striking how well such simplified models can capture real protein behaviour, which will be discussed in more detail in the following sections.

4.3 EQUILIBRIUM PROPERTIES OF THE 2-d GÖ MODEL

As an introduction into the world of lattice proteins, we will review a well studied 2-d lattice example, by presenting some of its thermodynamic properties. First, the case of a quadratic lattice will be considered, for which a native state needs to be defined in order to further assign interaction energies. The native state definition is motivated by the observation of real life proteins. In comparing with nature, it is often found that native states are generally very compact structures, so it seems to be an obvious choice to use a maximally compact random walker on the lattice. The length of the walker L_p indicates the number of amino acids along the chain. As the length increases the number of possible native structures to be used increases drastically [93, 94]. These enumerations are generally achieved via recursive algorithms and to the author's knowledge, no exact formula is available to give the number of possible configurations (maximally compact or not) for a chain of a given length. For simplicity, a chain with $L_p = 16$ (16mer) amino acids is considered here. An exemplary choice of a native state is depicted in figure 4.3, which was also used in [95]. **Native contacts** are defined as the set of nearest neighbour contacts in the native state, observed not along the backbone of the chain, also indicated in figure 4.3, by the dashed blue lines. The number of native contacts (N_c) in figure 4.3 is $N_c = 9$. Any arrangement of this chain is referred to as the protein's **configuration**.

4.3.1 Hamiltonian

The interaction energies defined by the Hamiltonian of the Gō model only considers nearest neighbour interactions, between native contacts. The Hamiltonian can be expressed as:

$$\mathcal{H} = \sum_{i < j}^{L_p} B_{ij} \Delta_{ij} , \quad (4.1)$$

where Δ_{ij} is a nearest neighbour contact map of the protein configuration and B_{ij} holds the information of the native contacts between residue i and j . The strength of the interaction potentials can be cho-

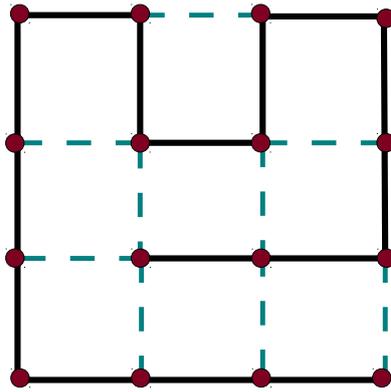


Figure 4.3: Native configuration of a 16mer. Black links represent the backbone, red circles the amino acids and the dashed blue lines represent the $N_c = 9$ native contacts.

sen arbitrarily, but will initially be set to -1 in arbitrary units. With the knowledge of the Hamiltonian, the partition function and equilibrium properties associated with it can be evaluated, as discussed in chapter 2. So, for example, the probability of being in the native state is known, if the enumeration of all configurations is possible:

$$P(E_g) = \frac{\exp(-\beta E_g)}{Z} . \quad (4.2)$$

$P(E_g)$ is the probability being native, i.e. being in a ground state energy and Z as usual the partition function, consisting of the sum of all possible configurations. In the regime of short chain lengths $L_p < 28$ on a square lattice, exact enumerations of all conformations is achievable on a standard work station. From this enumeration the exact partition function of the system can be deduced, and thus all thermodynamic properties can be calculated exactly. This is, for example, the probability of being in the native state at a given temperature, as given by equation (4.2). This has been studied by Lee and Chan [96, 91, 97] amongst others. For larger systems or the dynamic behaviour of these systems however, simulations will be necessary. In order to setup these simulations, it is important to define a moveset and define an acceptance function, such as the Metropolis criterion, for the proposed MC moves. With the simulations properties such as mean or median folding times at different temperatures can then be studied.

4.3.2 MC Simulations

Before being able to run a standard MC simulation much in the same way as presented in chapter 3, an equivalent of the single spin flip move needs to be established. Such a move set was previously used in lattice model studies by Shaknovich [98] and also others. It consists

Listing 4.1: Pseudocode: MC algorithm for lattice protein simulation

```

Step1: Compute the instantaneous energy of the system
       according to the hamiltonian
Step2: Choose with 0.8 probability to do a single bead move
       and 0.2 probability to do a double bead move
Step3a: If single bead move: Choose a random bead and random
        move direction.
Step3b: if the move is allowed (i.e., chain doesn't break or
        two beads occupy same lattice site), accept move
        according to the Metropolis criterion
Step4a: If double bead move: Chose a random neighbour and a
        move direction (up down).
Step4b: If the move is allowed (i.e. the chosen beads are in
        a correct crank shaft position), accept move according
        to the Metropolis criterion.

Step 5: Repeat 1-4 until the loop criterion is fulfilled.

```

of single residue and double residue moves which are depicted in figure 4.4. Most importantly it fulfils detailed balance.

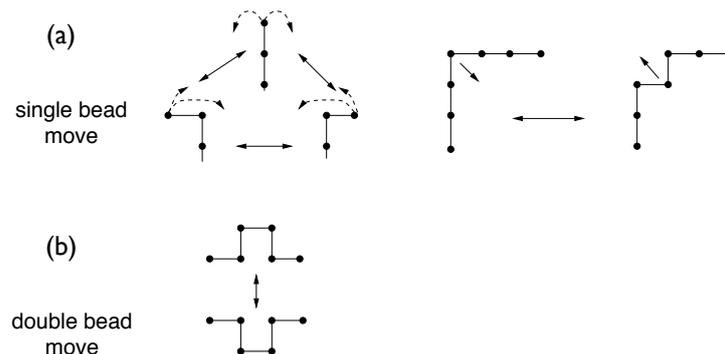


Figure 4.4: (a) (left) End flip of 90° (right) corner flip. (b) The double residue move is a crank-shaft move, where two residues are flipped by 180° .

The algorithmic setup closely follows [95] and as such aims to reproduce some of the results presented in [95], in order to get an idea about properties of interest as well as validating the algorithmic aspects of the work. The MC algorithm using the proposed moveset will follow the steps displayed in the listing 4.1. The used acceptance criterion is again chosen to be the Metropolis acceptance criterion i.e equation 3.3 in chapter 3. Ratios for attempted single and double bead flips were initially varied but, best results were achieved with a ratio of 20% to 80%. Generally, time units (such as a MC lattice sweep) are scaled according to the amino acid chain length L_p , in order to have comparable numbers between chains of different lengths.

Next a set of observables need to be defined, such as: the average energy $\langle E \rangle$ with respect to temperature T , average number of nativeness, which is just the fraction of native contacts $q = \frac{n_c}{N_c}$ formed at a given point in time (n_c "current number of native contacts" and N_c "the total number of possible native contacts") with respect to temperature, i.e. $q(T)$. This observable can then be averaged over a whole simulation $\langle q(T) \rangle = \frac{1}{t_{\text{obs}}} \sum_{t'=0}^{t_{\text{obs}}} q(t')$. Dynamic types of observables might be the mean first passage time (MFPT) to the native state [95], or contact appearance order (CAO) [82]. In the following some example simulations were carried out to illustrate how these observables within this model mimic some of the real world protein behaviour. These results follow some of those presented in [95] and were used to validate the Monte Carlo algorithm.

4.3.3 Equilibrium Simulation Results

As an example, the 16mer, with the depicted native state of figure 4.3 is simulated using the MC dynamics described in the listing 4.1 with the Gō model Hamiltonian of equation (4.3). First of all, the equilibrium

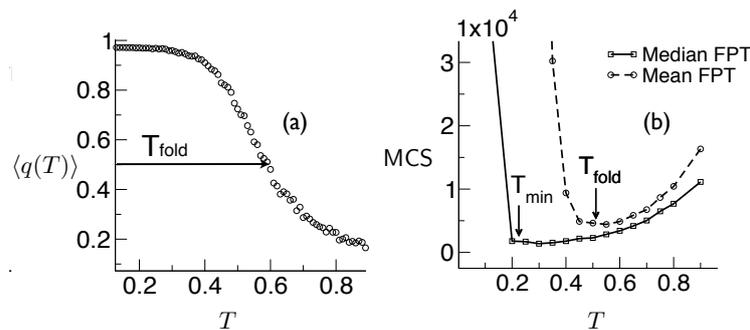


Figure 4.5: (a) Depicts the average nativeness as a function of temperature $\langle q(T) \rangle$. Errors are omitted for clarity. (b) Illustrates the results of a MFPT and median first passage time simulation with respect to temperature. Results are obtained from 10000 independent simulation runs at each temperature. Standard errors are within the size of the symbols.

observable of nativeness with respect to temperature $\langle q(T) \rangle$ is considered. Figure 4.5(a) shows a sigmoidal curve of having a temperature regime for which the polymer is mainly native and a regime where it is mainly unfolded. Sometimes this behaviour is said to be first-order like. The observed behaviour compares well to real life proteins which denature as their surrounding temperature is increased. Secondly, figure 4.5(b) shows the MFPT and median first passage time required for the native state to be reached as a function of temperature. Averages were taken over 10000 independent realisations. The size of the data points is representative of the size of the standard error, i.e. the stan-

standard deviation from the mean, divided by square root of the number of realisations:

$$\epsilon_{\text{std}} = \frac{\sigma_{\text{MFPT}}}{\sqrt{n}} .$$

All the results presented are in reduced units for $k_B = 1$ and interaction energies set to -1 in arbitrary units.

From figure 4.5 an optimal folding temperature can be extracted for which the first passage time is the shortest. Secondly the **folding temperature** T_{fold} can be defined as the temperature where half of the native contacts are formed on average, i.e. $\langle q(T_{\text{fold}}) \rangle = 0.5$. These are both observables which will be of interest in later parts of this chapter.

4.4 THE 3-d LATTICE MODEL

Next, a further dimension is added, moving to a lattice protein living on a cubic lattice. The main results of this chapter will be presented for the 3-d model. These will include some previously known equilibrium results as well as a novel investigation into the dynamics. This is done by employing the formalism as introduced in chapter 2.

4.4.1 Hamiltonian and Native State Definition

The native state of the system is again defined by a possible maximally compact structure of a self avoiding random walker. For the shortest lattice protein, an 8mer structure, one of the three possible native states and the corresponding native contacts are depicted in figure 4.6(a) and (b) respectively.

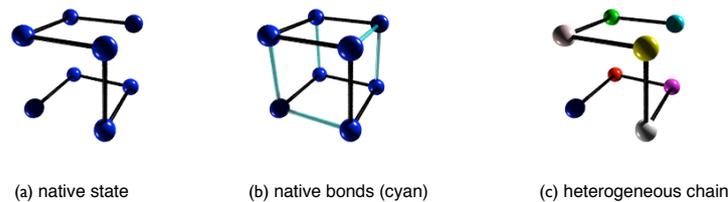


Figure 4.6: (a) The maximally compact state of an 8mer chain with beads as amino acids is depicted. (b) Native bonds are indicated in cyan and in (c) a heterogeneous sequence as in the HeG \bar{o} and Full case are illustrated.

As before, it is possible to define a Hamiltonian. Now not only the Hamiltonian of the G \bar{o} model will be considered, but also two different variants of it. For an easy distinction of the three variants the following nomenclature is introduced: The homogeneous G \bar{o} (G \bar{o}) represents the original case, where only native contacts are considered, associated with a uniform interaction energy between them. This could be also considered as an amino acid sequence consisting of a single monomer.

Secondly, a variant of the Gō potential with heterogeneous interaction energies, here referred to as the HeGō case, is considered. In this case again, only native contact interactions are used, but the underlying amino acid sequence is heterogeneous, i.e. a sequence consisting of a combination of the 20 natural amino acids. A pictorial example would be figure 4.6(c), where each amino acid is represented by a differently coloured bead. Lastly, a full interaction potential with native and non-native nearest neighbour interactions is considered based on the amino acid sequence of choice. This model will be referred to as the Full model.

Mathematically the Hamiltonian takes the form:

$$\mathcal{H} = \sum_{i=1}^{L_p-1} \sum_{j>i}^{L_p} U(r_{ij}) + \sum_{i=1}^{L_p-3} \sum_{j=i+3}^{L_p} N_{ij} B_{ij} \Delta(r_{ij} - a) , \quad (4.3)$$

where $r_{ij} = |r_i - r_j|$ is the distance between two amino acid residues. The potential $U(r)$ restricts the walker to be self avoiding, as it takes a value of ∞ for $r = 0$ and 0 for any value of $r > 0$. In the second term B_{ij} , is an energy interaction matrix as before in the 2-d case but is now determined by the sequence of the chain. The lattice spacing is given by a , as previously. Interaction values for different amino acids are drawn from the model of Miyazawa and Jernigan, which uses experimental values of protein interaction energies [99]. This provides a table of all amino acid interaction energies between each other and can be thought of as non-covalent bond interactions. These range from positive interaction energies to values of $\epsilon = -1.06\epsilon_0$ for interactions between cystein residues. This interaction energy scale also sets the reduced units for all following simulations. Thus, the temperature for this model is always given in the reduced scale of $T = \frac{\epsilon_0}{k_B}$.

The available permutations and thus the choice of native states and amino acid sequences rapidly becomes very large as the chain length increases. For simplicity, here mainly one particular native state of a 48mer will be the model of choice, with its 57 native contacts. This was partly motivated by previous studies, where this structure was used [82], such that a comparative analysis would be possible but also aided in the model verification. With the help of PyMOL, the native state structure can be depicted as seen in figure 4.7(a). Figure 4.7(b) shows an exemplary amino acid sequence in single letter code. This also presents the sequence used for most the results presented in this chapter. The appendix B.1 contains an explanation for the single letter code.

4.4.2 Equilibrium Simulation Results

Initially, some equilibrium properties and the MFPT behaviour are investigated. This was done in a similar manner as presented by Sha-

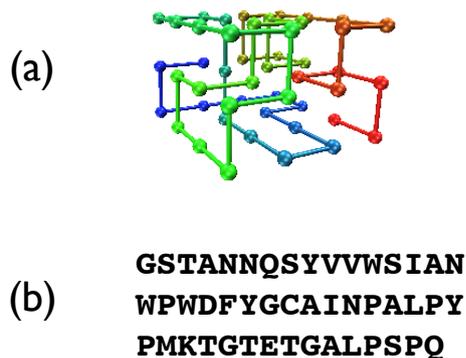


Figure 4.7: (a) Chosen native structure. Colour code aids the eye in following the chain from one end to the other. (b) Representative amino acid sequence in one letter code.

knovich and co-workers [80]. The motivation for this investigation in our case, is to scan the equilibrium parameter space to set the ground work for more complex investigations. The flexibility of the Hamiltonians allows a detailed investigation of the equilibrium and folding behaviour of the lattice protein and in particular considering the effect of non-native interactions have on the dynamics. The simulations are conducted in the same way as the 2-d model, discussed in section 4.3.2. The only difference is that the crank-shaft move now results in a 90° orientation change.

Once the native state is chosen, next the interaction potential needs to be established. Thus, for any non-G \bar{o} simulation, an amino acid sequence is defined. For most of the following work the sequence is chosen to be the one of figure 4.7(b) unless otherwise stated. This choice is supported by the fact that the sequence is a fast folder in the Full and HeG \bar{o} potential. The sequence was generated using an evolutionary algorithm in order to optimise folding efficiency for the particular native structure given. For more details on evolutionary sequence generation see [82].

For the simple G \bar{o} model the choice of sequence is also adjusted to the amino acid interaction energy scale and chosen to be all Valin amino acids, where a Valin-Valin interaction in the reduced energy scale is: $-0.29\epsilon_0$.

Initial configurations for the simulation setup are drawn from a high temperature ensemble and equilibrated at the simulation temperature for 1×10^5 simulation steps.

Initially, the temperature dependence on $\langle q(T) \rangle$ is evaluated in order to obtain a value for T_{fold} , from a series of independent temperature simulations. This was done for a temperature range of $T = [0.1, \dots, 0.3]$ for the G \bar{o} potential, the HeG \bar{o} potential and the Full potential.

Generally, from these models a first-order like sigmoidal curve for the dependence of nativeness with respect to the temperature is expected. This behaviour is clearly observed in figure 4.8(a). Results are drawn from 1000 independent simulation runs at each temperature. The simulation length, or observational time $t_{\text{obs}} = 1 \times 10^7$ for the Gō model and the Full model and $t_{\text{obs}} = 5 \times 10^6$ MC steps for the HeGō model. The HeGō (red) and Full (blue), which have the same amino acid se-

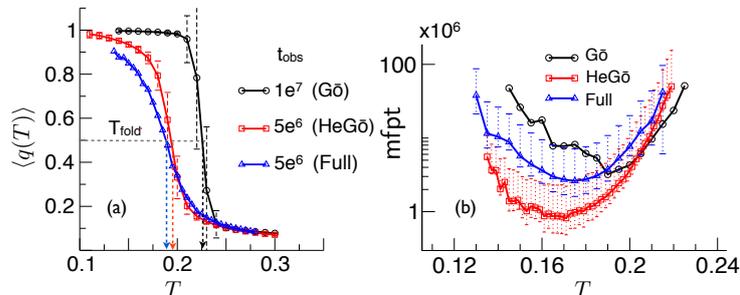


Figure 4.8: (a) Fraction $\langle q(T) \rangle$ of native contacts made with respect to the temperature for all three models (Gō black, HeGō red, Full blue) over 1000 independent realisations. T_{fold} is indicated in all cases, error bars shown are the standard deviation from the simulation and omitted in the case of the Full model for clarity. (b) Mean first passage time to the native state on a log scale for all interaction model with the same colour code as in (a). Error bars represent $1\text{-}\sigma$ confidence intervals from 1000 realisations. Error bars for the Gō model are omitted for clarity.

quence exhibit very similar behaviour, whereas the Gō structure is shifted, probably accounted for by the lower ground state interaction energy of $E_g(\text{Gō}) \sim -16\epsilon_0$ as opposed to the ground state interaction for the heterogeneous sequence of $E_g(\text{HeGō}) \sim -14\epsilon_0$. Furthermore, it becomes apparent that T_{fold} is higher for the Gō structure.

Next the MFPT will be investigated. Results of this investigation are shown in figure 4.8(b). Again all three variants of the model were considered. The HeGō model performs much better at folding to the desired compact structure, as a more specific folding pathway can be selected aided by the heterogeneity of the sequence. Also in the case of the Full interaction potential the folded state is efficiently reached and thus non-native contacts do not prevent the model from folding in this case. The slowest folder, with the longest MFPT is the Gō model despite its favourable potential. This can be seen in the log plot of the MFPT with respect to temperature in figure 4.8(b).

These two observables (q and MFPT) all give information about the average behaviour of all variants of the Gō model. When varying chain lengths and sequences, it can be shown that most of them will behave in a similar way (provided they are able to fold at all.) It is generally believed, that a nucleus will need to form in order for the folding process

to be a downhill event. The classification and formation of these nuclei has been done by Shaknovich and co-workers [94]. Furthermore, for specific sequences very distinctive folding pathways can be extracted, whereas for the G \bar{o} structure this is not the case. This was further established recently, by introducing an observable looking at the contact appearance order (CAO) in the folding process [82]. The CAO is defined in a way that once a contact is formed and stays formed until the native state is found, is a permanent contact and thus appears in the CAO.

4.5 THE LATTICE PROTEIN MODEL IN THE S-ENSEMBLE

The next part of this chapter will be concerned with applying the s-ensemble formalism formally defined in chapter 2 and applied to the Ising model in chapter 3, to the lattice protein dynamics in order to probe the behaviour of the trajectory space of these models further.

4.5.1 MC Simulations Setup

As discussed in 4.3.2, MC simulations were carried out for the lattice protein model. With respect to these simulations an s-ensemble bias can also be used with these simulations, much in the same way of the adapted TPS approach, as presented in chapter 3.

For this purpose, a set of dynamic observables needs to be defined. There are two obvious variables that come to mind, both scaled in system size and observational time. The first one is the general activity of the system (gActivity) $\kappa_g = \frac{K}{N t_{\text{obs}}}$, which just represents any accepted MC move of the chain resulting in a conformational change of the chain. The activity K is the number of accepted MC moves. The second one is a more restrictive one, looking at the native activity of the system (nActivity) $\kappa = \frac{K_n}{N t_{\text{obs}}}$, which represents the incremental number of broken or formed native contacts (K_n) during the observed trajectory. The latter observable will be the observable of choice in this thesis, as the primary interest is to look at the influence of the native contacts with respect to the dynamics. It should be noted that notation is the same as the activity of chapter 3. The generation of s-ensemble biased trajectories follows the same ideas as presented in chapter 3. The schematic that illustrates the TPS scheduled used for biasing (figure 3.6), is still valid, but now instead at each timestep saving a lattice spin configurations, the 3-d lattice protein configuration is saved, generating time dependent trajectories in this way. Along the trajectory, shooting points are chosen at random and from these, the new trajectory is evolved according to the MC dynamics. A new trajectory is accepted based on the s-ensemble acceptance criterion:

$$P_{\text{accpt}} = \min(1, \exp(-s\Delta K)) , \quad (4.4)$$

where K represents the unnormalised activity measure of choice. The choice of the observation time is a bit more limited than in the Ising

model. The idea is to generally use observation times which are longer than the mean first passage time to the native state of the system, which is assumed to be the slowest time scale of the system. This means that already initial observation times are long and finite time scaling becomes non-trivial, as the TPS algorithm will be more likely to reject newly proposed trajectory, purely due to the length of newly generated trajectory parts. Due to this, we will mainly use observation times between 1×10^6 and 1×10^7 MC steps, as these observation times are of the order of the mean first passage time or longer. None the less, with these ideas, it is hoped to uncover the rare dynamic fluctuations of the system.

4.5.2 $G\bar{o}$ Results

First the homogeneous $G\bar{o}$ chain will be discussed. The dynamic behaviour is very straight forward which is not at all surprising, as the Hamiltonian contains no heterogeneity.

4.5.2.1 $G\bar{o}$: Dynamic Behaviour in the Equilibrium Ensemble

The average native activity $\langle \kappa \rangle_0$ (the 0 subscript means we are considering the system at equilibrium, i.e $s = 0$) with respect to temperature will be considered first in order to have an idea about the scale of the of activities accessible at equilibrium. Figure 4.9(a) shows the average native activity of the system with respect to temperature. A fairly sharp sigmoidal increase is observed, for observation times of the order of the MFPT of the system. In fact, by looking at different observation times it was observed that altering the overall observation time of the system, the qualitative behaviour of figure 4.9(a) did not change. The T_{fold} temperatures corresponds to the sharp changeover in activity. Looking at the probability distribution of observed native activity κ , two clear peaks are observed for different native activities, as seen in figure 4.9(b).

This suggests a first-order like transition between two activity regions. In this simplified potential the low and high activity regions correspond to trajectories mainly containing native or near native states. In this case the dynamics gets restricted by the energetically favourable compactness of the system. The more active regions correspond to trajectories with mainly unfolded configurations. This can also be illustrated with a plot of κ with respect to the integrated nativeness integrated, for each trajectory, as seen in figure 4.10. The integrated nativeness is thus given by:

$$Q = \frac{1}{t_{\text{obs}}} N_c \sum_{t'=0}^{t_{\text{obs}}} q(t') . \quad (4.5)$$

The observed dependence of Q and κ is linear. A linear fit of the data suggests that the system will always retain a base activity of

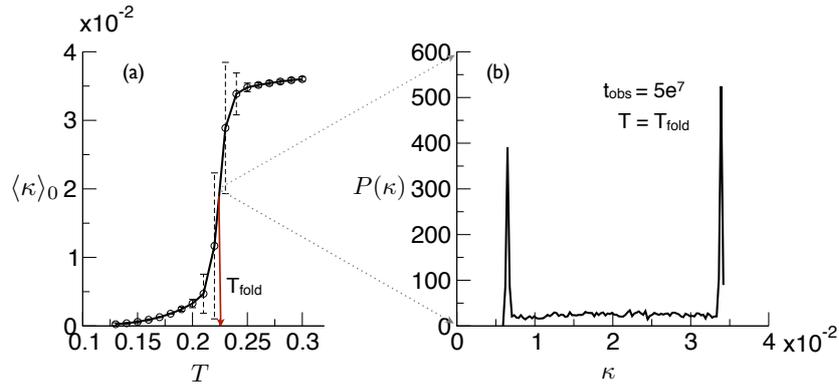


Figure 4.9: (a) Shows $\langle \kappa \rangle_0$ from averages of 20000 simulations. The standard deviation is shown as the error bars, which means that standard errors will actually fall into the size of the data symbols. (b) Displays $P(\kappa)$ (a histogram of all values of κ observed) from simulations of length $t_{\text{obs}} = 5 \times 10^7$ and 20000 independent runs.

$\kappa \approx 0.0054$, the case where all native contacts are formed at the given temperature $T = 0.225$. This base activity will be lower for lower temperatures. It should be noted that if the observational time is sig-

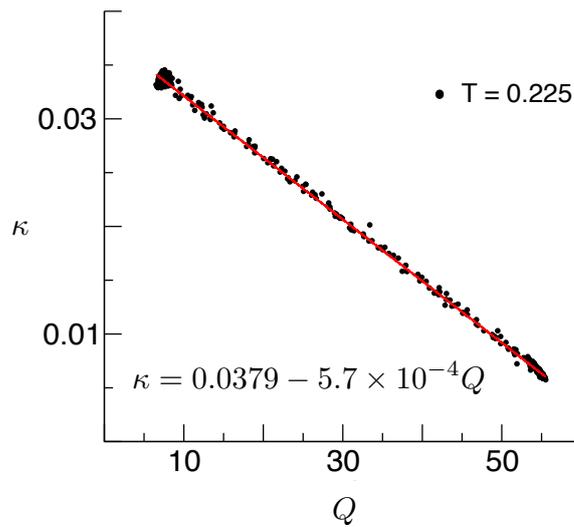


Figure 4.10: Scatter plot of κ with respect to Q for a subsample of 5000 of the 20000 trajectories. The linear fit in red suggests a minimal activity of $\kappa \approx 0.0054$

nificantly increased, the bimodal behaviour of the system is lost, thus confirming the notion of a finite system in which no phase transition can be observed.

4.5.2.2 *Gō model: s-ensemble Biasing*

Up to this point the dynamic observable κ was only considered in the equilibrium ensemble, where $s = 0$. Now the idea is to bias the system in such a way that even at temperatures below the folding temperature and above the folding temperature, the dynamic phases between high and low activity can be accessed. Taking $T_{\text{fold}} \approx 0.225$, then in each case two temperatures below ($T = 0.18$ and $T = 0.19$) and above that folding temperature ($T = 0.25$ and $T = 0.26$) were chosen for the s -ensemble biasing. Averages were taken over 2×10^4 realisations which was relatively easily achievable on computer cluster and resulted in sufficient data for the error analysis.

In case of the Ising model, a critical value of s was known for any temperature due to the analytical result available for the system as discussed in chapter 3. Here the space of s is initially scanned blindly in order to obtain an estimate for a critical value. Figure 4.11 shows the results of the analysis from the activity behaviour $\langle \kappa \rangle_s$ (now the subscript s indicates that we are taking averages over activity values with a biasing value of s) with respect to the biasing parameter s . For

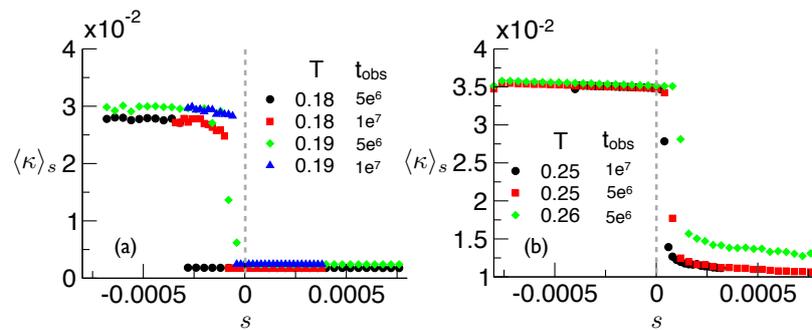


Figure 4.11: (a) Shows the results of $\langle \kappa \rangle_s$ with respect to s for $T = 0.18$ and $T = 0.19$ at two different t_{obs} . Standard errors are within the data points as averages were taken over 2×10^4 realisations (b) Results are shown for the two high temperatures $T = 0.25$ and $T = 0.26$ and different t_{obs} . Standard errors are within the data points as averages were taken over 2×10^4 realisations. The unbiased value of s is indicated by a grey vertical line in (a) and (b).

both the low temperature (figure 4.11(a)) and high temperature (figure 4.11(b)) case with a negative and positive s value respectively, the high or low activity trajectories can both be reached (i.e. mainly native and mainly unfolded trajectories, non-typical for the observed temperatures).

As was done in the Ising model, a finite time scaling approach was used. For this purpose the investigation was carried out with different observation times, which are: $t_{\text{obs}} = 5 \times 10^6$ and $t_{\text{obs}} = 1 \times 10^7$. Trajectories with even longer observational times became numerically challenging due to a high number of rejections of proposed trajectories, thus making the simulations very slow even with the adapted TPS

method. However, in figure 4.11(a) and (b) it can be seen, that the dynamic crossover at $s = 0$ can be shifted with the application of the external biasing field s to lower and higher temperatures and scaling is observed to the extent, that the critical value of s_c decreases as the observational time is increased. It should be remembered that here the chain length is a further limiting factor, as no periodic boundary conditions in the spatial dimension can be applied.

4.5.2.3 *Gō Model: Phase Diagram*

The analysis from the results in the previous section leads to an idealised phase diagram in the context of this lattice model, where a first-order like dynamic phase transitions can be assigned to a set of critical biasing parameters s with respect to temperature. The critical values for s are extrapolated from figure 4.11. From this, a phase diagram in parameter space $\{T, s\}$ can be constructed as seen in figure 4.12.

It becomes clear, that there is a **static** as well as **dynamical** separa-

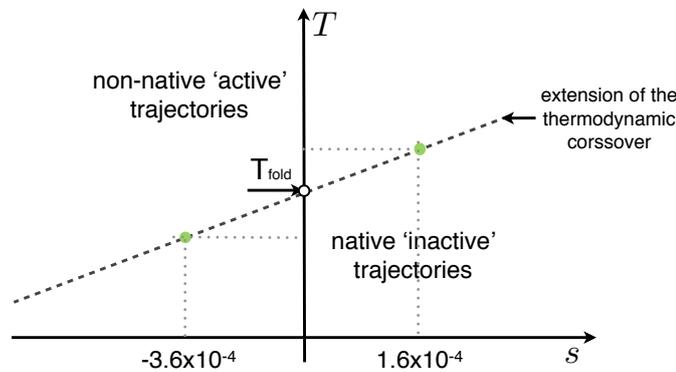


Figure 4.12: Schematic of the phase diagram of the Gō model. Active and inactive phases are indicated and separated by the first-order like dynamic crossover (grey dashed line). The green points are values taken from figure 4.11, where the exact location of the critical value is much harder to determine than in the Ising model.

tion between the folded and unfolded state of this system, which can be described by a first-order like transition between these states. The active phase, with many native contacts being broken and formed, always corresponds to trajectories which on average are more unfolded than folded. The less active phase however, consists of trajectories on average found to be in the native state. The artificial biasing with the conjugate field s can bring this interplay between folded and unfolded trajectories back into balance even at lower and higher temperatures than those, where the transition is observed in equilibrium.

The same behaviour as observed for the 48mer can also be observed for different chain lengths and different compact native structures. The

same investigation was carried out for a 27mer and different target native structures. In all cases there was a clear dependence on native activity κ with the number of native contacts present on average along the trajectory. Therefore $G\bar{o}$ models in general can be associated with a first-order like dynamic phase transition, associated to native and non-native trajectories.

Lastly, one might consider also using the general activity as an observable for biasing. Again a clear dependence on the average nativeness and activity for a given trajectory is found leading to a very similar overall dynamic behaviour. Despite the general activity indicating the system to be more active than when considering native activity, the overall qualitative observation of the dynamic phase behaviour is the same. A discussion of the thorough analysis of these was omitted for brevity.

4.5.3 *HeG \bar{o} and Full Interaction Potential Results*

Having discussed the dynamic biasing of the $G\bar{o}$ model, the next step is to look at the dynamic dependence of the HeG \bar{o} and Full interaction potentials.

4.5.3.1 *HeG \bar{o} and Full Model: Dynamic Observable in the Equilibrium Ensemble*

As was done for the case of the $G\bar{o}$ potential, the average native activity is considered as a function of temperature. Figure 4.13(a) shows this behaviour for both the HeG \bar{o} model (red line) and the Full model (blue line). Already in the dynamics at equilibrium a sigmoidal behaviour

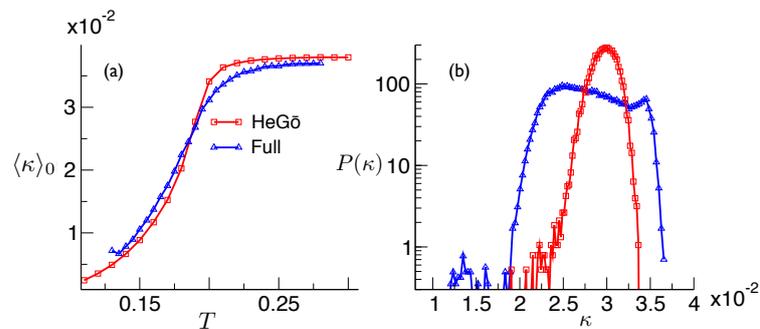


Figure 4.13: (a) Shows $\langle \kappa \rangle_0$ for the HeG \bar{o} (red) and Full (blue) model with respect to the simulation temperature from averages of 2×10^4 realisations. Standard errors are within the size of the data point symbols. (b) Displays the probability distribution of the native activity κ at the folding temperature of around $T \sim 0.19$ on a log-scale. All data is obtained from simulation trajectories of length $t_{\text{obs}} = 2 \times 10^6$, near the mean first passage time for both potentials.

is observed. From the log plot of the probability distribution of the

native activity at T_{fold} a long tail of low activity is found, hinting at a low activity dynamic behaviour becoming more accessible through s-ensemble biasing. This is depicted in figure 4.13(b).

4.5.3.2 HeGō Model: s-ensemble Biasing for 3 Temperatures

First the HeGō potential will be considered. In this case, the dynamic behaviour is more interesting in comparison to the Gō model, yet still conclusive. The HeGō simulations of figure 4.14 constitute the main result of this investigation. First we will discuss results for the observation temperature $T = 0.19$, which is close to the T_{fold} temperature. The main result is the observation of a transition from an active set of trajectories to a less active, or low activity set of trajectories, as the biasing factor s is increased. A coexistence region for this behaviour can be found in the TPS trajectories, where both dynamic states are roughly equally populated, which corresponds to the critical value of s_c .

From such a coexistence region a histogram of the native activity can be evaluated. This was done from 40000 critical trajectories each of observation length $t_{\text{obs}} = 5 \times 10^6$. Coexistence is assured by calculating the area under the histogram, making sure that the histogram of s_c has the same number of highly active trajectories as low activity trajectories. With reweighting of the histogram, the dependence of $\langle \kappa \rangle_s$ and s can be extracted. For this purpose methods such as weighted histogram analysis method (WHAM) can be used [100]. In this way, only good sampling around the critical value s_c is required. From the reweighting of the histograms, the dependence on activity and s can be extracted. This is particularly advantageous for large values of s , where simulations become tricky, as newly proposed trajectories become rare even in the TPS simulation scheme. The reweighting scheme is employed here, as opposed to direct sampling for many different values of s . The reason for this choice is that the protein trajectories are significantly more costly than those of the Ising model, thus reducing the overall required amount of simulations for this model. The reweighting scheme had also been employed in previous s-ensemble studies [4, 5].

All results displayed in figure 4.14, use the histogram reweighting method in order to extract the activity behaviour in terms of the biasing variable. Let us look at the T_{fold} results first found in figure 4.14(d)-(f). The reweighting of histograms was applied to trajectories at three different observational times $t_{\text{obs}} = 1 \times 10^6$ (black), $t_{\text{obs}} = 5 \times 10^6$ (red) and $t_{\text{obs}} = 1 \times 10^7$ (blue). From the finite time scaling using the observation time as the scaling parameter, from figure 4.14(d) and (e) it becomes clear, that a first-order like dynamic transition in trajectory space is observed. As t_{obs} increases, the critical value of s_c decreases as seen by the moving peaks and increased sharpness of the susceptibility (χ_κ) of the activity κ in figure 4.14(e). Here κ serves again as

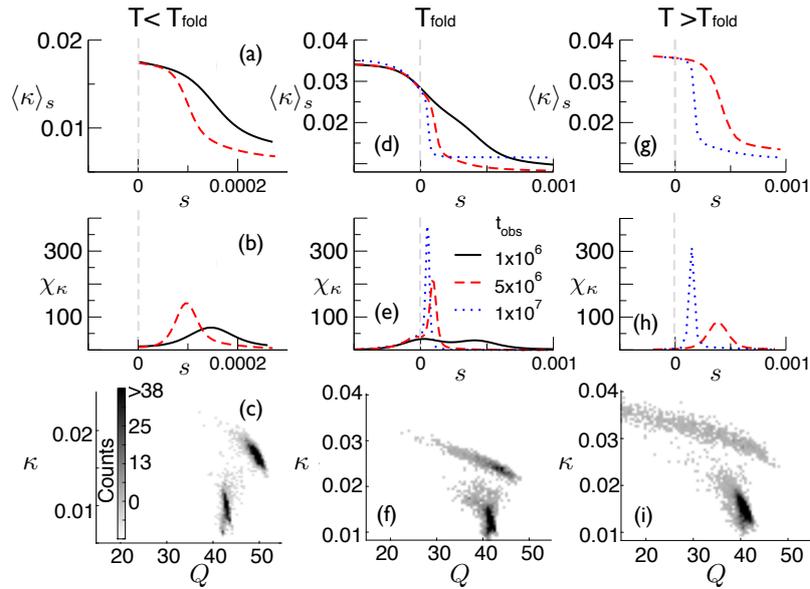


Figure 4.14: (a)-(c) Simulations for $T = 0.18 < T_{\text{fold}}$, with (a) showing the dependence of $\langle \kappa \rangle_s$ on s , (b) the susceptibility χ_κ with respect to s and (c) a density plot of the joint distribution of κ and integrated nativeness Q from a subsample of 5000 representative trajectories at a critical value of $s_c \approx 0.0001$.

(d)-(f) Simulations for $T = 0.19 \approx T_{\text{fold}}$, with (d) showing the dependence of $\langle \kappa \rangle_s$ on s , (e) the susceptibility χ_κ with respect to s and (f) a density plot of the joint distribution of κ and integrated nativeness Q from a subsample of 5000 representative trajectories at a critical value of $s_c \approx 0.00005$.

(g)-(i) Simulations for $T = 0.2 > T_{\text{fold}}$, with (g) showing the dependence of $\langle \kappa \rangle_s$ on s , (h) the susceptibility χ_κ with respect to s and (i) a density plot of the joint distribution of κ and integrated nativeness Q from a subsample of 5000 representative trajectories at a critical value of $s_c \approx 0.0002$.

The Colour-code according to the observation time is indicated by the legend.

the order parameter. This is very similar to what has been observed in the Gō model, provided that the 'less active' state corresponds to mainly native trajectories and the 'active' state corresponds to mainly non-native trajectories. For this a density plot of the native activity with respect to the integrated nativeness is provided in figure 4.14(f). Instead of observing a one to one correspondence between the average number of native contacts formed with respect to the activity (as for the Gō model), different activities can be observed for the same number of integrated native contacts, thus hinting at a trapping state. This behaviour will be looked at in more detail in a subsequent section, but first we want to comment on the behaviour observed for different temperatures as shown in figure 4.14(a)-(c) for $T = 0.180$ and in figure 4.14(g)-(i) for $T = 0.2$. For both temperatures a transition with

increasing sharpness as the observation time is increased is observed, further confirming the notion of a first-order transition in dynamic trajectory space. This is illustrated by figure 4.14(a) and (b) for the lower temperature and figure 4.14(g) and (h) for the higher temperature, again reweighting the histograms of trajectories with critical s . Furthermore, the dependence of the density distributions of the joint space of κ and Q is interesting as in all three cases the low activity state corresponds to trajectories displaying an average nativeness of $Q > 40$, yet by no means close to the expected number of native contacts where $Q \approx N_c = 57$. This is shown in figure 4.14(c), (f) and (i). Therefore, the obvious next step is to look at the nature of the observed dynamic phases with respect to their integrated nativeness Q in more detail.

4.5.3.3 HeGō Model: Classifying the Observed Dynamic Phases According to Nativeness

Looking at figure 4.14(f) in more detail, we show the joint density distribution at $s = 0$ and $s \approx s_c$ in figure 4.15(a). It can be observed that for $s = 0$ a linear cigar like dependence of Q and κ is observed, whereas for $s \approx s_c$ this is lost.

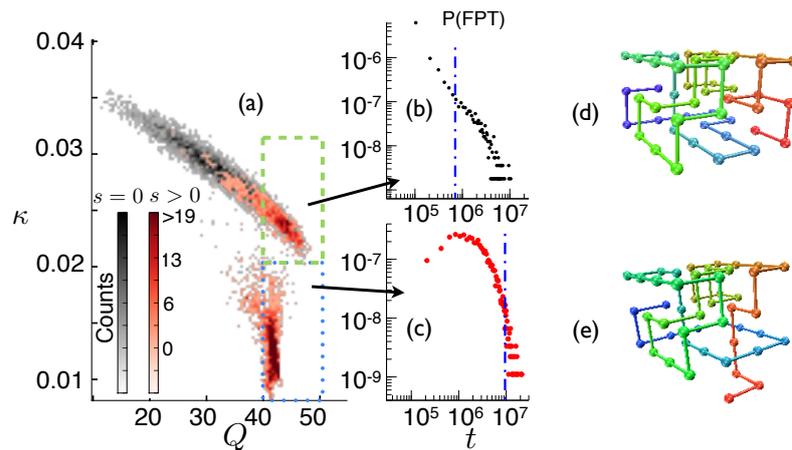


Figure 4.15: (a) Joint density distribution of κ and Q for $s = 0$ in black and $s \approx s_c$ in red. Distribution from a subsample of 5000 trajectories. (b) Shows the first passage time to the native state distribution for initial configuration drawn from the biased ensemble $\kappa < 0.003$ and $\kappa > 0.002$ and $Q > 40$, quenched to $s = 0$ (indicated by the blue rectangle). (c) Shows the first passage time distribution for initial configurations drawn from trajectories fulfilling conditions of $\kappa < 0.002$ and $Q > 40$ (indicated by the blue rectangle). (d) Native state and (e) representative trapped configuration. The colour code is again used for easy guidance along the chain.

Evidently, there is a highly populated non-native state in this dynamic ensemble, which will also have a contribution to the equilibrium

properties of the system. This arises solely from the introduction of a heterogeneous sequence and does not involve any non-native contact energies.

The next step will be to classify this set of configurations further, which was done by looking at the first passage time distribution to the folded state of representative conformations taken from trajectories found in the 'active' and 'less active' dynamic states respectively. The results from these are seen in figure 4.15(b) and (c). For this purpose configurations were chosen at random from a trajectory that would fulfil the respective constraints ($\kappa < 0.003$ and $\kappa > 0.002$ and $Q > 40$ for figure 4.15(b) and $\kappa < 0.002$ and $Q > 40$ for figure 4.15(c)). From the chosen configurations dynamics are run at $s = 0$ until the folded state was reached, for 50 representative structures with each 100 folding trajectories from which figure 4.15(b) and figure 4.15(c) were constructed. The respective dotted dashed lines in blue indicate the mean of each of the distributions. Further it should be noted, that configurations drawn from high activity trajectories follow a power law first passage time distribution, whereas those drawn from low activity configurations follow an exponential first passage time distribution. A representative initial configuration from the low activity trajectories is shown in figure 4.15(e).

In order to investigate this behaviour and classify a potential trapping state further, a root mean square distance (RMSD) analysis of a subset of representative trajectories was carried out. For this purpose, 100 trajectories from each of the two dynamic phases were chosen at random. To be precise, from simulations at $T = 0.19$ and $s = 0.00012$ and $t_{\text{obs}} = 5 \times 10^6$, trajectories for which $\kappa > 0.02$ and $Q > 40$ were used for highly native and more active trajectories and $\kappa < 0.02$ and $Q > 40$ for highly native and less active trajectories. Snapshots at intervals of 5000 MC steps were taken and written out to trajectory files. The RMSD for each of the trajectories was then computed by means of a least square fit of each structure along the trajectory in comparison to the native state. Thus large RMSD values mean dissimilar structures from the native state and an $\text{RMSD} = 0$ is that of the native structure. The analysis was done using the Gromacs tool `g_rms` [85]. The tool `g_rmsf` computes the root mean square (RMS) fluctuations for each residue and returns its fluctuations along the trajectory.

The results from the s-ensemble biased RMSD analysis are seen in figure 4.16. Here the probability of the RMSD from the 100 representative trajectories is shown in figure 4.16(a). The low activity trajectories (black) are highly peaked around an $\text{RMSD} = 1$ lattice unit and have practically no probability of being native. The higher activity trajectories have a sharp peak at the native state and the rest of the RMSD distribution is much broader (red). This shows that in higher activity trajectories, the native state itself is responsible for the high nativeness

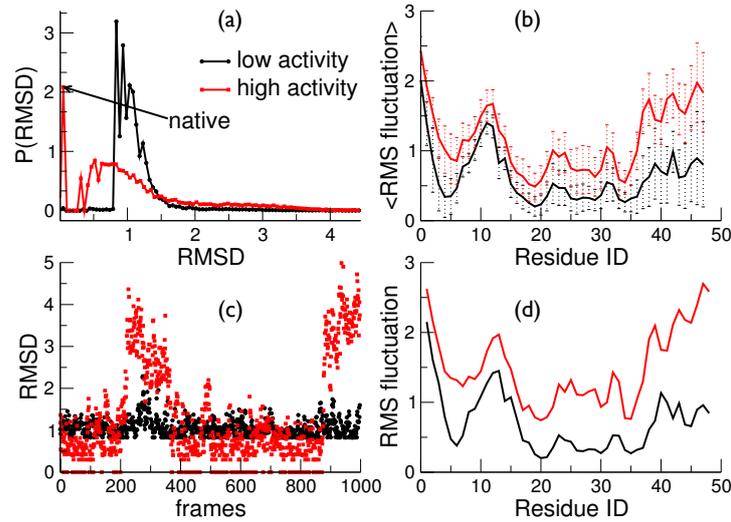


Figure 4.16: All data are from s -ensemble simulations at $s = 0.00012$, with snapshots taken every 5000 MC steps. (a) Shows the RMSD distribution from 50 independent trajectories, where high activity means that $\kappa > 0.02$ (red) and low activity $\kappa < 0.02$ (black). (c) Shows a representative trajectory from the distribution of the RMSD over the trajectory. (b) Shows the RMS fluctuations of each residue in the 48mer for both activity phases, errorbars correspond to one standard deviation. The corresponding single trajectory representation is found in (d).

along the entire trajectory, whereas the lower activity trajectories are trapped in a state, in misfolded configurations, which are close to being native, i.e. similar to the one seen in 4.15(e).

The average RMS fluctuations over these trajectories are shown in figure 4.16(b) for each amino acid (residue ID). As expected, the ends of the chain on average move more, but also as expected the more active trajectories (red) have on average a higher fluctuation per residue ID. The errorbars indicated are the standard deviations from the calculated averages of the 100 trajectories. In figure 4.16(c) a single trajectory with its 1000 frames for both a high activity and low activity trajectory are shown. As expected, the RMSD fluctuates more in the active trajectory, which is also confirmed by a single realisation of the RMS fluctuations in 4.16(d). This clearly supports the concept of a trapping state dominating inactive trajectories.

The next question is, to what extent such trapping states in the inactive trajectory influence non-biased equilibrium dynamics of the system. This was already addressed in figure 4.15, where it was shown that MFPT from inactive states are by an order of magnitude slower, than from the active counterparts with the same integrated nativeness of $Q \sim 40$.

To investigate the folding behaviour a bit further, the same RMSD anal-

ysis as carried out for the biased trajectories is also applied to folding trajectories. For this purpose, configurations along the biased trajectories are chosen at random as the initial starting position for the folding simulation. The following conditions were used: The chosen configuration must have a nativeness of $Q > 40$ for both activity phases and the activity must be below $\kappa < 0.02$ for the low activity trajectory, from which the configuration is picked. For both phases 50 initial configurations are chosen at random, from which equilibrium folding simulations are carried out. The folding trajectories are then analysed according to their [RMSD](#), the results for this can be found in figure 4.17.

The distribution of [RMSDs](#) along the 50 folding trajectories for both phases is shown in figure 4.17(a) with a logarithmic scale on the y-axis. Folding trajectories started from configurations along low activity trajectories, exhibit a double peak behaviour in the [RMSD](#), where one peak corresponds to an $\text{RMSD} = 1$ and the second broader peak to an $\text{RMSD} = 3$ lattice units. This suggests that at first trajectories continue to be trapped, then unfold until they finally refold. This picture is supported by the [RMSD](#) of a single folding trajectory as seen in figure 4.17(c). On the other hand, starting configurations taken from active trajectories, the [RMSD](#) is only peaked at a value of $\text{RMSD} = 1$ lattice units (red).

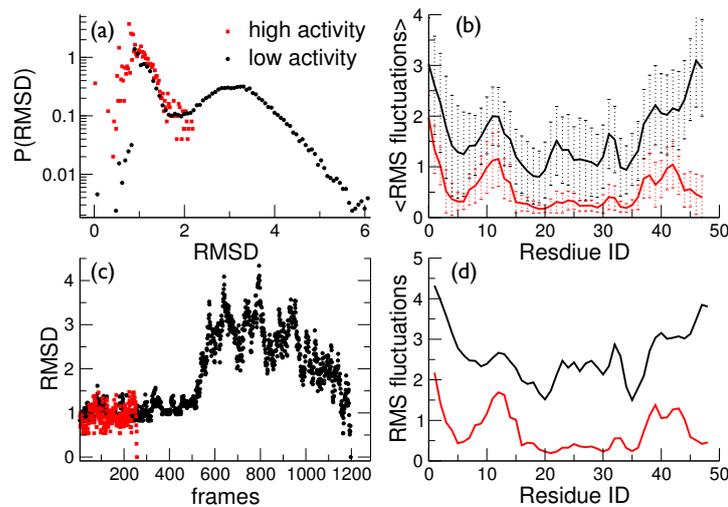


Figure 4.17: Folding trajectories from representative starting structure for $s = 0$ from both phases ($\kappa > 0.02$ and $\kappa < 0.02$). (a) Shows the distribution of [RMSD](#) from 50 folding trajectories with initial structures $Q > 40$, for low activity structures (black) and high activity structures (red). Representative folding trajectories for both starting structures are shown in (c). (b) Shows the average [RMS](#) fluctuations for each contact, with error bars representing the standard deviation. For both activity phases representative trajectories are shown in (d).

From figure 4.17(c) faster folding from high activity configurations can be observed. The folding event occurs in a much smaller number of frames, than for the low activity configurations. Here frames are snapshots of the protein configuration written out at a given interval along the folding trajectory. Frames for the low activity configurations were written out every 100 MC steps, whereas for the high activity configuration frames were written out every 50 MC steps supporting the fast folding behaviour from active configurations.

Also the fluctuations per residue were looked at along the folding trajectories. The average from 50 trajectories is seen in figure 4.17(b) with the error bars indicating standard deviations from the average. In figure 4.17(d), the RMS fluctuations for representative trajectories are shown. This clearly indicates that the fluctuation in the folding trajectories started from trapping configurations, are much larger than those started from 'on folding pathway' configurations.

4.5.3.4 HeGö Model: General Activity κ_g

One might now argue that this is not surprising as an intrinsic bias was introduced with the native activity as the observable, which is biased according to s . In the following, it can be shown that using the general activity as the dynamic observable to which s is the conjugate biasing field, the same inactive non-native state is found, as it just represents a trapping state for this type of lattice protein, with the given sequence. To investigate this, similar simulations were carried out as found in section 4.5.3.2. The temperature is again set to $T = 0.19$ and the activity measure now consists of all MC moves carried out. Results from this simulation can be found in figure 4.18.

Figure 4.18(a) displays the behaviour of κ_g with respect to s , from

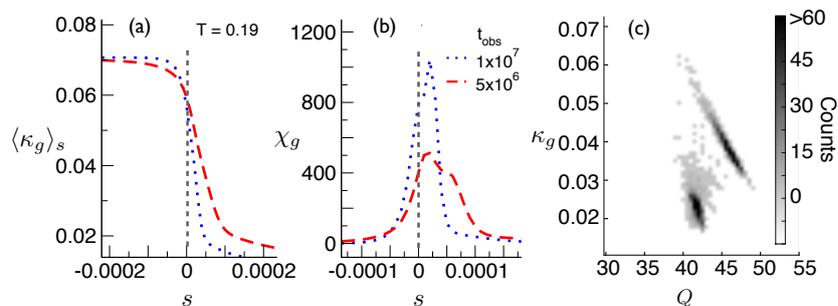


Figure 4.18: The behaviour of the general activity with respect to s is shown. (a) $\langle \kappa_g \rangle$ with respect to s (b) fluctuations in the activity with respect to s . A critical s can be established for both depicted observational times $t_{\text{obs}} = 5 \times 10^6$ and $t_{\text{obs}} = 1 \times 10^7$. A density of κ_g with respect to the integrated nativeness is found in (c).

the reweighted histograms. In figure 4.18(b) the susceptibility of κ_g is shown and (c) represents the joint density of κ_g and the integrated nativeness for $t_{\text{obs}} = 1 \times 10^7$. The critical value for s needed for a co-existence of the phases is much smaller in the general activity, suggest-

ing that the transition to the low activity trajectories is not far from the equilibrium trajectories. For trajectories of an observational time $t_{\text{obs}} = 1 \times 10^7$, $s_c \sim 2 \times 10^{-5}$. This means there is a clear influence of these rare event trajectories on the equilibrium behaviour.

4.5.3.5 Full model: s-ensemble biasing at T_{fold}

Having exhausted the analysis for the HeG \bar{o} model, the next step is to look at the dynamic behaviour of the Full interaction potential model given by equation (4.3). For this purpose again the same sequence at $T = 0.19$ was used, following the same approach as before, i.e. varying s according to the fixed temperature with two different simulation lengths and find a dynamic phase coexistence region, from which an estimate for the critical value s can be found. From this data it is then possible to construct histograms and reweight the data in order to extract the behaviour of the average native activity $\langle \kappa \rangle$ versus the biasing field s as seen in figure 4.19. In general the behaviour of the Full inter-

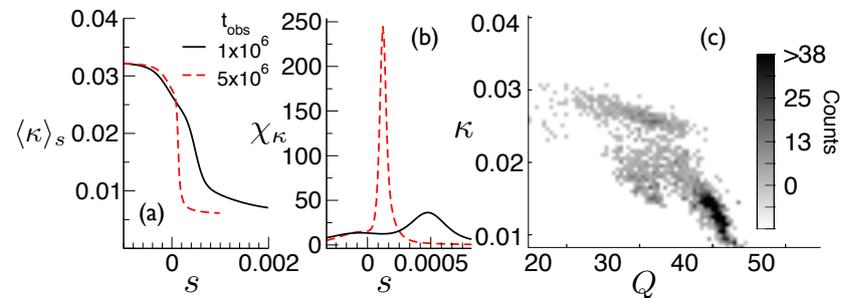


Figure 4.19: (a) Shows $\langle \kappa \rangle_s$ with respect to s , (b) the fluctuations (χ_κ) in κ with respect to s and (c) density plot of the joint distribution of Q and κ for $t_{\text{obs}} = 5 \times 10^5$. All data is taken from 20000 trajectories at $T = 0.19$

action potential is very similar to that of the HeG \bar{o} potential. Again the same trapping state is found and has similar effects on the overall dynamics of the system. This is seen from figure 4.19 (a)-(c), where the finite size scaling shows the transition from the active to the less active trapped trajectories at some critical value of s . However, the density plot (c) hints at a more complex behaviour of multiple low activity states with different average nativeness. This is not surprising, as more trapping states with more complex potentials would be expected. The actual dynamical interplay arising from the higher complexity would need to be studied in greater detail, but is beyond the scope of this thesis.

4.5.3.6 Dynamic Phase Diagram

Using all the data displayed in the previous sections collectively, a phase diagram in the parameter space of $\{s, T\}$ can be constructed. A

schematic representation of this is found in figure 4.20. In the phase

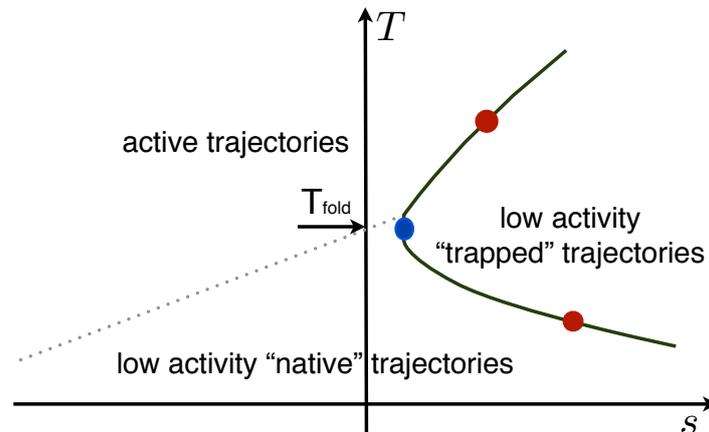


Figure 4.20: Schematic of the dynamic phase diagram in parameter space $\{s, T\}$. Dynamic first-order transition between inactive trapped trajectories and active trajectories is indicated by the green continuous line. Critical values for different T and s are indicated by circles. The grey dotted line is an extension of the thermodynamic crossover between active and inactive (now native) trajectories. The folding temperature T_{fold} is indicated. For the Full model, some more complexity could be added to the low activity phase but is omitted for clarity.

diagram a first-order dynamic transition induced by the biasing field s is shown (green continuous line). Red points are those observed from figure 4.14(b) and (h) at least qualitatively. The blue point represents the dynamic transition at T_{fold} between unfolded, active and inactive, trapped trajectories. At $s = 0$ the thermodynamic equilibrium at T_{fold} exhibits a first order like dynamic behaviour and can be extended in to the space of negative s values as indicated by the dotted grey line. This observed line is analogous to the behaviour observed for the Gō model, where the thermodynamic crossover falls together with the dynamic transition and extends from the coexistence at T_{fold} into values of negative biasing values. Establishing exact values for the transition is not straightforward for the HeGō and Full model, thus only the qualitative behaviour of the models is depicted here. It can be said, that the HeGō and Full model share most dynamic features, but the Full model, due to its broader interaction possibilities will exhibit even further complexity, which is neglected here, as further computer simulations would be necessary.

4.5.3.7 Extending Ideas to Different Sequences and Structures

This concept can of course be extended to a pool of infinite sequences and native states obtained from constructing maximally compact walkers of length 48. Most combinations of models will result in some

kind of dynamic phase-space behaviour ranging from multiple very complex transition to very clear single transitions to unique trapping states. The analysis of further sequences within the HeGō potential is discussed in the appendix B.2 in order to further the notion of the broad validity of such dynamic transitions.

The s-ensemble as a robustness test

The s-ensemble biasing allows to probe the robustness of a given sequence. Sequences with more homogeneous interaction energies are by construction much more robust in folding to the target structure. The s-ensemble biasing gives an idea of how close to the equilibrium potential misfolding, or trapping states are and thus how likely it is for a sequence to get trapped in those. If only small values of s are required to reach such dynamically arrested non-native states, such states will have a much stronger influence on the folding behaviour than sequences where much stronger biasing, i.e. larger values of s are required in order to access these dynamically arrested states.

Chain lengths

Another obvious step might be a scaling in the chain length. This was done only for the Gō structure as a sequence is readily available. Here the behaviour is very much the same as found for the 48mer long structure, meaning that at high temperatures the inactive native state is recovered with the dynamic biasing as well as the active dynamic non-native trajectories at low temperatures can be recovered as presented in 4.14. This will break down for shorter observation times, when the length of the sequence is shorted, as the finite spatial size starts influencing the system sooner.

4.5.4 *Summary*

In summary, s-ensemble biasing in lattice proteins uncovers rich dynamic behaviour in heterogeneous sequence models. By means of an RMSD analysis two dynamic phases could be classified: An active phase where most trajectories within correspond to equilibrium trajectories and a low activity phase where the trajectories contained are mainly populating trapping states of the lattice polymer. The dynamic phases are connected through a first-order (rounded, due to the finite system size) phase transition at a critical value of the biasing parameter s. Scanning the parameter space in $\{s, T\}$ made it possible to establish a clear phase behaviour for the main sequence of investigation. Furthermore, probing different sequences robustness of folding pathways can be tested and potential trapping states identified. This method could in the future be applied to more complex systems trying to probe the dynamic behaviour of real life proteins.

4.6 CONCLUSION

Lattice models have in the past given valuable clues about real protein dynamics and are able to mimic some real world protein properties well (such as the behaviour of nativeness with respect to temperature.) The type of lattice model studied here was developed amongst other reasons to probe the two state folding behaviour of real proteins. It was recently shown that many proteins exhibiting two state behaviour may in fact show hidden complexity arising from dynamics [101]. In this study it was shown that near native states give rise to metastable states, which would have previously been attributed to the native state itself. This ties well with what has been found here, as even very simple models, seemingly having a two state (unfolded-folded) folding behaviour with no off pathway diversions, do in fact exhibit more complex dynamic behaviour than expected. This arises from far from equilibrium fluctuations within the dynamic trajectory space.

This chapter will give an introduction to general techniques used for estimating the stationary behaviour of complex stochastic processes. These ideas are then seized to introduce a novel estimator which combines multi ensemble techniques with Markov state models, allowing very efficient estimation of equilibrium properties. All results presented in this chapter can be found in a publication in collaboration with F. Noé currently under review in *Phys. Rev. X* [102].

5.1 MOLECULAR DYNAMICS: SAMPLING THE EQUILIBRIUM

In section 2.2.5 of chapter 2, we assumed that if a system of interest becomes too complex and it is not possible to evaluate the partition function exactly, it is straight forward to use computational sampling instead in order to estimate stationary behaviour. In fact, we used this assumption on a "trivial" example of a single particle diffusing around a potential with two minima. We even applied a quite complex Markov model analysis to the dynamics, without really verifying that the underlying algorithm does indeed capture the stationary properties of the system. The aim of this section is to show that such a computational approach is valid and does capture the correct equilibrium behaviour of the system.

When modelling a stochastic process, commonly one considers either Brownian or Langevin dynamics. The latter is based on a stochastic differential equation called the Langevin equation, used previously in chapter 2, given by:

$$m \frac{d^2 \mathbf{x}}{dt^2} = -\mathbf{F}(\mathbf{x}) - \zeta m \frac{d\mathbf{x}}{dt} + \boldsymbol{\eta}(t) . \quad (5.1)$$

The parameters are as before: m is the mass of the system, \mathbf{F} the force due to the potential, ζ is a friction constant and $\boldsymbol{\eta}(t)$ a Gaussian random noise such that:

$$\langle \boldsymbol{\eta}(t) \rangle = 0 , \quad (5.2)$$

and the time correlation is given by:

$$\langle \boldsymbol{\eta}(0) \boldsymbol{\eta}(t) \rangle = 2\zeta k_B T_{\text{ref}} m \delta(t) . \quad (5.3)$$

The temperature T_{Ref} is the reference temperature of the system. The Langevin equation can then be used to sample a diffusive particle system, such as a single particle living in a potential. Of course, if the

potential is not very complex, i.e. $d < 5$, evaluating the stationary probability according to the potential with a grid based finite elements method, is exact and therefore a better approach. However, if the low dimensionality of the system is not given, we do have to turn to the computational approach and thus evaluate the Langevin equation according to the potential of the system.

Let us look at how we can get a computational representation of the Langevin equation. First of all a discrete timestep is defined given by Δt . Then an updating scheme used for updating the particles position and velocity according to the potential and the Langevin dynamics needs to be found. Here, ideas from deterministic dynamics are borrowed, were often a leap frog algorithm is used to computationally integrated Newton's equations of motion. As the name suggests, the leap frog algorithm updates positions and velocities with a half a time step lag between them. As an example, we can denote the particles position by \mathbf{x} in state space $\mathbf{x} \in \Omega \in \mathbb{R}^d$ and its corresponding velocity by $\dot{\mathbf{x}}$ ($\dot{\mathbf{x}} = \frac{d\mathbf{x}}{dt}$). The acceleration is provided by the potential according to:

$$\ddot{\mathbf{x}} = -\nabla U(\mathbf{x}) . \quad (5.4)$$

The set of update equations according to the leap frog algorithm are then given by:

$$\begin{aligned} \mathbf{x}_t &= \mathbf{x}_{t-\Delta t} + \dot{\mathbf{x}}_{t-\frac{1}{2}\Delta t} \Delta t \\ \dot{\mathbf{x}}_{t+\frac{1}{2}\Delta t} &= \dot{\mathbf{x}}_{t-\frac{1}{2}\Delta t} + \ddot{\mathbf{x}}_t \Delta t . \end{aligned} \quad (5.5)$$

More details and a derivation of this algorithm can be found in the textbook by Daan Frankel and Berendt Smit, "*Understanding Molecular Simulations*" [103]. This can now be applied to the stochastic Langevin equation resulting in a more complex updating scheme, as proposed and derived in [29]. Generally it is assumed that initial positions and velocities of all particles in the potential are known. If the velocities are not know, they can be drawn uniformly from a Boltzmann distribution at the reference temperature T_{ref} . Secondly, the force needs to be evaluated from the known potential such that:

$$\mathbf{F}(\mathbf{x}) = -\nabla U(\mathbf{x}) . \quad (5.6)$$

First, velocities are updated by half a time step ($\frac{1}{2}\Delta t$), then the positions are updated followed by another velocity update of half a time step, resulting in the set of updating equations given by equation (5.7).

$$\begin{aligned} \dot{\mathbf{x}}_{t+\frac{1}{2}\Delta t} &= \dot{\mathbf{x}}_t \exp\left(-\zeta \frac{\Delta t}{2} - \frac{\mathbf{F}}{m} \Delta t \left(\frac{1 - \exp(-\zeta \Delta t/2)}{\zeta \Delta t/2}\right)\right) + \\ &\quad + \frac{\sqrt{k_B T_{\text{ref}} \mathcal{N}}}{\sqrt{m}} \sqrt{1 - \exp(-\gamma \frac{\Delta t}{2})} \\ \mathbf{x}_{t+\Delta t} &= \mathbf{x}_t + \Delta t \dot{\mathbf{x}}_{t+\frac{1}{2}\Delta t} \end{aligned}$$

$$\begin{aligned} \dot{\mathbf{x}}_{t+\Delta t} = & \dot{\mathbf{x}}_{t+\frac{1}{2}\Delta t} \exp\left(-\zeta\frac{\Delta t}{2} - \frac{\mathbf{F}}{m}\Delta t\left(\frac{1-\exp(-\zeta\Delta t/2)}{\zeta\Delta t/2}\right)\right) + \\ & + \frac{\sqrt{k_B T_{\text{ref}}}\mathcal{N}}{\sqrt{m}}\sqrt{1-\exp\left(-\zeta\frac{\Delta t}{2}\right)}. \end{aligned} \quad (5.7)$$

The acceleration is given by $\ddot{\mathbf{x}}_t = \mathbf{F}(\mathbf{x}_t)$. The next question is, does this set of update equations given by equation (5.7), generate an equilibrium ensemble, i.e. do detailed balance (2.19) and ergodicity hold? In order to address this, we will revisit the double well potential from chapter 2, of the form:

$$U(x) = \begin{cases} -10 + 5(x+2)^2 & \text{if } x < -1 \\ -5x^2 & \text{if } x \geq -1 \text{ and } x < 0 \\ -7.5x^2 & \text{if } x \geq 0 \text{ and } x < 1 \\ -15 + 7.5(x-2)^2 & \text{if } x \geq 1 \end{cases}. \quad (5.8)$$

A pictorial representation of the potential is seen in figure 5.1(a). For the toy potential the position variable x is in arbitrary units of distance. In an actual molecular system typical distances would be measured in nm. The goal is to show that a computer simulation of a particle within this potential will actually sample the equilibrium and thus converge to the stationary distribution. For the potential of choice, the exact solution for the stationary probability is readily obtained through numerical integration, as it is given by:

$$\pi(x) = \frac{\exp(-\beta U(x))}{\int_{-\infty}^{+\infty} dx \exp(-\beta U(x))}. \quad (5.9)$$

Let us therefore compare the exact solution of the stationary probability obtained from the integration of the potential with the computer simulations. The results of this are shown in figure 5.1(b) and (c). Overlaid onto the exact stationary distributions are the histograms from the simulations using the Langevin leap frog integrator from equation (5.7) for 10000 simulation steps evolving a single particle of mass 1 with a time step of $dt = 0.01$ within the potential. The system is simulated at two different temperatures, a high temperature of $T = 15$ in reduced units and a low temperature of $T = 1$. The temperatures are given in reduced units of the Boltzmann constant, which was set to $k_B = 1$ for convenience. From figure 5.1(b), the high temperature simulation, it can be seen that the exact solution (dashed black curve) agrees well with the simulation (continuous red curve). This confirms what has previously been assumed to be true, that Langevin dynamics can capture the equilibrium dynamics of a stochastic system well and the set of leap frog equations for this process are valid.

In contrast however, looking at figure 5.1(c), the simulation only samples one of the wells and therefore the stationary distribution from the sampled data, does not agree with the theoretical solution. Why can the simulation not capture the equilibrium behaviour correctly at

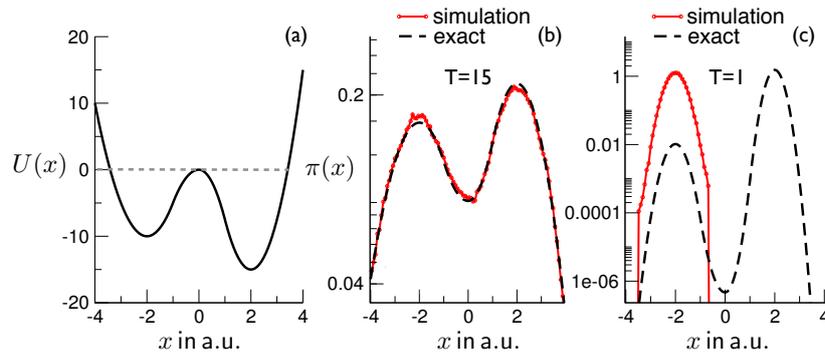


Figure 5.1: (a) Shows the potential and (b) the exact stationary distribution (black) and a histogram (red) from an MD simulation with 10000 steps at $T = 15$. (c) Shows the exact stationary (black) distribution at $T = 1$ and a simulation histogram (red).

low temperatures? The reason is not that the sampling algorithm – the Langevin leap frog algorithm – is not valid at low temperatures, but the fact that the system is metastable. Overcoming the high potential barrier of $5k_B T$ becomes very unlikely at low temperatures and is therefore very rarely observed. Waiting an infinite amount of time, would also recover the stationary distribution. Though, it is of obvious interest to accelerate this process. This leads to the formulation of some well established multi ensemble algorithms, which allow faster sampling of state space, while reproducing the equilibrium ensemble, still using Langvin dynamics. Two commonly used multi temperature algorithms will be introduced in the subsequent section 5.2. These will be referred to as parallel tempering (PT) and simulated tempering (ST).

5.2 ADVANCED SAMPLING TECHNIQUES

The toy example of a metastable potential in the previous section is an idealised version of a very common problem in more complex systems. An example of such a metastable system are spin glasses. These models represent frustrated magnets and exhibit many metastable states. A standard MC simulation can easily get "stuck" in these metastable states, with a simulation temperature below a certain freezing temperature. In order to overcome this, in the 1980's a multiple temperature ensemble algorithm was developed by Swendsen *et al.* [14]. This algorithm is called PT in this thesis. The algorithm has seen been used in many different systems, such as quantum frustrated spin systems [104, 105], quantum chromo dynamics (QCD) [106, 107, 108] and simulations of biomolecules [109, 110]. This chapter will focus on the latter of these. The naming of these multi temperature algorithms can be inconsistent in the literature and sometimes the term replica exchange is used in the context of molecular dynamics simulations, referring to the same algorithm as PT.

5.2.1 Parallel Tempering

The **PT** algorithm utilises the fact that at higher temperatures energetic barriers are more likely to be overcome. Let us consider a particle evolving according to the diffusive process \mathbf{X}_t in the potential (5.8) at M different temperatures in parallel. These are referred to as **replicas**. At higher temperatures the barriers will easily be crossed, whereas at lower temperatures this is not the case. After a certain number of discrete time evolution steps n , exchanges between neighbouring replicas are attempted. This means the temperatures of each replica is swapped. Generally this is done in a way that on each alternating swapping step even replicas and odd replicas are exchanged, as illustrated in figure 5.2.

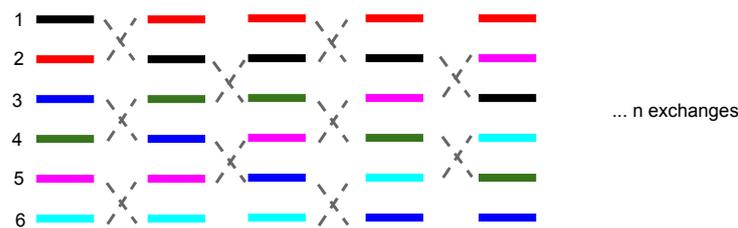


Figure 5.2: Illustration of replica swapping schedule

In order to retain detailed balance in the system, the swapping is done according to a simultaneous Metropolis Hastings acceptance criterion.

$$P_{\text{PT}}(\mathcal{H}_1(\mathbf{x})\beta_1 \rightarrow \mathcal{H}_2(\mathbf{x})\beta_2) = \min\{1, \exp(\Delta\beta\Delta\mathcal{H})\} , \quad (5.10)$$

where \mathcal{H} is, as usual, the Hamiltonian of the system. The temperature difference of the two states is given by their inverse temperatures $\Delta\beta = \frac{1}{k_B T_2} - \frac{1}{k_B T_1}$ and the difference in energy is evaluated according to the potential. If the set of temperatures is chosen in such a way that the energy distributions overlap with significant probability, acceptance is guaranteed. Generally, a 20% acceptance is desired [111] and an exponential spacing of temperatures is preferred, as the energy distributions are narrower at lower temperatures. The potential distribution overlap required to obtain the desired acceptance is around 60 – 70%. The goal is that each replica spends roughly an equal amount of time at each temperature, to ensure good mixing of the system. The idea of the **PT** concept can be illustrated using the potential from the previous section, defined by equation (5.8). It can now be shown that the dynamics using the **PT** simulation schedule will converge to the stationary distribution even at the low temperature in finite simulation time, through the additional information gained from the higher temperatures. The simulation results for this are seen in figure 5.3. The drawback of this method is that it creates a computational overhead, as the system needs to be simulated M times, i.e. the

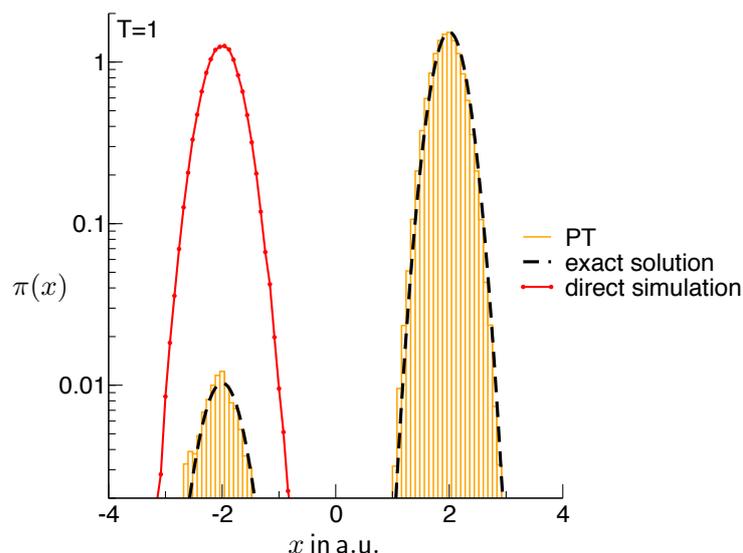


Figure 5.3: The recovery of the stationary distribution at $T = 1$ after $n_{\text{ex}} = 10000$ replica exchanges in orange, for an exponential spacing of 6 temperatures between $T = [1, \dots, 15]$ is shown. In red line originates from data using a single temperature ($T = 1$) simulation, with the particles initial position being in the left well.

number of temperatures desired. But in this way allows the system to escape from metastable states such as the left basin of the chosen example from figure 5.1. This becomes very important in higher dimensional systems, where much more of the conformational space can be explored in this way and the computational overhead is acceptable. Furthermore, this method can easily be extended to different ensembles other than the temperature ensemble. Many MD software packages, developed for simulating biomolecules incorporate this method as a REMD option. Amongst these are GROMACS and NAMD [85, 86].

5.2.2 Simulated Tempering

The second example of a multi temperature method is ST [112]. This method is closely related to PT, but only uses a single replica which is allowed to diffuse in temperature space. To be more precise, during a simulation after a certain number of simulation steps (MD or MC steps), an attempt is made to raise or lower the temperature. The temperature space is chosen in advance and generally makes use of the same spacing as one would choose for a PT simulation. This is a valid approach, as the same underlying energy distributions are not affected by the multi temperature algorithm used, but the underlying potential. A change in temperature is again accepted by a Metropolis criterion to ensure detailed balance:

$$P_{\text{ST}}(\mathcal{H}_1(\mathbf{x})\beta_1 \rightarrow \mathcal{H}_1(\mathbf{x})\beta_2) = \min\{1, \exp(-\mathcal{H}(\mathbf{x})\Delta\beta + \Delta g)\}, \quad (5.11)$$

where Δg is the difference in the weight factors. The weight factor g_m ensures equal probability of all temperatures and is given by:

$$g_m = -\ln \int_{\Omega} dx \exp(-\beta_m \mathcal{H}(\mathbf{x})) = -\ln Z_m, \quad (5.12)$$

with Z_m being the partition function of the system at temperature T_m , with $m \in M$. This evidently poses a problem, as Z_m generally is a property of the system the simulation aims to approximate in the first place. This resulted in an initial unpopularity of the algorithm over its multi replica counterpart, but now different approaches are available in order to approximate the weight factor. One option is to evaluate energy averages at each simulation temperature, prior to the simulation run and used average differences in the simulation temperatures as an approximation [112]. This method is particularly useful, if available computational power is limited. Another option is to use adaptive update techniques, with an initial guess which is improved, as the simulation progresses [113].

5.3 MULTIPLE STATE EQUILIBRIUM ESTIMATION

The work presented in the following sections is new and was developed and implemented by the author. A novel estimation method for stationary probabilities will be introduced. This combines the ideas of an extended temperature ensemble with those of a Markov model as already introduced in chapter 2 section 2.2.

5.3.1 Defining the Extended Ensemble

Generally, the problem in molecular systems is that the state space is too large to work directly on this space and therefore defining a set of discrete states is desirable, as long as these states still capture the essential dynamical and equilibrium properties of interest. Therefore using Markov models is a good approximation of a real system. The discretisation of the state space in itself can be challenging, but is assumed to be optimal for the purpose of this chapter. As already discussed in chapter 2, all dynamical and static information is retained in a transition matrix which can be reversibly estimated from a discrete Markov chain. In order to construct correct timescales for the Markov jump process, the Markov property must hold. However, for the purpose of the equilibrium estimates this is not necessary. Therefore a lagtime of the order of the simulation timestep can be used for the equilibrium estimator. In general, one can define a discrete jump process, for a discrete set of states $S = \{S_1, S_2, \dots, S_k\}$ at a given temperature T_m . From this a transition matrix \mathbf{T}_m , at temperature m , can be estimated reversibly. This transition matrix is spanned by the space of discrete states $K = |S|$ such that $\mathbf{T}_m \in \mathbb{R}^{K \times K}$. The subscript m indicates that this transition matrix is evaluated at a particular tem-

perature m . Generally this subscript is not written explicitly, but here it should help to remind the reader that we will consider an extended ensemble in temperature and state space.

Let us now move to a different transition matrix which spans the space of temperatures and states, such that $\tilde{\mathbf{T}} \in \mathbb{R}^{KM \times KM}$, with K different states and M different temperatures. The extended transition matrix represents the proposed estimator and is what is sought after to be approximated from simulation data. We will now discuss some properties of the extended transition matrix and explain, why it can be used as an estimator for stationary properties.

This extended matrix $\tilde{\mathbf{T}}$ will also hold all the properties of a transition matrix, discussed in section 2.2.3. This also includes the property, that the largest eigenvalue is $\lambda_1 = 1$ and the corresponding right eigenvector represents the stationary probability of the system.

$$\mathbf{1} = (\alpha_1 \boldsymbol{\pi}_1, \dots, \alpha_m \boldsymbol{\pi}_m) \tilde{\mathbf{T}}, \quad (5.13)$$

where $\boldsymbol{\pi}_m$ corresponds to the stationary distribution at temperature m and α_m is a normalisation constant, which ideally gives equal probability for each temperature. An equal probability for temperatures is desirable as it avoids numerical issues in the estimation of the stationary distribution at different temperatures. In the following the transition matrix and its respective entries will be quantified a bit further. The extended ensemble is defined as:

$$\tilde{\mathbf{T}} = \tilde{\mathbf{C}} - \mathbf{S}, \quad (5.14)$$

with $\tilde{\mathbf{C}}$ given by:

$$\tilde{\mathbf{C}} = c \begin{bmatrix} \mathbf{T}^{0,1} & \mathbf{U}^{1,1} & \dots & \mathbf{U}^{1,m-1} \\ \mathbf{D}^{2,1} & & \ddots & \vdots \\ \vdots & \ddots & & \mathbf{U}^{m-1,1} \\ \mathbf{D}^{m,m-1} & \dots & \mathbf{D}^{m,1} & \mathbf{T}^{0,m} \end{bmatrix}, \quad (5.15)$$

where c is normalisation constant and $\mathbf{T}^{0,m}$ are transition matrices at the observed temperature m with all diagonal entries set to zero. The diagonal matrices \mathbf{D} and \mathbf{U} describe temperature jump probabilities from each state S_k up and down in temperatures. The first superscript indicates the starting temperature and the second one the difference in the temperature jump. As in the case for **PT** and **ST**, for the detailed balance condition to hold, Metropolis Hastings probabilities are used to evaluate entries. Entries for the diagonal matrices of the up jumps are given by:

$$u_{m,i}^{\Delta m} = \frac{Z_m}{Z_{m+\Delta m}} \int_{\mathbf{x} \in S_i} d\mathbf{x} \mu_i(\mathbf{x}) \min\{1, \exp[-\mathbf{U}(\mathbf{x})\Delta\beta_{\Delta m}]\}, \quad (5.16)$$

with $\Delta\beta_{\Delta m} = \frac{1}{k_B T_{m+\Delta m}} - \frac{1}{k_B T_m}$ is the difference in inverse temperatures and $\mu_i(\mathbf{x}) = \pi_{m,i}^{-1} \exp(-\mathbf{U}(\mathbf{x})/k_B T_m)$ the local stationary density in S_i . With $T_m < T_{m+\Delta m}$ entries are given by the ratio of the

partition functions of temperature m and $m + \Delta m$: $u_{m,i} = \frac{Z_m}{Z_{m+\Delta m}}$. Analogously, we can define the entries for the down jumps:

$$d_{m,i}^{\Delta m} = \frac{Z_{m+\Delta m}}{Z_m} \int_{\mathbf{x} \in S_i} d\mathbf{x} \exp [U(\mathbf{x})\Delta\beta_{\Delta m}] . \quad (5.17)$$

In order to normalise $\tilde{\mathbf{T}}$ to a transition matrix the diagonal matrix \mathbf{S} is used, given by:

$$\mathbf{S} = \text{diag}\{1 - \text{const} \sum_i C_i\} . \quad (5.18)$$

Now a non-reversible transition matrix $\tilde{\mathbf{T}}$ is fully defined. As a reversible estimate is much more desired, this non-reversible matrix can be used as an input countmatrix for a reversible estimate according to [28].

5.3.2 The Choice of Scaling Variable

The definition of the extended ensemble from equation (5.15), is ideal for the following to reasons. Firstly, the down jump probabilities can be easily approximated from simulation data by summing over all simulation frames, as will be discussed in section 5.3.3. Secondly, the approximation of the scaling variable, which is the ratio of partition functions at the different temperatures, is also possible from a good set of simulation data. This in turn allows equal probabilities for all temperatures of the extended ensemble. Let us simplify the extended ensemble by only considering nearest neighbour temperatures. This has two advantages. Firstly the estimated transition matrix $\tilde{\mathbf{T}}$ will be sparsely populated as it now has a tri-diagonal block form, making the diagonalisation even for larger systems possible. Secondly, nearest neighbour temperatures will have the largest overlap between the energy distribution functions, meaning that their probability contribution to the up and down jumps will be largest and thus sufficient for a good estimate. In this way a small $\mathbb{R}^{4 \times 4}$ system can serve as an example in order to illustrate how the scaling variable can be optimally chosen. The transition matrix can be defined as a rate matrix

$$\tilde{\mathbf{T}} = \mathbf{K} - \mathbf{I} = \begin{pmatrix} 1 & T_2^0 & \gamma & 0 \\ T_3^0 & 1 & 0 & \gamma \\ \Xi d_1 & 0 & 1 & T_2^1 \\ 0 & \Xi d_2 & T_3^1 & 1 \end{pmatrix} - \begin{pmatrix} T_2^0 + \gamma & 0 & 0 & 0 \\ 0 & T_3^0 + \gamma & 0 & 0 \\ 0 & 0 & \Xi d_1 + T_2^1 & 0 \\ 0 & 0 & 0 & \Xi d_2 + T_3^1 \end{pmatrix} , \quad (5.19)$$

with $\Upsilon = \frac{Z_m}{Z_{m+1}}$ is the scaling variable and d_1 and d_2 are the estimated down jumps from equation (5.17), for each of the two discrete states. The goal is to chose Υ in such a way that for the overall transition matrix $\tilde{\mathbf{T}}$ the normalisation constants $\alpha_1 = \alpha_2$, i.e.:

$$\alpha_1 = \left| \frac{\pi_1 + \pi_2}{\sum_{i=1}^4 \pi_i} \right| \quad \text{and} \quad \alpha_2 = \left| \frac{\pi_3 + \pi_4}{\sum_{i=1}^4 \pi_i} \right|, \quad (5.20)$$

where $\boldsymbol{\pi}$ is the stationary vector of $\tilde{\mathbf{T}}$. This means there must be a stationary vector $\boldsymbol{\alpha}$ for a transition matrix projected only on the temperature states. This projected transition matrix can be defined as:

$$\tilde{\mathbf{T}}^* = \mathbf{P} \tilde{\mathbf{T}}^* \mathbf{P}^\top, \quad (5.21)$$

where \mathbf{P} is a projection matrix, such that:

$$\mathbf{P} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix}. \quad (5.22)$$

This leads to a projected transition matrix $\tilde{\mathbf{T}}^*$

$$\tilde{\mathbf{T}}^* = \frac{1}{2} \begin{pmatrix} 2(1-\Upsilon) & 2\Upsilon \\ \Xi(d_1 + d_2) & 2 - \Xi(d_1 + d_2) \end{pmatrix}, \quad (5.23)$$

with the condition $\alpha_1 = \alpha_2$ in hindsight and the matrix equation:

$$\tilde{\mathbf{T}}^{*\top} \boldsymbol{\alpha} = \mathbf{1} \boldsymbol{\alpha}. \quad (5.24)$$

The scaling parameter Υ can be found to be:

$$\Upsilon = \sqrt{\frac{\tilde{Z}}{2}}, \quad (5.25)$$

where $\tilde{Z} = \int_{\mathbf{x} \in S_1} \exp(\mathbf{U}(\mathbf{x})\Delta\beta) + \int_{\mathbf{x} \in S_2} \exp(\mathbf{U}(\mathbf{x})\Delta\beta)$. The variable Ξ , follows from this, as Ξ and Υ are related by their inverse. This information is easily obtained from the simulation data and can readily be extended to a large number of spatial as well as temperature states. To be precise for any number of k temperatures Υ is given by:

$$\Upsilon = \sqrt{\frac{\tilde{Z}}{M}}, \quad (5.26)$$

and \tilde{Z} is the integral over all discrete states K . In practice this integral is approximated by the sum over all available simulation data. If the sampling is good then the estimate for Υ becomes accurate. However, it is always useful to approximate Υ , by a sum over all observed data for each discrete state, as it will tend towards a more equal distribution of the temperatures and therefore make the extended matrix numerically more stable.

5.3.3 Estimating $\tilde{\mathbf{T}}$ from Simulation Data

Having established a way of optimally assigning the scaling parameter Υ , a recipe of how to use different types of simulation data for the estimation of $\tilde{\mathbf{T}}$ will be discussed. Initially, the state space needs to be discretised according to some relevant states. This will vary depending on the system of interest. The discretisation can for example be obtained using the Markov model building software EMMA [32]. Alternatively it can also be done by hand through some obvious coarse graining of an order parameter for example. The Markov model can be build from each recorded frame, meaning the lagtime τ for the transition matrix estimate is that of the recorded data. Again, only neighbouring temperature jumps will be taken into account in this case, even large systems can be considered. Any type of multi temperature simulation can be used to feed into the estimation of $\tilde{\mathbf{T}}$, such as a *PT*, *ST* or individual simulation runs at different temperatures. Ideally a random swapping (*RS*) protocol should be employed, in which a single replica diffuses through temperature space, but each temperature move is always accepted. The resulting out of equilibrium state after the temperature swap is compensated by discarding some initial simulation steps at the new temperature to ensure that a local equilibrium is reached. This is a feasible approach, as the transition matrix only requires local equilibrium and not a global equilibrium, generated by algorithms such as *PT* and *ST*. However, a good theory of how to determine whether a system is in local equilibrium or not, is not yet available, therefore such an *RS* schedule should be used with caution, if at all. From the discrete trajectories the transition counts at each temperature are obtained and give an estimate of the transition matrix \mathbf{T}_m at each temperature. Entries for \mathbf{U} are set to Υ as defined in equation (5.25). The down transitions are estimated according to the metropolis Hastings criterion and result in a discrete version of equation (5.17):

$$d_{m,i}^1 = \frac{1}{\Upsilon} \sum_{\mathbf{x} \in S_i} \exp[\mathbf{U}(\mathbf{x})\Delta\beta_1] . \quad (5.27)$$

Following these steps, allows for a non-reversible estimate of $\tilde{\mathbf{T}}$, which can then be used as an input for a reversible estimate according to [28]. The transition matrix then holds all the relevant information about the stationary probabilities of the system and in such represents the proposed estimator.

5.4 RESULTS

In order to illustrate the validity and superiority of the estimator over standard *PT* and *ST* techniques, three examples have been considered. First, the previous example of the double well potential which was used to illustrate the validity of the *PT* method is reused. Then, a d-

dimensional artificial folding potential is used and lastly REMD simulation data from MD simulations of alanine dipeptide generated with the MD toolkit OpenMM [114].

5.4.1 Double Well Potential

The simulation set up for the double well given by equation (5.8) is as follows. As previously described a single particle is integrated using a Langevin leapfrog algorithm from equation (5.7). Additionally, solvent particles are added to perturb the system and increase the number of degrees of freedom. The solvent particles diffuse in a harmonic potential of $U(x) = x^2$. Three different simulation protocols are used: a PT, ST and RS simulation. This defines the simulation techniques and the potential. Next, the temperature space needs to be defined. Six temperatures are exponentially spaced between $T = [1, \dots, 15]$. The potential is discretised into a set of two states $S = [S_1, S_2]$ each of the states corresponding to one of the wells, with the state boundary at $x = 0$, as seen in figure 5.4(a). The stationary probability distribution at $T = 1$ is shown in figure 5.4(b). This means the system is now defined in an extended space of $\tilde{\mathbf{T}} \in \mathbb{R}^{12 \times 12}$. Initially the system was perturbed with 2 solvent particles. The idea is now to quantify the error of the estimate of the stationary probability of being in any of the two given states. For illustration purposes, the relative error of the estimate of the stationary probability for state S_1 at temperature $T = 1$ is chosen. The exact probability of being in state S_1 , $P(S_1) = 0.008$, readily

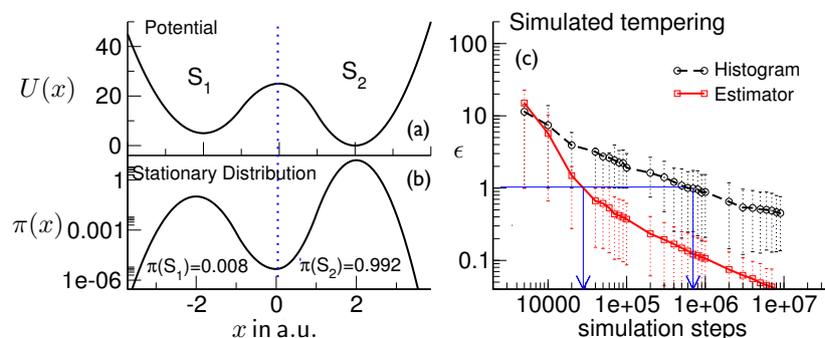


Figure 5.4: (a) Shows the potential and (b) the corresponding stationary distribution at $T = 1$. The two discrete states and their probabilities are also indicated. (c) Shows the the results obtained from PT simulation data in a log-log plot of the number of simulation steps taken versus the relative error of the estimate of the probability of being in S_1 . The black broken line is obtained through direct summing of the histogram of the ST data and the solid red line is the same data used in conjunction with the estimator. Standard errors are indicated.

evaluated by means of numerical integration. The relative error in the estimate is given by:

$$\epsilon = \left| \frac{\pi(S_1)_{\text{exact}} - \pi(S_1)_{\text{estimate}}}{\pi(S_1)_{\text{exact}}} \right|. \quad (5.28)$$

First, data obtained from a **ST** simulation is considered. As the potential can be evaluated exactly, the exact weightfactor g_m , presented in equation (5.11) was used, in order to get an optimal simulation schedule. The histogram obtained from the **ST** simulation is summed according to the two discrete states and a direct estimate of $P(S_1)$ is obtained. The same data is used to reversibly estimate entries for $\tilde{\mathbf{T}}$. The stationary vector $\boldsymbol{\pi}$ is calculated by solving the eigenvalue problem as given by equation (5.13 and the stationary probability at temperature $T = 1$ of being in state S_1 evaluated. In order to obtain good statistics averages are taken over 1000 trajectories, where each individual simulation is initiated in S_1 .

The convergence of the relative error with respect to the number of simulation steps taken is evaluated from histogramming the data and compared to an estimate of the stationary probability obtained from the estimator. This is depicted in figure 5.4 for the results of a simulated tempering simulation. Here a log-log plot of the relative error with respect to the number of simulation steps is shown. The stan-

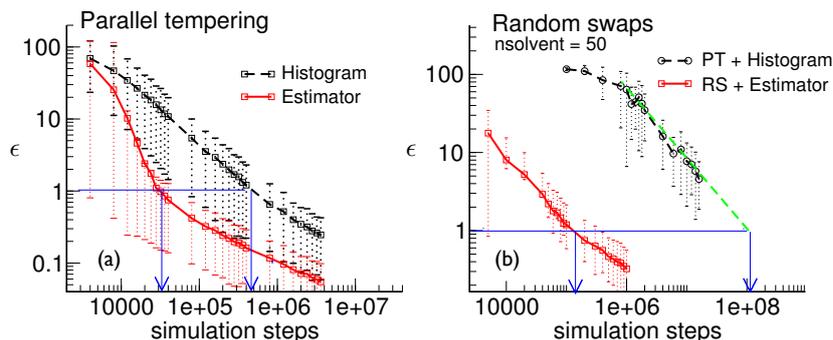


Figure 5.5: (a) log-log plot of relative error with respect to number of simulation steps from **PT** data. In red data obtained from the estimator is shown in black direct histogram counts from the **PT** simulation are shown. (b) Shows the results of the **RS** simulation on a log-log plot relative error with respect to simulation steps for a **PT** simulation using 16 replicas (black broken line) and 50 solvent molecules, in comparison with the **RS** simulation using only 6 replicas (red continuous line). The error level of $\epsilon = 1$ had to be extrapolated for the **PT** simulation and is shown in green. Blue arrows indicate the number of simulation steps needed for a given error level of e.g. $\epsilon = 1$

standard **ST** simulation, as well as the results from the estimator converge sublinearly following a powerlaw of $O(t^{-0.5})$. However, in order to reach the same relative error of $\epsilon = 1$ for the direct **ST** estimate in comparison with the estimator estimate, the simulation needs to be

run an order of magnitude longer. This is indicated by the blue arrow in figure 5.4(c).

The same set up as for the *ST* simulation was also used for the *PT* simulation. Figure 5.5(a), shows the results of the simulation using 1000 independent realisations. Again, an order of magnitude more simulation steps were required in order to achieve a relative error $\epsilon = 1$ from the *PT* histograms. The third simulation set up used was a *RS* simulation, which is compared to a parallel tempering simulation. The results of this are depicted figure 5.5(b).

For this simulation set up the number of solvent molecules perturbing the system is increased to 50. For this system, if the same number of replicas i.e. 6 was retained, the acceptance for exchanges would on average be: $P_{\text{accept}} = 0.038 \pm 0.08$ for a *PT* simulation, which is not sufficient for an effective simulation. Therefore the number of replicas needs to be increased for which the acceptance is $P_{\text{accept}} = 0.218 \pm 0.10$, which is much more desirable. As mentioned before, acceptance is governed by the overlap of the energy distribution functions. Through the additional degrees of freedom the overlap was reduced and therefore more replicas needed. Figure 5.6(b) shows the overlap of the energy distribution function for 6 exponentially spaced temperatures and figure 5.6(a) for 16 exponentially spaced temperatures, with 50 solvent molecules in the system. It is clearly seen that the overlap is much smaller for the 6 replicas.

Figure 5.5(b) compares the relative estimation error of the direct histogram from the *PT* simulation using 50 solvent molecules, to the simulation of only 6 temperature and using *RS* data in conjunction with the estimator. Here the convergence to a relative error of $\epsilon = 1$ needs to be extrapolated from the *PT* simulation, according to a power law distribution fit. This means that the estimator outperforms the *PT* simulation by three orders of magnitude.

It can be concluded, that for the simple 1-d potential well the estimator on average improves the estimate by up to three orders of magnitude using the same amount of simulation data and therefore shows to be a promising tool for large system which can be approximated well using Markov models.

5.4.2 Folding Potential

In the following, a more complicated, yet exact potential will be considered. The potential mimics a folding funnel, with an energetically favourable *native* state and an entropically favourable *unfolded* state. The potential is represented by a vector \mathbf{r} in d dimensions and is given by:

$$U(|\mathbf{r}|) = \begin{cases} -2.5|\mathbf{r}|^2 & \text{if } |\mathbf{r}| < 3 \\ 0.5|\mathbf{r}|^3 - |\mathbf{r}|^2 & \text{if } |\mathbf{r}| \geq 3 \end{cases} . \quad (5.29)$$

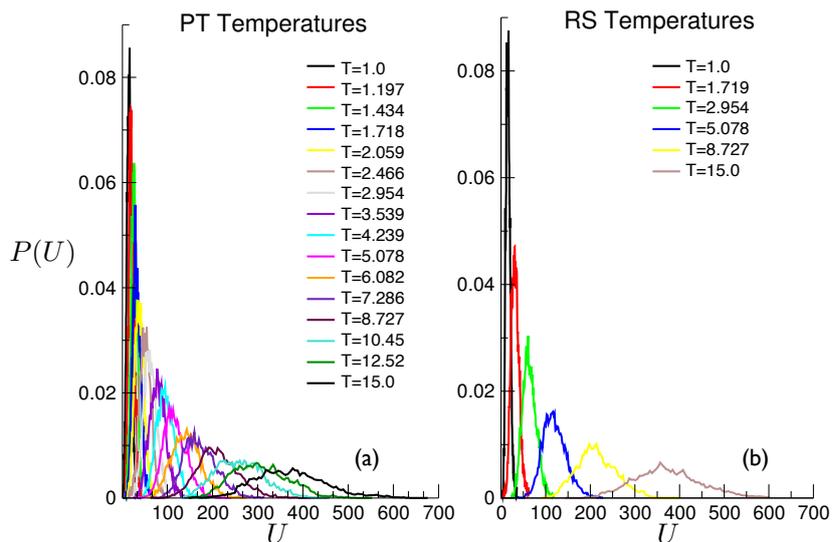


Figure 5.6: (a) The potential energy distribution for 16 replica temperatures between 1 and 15. (b) Potential energy distribution for 6 replica temperatures. In both cases the system is perturbed with 50 solvent molecules.

Again, the potential allows the exact evaluation of the stationary dis-

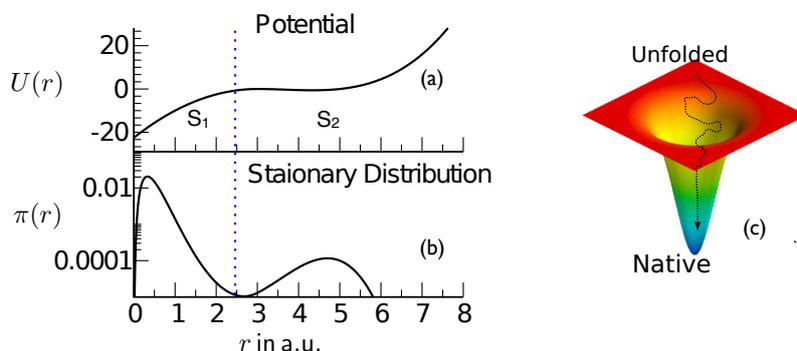


Figure 5.7: (a) Shows the potential and (b) the corresponding stationary distribution at $T_0 = 1.1$ and the two discretely chosen states. (c) Is a schematic of an idealised folding funnel.

tribution, which is done for a system where $\mathbf{r} \in \mathbb{R}^5$. However, the dimensionality is reduced to two distinct states for the Markov model. S_1 corresponds to the native state and S_2 corresponds to the unfolded state, with the state boundary at $|\mathbf{r}| = 2.5$, the radius of lowest probability density. The potential and the stationary distribution evaluated through numerical integration of the potential is seen in figure 5.7(a) and (b) respectively. The same simulation set up was conducted, as was done for the previous section 5.4.1, i.e a set of *ST* and *PT* simulations, as well as the *RS* schedule, with all simulations initialised in

state S_1 (the 'native' state). A single particle is integrated in the potential 5.29 using the same Langevin dynamics as previously described. A total of six temperatures were again spaced exponentially between $T = [1.1, \dots, 1.7]$. The temperature spacing is chosen in such a way, that at the low temperature the folded state (S_1) is the most populated state and at the highest temperature the unfolded entropic state (S_2) is the most populated one. The results of all simulations are presented in figure 5.8. Figure 5.8(a) shows the results from *ST* simulations

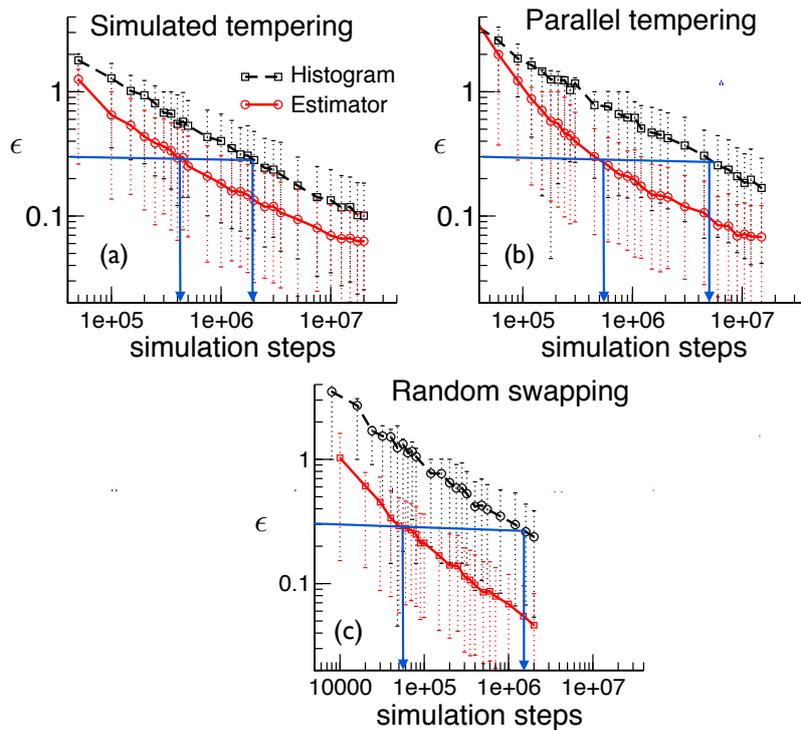


Figure 5.8: Shows the results of *PT*, *ST* and *RS* simulations in a log-log plot of the relative error of the stationary probability of being in the folded state ($P(S_1)$). (a) Shows plots the *ST* data obtained from the direct histogramming (black) as and the estimator used with the same data. (b) Shows the same plot but using *PT* data and panel (c) compares *PT* data with *RS* data used in conjunction with the proposed estimator. The number of simulation steps to reach a relative error of $\epsilon = 0.3$ is indicated in all three cases.

and shows that the estimator (red) outperforms the direct histogramming (black). The same applies to data obtained from *PT* and *RS* simulations, where results are depicted in 5.8(b) and (c) respectively. With the folding potential it can again be seen, that the relative error of the stationary estimate is up to over a magnitude better than with the direct histogram evaluation, meaning that fewer simulation steps are required for a good estimate of the stationary probabilities of the system.

5.4.3 Alanine Dipeptide

As a last example we consider a set of replica exchange molecular dynamics (Monte Carlo dynamics) (REMD) simulations of alanine dipeptide. Replica exchange here means nothing other than the previously defined PT, as we only consider temperature replicas. The name REMD is more broad and would allow for replicas of a different ensemble other than temperatures. Here we shall only consider temperature, yet use the name REMD, to be consistent with many MD software packages. The choice of this molecule is evident, as it has been studied before and thus validation of the simulation can be achieved through the comparison of existing studies, such as [115, 116]. Figure 5.9(a) shows the peptide with the two dihedral angles indicated in its α -helical configuration. This will also serve as the initial structure for all simulations (State IV in figure 5.9(b)).

The dominant conformations of this system are the different rotamers set by the dihedral angles ψ and ϕ . At a temperature of $T = 300$ K, these conformations are separated by relatively low energy barriers, such that a free energy surface at $T = 300$ K can be readily estimated, as seen in 5.9(b). Before the start of the simulation two choices are made: the MD software package used for the simulation and the desired force field.

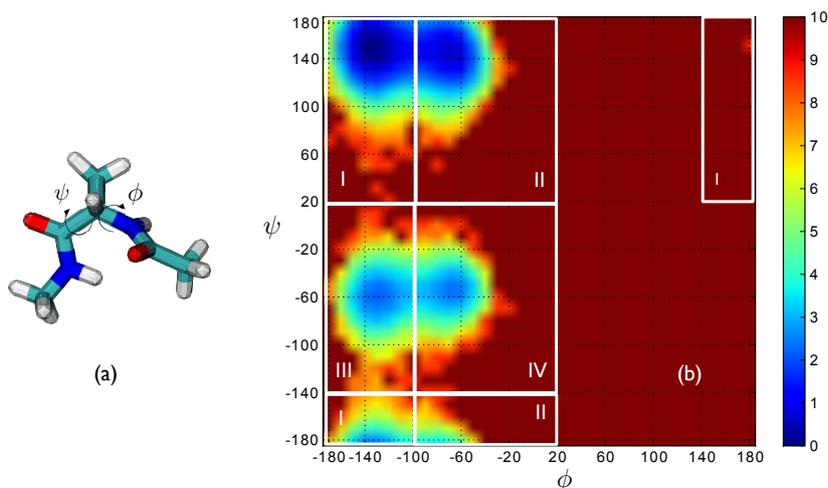


Figure 5.9: (a) Graphical representation of Alanine dipeptide: ACE-ALA-NME. (b) Free energy surface as projected onto the torsional angles ϕ and ψ for Alanine dipeptide, from 75 ns REMD simulation at 300 K. The energy scale is given in $k_B T$.

OpenMM, with its easy to use Python API was the chosen MD simulation toolkit, whose advantage is that it can be run on GPUs [114]. We used three GTX 580 GPUs, using the Amber99-sb forcefield [117]. In fact this force field was altered slightly by adding a set of circular van Mises potentials to the observed minima to further increase the

metastability of the system in order to demonstrate the clear superiority of the estimator. The van Mises potential is given by:

$$U(\Theta) = e \exp\left(\frac{q \cos(\Theta - \Theta_0)}{2\pi I_0}\right), \quad (5.30)$$

where Θ is the angle (i.e. ϕ or ψ to which the potential is added and I_0 is a zeroth order Bessel function). With the additional potential in place, the peptide was solvated in 645 water molecules and minimised with a steepest decent algorithm for 200 steps. Then an equilibration with modified a Berendsen thermostat and position restraints on the peptide was carried out for 100 ps.

This system was then cloned for the REMD setup, with all 32 replicas only differing in their reference temperature. The lowest temperature is set to $T_1 = 300$ K and the highest to $T_{32} = 600$ K. All other temperatures are spaced such that a roughly equal exchange probability for all replicas is given. This was achieved by short trial runs readjusting the temperature until a desired optimal spacing was reached. Each system (i.e. at each temperature) was equilibrated further at the given temperature for 500 ps, followed by a REMD production run.

For production runs the time step was chosen to be 2 fs, the saving interval for coordinates was 0.1 ps. Long range electrostatic interactions were evaluated with Particle Mesh Ewald and a bonded cut-off of 1 nm. All hydrogen bonds were constrained. The production run was carried out using a Langevin integrator with a collision rate of 1 ps^{-1} and replicas being exchanged every 1 ps. All parameter choices are commonly used and generally recommended in the user manual of OpenMM [114].

Accepted exchange attempts were observed to occur around 15 – 20% of the time. Exchanges were attempted every 1 ps. For the REMD simulations 15 independent realisations each of 5 ns were carried out. The minima for the van Mises potential were evaluated from a short unperturbed equilibration simulation, from which choices of Θ_0 for each of the torsional angles were made, as given by table 5.1:

Minimum	ϕ in $^\circ$	ψ in $^\circ$
I	-150	150
II	-70	135
III	-150	-65
IV	-70	-50

Table 5.1: Angle minima Θ_0 for von Mises potential

Other values were chosen such that: $e = -40$ KJ/mol the angular deviation $\sigma = 45^\circ$ and $q = \sigma^{-2}$. Using all the data available from all 15 realisations at 300K, the free energy surface in ϕ, ψ space was constructed, as seen in figure 5.9(b). From the simulations a set of uncorrelated data with a time series analysis was extracted which was then

used to generate the free energy surface. This was achieved by using bins of 10° width, with a cut-off of $10 k_B T$ for the free energy. From this surface 4 distinct states can be established as metastable states on this 2-dimensional projection. The states have been distinguished by eye according to the free energy surface and are indicated by white boxes and numbered accordingly and displayed in figure 5.9(b).

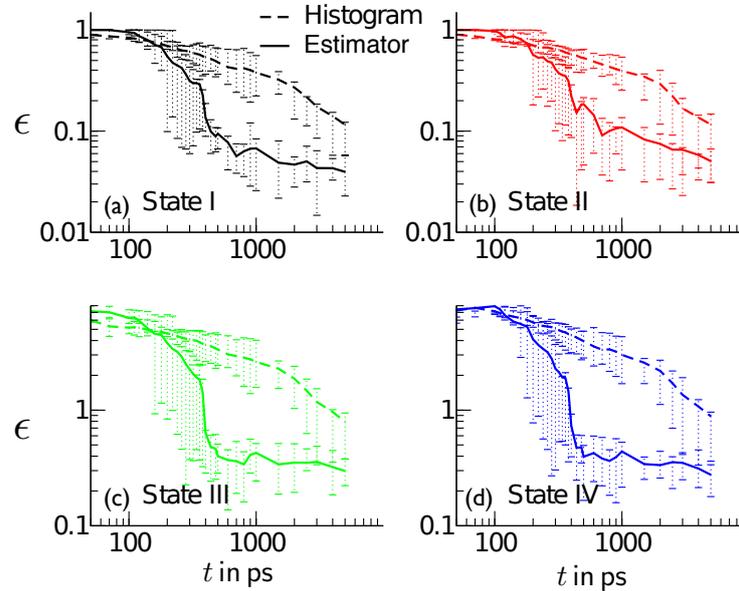


Figure 5.10: (a)-(d) Relative error convergence of the estimate using histogram counts (dashed lines) and estimator (continuous line). Reference value are all histogram counts from all realisations, with the initial 3 ns data discarded toward equilibration. Error bars are from averages of the 15 realisations.

From the generated data, the following data analysis was carried out. The generated trajectories were discretised according to the 4 defined states. As a reference (i.e. π_{exact} before, where the exact stationary probabilities were known.) histogram counts for each state from the last 2 ns of all realisations was used. This clearly is not the right value, but provides a "worst case estimate" scenario. As before for the double well potential and the folding potential the relative error over progressing simulation time is estimated, now for all 4 states. The estimate is again achieved in two ways. Firstly, by the proposed estimator, where the extended transition matrix (now in $\mathbb{R}^{132 \times 132}$) is evaluated and the obtained estimate is used in equation (5.28) as the estimated value. The results obtained this way for each state are shown in figure 5.10(a)-(d) respectively denoted by the continuous line, with the relative error plotted with respect to the simulation time in ps. Secondly, an estimate from direct counting of how many times each state is observed in each realisation, gives an estimate on the stationary probability, which is then used in equation (5.28) and plotted as the histogram estimate in figure 5.10(a)-(d), denoted by the dashed lines. All errors are obtained

as standard errors from the mean over the 15 different realisations. It can be seen again, that the novel estimator outperforms standard histogram methods, even in this much more complex molecular system. Lastly, we would like to look at how the actual stationary probability of each state slowly converges with respect to the simulation time spent. This is shown in figure 5.11(a)-(d) for all 4 states. All estimates using the novel estimator already reach a stationary plateau for the estimate with only small observed changes and only small deviations within the different realisations, whereas the direct counting estimate shows no plateaued convergence yet. This is true for all four states as seen in figure 5.11(a)-(d), where again continuous lines represent the novel estimator and dashed lines the histogram method. This results in a converged estimate for the novel estimator after around 500 ps already – indicated by the grey vertical line in 5.11 – which the direct count does not even reach after 5 ns simulation time.

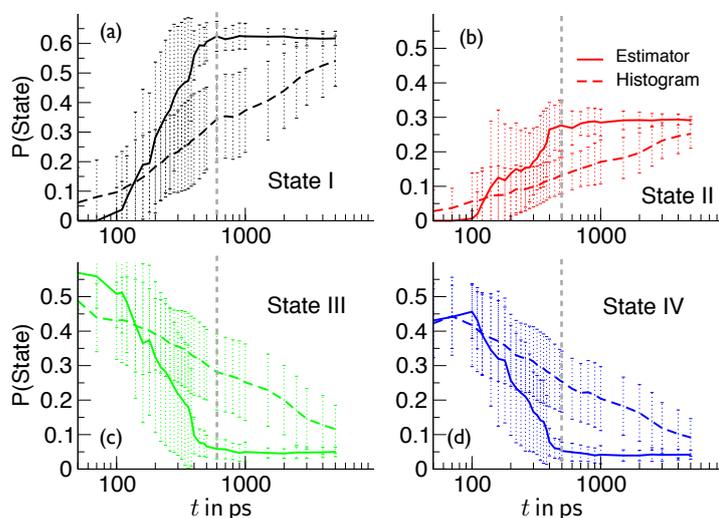


Figure 5.11: The stationary probability for each state (a)-(d) with respect to the simulation time is shown for both direct counting estimates (broken line) and estimator estimates (solid lines). Convergence using the estimator is significantly better than from histograms.

5.5 DISCUSSION AND CONCLUSION

The advantage of the novel estimation method is that it can be used with standard generalised ensemble simulations such as [PT](#) or [ST](#) simulations and in comparison to obtaining probabilities of the configurations of interest via direct counting or histogramming, it exhibits a much faster convergence of the error in the estimated stationary probabilities. Due to its much faster convergence, an order of magnitude fewer simulation steps are on average needed in order to achieve good estimates of the stationary behaviour. Thus it poses

an ideal tool for complex simulation where often the problem is that required simulation times for good stationary estimates are too long to be reached easily with current computational tools. This estimator aims to remedy this.

Here we can use the justification for the speed up, in a slightly adapted version as presented in the publication related to this estimator [102]: At least two aspects are responsible to achieve this significant speedup.

Firstly, the estimator achieves a reduction of the burn-in time after which the bias set by the starting configurations in a *PT* or *ST* simulation is lost. Traditional counting of configuration frequencies relies on the overall simulation (*PT* or *ST*) to be in global equilibrium, i.e. after the burn-in time, when sufficiently many replicas have moved to other configurations and mixed between temperatures, it can be expected that the temperature of interest samples from its stationary distribution in an unbiased fashion. As the novel estimate uses ideas from Markov modelling and constructs an estimate of the stationary probabilities from a transition matrix in the joint space of configurations and temperatures, global equilibrium is no longer a requirement and finding the system in a local equilibrium of the state and that temperature is sufficient. In *ST*, *PT* and *REMD* simulations, local equilibrium is maintained by the appropriate Metropolis trials conducted before any temperature transition. As a result of discarding the global equilibrium requirement, the burn-in phase is significantly reduced and in contrast to direct counting estimations, no simulation data must be discarded.

Secondly, the estimator achieves a reduction of the correlation time required to draw new uncorrelated samples at the temperature(s) of interest. It estimates transition probabilities with all available data, therefore maximising the use of all available data for analysis, and thus profits from the knowledge of how likely a transition between temperatures is, but not that the transition probability is large in order for the temperature step to be accepted. Thus, even in situations when the underlying *PT* or *ST* simulation does not mix well, the stationary probabilities can still be recovered.

In order to make the estimator widely available and readily usable, an implementation of it will be included in the next release of the Markov model estimation tool EMMA [32].

The fast and efficient estimation of stationary behaviour of complex systems will in general be useful for many different systems. The versatility of the estimator makes it a perfect tool for a broad range of applications and not only molecular systems which served here as a demonstration. One could think of other situations such as spin glasses or magnets, where a discretisation of the order parameter can be used to achieve very accurate estimates of stationary probabilities, from multi ensemble simulations. In particular bearing in mind, that this type of estimator is not restricted to the temperature ensemble.

SUMMARY AND OUTLOOK

This thesis looked at the equilibrium and dynamic behaviour of different stochastic models varying in their complexity. After a general introduction to the theoretical framework used to present this work, chapter 3 discusses the results of an exhaustive computational study of the dynamic properties of the 1-d Ising model. A rich dynamics is uncovered in the context of trajectories. The work was initially motivated in order to computationally verify the dynamic second-order phase transition predicted by Jack *et al.* [45]. The transition in the model was computationally observed by using a biasing ensemble (s -ensemble), in which the activity of trajectories was restricted. The verification of the second-order transition between a ferromagnetic and low activity phase and a paramagnetic high activity phase was achieved by means of finite time scaling. In the case of finite time scaling the observation time t_{obs} served as the scaling variable, instead of the generally used spatial system size. The activity of the observed trajectories served as the order parameter.

The theoretical model by Jack has limitations and did not reach beyond the prediction of the line of critical points as given by equation (3.28), i.e. the second-order phase transition. This is where the computational model was able to contribute further results to the system, by considering the influence of an external magnetic field on the 1-d Ising chain in trajectory space. In this way a phase space spanned by the parameters of $\{s, T, h\}$ could be described and resulted in the phase diagram given by 3.5. With the introduction of the external magnetic field, a surface of first-order transitions below the line of critical points for the case with no external magnetic bias was observed. The magnetic behaviour was investigated further by means of hysteresis loops, with a particular interest in how the biasing field s influences the size of these loops.

Lastly, the dynamics of the 1-d Ising model were changed which allowed to show that the dynamic phase behaviour is drastically altered by the choice of dynamics. This change meant moving from a single spin flip algorithm to a spin exchange algorithm (Kawasaki dynamics). Now the overall magnetisation of the system was conserved and for the purpose of the study set to zero. The line of critical points in the s -ensemble was no longer valid. Instead for no external bias (i.e. $s = 0$), a dynamic coexistence of an active paramagnetic and inactive ferromagnetic phase was observed. This observation was not unexpected, as similar behaviour had previously been observed in kinetically con-

strained models [5, 2].

The Ising model served as a starting point in the study of ensembles of trajectories. In this case a theoretical model was available and the phase behaviour was readily observed. Next, we turned to a more challenging system; that of a lattice protein. The results from the computational investigation of the s -ensemble biasing in the lattice protein model were presented in chapter 4. The choice of this model represented a good compromise between increased complexity and manageability. By this we mean that this model displays more complexity than many of the previously studied models in the s -ensemble framework. It is a simple lattice model, yet it displays features of real world proteins. In this context a general overview into protein dynamics and folding behaviour was given. The model was initially introduced in 2-d discussing some of its well known equilibrium properties. The 3-d case was considered next and served as the main model of study in the dynamical s -ensemble framework.

Three different variants of the lattice model were considered, distinguishable according to their different interaction potential. First a variety of the Gō model was studied, where all nearest neighbour interactions are uniform and defined according to the native state of the protein (here that of a 48mer). It was demonstrated that static properties, such as the nativeness $q(T)$ with respect to temperature are separated by a first-order like crossover. Furthermore, it was shown that this crossover is also present in the dynamics, meaning that dynamic trajectories are separated according to their activity. Low activity trajectories were mainly found to be native, whereas high activity trajectories were found to be unfolded. With the s -ensemble it was possible to reweight the trajectory dynamics, such that a first-order like crossover away from the folding temperature T_{fold} persisted at other temperatures as well. This was illustrated in a schematic phase diagram of figure 4.20.

As real life proteins build on heterogeneous amino acid sequences, the specific example of an amino acid sequence for the heterogeneous Gō potential (HeGō) and Full interaction potential was studied. In both cases the dynamic behaviour observed with the s -ensemble biasing was very similar, showing more complexity in the Full model than the HeGō model. In both models a first-order like dynamic transition between trapped and equilibrium trajectories was observed. The observed low activity trajectories were looked at in detail by means of MFPT and RMSD analysis, and it was possible to show that indeed the protein conformations belonging to low activity trajectories are trapping conformations, slowing down the overall dynamics of the system. The analysis also made use of finite time scaling with t_{obs} as the choice of the scaling variable. In conjunction with this investigation, different

dynamic observables and different sequences were considered, proposing the general argument that a s -ensemble analysis can serve as a good tool in order to probe the robustness of equilibrium dynamics.

What we have learnt so far from the s -ensemble is that rare fluctuations in the dynamics can influence the equilibrium behaviour of systems. These fluctuations are assumed to be the root cause of the slow relaxation times observed in glasses, but also seem to play a role in the lattice protein model. So far mainly very idealised systems have been considered, therefore the next step would be to try and employ the s -ensemble formalism to e.g. an all atom model of a protein using the full biased dynamics. This poses a computational challenge as significant timescales cannot be reached easily in an all atom protein simulation. Hopefully such challenges can be overcome, as more sophisticated algorithms are developed and faster computers become available.

So far the Ising model was the first toy system of a magnet whose non-equilibrium dynamic behaviour was considered. At this point it is unclear what implications such rare dynamic fluctuations have on real life magnets. Therefore, it would be desirable to find a way of experimentally probing observations due to the s -ensemble biasing. At this point in time, we are not aware of the availability of such methods, but this problem may be resolved in the future.

Chapter 5 looked at effective ways of estimating equilibrium probabilities in systems exhibiting metastability. Metastability often hinders the system to explore configurational space efficiently, thus making it very difficult to obtain estimates of the equilibrium behaviour, due to oversampling of the metastable states. It was shown that combining Markov model tools with multi-ensemble algorithms allows fast and efficient estimation of stationary properties. The main advantage of this method is its use of simulation data, meaning that on average an order of magnitude fewer simulation data is needed to get a sufficiently accurate estimate for the stationary probability of a given system. The direct consequence of this is that estimates of the stationary behaviour of even large and complex biological systems are no longer unattainable. Before, lacking sufficiently long simulations, only poor equilibrium estimates are often achievable. Furthermore, one should remember that these ideas are by no means restricted to biological systems, but the extended ensemble estimator can also be applied to metastable systems such as glasses, where the discretisation of the respective order parameter would be necessary.

In the near future this estimator can be applied to more complex simulations of protein systems generated with software tools such as GROMACS [85]. It is planned to incorporate this estimation method

into future releases of the Markov model analysis software EMMA [32]. It is not far fetched to assume that this analysis tool could also be useful for simulations in entirely different fields of physics, as long as systems exhibiting metastability are investigated. Another interesting aspect is that generally a good knowledge of the stationary behaviour allows a much more efficient and accurate estimation of dynamic properties of metastable systems. One could envision that such an estimation method forms part in a broader adaptive sampling schedule in order to optimally find states in configuration space that would require further sampling for the improvement in accuracy of timescales of processes. Therefore this would be a valid approach towards lowering the timescale gap which currently plagues simulations of many complex systems, by not sampling for longer but more cleverly.

X-ray diffraction

X-Ray diffraction is a method that had been known for a while but in 1962 it was possible to create a crystal of myoglobin and haemoglobin. From this crystal it was possible for the first time to resolve the 3 dimensional structure of a protein through the efforts of Kendrew and Perutz [118, 119]. The difficulty with this method is, that it requires the production of a perfect protein crystal. This can be a challenge, and may also lead to the question how representative a crystal structure is of the actual three dimensional structure of the protein. Due to the strong interactions through hydrogen bridges and disulphide bridges it has been experimentally been tested that the crystal structure and the structure in solution is very similar, which justifies the use of x-ray diffraction as a method to probe secondary structure of proteins.

Nuclear magnetic resonance – NMR

NMR was developed in the late 1930's by Rabi et al, who obtained a Nobel price for his work [120]. As the method was developed further it soon found application in determination of protein structures [121]. Most commonly used is proton NMR where the spin relaxation of a hydrogen isotope in a magnetic field is measured. Measurements give information about the distance between hydrogen atoms within the different amino acids, from which then the underlying structure can be deduced. The signals for α helices and β sheets is very distinctive, which makes the identification of such structures easy using this technique. Furthermore for NMR experiments the protein can be looked at in solution. This is an advantage for small proteins, but with larger ones getting distinctive signals from the proteins core can be difficult. Another limiting factor is the high protein concentration needed within the solution, which favours aggregation of proteins, which needs to be evaded.

Single molecule pulling

Single molecule pulling experiments were developed from the early 1990's onwards, where a molecule is set up in such a way that it has anchors at two different sides from which can be pulled [122, 123]. This is possible with optical tweezers. The setup is as follows: A Styrofoam ball is held in an optical trap, to which a DNA or RNA handle (a strand of around 1000 base pairs for example) is attached, via an antibody in-

teraction. Onto the DNA handle the molecule of interest is attached (e.g.) a protein. The setup is symmetric, as in on the other side of the protein a second DNA handle attaches which is then again held by a Styrofoam ball in an optical trap. Generally the second trap is moveable so that a force can be exerted onto the molecule of interest, allowing either the molecule to be studied under a constant force, or a constantly increasing force pulling the molecule apart. From these experiments it is possible to study unfolding and refolding behaviour [76, 77].

Förster resonance energy transfer – FRET

This technique works between two molecules with fluorescently active chromophores. One of them will be the donor the other one the acceptor. The donor is in an excited electronic state, transferring energy to the acceptor via non-radiative dipole-dipole interaction. From the efficiency of this energy transfer depends on the two acceptor and donor molecules, thus it is possible to obtain information about the location of molecules within the experimental setup. As the efficiency depends on the 6th power of the distance, even small changes within the distance will alter the FRET efficiency drastically, which makes this method so viable. [124]

dSTORM

dSTORM stands for direct stochastic optical reconstruction microscopy and is a technique of getting high resolution images of up to 20nm within living cells [74, 75], in comparison to conventional methods, which achieve a resolution of around 200nm. The technique relies on standard fluorescent dyes, such as ATTO, covering the entire spectral range. The high resolution is possible, as in this method the exact location from which the fluorescent signal is emitted is reconstructed. The signal obtained from a sensitive CCD camera, is taken as a position function and with the help of fitting Gaussians the most likely location of the light emitting molecule is determined. This allows to study molecular movement within cells.

APPENDIX: AMINO ACIDS AND ADDITIONAL RESULTS

B.1 AMINO ACID LETTER CODE

Table B.1 contains all amino acids with their 1 and 3 letter code. This is presented so that the single letter amino acid sequences used in the thesis can be associated to an amino acid sequence.

Alanine	Ala	A
Cysteine	Cys	C
Aspartic Acid	Asp	D
Glutamic Acid	Glu	E
Pheylalanine	Phe	F
Glycine	Gly	G
Histidine	His	H
Isoleucine	Ile	I
Lysine	Lys	K
Leucine	Ley	L
Methionine	Met	M
Asparagine	Asn	N
Proline	Pro	P
Glutamine	Glu	Q
Serine	Ser	S
Theronine	Thr	T
Valine	Val	V
Tryptophane	Trp	W
Tyrosine	Tyr	Y

Table B.1: Amino acid 3 letter and single letter code

B.2 ADDITIONAL RESULTS OF THE 3-d LATTICE PROTEIN IN THE S-ENSEMBLE

In this section we present some additional results of the dynamic analysis of the 3-d lattice protein model from chapter 4 in section 4.5 and following. The same native state as seen in figure 4.7(a) is used, but

this time the sequence is changed to the one seen in figure B.1. The sequence was obtained using the previously mentioned evolutionary dynamics algorithm [82]. The ground state energy of this sequence is: $E_g = -12.91\epsilon_0$ and its folding temperature is $T_{\text{fold}} \approx 0.155$. The ground state energy can be calculated straight away. The dependence of the nativeness $q(T)$ and T was established through a series of simulations similar to those for the original sequence, but all equilibrium results will be omitted here. Instead, we will concentrate on the s-ensemble dynamic analysis, as presented in figure B.1. In figure B.1(a)-

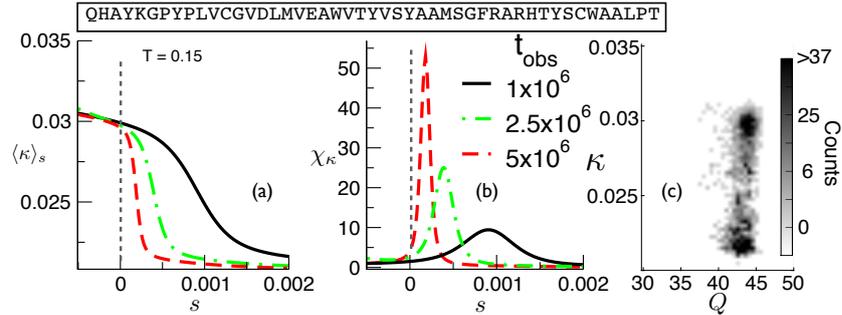


Figure B.1: The single letter code for the sequence used for the analysis is displayed on top of the figure. All data is taken from 2×10^4 trajectories unless otherwise stated. (a) Shows $\langle \kappa \rangle_s$ with respect to s , (b) the fluctuations of κ and (c) the density of the joint distribution of κ and Q from a subsample 5000 critical trajectories.

(c) the results of the investigation of the second protein sequence is given. Results are displayed from a series of simulations at temperature $T = 0.15$, slightly below T_{fold} for this sequence. Finite time scaling is achieved with observational times between $t_{\text{obs}} = 1 \times 10^6$ and $t_{\text{obs}} = 5 \times 10^6$, which are times near the mean first passage time of that sequence for the given temperature. The native activity κ serves as the main observable and order parameter. Figure B.1(a) is evaluated from reweighted coexistence histograms. The fluctuations in κ are evaluated and displayed in (b). The density plot of the joint distribution of κ and the integrated number of native contacts is found in figure B.1(c), from a trajectory of $s \approx s_c$. The low activity biased trajectories occupy states with very similar number of native contacts as in the equilibrium trajectories. The cigar like shape as observed in the dynamics of the original sequence as seen in figure 4.15, is not observed at all for this sequence. Instead, two highly populated regions in the density plot are observed. Both have different activity but similar intertged nativeness of $Q \sim 42$. In fact the two dynamic states are very closely connected and the lower activity state already sees a significant population in the equilibrium dynamics. The critical biasing parameter $s \sim 0.0001$ however is of the same order of magnitude as found for sequence presented in the main text. To what extend the MFPT is affected by these lower activity trajectories would need to be

probed further. It can be said however, that the equilibrium MFPT to the native state at $T = 0.15$ is around 4×10^6 MC steps, whereas for the sequence of the main text this is four times smaller, i.e. 1×10^6 MC steps on average.

It can be conjectured that the set of low activity trajectories sample a similar set of trapping configurations as those of the sequence of the main text. This however, would have to be investigated further. From figure B.1(c) it would seem that the nativeness of the trapping states are similar for both sequences in observation. The increased mean first passage time of the second sequence to the native state would mean that it is more likely for this sequence to go off route into the set of trapping states first, before being able to fold. Nonetheless, the sequence is by no means a slow folder, yet easily accessible trapping states slow the overall dynamics in equilibrium.

BIBLIOGRAPHY

- [1] R.L. Jack, J.P. Garrahan, and D. Chandler. Space-time thermodynamics and subsystem observables in a kinetically constrained model of glassy materials. *J. Chem. Phys.*, 125:184509, 2006.
- [2] 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(19):195702, May 2007.
- [3] 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.
- [4] 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):1309–1313, 2009.
- [5] Y.S. Elmatad, R.L. Jack, D. Chandler, and J.P. Garrahan. Finite-temperature critical point of a glass transition. *Proc. Natl. Acad. Sci.*, 107(29):12793–12798, 2010.
- [6] Francois Lelord. *Le voyage d’Hector*. Éditions Odile Jacob, 2004.
- [7] Johannes Voit. *The Statistical Mechanics of Financial Markets*. Springer, 3 edition, December 2005.
- [8] Ralf Blossey, editor. *Computational Biology: A Statistical Mechanics Perspective*. Springer & Hall, 2006.
- [9] N. Bose. Plancks Gesetz und Lichtquantenhypothese. *Z. Phys.*, 26(1):178–181, 1924.
- [10] P.A.M. Dirac. On the theory of quantum mechanics. *Proceedings of the Royal Society of London. Series A*, 112(762):661–677, 1926.
- [11] E. Fermi. Zur Quantelung des idealen einatomigen Gases. *Z. Phys.*, 36(11-12):902–912, 1926.
- [12] R.M. Gray. *Probability, Random Processes, and Ergodic Properties*. Springer, 2nd edition, 2009.
- [13] D.P. Landau and K. Binder. *A Guide to Monte Carlo Simulations in Statistical Physics*. Cambridge University Press, Cambridge, UK, 2005.
- [14] R.H. Swendsen and J.-S. Wang. Replica Monte Carlo simulation of spin-glasses. *Phys. Rev. Lett.*, 57:2607–2609, Nov 1986.

- [15] D.J. Earl and M.W. Deem. Parallel tempering: Theory, applications, and new perspectives. *Phys. Chem. Chem. Phys.*, 7:3910–3916, 2005.
- [16] A. Laio and F.L. Gervasio. Metadynamics: a method to simulate rare events and reconstruct the free energy in biophysics, chemistry and material science. *Rep. Prog. Phys.*, 71(12):126601, 2008.
- [17] D. Granata, C. Camilloni, M. Vendruscolo, and A. Laio. Characterization of the free-energy landscapes of proteins by NMR-guided metadynamics. *Proc. Natl. Acad. Sci.*, 110(17):6817–6822, 2013.
- [18] H. Xu, Y.N. Osetsky, and R.E. Stoller. Self-evolving atomistic kinetic Monte Carlo: Fundamentals and applications. *J. Phys.-Condens. Mat.*, 24(37):375402, 2012.
- [19] P.G. Bolhuis, D. Chandler, C. Dellago, and P.L. Geissler. Transition path sampling: Throwing ropes over rough mountain passes in the dark. *Ann. Rev. Phys. Chem.*, 53:291, 2002.
- [20] C. Dellago, P.G. Bolhuis, and P.L. Geissler. Transition path sampling. *Adv. Chem. Phys.*, 123, 2002.
- [21] J.P. Garrahan and I. Lesanovsky. Thermodynamics of quantum jump trajectories. *Phys. Rev. Lett.*, 104(16):160601, Apr 2010.
- [22] E.S. Loscar, A.S.J.S. Mey, and J.P. Garrahan. Thermodynamics of trajectories of the one-dimensional Ising model. *J. Stat. Mech.-Theory E.*, 2011(12):P12011, 2011.
- [23] S. Ross. *A First Course in Probability*. Prentice Hall press., 8th edition, 2009.
- [24] J.R. Norris. *Markov Chains*. Cambridge University Press, Cambridge, UK, 1998.
- [25] A.M. Glazer and J.S. Wark. *Statistical Mechanics: A Survival Guide*. Oxford Univ Press, Oxford, 2001.
- [26] W. Huisinga. *Metastability of Markovian Systems*. PhD thesis, Free University Berlin, 2001.
- [27] D. Boffi. Finite element approximation of eigenvalue problems. *Acta Numerica*, pages 1–120, 2010.
- [28] J.-H. Prinz et al. Markov models of molecular kinetics: Generation and validation. *J. Chem. Phys.*, 134(17):174105, 2011.
- [29] W.F. van Gunsteren and H.J.C. Berendsen. A leap-frog algorithm for stochastic dynamics. *Mol. Simulat.*, 1(3):173–185, 1988.

- [30] T. Kanungo, D.M. Mount, N.S. Netanyahu, C.D. Piatko, R. Silverman, and A.Y. Wu. An efficient k-means clustering algorithm: Analysis and implementation. . *IEEE Trans. Pattern Analysis and Machine Intelligence.*, 24:881, 2002.
- [31] P. Deuffhard and M. Weber. Robust Perron cluster analysis in conformation dynamics. *Linear Algebra Appl.*, 398(o):161 – 184, 2005.
- [32] M. Senne, B. Trendelkamp-Schroer, A.S.J.S. Mey, C. Schütte, and F. Noé. EMMA: A software package for Markov Model building and analysis. *J. Chem. Theo. Comput.*, 8(7):2223–2238, 2012.
- [33] F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T.R. Weigl. Constructing the full ensemble of folding pathways from short off-equilibrium simulations. *Proc. Natl. Acad. Sci.*, 106:19011–19016, 2009.
- [34] H.B. Callen. *Thermodynamics*. Wiley and Sons, New York, 1960.
- [35] K. Binder. Theory of first-order phase transitions. *Rep. Prog. Phys.*, 50:783–859, 1987.
- [36] L.D. Landau and E.M. Lifshitz. *Statistical Physics*. Oxford: Pergamon Press, 1958.
- [37] D.J. Amit and V. Martin-Mayor. *Field Theory; The Renormalization Group and Critical Phenomena*. World Scientific Publishing, 3 edition, 2005.
- [38] M.D. Donsker and S.R.S. Varadhan. Asymptotic evaluation of certain Markov process expectations for large time. II. *Comm. Pure Appl. Math.*, 28:279–301, 1975.
- [39] M.D. Donsker and S.R.S. Varadhan. Asymptotic evaluation of certain Markov process expectations for large time. III. *Comm. Pure Appl. Math.*, 29:389–461, 1976.
- [40] M.D. Donsker and S.R.S. Varadhan. Asymptotic evaluation of certain Markov process expectations for large time. IV. *Comm. Pure Appl. Math.*, 36:183–212, 1983.
- [41] H. Cramér. Sur un nouveau théorème limite dans la théorie des probabilités. *Colloque consacré a la théorie des probabilités*, 3:2–23, 1938.
- [42] R. S. Ellis. The theory of large deviations: From Boltzmann’s 1877 calculation to equilibrium macrostates in 2D turbulence. *Physica D*, 133:106–136, 1999.
- [43] H. Touchette. The large deviation approach to statistical mechanics. *Phys. Rep.*, 478(1-3):1 – 69, 2009.

- [44] R.L. Jack and J.P. Garrahan. Metastable states and space-time phase transitions in a spin-glass model. *Phys. Rev. E*, 81(1):011111, Jan 2010.
- [45] R.L. Jack and P. Sollich. Large deviations and ensembles of trajectories in stochastic models. *Prog. Theor. Phys. Supp.*, 184:304–317, 2010.
- [46] N. Goldenfeld. *Lectures on Phase Transitions and the Renormalization Group*. Addison-Wesley., 1993.
- [47] M.E.J. Newman and G.T. Barkema. *Monte Carlo Methods in Statistical Physics*. Oxford Univ Press, 1999.
- [48] D. Chandler. *Introduction to Modern Statistical Mechanics*. Oxford University Press, USA, 1987.
- [49] E. Ising. Beitrag zur Theorie des Ferromagnetismus. *Z. Phys.*, 31:253–258, 1925.
- [50] L. Onsager. Crystal statistics. I. A two-dimensional model with an order-disorder transition. *Phys. Rev.*, 65:117–149, Feb 1944.
- [51] I. Herbut. *A Modern Approach to Critical Phenomena*. Cambridge University Press, Cambridge, UK, 2007.
- [52] N.D. Mermin and H. Wagner. Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic Heisenberg models. *Phys. Rev. Lett.*, 17:1133–1136, Nov 1966.
- [53] N. Metropolis and S. Ulam. The Monte Carlo method. *J. Am. Stat. Assoc.*, 44:335–341, 1949.
- [54] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21(6):1087–1092, 1953.
- [55] K. Kawasaki. Diffusion constants near the critical point for time-dependent ising model. *Phys. Rev.*, 145:224, 1966.
- [56] H. Rieger and N. Kawashima. Application of a continuous time cluster algorithm to the two-dimensional random quantum Ising ferromagnet. *Eur. Phys. J. B*, 9:233–236, 1999.
- [57] M.P. Nightingale and H.W.J. Blöte. Transfer-matrix Monte Carlo estimates of critical points in the simple-cubic Ising, planar, and Heisenberg models. *Phys. Rev. B*, 54:1001–1008, Jul 1996.
- [58] J.G. Amar, F.E. Sullivan, and R.D. Mountain. Monte Carlo study of growth in the two-dimensional spin-exchange kinetic Ising models. *Phys. Rev. B*, 37:196–208, Jan 1988.
- [59] B.D. Cullity and C.D. Graham. *Introduction to Magnetic Materials*. Wiley, 2009.

- [60] S. Sachdev. *Quantum Phase Transitions*. Cambridge University Press, Cambridge, UK, 1999.
- [61] R.J. Glauber. Time-dependent statistics of the Ising Model. *J. Math. Phys.*, 4:294, 1963.
- [62] D.C. Jiles and D.L. Atherton. Theory of ferromagnetic hysteresis. *J. Magn. Mater.*, 61(1):48 – 60, 1986.
- [63] T. Bodineau, B. Derrida, V. Lecomte, and F. van Wijland. Long range correlations and phase transitions in non-equilibrium diffusive systems. *J. Stat. Phys.*, 133:1013–1031, 2008.
- [64] A.S.J.S. Mey, P.L. Geissler, and J.P. Garrahan. Rare-event trajectory ensemble analysis reveals metastable dynamical phases in lattice proteins. *arXiv:1305.5748*, 2013.
- [65] A. Sanger. The arrangement of amino acids in proteins. *Adv. Protein Chem.*, 7:1–67, 1952.
- [66] L. Pauling and R.B. Corey. Atomic coordinates and structure factors for two helical configurations of polypeptide chains. *Proc. Natl. Acad. Sci.*, 37:235, 1951.
- [67] C.B. Anfinsen. Principles that govern the folding of protein chains. *Science*, 181:223–230, 1973.
- [68] F.C. Bernstein, T.F. Koetzle, G.J.B. Williams, E.F. Meyer Jr., M.D. Brice, J.R. Rodgers, O. Kennard, T. Shimanouchi, and M. Tasumi. The protein data bank: A computer-based archival file for macromolecular structures. *J. Mol. Biol.*, 112(3):535 – 542, 1977.
- [69] J.M. Berg, J.L. Tymoczko, and L. Stryer. *Biochemistry*. Freeman, W. H., 2011.
- [70] A. Miranker, S.E. Radford, M. Karplus, and C.M. Dobson. Demonstration by NMR of folding domains in lysozyme. *Nature*, 349:633–636, 1991.
- [71] G.N. Ramachandran, C. Ramakrishnan, and V. Sasisharan. Stereochemistry of polypeptide chain configurations. *J. Mol. Biol.*, 7:95, 1963.
- [72] J.K. Meyers and T.G. Oas. Mechanisms of fast protein folding. *Annu. Rev. Biochem.*, 71:783–815, 2002.
- [73] A.M. Leske. *Protein Science*. Oxford University Press, 2004.
- [74] M. Heilemann et al. Subdiffraction-resolution fluorescence imaging with conventional fluorescent probes. *Angew. Chem. Int. Ed.*, 47:6172, 2008.

- [75] S. van de Linde, U. Endesfelder, A. Mukherjee, M. Schuttpelz, G. Wiebusch, S. Wolter, M. Heilemann, and M. Sauer. Multicolor photoswitching microscopy for subdiffraction-resolution fluorescence imaging. *Photochem. Photobiol. Sci.*, 8:465–469, 2009.
- [76] M.S.Z. Kellermayer, S.B. Smith, et al. Folding-unfolding transitions in single titin molecules characterized with laser tweezers. *Science*, 276(5315):1112–1116, 1997.
- [77] O.K. Dudko, G. Hummer, and A. Szabo. Intrinsic rates and activation free energies from single-molecule pulling experiments. *Phys. Rev. Lett.*, 96:108101, Mar 2006.
- [78] K.F. Lau and K.A. Dill. A lattice statistical mechanics model of the conformational and sequence spaces of proteins. *Macromolecules*, 22(10):3986–3997, 1989.
- [79] H. Abe and N. Gō. Noninteracting local-structure model of folding and unfolding transition in globular proteins. II. Application to two-dimensional lattice proteins. *Biopolymers*, 20(5):1013–1031, 1981.
- [80] V.I. Abkevich, A.M. Gutin, and E.I. Shakhnovich. Specific nucleus as the transition state for protein folding: evidence from the lattice model. *Biochemistry*, 33(33):10026–10036, 1994.
- [81] K. Yue. A test of lattice protein folding algorithms. *Proc. Natl. Acad. Sci.*, 92(1):325–329, 1995.
- [82] B.C. Gin, J.P. Garrahan, and P.L. Geissler. The limited role of nonnative contacts in the folding pathways of a lattice protein. *J. Mol. Biol.*, 392(5):1303 – 1314, 2009.
- [83] C. Clementi. Coarse-grained models of protein folding: Toy models or predictive tools? *Current Opinion in Structural Biology*, 18(1):10 – 15, 2008.
- [84] D.H. de Jong, G. Singh, W.F. Drew Bennett, C. Arnarez, T.A. Wassenaar, L.V. Schäfer, X. Periole, D.P. Tieleman, and S.J. Marrink. Improved parameters for the Martini coarse-grained protein force field. *J. Chem. Theo. Comput.*, 9(1):687–697, 2013.
- [85] B. Hess et al. GROMACS 4: Algorithms for highly efficient, load-balanced, and scalable molecular simulation. *J. Chem. Theo. Comput.*, 4(3):435–447, 2008.
- [86] J.C. Phillips, R. Braun, W. Wang, J. Gumbart, E. Tajkhorshid, E. Villa, C. Chipot, R. D. Skeel, L. Kalé, and K. Schulten. Scalable molecular dynamics with NAMD. *J. Comput. Chem.*, 26(16):1781–1802, 2005.

- [87] A. Cheng and K.M. Merz. Application of the Nosé-Hoover chain algorithm to the study of protein dynamics. *J. Phys. Chem.*, 100(5):1927–1937, 1996.
- [88] H.J.C. Berendsen, J.P.M. Postma, W.F. van Gunsteren, A. DiNola, and J.R. Haak. Molecular dynamics with coupling to an external bath. *J. Chem. Phys.*, 81(8):3684–3690, 1984.
- [89] V.S. Pande, I. Baker, J. Chapman, S.P. Elmer, S. Khaliq, S.M. Larson, Y.M. Rhee, M.R. Shirts, C.D. Snow, E.J. Sorin, and B. Zagrovic. Atomistic protein folding simulations on the submillisecond time scale using worldwide distributed computing. *Biopolymers*, 68(1):91–109, 2003.
- [90] C. Young, J.A. Bank, R.O. Dror, J.P. Grossman, J.K. Salmon, and D.E. Shaw. A $32 \times 32 \times 32$, spatially distributed 3D FFT in four microseconds on Anton. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, SC '09*, pages 23:1–23:11, New York, NY, USA, 2009. ACM.
- [91] H.S. Chan and K.A. Dill. Compact polymers. *Macromolecules*, 22:4559, 1989.
- [92] C.A. Rohl, C.E.M. Strauss, K.M.S. Misura, and D. Baker. Protein structure prediction using rosetta. *Methods Enzymol.*, 383:66 – 93, 2004.
- [93] E.I. Shakhnovich and A.M. Gutin. Enumeration of all compact conformations of copolymers with random sequence of links. *J. Chem. Phys.*, 93:5967–5971, 1990.
- [94] E.I. Shakhnovich and A.M. Gutin. Engineering of stable and fast-folding sequences of model proteins. *Proc. Natl. Acad. Sci.*, 90:7195–7199, 1993.
- [95] M. Cieplak, T.X. Hoang, and M.S. Li. Scaling of folding properties in simple models of proteins. *Phys. Rev. Lett.*, 83(8):1684–1687, Aug 1999.
- [96] J. Lee. Exact partition function zeros of two-dimensional lattice polymers. *J. Korean Phys. Soc.*, 44:617–620, 2004.
- [97] P. Stolorz. Recursive approaches to the statistical physics of lattice proteins. In *System Sciences, 1994. Proceedings of the Twenty-Seventh Hawaii International Conference*, volume 5, pages 316 – 325, Jan. 1994.
- [98] E. Shakhnovich, G. Farztdinov, A.M. Gutin, and M. Karplus. Protein folding bottlenecks: A lattice Monte Carlo simulation. *Phys. Rev. Lett.*, 67:1665–1668, Sep 1991.

- [99] S. Miyazawa and R.L. Jernigan. Estimation of effective inter-residue contact energies from protein crystal structures: Quasi-chemical approximation. *Macromolecules*, 18:534–552, 1985.
- [100] S. Kumar, J.M. Rosenberg, D. Bouzida, R.H. Swendsen, and P.A. Kollman. The weighted histogram analysis method for free-energy calculations on biomolecules. I. The method. *J. Comput. Chem.*, 13(8):1011–1021, 1992.
- [101] K.A. Beauchamp, R McGibbon, Y.S. Lin, and V.S. Pande. Simple few-state models reveal hidden complexity in protein folding. *Proc. Natl. Acad. Sci.*, page PMID: 22778442 (ahead of publication), 2012.
- [102] A.S.J.S. Mey and F. Noé. Efficient estimation of equilibrium expectations from generalized-ensemble simulations. *Phys. Rev. X*, 2012. under review.
- [103] B. Smit and D. Frenkel. *Understanding Molecular Simulation*. Academic Press, 2001.
- [104] R.G. Melko. Simulations of quantum XXZ models on two-dimensional frustrated lattices. *J. Phys. Condens. Matter*, 19:145203, 2007.
- [105] A. de Candia and A. Coniglio. Spin and density overlaps in the frustrated Ising lattice gas. *Phys. Rev. E*, 65:016132, Dec 2001.
- [106] E.-M. Ilgenfritz et al. Parallel tempering in full QCD with Wilson fermions. *Phys. Rev. D*, 65:094506, 2002.
- [107] G. Burgio, M. Fuhrmann, W. Kerler, and M. Müller-Preussker. Modified SO(3) lattice gauge theory at $T \neq 0$ with parallel tempering: Monopole and vortex condensation. *Phys. Rev. D*, 75:014504, Jan 2007.
- [108] B. Joó, B. Pendleton, S.M. Pickles, Z. Sroczynski, A.C. Irving, and J.C. Sexton. Parallel tempering in lattice QCD with O(a)-improved Wilson fermions. *Phys. Rev. D*, 59:114501, Apr 1999.
- [109] U.H.E. Hansmann. Parallel tempering algorithm for conformational studies of biological molecules. *Chem. Phys. Lett.*, 281(1-3):140 – 150, 1997.
- [110] Y Sugita and Y. Okamoto. Replica-exchange molecular dynamics method for protein folding. *Chem. Phys. Lett.*, 314:141, 1999.
- [111] M.J. Abraham and J.E. Gready. Ensuring mixing efficiency of replica-exchange molecular dynamics simulations. *J. Chem. Theo. Comput.*, 4(7):1119–1128, 2008.
- [112] S. Park and V. S. Pande. Choosing weights for simulated tempering. *Phys. Rev. E*, 76:016703, Jul 2007.

- [113] S. Rauscher, C. Neale, and R. Pomřšs. Simulated tempering distributed replica sampling, virtual replica exchange, and other generalized-ensemble methods for conformational sampling. *J. Chem. Theo. Comput.*, 5(10):2640–2662, 2009.
- [114] P. Eastman et al. Openmm 4: A reusable, extensible, hardware independent library for high performance molecular simulation. *J. Chem. Theo. Comput.*, 9(1):461–469, 2013.
- [115] W.-N. Du, K.A. Marino, and P.G. Bolhuis. Multiple state transition interface sampling of alanine dipeptide in explicit solvent. *J. Chem. Phys.*, 135(14):145102, 2011.
- [116] J.-H. Prinz, J.D. Chodera, V.S. Pande, W.C. Swope, J.C. Smith, and F. Noé. Optimal use of data in parallel tempering simulations for the construction of discrete-state Markov models of biomolecular dynamics. *J. Chem. Phys.*, 134(24):244108, 2011.
- [117] V. Hornak, R. Abel, A. Okur, B. Strockbine, A. Roitberg, and C. Simmerling. Comparison of multiple Amber force fields and development of improved protein backbone parameters. *Proteins*, 65(3):712–725, 2006.
- [118] J.C. Kendrew. The structure of globular proteins. *Comp. Biochem. Physiol.*, 4(2):249 – 252, 1962.
- [119] M.F. Perutz et al. Structure of haemoglobin: A three-dimensional fourier synthesis at 5.5-Å. resolution, obtained by x-ray analysis. *Nature*, 185:416–422, 1960.
- [120] I.I. Rabi, J.R. Zacharias, S. Millman, and P. Kusch. A new method of measuring nuclear magnetic moment. *Phys. Rev.*, 53:318 – 318, Feb 1938.
- [121] Kurt Wřthrich. *NMR of Proteins and Nucleic Acids*. New York (NY), USA: Wiley-Interscience, 1986.
- [122] E. Evans and A. Berk, D.and Leung. Detachment of agglutinin-bonded red blood cells. I. Forces to rupture molecular-point attachments. *Biophys. J.*, 59(4):838 – 848, 1991.
- [123] E.-L. Florin, V.T. Moy, and H.E. Gaub. Adhesion forces between individual ligand-receptor pairs. *Science*, 264:415–417, 1994.
- [124] L. Stryer and R.P. Haugland. Energy transfer: A spectroscopic ruler. *Proc. Natl. Acad. Sci.*, 58:719 –726, 1967.

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". `classicthesis` is available for both \LaTeX and \LyX :

<http://code.google.com/p/classicthesis/>