Fast and Slow Dynamics in Kinetically Constrained Models of Glasses

Douglas James Ashton

School of Physics & Astronomy University of Nottingham

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

 \cdot November 2008 \cdot

Abstract

Kinetically constrained models (KCMs) are able to account for many of the slow dynamical properties of glass forming systems such as dynamic heteorgeneity and Stokes-Einstein breakdown using simple models with simple dynamical rules. In this thesis we study several KCMs and extend them to include fast degrees of freedom. We show how the method of Monte Carlo with absorbing Markov chains can be applied to a particular class of KCMs, the facilitated spin models, to create an efficient numerical algorithm that can speed up simulations by several orders of magnitude. Another branch of KCMs, the constrained lattice gases, are studied and new results for a version on an FCC lattice in three dimensions are presented. This model is necessary when fast dynamics are studied and dimension plays an important role. To establish how fast degrees of freedom can be introduced without changing the character of the underlying KCMs we introduce coupled Ising spins to several existing models. We find that these models can reproduce much of the fast behaviour seen in the β -relaxation of real supercooled liquids without changing the slow behaviour that is already well described by KCMs. Lastly, by considering harmonic interactions between particles we study the relation between short-time vibrational modes and long-time relaxational dynamics in two constrained lattice gas models. We find an excess in the vibrational density of states similar to the 'Boson peak' of glasses and we find a correlation between the location of these low (high) frequency vibrational modes and regions of high (low) propensity for motion in agreement with recent results from atomistic simulations.

Acknowledgements

This thesis would not exist were it not for the support of a number of people. First and foremost is my supervisor, Juanpe. Going all the way back to our lively tutorials at Balliol College, no one has had a bigger influence on how I think about science. Thanks for all your support and guidance.

I've made many great friends in my time in Nottingham. Thanks to Lester, my physics brother, for three great years living and working together. I always enjoyed chatting physics and I hope we can keep working together in the future. Thanks also to my other housemates: Myanna, Alec and Kuldeep, for all the good times. To Jara, thanks for being such a great friend and for sharing all those cups of tea. Working late won't be the same again. Thanks as well to the rest of C18 and the physics department - 10:45 will always be tea time.

I'm much indebted to David Chandler and the rest of the Chandler group for their hospitality during my three months in Berkeley. I learned so much from all of you. Particular thanks to Rob for answering all my silly questions and to Yael for being so welcoming.

Most importantly, thank you to my family, especially my parents. Without your unwaivering support and friendship I would never have made it this far.

Citations to Previously Published Work

All of the work in this thesis is original except for relevant background material in chapter 1 and elsewhere - this is always clearly cited. Any figures that use data from other sources are labelled appropriately. Some figures, for example in chapter 4, reproduce results that have already been performed. However, unless otherwise stated the data for these figures has been produced by the author.

Chapter 2 is the fruit of collaboration with Lester Hedges and has been published in

• Fast Simulation of Facilitated Spin Models, Douglas J. Ashton, Lester O. Hedges and Juan P. Garrahan, J. Stat. Mech. P12010 (2005).

Chapters 3 and 4 are unpublished work carried out by the author. Chapter 5 is available on pre-print servers as

• Relationship between vibrations and dynamical heterogeneity in a model glass former: extended soft modes but local relaxation, Douglas J. Ashton and Juan P. Garrahan, arXiv:0808.2412, submitted to Phys. Rev. Lett.

Contents

Acknowledgements i								
$\mathbf{C}\mathbf{i}$	tatic	ons to F	Previously Published Work	ii				
1	Intr	on	1					
	1.1	Formers	2					
	1.2	Dynan	nic Heterogeneity	4				
	1.3	Fast D	Dynamics	5				
	1.4	Kineti	cally Constrained Models	8				
	1.5	Aim o	f the thesis	8				
2	Adv	vanced	Simulation Techniques	9				
	2.1	Numer	rical Methods	9				
	2.2	Facilit	ated Spin Models	10				
	2.3	Contin	uous Time	13				
	2.4	Monte	Carlo with Absorbing Markov Chains	14				
		2.4.1	Introduction to MCAMC	15				
		2.4.2	Application to the East Model	16				
		2.4.3	Approximations for the Update Time	18				
		2.4.4	Generalisation to Any Dimension	19				
		2.4.5	FA-East Crossover model	20				
		2.4.6	Higher order MCAMC	21				
		2.4.7	Speed Tests	23				
		2.4.8	Example of results	25				
	2.5	Discus	sion	27				
3	A KCM in Three Dimensions: The FCC Constrained Lattice Gas 28							
	3.1	The co	onstrained triangular lattice gas	29				
	3.2	The F	CC lattice gas	31				
		3.2.1	Definition of the model	31				
		3.2.2	Backbones	31				
		3.2.3	Two-point correlation functions	34				
		3.2.4	Stokes-Einstein breakdown	36				
		3.2.5	Persistence and exchange times	38				
	33	3.3 Conclusion						

4	Fast	East Degrees of Freedom								
	4.1	Theories of fast dynamics								
		4.1.1 The β -relaxation								
		4.1.2 Correlated dynamics								
		4.1.3 Schemes for fast KCMs 43								
	4.2	One dimension								
		4.2.1 Thermodynamics								
		4.2.2 Dynamics								
		4.2.3 Spatial correlation								
		4.2.4 Summary								
	4.3	Two dimensions								
		4.3.1 Diluted Ising model								
		4.3.2 Blume-Capel model								
		4.3.3 Thermodynamics								
		4.3.4 Dynamics								
		4.3.5 Dynamic Heterogeneity								
	4.4	Concluding Remarks								
5	Vib	rations 65								
	5.1	Vibrations and rigidity								
	5.2	Dynamics of the constrained lattice gases								
	5.3	Central Force Network								
	5.4	Comparison of the Two Models								
	5.5	Nonaffine Displacement								
	5.6	Conclusion								
6	Cor	aclusion 81								
-										
\mathbf{A}	Tra	nsfer Matrices 82								
	A.1	General results								
		A.1.1 Bulk properties								
		A.1.2 Magnetisation								
		A.1.3 Spatial Correlation								
	A.2	Results for the ns-model								
		A.2.1 Magnetisation								
		A.2.2 Spatial Correlations								
	A.3	Equilibrium configurations								
р	P. Chuston Almonithma									
D		1 Detailed Palance								
	В 0 Г.Т	The Metropolic algorithm								
	D.2 Рэ	Incomposition 09 Wolff algorithm 00								
	D.) Д 4	Wom agonum								
	D.4	invaded cluster algorithm								

Chapter 1

Introduction

Since its birth the field of statistical mechanics has made remarkable achievements in connecting the microscopic world of atoms and molecules to the macroscopic world that we can perceive. We now understand how, for a tiny change in conditions, a liquid can freeze into a solid or how a material can spontaneously magnetise. Perhaps at its most elegant it says that near a phase transition the exact details of a system can become unimportant. Substances can be grouped together into broad universality classes that depend only on general features such as dimensionality and symmetry. This emergence of simplicity from the seemingly overwhelming complexity is one reason why the study of simple, idealised models, has been so successful in developing our understanding of what drives matter to behave in the way it does.

Glasses and glass forming liquids represent a significant outstanding problem in statistical physics. They are very common materials in our everyday life that have been exploited for millennia. While much is known experimentally we still do not have a full understanding, on the microscopic level, of how a glassy liquid flows or why it eventually becomes a glass. Advances in computing power have allowed for large scale simulations at the atomic level that provide an unprecedented level of detail covering a growing range of time scales. There is still a place, however, for the simple model. By stripping the level of detail down to its bare minimum simple models help to build an intuition and, depending on their success or failure, can sort between which features are really important.

A challenge for any simplified model of glassy behaviour is that there is a lot of behaviour to explain. Depending on what time scale is of interest the same material can have properties that are usually assigned separately to liquids and solids. The aim of this thesis is to try and explain as much of this unusual behaviour as possible using as few ingredients as possible. The emphasis is placed on understanding the model first and then seeing how much behaviour it can reproduce rather than best reproducing the behaviour and then trying to understand the model. While this can sometimes lead to less successful models it is our opinion that it also leads to a better understanding.

1.1 Glass Formers

Glasses are all around us. There are the familiar uses in our everyday lives: building materials, containers and so forth. There are also many more not-so-familiar uses throughout technology. From optical fibres, to plastics and photovoltaic cells [1-3], many of the materials in use today are in a glassy state. Modern pharmocology can make drugs that are absorbed much faster when delivered as a glass [4] and it is even thought that proteins have many features in common with glasses [3,5]. Gaining a better understanding of how and why glasses form, and how to control it, is therefore of great relevance as well as great interest.

To describe exactly what a glass is it is easier to describe how one is made. If a liquid is cooled slowly to its melting point, T_m , then it will undergo a phase transition and crystallise into a solid. If it is cooled quickly enough then the molecules do not have time to form a crystal and the phase transition can be avoided forming a supercooled liquid. As Figure 1.1 shows, there is no sudden change on passing through T_m and the supercooled liquid has the same structure as the normal liquid. In this state the liquid can be thought of as being in a kind of metastable equilibrium because the relaxation times are short compared to the time to crystallise.

While the structure does not noticeably change [6], the dynamical properties do begin to deviate quite rapidly. As the liquid is cooled further below T_m the viscosity starts to increase. This increase quickly crosses several orders of magnitude until eventually it can no longer be measured. At this point the system has fallen out of any kind of equilibrium and becomes a glass, it will no longer flow and is solid. The temperature that this occurs, T_g , depends on the rate of cooling and is usually defined as the temperature where the viscosity has reached 10^{13} Poise, the largest that can be measured.

The exact nature of the glass transition is not widely agreed upon. It is certainly not a phase transition in the thermodynamic sense [7], there is no observed diverging static lengthscale and no well defined transition temperature. Neither does it appear to be a dynamic phase transition, where the timescales diverge at a finite temperature, although this is somewhat contested.

On approach to the glass transition different liquids slow down in different ways. Broadly speaking they can be broken down into two categories, 'strong' and 'fragile'. The timescales of strong glass formers, such as silica (SiO₂) or GeO₂, change with temperature in an Arrhenius way. This means that the viscosity, η , can be fitted quite well with the form $\eta \sim \exp(E/k_BT)$, where E is a constant that does not depend of temperature.

Fragile glass formers, such as O-Terphenyl (OTP), have a much greater dependence on temperature. In Figure 1.2 the fragile liquids are the ones that follow the steepest curve, with an increase in viscosity near to the glass transition that is quite spectacular. The relationship is often fitted by the Vogel-Tammann-Fulcher (VTF) equation, where $\eta \sim \exp(B/(T - T_0))$, although this fit does not work so well for the most fragile liquids and it contains a divergence at T_0 . The most important thing to note at this stage is that, in all glass formers, there is an enormous slowing down without any significant changes in the structure.



Figure 1.1: A representation of specific volume for a liquid cooled at different rates. For faster cooling there is no noticeable change around the melting point, T_m , and moving from the supercooled liquid to the glass is seen with a smooth variation that depends on the cooling rate. In this plot glass 2 was cooled more slowly than glass 1. Taken from Ediger [2].



Figure 1.2: Plot of log viscosity, $log(\eta)$, against scaled inverse temperature and fitted with the Vogel-Tamman-Fulcher equation (see text). Strong glass formers form a straight line (Arrhenius) while fragile glass formers have a more drastic dependence on temperature. The viscosity can jump by up to 14 orders of magnitude in a relatively narrow temperature range. Taken from Angell [8].



Figure 1.3: A supercooled colloidal suspension where the fastest particles have been tagged with large spheres and other particles with small spheres. Particles are $\sim 3\mu$ m in size. Figure taken from Weeks *et al.* [9].

1.2 Dynamic Heterogeneity

If the structural properties do not show any remarkable changes on approach to the glass transition then the dynamical properties more than make up for it. Two point correlation functions can tell us how the system relaxes after a perturbation. For a normal liquid the correlations in time decay as an exponential, in supercooled liquids this form breaks down and the correlations are often fitted with a stretched exponential of the form

$$C(t) \approx e^{-(t/\tau)^{\beta}} \tag{1.1}$$

where τ is the average relaxation time and β is a parameter that quantifies the stretching.

The reason for this unusual scaling form is not that all parts of the liquid are relaxing in an homogenous, non-exponential manner, but rather that different parts of the liquid are relaxing in very different ways. This is dynamic heterogeneity. At any given time scale one region of the liquid could appear quite fluid whereas another region, possibly very close by, hardly moves at all. These particles are locked in place until such a time that the collective movements of all the particles nearby, and not so nearby, conspire to release them. A vivid demonstration of this breaking up into fast and slow regions is given in an experiment done by Weeks *et al.* on a colloidal system [9]. Colloids demonstrate glassy behaviour but are also large enough to see. By tracking individual particles using conformal microscopy they were able to label each one by how much it was moving. Figure 1.3 shows that in the supercooled regime the fast particles are tightly correlated in space and clusters can involve many particles.

Apart from colloidal systems dynamic heterogeneities have been seen in experiments [10] and in computer simulations of glass forming materials [11] and they are now at the core of our current understanding of glassy dynamics. Instead of a picture of a growing static lengthscale that is familiar from thermodynamic phase transitions, we here see a growing dynamic length scale.



Figure 1.4: The self-intermediate scattering function, $F_s(k = k_0, t)$, for a variation on the binary Lennard-Jones mixture. k_0 is the wavevector corresponding to the first peak in the structure factor. Simulations use a WCA cut-off potential [13] for temperatures ranging from high liquid like, T = 5, to low supercooled, T = 0.36, (temperature units relative to the interaction energy of the particles). Data courtesy of Lutz Maibaum.

1.3 Fast Dynamics

The full structural relaxation that goes with viscosity is not the only process of interest in a glass former. As the temperature drops and structural relaxation time increases, other faster processes begin to appear as distinct. This is best seen from a microscopic viewpoint using two-point correlation functions as described above. Figure 1.4 shows the results from computer simulations on binary Lennard-Jones mixtures measuring the self-intermediate scattering function defined as

$$F_s(\vec{k},t) = \frac{1}{N} \sum_i e^{i\vec{k} \cdot (\vec{r}_i(t) - \vec{r}_i(0))}$$
(1.2)

where $r_i(t)$ is the position of particle *i* at time *t* and \vec{k} is a wave vector that sets the lengthscale over which we are interested. In the figure \vec{k} is set such that the correlation decays when most particles have moved a particle diameter from their starting position. For high temperatures the mixture relaxes exponentially with only one time scale, τ_{α} . As temperature is decreased the α -relaxation moves to longer and longer times, fitting a stretched exponential of Eq. 1.1, and a two-step picture begins to emerge. A shorter, secondary process, known as the β -relaxation [12] occurring over a characteristic time, τ_{β} , begins to appear as distinct. At low temperatures it becomes widely separated from the α -relaxation with an extended plateau separating them.

A simplified picture of this time scale separation goes something like this: Particles are locally restricted, or caged, by their neighbouring particles preventing them from making large movements away from their starting position. At short times particles can 'rattle' in these cages (and indeed so can the particles forming the cage) without making any significant change to the structure. It is this movement within a fixed local structure that is the β -relaxation. After a while of doing this a particle might be able to escape from its cage signifying the start of the α -relaxation.

The crossover from β behaviour to α behaviour is by no means cleanly defined. One of the few predictions in this regime comes from mode coupling theory (MCT). MCT is a theory that uses approximations to solve the microscopic equations of motion and is very successful for intermediate temperatures [14]. At low temperatures it predicts a dynamical critical point, T_c , that does not match with experiment.

Around the plateau region MCT says that the correlation function should approach with a powerlaw decay, $F_s(\vec{k},t) = f + A_1 t^{-a}$, and that when leaving the plateau it should follow another powerlaw, $f - A_2 t^{-b}$. At later times it will go to the stretched exponential described above. Experimental and numerical data appears to confirm this functional form [15]. Even if this is not an exact result it is clear that the relaxation to the plateau is not a simple independent rattling, but a highly collective motion.

At time scales a little shorter and overlapping with the early β -relaxation, when the structure can be said to be frozen in place, we have vibrational dynamics. Vibrations are very interesting because they can be studied right through the glass transition and show some curious properties of their own. In a continuous elastic material the dynamics can be solved exactly and give homogenous plane wave solutions. A linear dispersion relation gives a vibrational density of states (DoS) that depends on frequency, ω , as

$$D(\omega) \propto \omega^{d-1} \tag{1.3}$$

where d is the dimension. This is the Debye law. For a crystal the Debye law works for low frequencies but eventually breaks down at a frequency ω_D that relates to a length scale where the microscopic structure of the crystal becomes visible. The Debye law is responsible for the low temperature T^3 dependency of the heat capacity in crystals.

In glasses the vibrational density of states can be measured using Raman and neutron scattering and it is found that the Debye law breaks down at a frequency lower than ω_D . This breakdown comes in the form of an excess of low frequency modes in the THz range that is known as the "Boson peak" (BP). The exact cause of the Boson peak is not agreed upon, partly because it is possible to account for an excess of modes using a number of different theoretical techniques [16–18] and so the detail appears to be in the structure of the anomalous modes.

The experiments shown in Figure 1.5 show that as the glass is heated up the BP moves back in frequency before eventually going to zero. This result motivates the view of Wyart *et al.* [19] that the Boson peak is the result of the glass being a marginally rigid solid. When a solid loses its rigidity it develops floppy modes of motion that are not resisted by any kind of restoring force. In a marginally rigid solid the anomalous low frequency modes are linked to beginnings of these floppy or 'soft' modes. Extrapolating from this view it has been suggested that the onset of glassiness and the onset of rigidity are closely linked [20]. This naturally raises the question of how the anomalous modes that are crucial to the onset, or breaking up, of rigidity are related to the dynamical heterogeneities that are crucial to the slowing down of the supercooled liquid.



Figure 1.5: The reduced density of states of toluene, ethylbenzene, dibutylphthalate, and glycerol. Debye behaviour would be a flat horizontal line in this plot. Arrows indicate the lowest temperature position of the Boson peak which, on heating, moves back in frequency towards zero. Taken from Chumakov *et al.* [21].

1.4 Kinetically Constrained Models

The aim of kinetically constrained models (KCMs) is to simplify the complicated dynamics of glasses by gaining as much understanding as possible with the minimum ingredients. As previously discussed there does not appear to be any great change in the thermodynamic properties upon supercooling, so KCMs are models with very simple, sometimes trivial, equilibrium states. They achieve glassiness by restricting the transitions between them. These restrictions, the 'kinetic constraints', are usually chosen with some physical motivation, although there is a lot of freedom for creativity.

Of the many varieties of KCM (see Ref [22] for a good review), the two types that are focused on in this thesis are spin facilitated models and constrained lattice gases. The former category, explained in detail in chapter 2, are coarse grained models that have dynamic heterogeneity at their heart. They are based on the idea that a region can only become dynamically active if facilitated by another active region nearby. The lattice gas models, thoroughly discussed in chapter 3, are less coarse grained and are constrained by particles locally blocking one another.

Both classes of KCM have been very successful in recreating many of the features described above. Both have heterogeneous dynamics, they display rapid slowing down and can make interesting observations on fragility. All this is done with very simple rules and dynamics that have coarse grained over the fast processes. Numerically, this extra simplicity allows us to explore deeply supercooled states with time scales that would be very difficult to access with more realistic simulations. The downside is that KCMs do not tell us anything about the β -relaxation or anomalous vibrations. Or perhaps more importantly, how the fast processes might relate to the structural α -relaxation that they do describe.

1.5 Aim of the thesis

The aim of this thesis is to make a connection between fast and slow dynamics using the kinetically constrained models as a basis. Continuing in the spirit of the simple model this will be done by adding the minimum of extra ingredients. The thesis is laid out as follows: In chapter 2 we introduce in detail the facilitated spin models and demonstrate how an advanced Monte Carlo algorithm can be applied to dramatically improve simulation times. In chapter 3 we turn to the constrained lattice gases and introduce a three dimensional model showing how it reproduces glassy behaviour. Chapter 4 sees the introduction of fast dynamics; we look at how to add fast degrees of freedom to the KCMs and how this affects the equilibrium properties. We finish in chapter 5 by studying the anomolous vibrations of an elastic model based on the lattice gases allowing direct comparison with the dynamic heterogeneity.

Chapter 2

Advanced Simulation Techniques

Numerical simulations are a hugely important tool in the study of supercooled liquids and glasses. They are vital for testing predictions and informing the theory, while at the same time acting as a kind of bridge between theory and experiment. As supercooled liquids are in their nature so slow, this can make numerical studies very difficult to perform. Kinetically constrained models (KCMs) are very simple models of glass formers and as such are much easier to treat numerically but still suffer from the slowness problem. Much of the time is spent trying to escape from deep energy or free energy traps.

In this chapter we will review the facilitated spin models (FSMs) as a class of KCM and some techniques that are available to improve the efficiency of numerical simulations. We introduce an algorithm based on the absorbing Markov chain that is specially designed for systems that are often stuck in such traps and apply it to a particular FSM, the 'East model'.

The work in this chapter, from section 2.4 onwards, was undertaken in collaboration with Lester Hedges and has been published as *Fast Simulation of Facilitated Spin Models*, Douglas J. Ashton, Lester O. Hedges and Juan P. Garrahan, J. Stat. Mech. P12010 (2005) [23].

2.1 Numerical Methods

Much of our understanding of glass formers has come from detailed microscopic simulations. These techniques calculate continuous forces from effective pair potentials that can either be based on real molecular interactions or on more generic Lennard-Jones potentials. The most common method, molecular dynamics (MD), is simply to integrate the equations of motion using a suitably chosen time increment, δt . Connection to the relevant statistical ensemble (usually the canonical ensemble) is achieved either from carefully chosen initial conditions or through the use of thermostats. MD is very computationally intensive as it requires the constant recalculation of the forces for each small increment in time. Of course the benefit is that, for a short time, one essentially has a laboratory inside the computer.

There are alternatives to MD, real liquids can also be modelled using stochastic dynamics. In this case connection to the statistical ensemble is made through the use of noise and friction terms in the microscopic equations of motion. With Brownian dynamics it is simplified even further by not considering acceleration. In as far as the thermodynamic properties are concerned all of these techniques are exactly the same because they are evolved with respect to the same ensemble and consider the same interactions. In terms of efficiency there is no improvement over MD but it is interesting to note that the dynamic properties, while different, have been shown to be qualitatively the same, particularly around the α -relaxation [24].

As a means of exploring the full phase space to measure thermodynamic quantities, all of the above methods are quite inefficient because they are limited to making tiny changes in configuration. Monte Carlo (MC) is a very general method that can sample the phase space in an arbitrary manner provided it preserves the Boltzmann distribution [25]. This can be incredibly useful, for example, near critical points. A system around its critical point becomes very slow due to the growing size of fluctuations and so algorithms have been developed, such as Wolff algorithm [25], that make large changes to clusters on the same length scale as the fluctuations allowing it to move efficiently through the phase space.

If it is the dynamics that we are interested in (and it is here) then it isn't much use making enormous changes that lose all the detail of how we got there. It is possible, however, to make small steps to nearby points in phase space that are close enough to give a meaningful picture of the dynamics, but far enough away to make a big difference to the efficiency. This approach was found to be very effective in a standard Lennard-Jones mixture showing that MC could reproduce the long-time stochastic dynamics exactly [26].

Kinetically constrained models have already been introduced in chapter 1 as an attempt to simplify our understanding of supercooled liquids [22]. They do this by coarse graining the microscopic degrees of freedom and by moving onto a lattice thereby greatly reducing the size of the phase space. KCMs use stochastic dynamics and, in a manner similar to above, make small changes to make these dynamics meaningful. A crucial consequence of the coarse graining is that the number of neighbouring points in phase space is small enough that we can begin to employ some clever algorithms to improve performance. In the next section we will give some concrete examples as to how this works in a KCM.

2.2 Facilitated Spin Models

Possibly the simplest of the KCMs is the Fredrickson-Andersen (FA) model [27]. In this model the liquid is coarse grained onto a lattice of Ising spins, n_i , where sites can either be mobile, $n_i = 1$, or immobile, $n_i = 0$. There is an energy cost for being in the mobile state so the Hamiltonian is simply defined as

$$H = J \sum_{i} n_i \tag{2.1}$$



Figure 2.1: Schematic representation of the MCAMC algorithm. Standard Monte Carlo makes many iterations with most resulting in failed moves. Continuous time jumps between successful moves but can still take a long time to escape a free energy trap. MCAMC jumps to a point where two or more steps are made away from the trap thus increasing efficiency.

where J = 1 from now on. The equilibrium configuration has no static correlations and sites are mobile (excited) with a concentration, $c = 1/(1 + e^{\beta})$, where β is the inverse temperature, $\beta \equiv 1/T$.

The dynamics proceed by randomly attempting to flip single spins. Glassiness in the FA model comes about through the constraint that a site can only change its spin if it is *facilitated* by a neighbouring mobile site. This simple constraint is based on the idea that a jammed region cannot become spontaneously unjammed, it requires some nearby mobility to get things started. If we define the facilitation of site i as $f_i = \sum_j n_j$, where the sum is over nearest neighbours, then for the standard FA model a move is allowed if $f_i \geq 1$. In higher dimensions this can be generalised to the f-spin facilitated FA model where the restriction is $f_i \geq f$.

In order to ensure that the equilibrium state is always that of the Boltzmann distribution the dynamics must always obey detailed balance. This is a simple restriction that says that, at equilibrium, there should be no net current between any two states. This can be expressed as

$$P_{\mu}W(\mu \to \nu) - P_{\nu}W(\nu \to \mu) = 0 \tag{2.2}$$

where P_{μ} is the equilibrium probability of being in state μ and $W(\mu \to \nu)$ is the rate of going from μ to ν . The ratio of the equilibrium probabilities can be fixed in terms of the Boltzmann weights giving

$$\frac{W(\mu \to \nu)}{W(\nu \to \mu)} = \frac{P_{\nu}}{P_{\mu}} = e^{-\beta(E_{\nu} - E_{\mu})}$$
(2.3)

The transition rates can be broken up into a selection rate and an acceptance rate such that $W(\mu \rightarrow \nu) = S(\mu \rightarrow \nu)A(\mu \rightarrow \nu)$. At this point it is possible to play with different schemes to create very efficient algorithms depending on what is being studied. For example the Wolff algorithm mentioned earlier uses selection rates to obey detailed balance and sets acceptance rates to one. The FA model has symmetric selection probabilities such that the ratio of transition rates is the same as the ratio of acceptance rates. Given it is

only the *ratio* of the rates that is restricted, it makes sense numerically to make each rate as large as possible. This is known as the "Metropolis algorithm" and is done by setting

$$A(\mu \to \nu) = \min\{1 : e^{-\beta(E_{\nu} - E_{\mu})}\}$$
(2.4)

The kinetic constraint does not violate detailed balance because transitions are not allowed in either the forward or reverse move so equation (2.2) is satisfied. The basic algorithm (from here on referred to as Monte Carlo (MC)) for the FA model therefore proceeds as follows:

- 1. From a starting configuration, $\{n_i\}$, choose a site *i* at random.
- 2. If none of its nearest neighbours are excited, $f_i = 0$, then jump to step 5.
- 3. Calculate the energy difference, ΔE , if the move was accepted.
- 4. Generate a random number, r, between zero and one. If $r < \min\{1 : e^{-\beta \Delta E}\}$ then accept the proposed move and let $n_i \to 1 n_i$.
- 5. Increment the clock and repeat from step 1.

At low temperatures the FA model becomes very slow because the concentration of excitations, $c \approx e^{-\beta}$, is very small and so most sites cannot change most of the time. In this regime movement is dominated by the diffusion of excitations through a thermally activated process

$$100 \xrightarrow{\epsilon} 110 \xrightarrow{1} 010 \tag{2.5}$$

where the numbers above the arrows indicate the relative transition rates between the states and $\epsilon \equiv e^{-\beta}$.

By slightly changing the kinetic constraint it is possible to greatly affect the dynamics. A variant of the FA model in one dimension, the "East model" [28], is defined in the same way as the FA model and shares the Hamiltonian from equation (2.1). The difference comes in the kinetic constraint. Instead of sites being facilitated by any of their nearest neighbours in the East model they are only facilitated by their neighbour to the left. This has the consequence that excitations propagate to the 'east'. The rates of the allowed moves are therefore

$$10 \xrightarrow{\epsilon} 11, \ 11 \xrightarrow{1} 10$$
 (2.6)

Diffusion cannot occur in the East model because the final move in (2.5) is not allowed. So whereas in the FA model movement is restricted by one energy barrier, in the East model it turns out that movement is restricted by a hierarchical series of increasing energy barriers (to be discussed later).

Despite their simplicity both models can reproduce much of the behaviour of a glass former described in chapter 1. Dynamical heterogeneity, arising from the excitations, appears naturally in the models. Excitations are distributed throughout the system and as temperature drops the distance between them increases, diverging as $T \rightarrow 0$. At low temperatures the decay of two-point correlation functions become stretched exponentials and the characteristic time scales increase rapidly with decreasing temperature. Fig.



Figure 2.2: (a) Persistence functions, P(t), for the FA model at inverse temperatures $\beta = 1$ to 6 (left to right) developing stretched exponentials at low temperature. (b) Angell plot for the FA and East models showing strong and fragile behaviour respectively. Persistence times, τ_{α} , are plotted against scaled inverse temperature. Here T_g is defined as the temperature where $\tau_{\alpha} = 3 \times 10^{12}$ Monte Carlo sweeps.

2.2 shows this with the persistence function, P(t). This two-point correlation function measures the fraction of sites have never changed up to time t.

The time scales for the FA model increase in an Arrhenius manner, $\tau_{\text{FA}} \propto e^{A/T}$ (A a constant), whereas the East model slows down much more rapidly with a form, $\tau_{\text{East}} \propto e^{f(1/T)}$, where f(1/T) can be fitted with a quadratic. In this way the FA and East models can be thought of as being strong and fragile glass formers respectively.

2.3 Continuous Time

The MC algorithm described above is extremely inefficient. Even at a modest temperature of T = 0.5 only about one in a hundred attempted moves would be successful (even less for the East model) and for lower temperatures it quickly becomes completely useless. Fortunately the simplicity of the models allow us to employ some clever tricks.

For any given configuration we can work out which sites are facilitated and which are not. If we know the number of facilitated sites in the excited state, N_1 and the number of facilitated sites that are not excited, N_0 , then we can calculate the probability that an attempted move will be successful.

The probability of choosing a facilitated excited spin from the N possible sites is N_1/N . The acceptance rate for flipping excited spins is 1 so at each iteration the probability that we will successfully change a spin from 1 to 0 is $p_1 = N_1/N$. For the unexcited spins we have an acceptance rate of $e^{-\beta}$ due to the energy penalty and so the probability of making one of these moves is $p_0 = N_0 e^{-\beta}/N$. The probability of successfully making any move is thus

$$p_s = p_1 + p_0 = \frac{N_1 + N_0 e^{-\beta}}{N} \tag{2.7}$$

If we were to attempt to do this a number of times, the probability that we have still not successfully made a move after m iterations is simply $P_m = (1 - p_s)^m$, where P_m starts at one and decays to zero at large m. If we rearrange for m we get

$$m = \frac{\log P_m}{\log(1 - p_s)} \tag{2.8}$$

By putting a random number into P_m we can generate time steps with the correct distribution. Once we have a time we need to decide what type of move we made. The probability that we flipped an excited spin is p_1/p_s and the remaining probability goes to flipping an unexcited spin. Provided we keep track of where the facilitated sites are, taking care to update as we go, then it is possible to advance the dynamics without ever rejecting a move. The procedure is then as follows

- 1. Generate a random number between zero and one and substitute into P_m in equation (2.8) to produce a time step, m, rounding up to the nearest integer. Increment the clock, $t \to t + m$.
- 2. Using another random number choose whether to flip an excited spin (probability p_1/p_s) or an unexcited spin (p_0/p_s) .
- 3. Make the chosen move and update the tables for the next move. Repeat from 1.

This sort of algorithm is often referred to as "rejection free" or "the *n*-fold way" [29]. It should not be possible to tell apart the results of this rejection free simulation and a MC simulation. It is possible however to use an average time increment rather than picking from the distribution as the approximation does not have an effect except at short times and speeds the algorithm up even further. This is known as "continuous time" (CT) [25].

Because the algorithmic overhead is relatively small in this model the speed gain from using the CT algorithm is $\sim e^{2\beta}$ times faster than MC. At low temperatures this can be many orders of magnitude. In more complicated models that have many different types of move it can still prove worth while using a CT algorithm. Eventually, as the complexity increases, the cost of keeping track of all the possible moves outweighs the benefits.

2.4 Monte Carlo with Absorbing Markov Chains

Continuous time is very effective when the majority of attempted moves are very unlikely to be accepted due to energy or kinetic constraints. Instead the algorithm fast forwards to a time when a move is accepted. For many slow systems making this first unlikely move is only just the beginning. The overwhelmingly likely next move will be to undo the first one leaving you back where you started. For example consider the East model as described above. At low temperatures excitations are mostly isolated from one another in a configuration such as:

 $\cdots 100 \cdots 100 \cdots 100 \cdots$

where the dots represent a chain of zeroes. A continuous time algorithm would flip one of the three facilitated spins to an excited state. Because the rate of relaxing the excitation is 1 compared to the rate of creating a new one, $e^{-\beta}$, the next step will almost certainly return us to the isolated state. Another area where this is a problem is in the study of magnetic reversal. An Ising model below T_c has a metastable configuration where all spins are aligned in one direction. This state represents a deep minimum in the energy. Individual spins will occasionally flip over but the time for enough to do it simultaneously to gather the pace to make a full reversal can be extremely long.

The natural response to these problems is to try and create an algorithm that can fast forward to a time when not just one unlikely move is accepted, but two or more in a row (see Fig. 2.1). The general framework for doing this was developed by Novotny [30] (see [31] for a pedagogical review) and is known as Monte Carlo with absorbing Markov chains (MCAMC).

2.4.1 Introduction to MCAMC

In a MC algorithm any given move depends only on the two states that the system is moving between and not on any previous moves. This is by definition a Markov process and allows us to treat the MC algorithm as a Markov chain. A Markov chain is characterised by the matrix **M** which defines transition probabilities between states. If the vector $\vec{x}^T(m)$ indicates the probability distribution of the system after iteration m, the probability distribution at the next step m + 1 is given by $\vec{x}^T(m+1) = \vec{x}^T(m)\mathbf{M}$.

An absorbing Markov chain is defined by separating the available states into s transient states and r absorbing ones [30]. The system always starts in a transient state and by successive applications of the Markov matrix explores the transient subspace until it lands in an absorbing (or exit) state from where it cannot leave. We can divide the general state vector \vec{x}^T into absorbing and transient parts, to get the (r + s)-dimensional vector $\vec{x}^T = (\vec{u}^T, \vec{v}^T)$ where \vec{v}^T contains the transient states. The initial state in this form must obey $\vec{x}_I^T = (\vec{0}^T, \vec{v}_I^T)$. With this structure the Markov matrix can be written in the form,

$$\mathbf{M} = \begin{pmatrix} \mathbf{I}_{r \times r} & \mathbf{0}_{r \times s} \\ \mathbf{R}_{s \times r} & \mathbf{T}_{s \times s}, \end{pmatrix}$$
(2.9)

where \mathbf{I} is the identity matrix, $\mathbf{0}$ is the zero matrix and subscripts indicate the size of each matrix. The positions of the identity and zero matrices guarantee that if the system falls into an absorbing state then it does not leave. The transient matrix, \mathbf{T} , gives the probabilities for moving between transient states and the recursive matrix, \mathbf{R} , gives the probabilities for moving from the transient states to the absorbing states.

For a given starting vector \vec{v}_I^T the probability of still being in the transient subspace, $p_{\text{trans.}}$, after m steps is

$$p_{\text{trans.}} = \vec{v}_I^T \mathbf{T}^m \vec{e}, \qquad (2.10)$$

where \vec{e} is a vector with all elements equal to 1. The probability of absorbing to a particular state after m steps is given by summing over the probabilities of absorbing at each time step. This gives the vector $\vec{p}_{abs.}^{T}$,

$$\vec{p}_{\text{abs.after }m}^{T} = \vec{v}_{I}^{T} \big(\mathbf{I} + \mathbf{T} + \dots + \mathbf{T}^{m-1} \big) \mathbf{R}.$$
(2.11)

If the exit has taken place at step m, then the probabilities of absorbing into the different exit states is given by:

$$\vec{p}_{\text{abs.at }m}^{T} = \frac{\vec{v}_{I}^{T} \mathbf{T}^{m-1} \mathbf{R}}{\vec{v}_{I}^{T} \mathbf{T}^{m-1} \mathbf{R} \vec{e}}.$$
(2.12)

Here it is convenient to introduce the fundamental matrix

$$\mathbf{N} = (\mathbf{I} - \mathbf{T})^{-1} = \mathbf{I} + \mathbf{T} + \mathbf{T}^2 + \cdots, \qquad (2.13)$$

which can be used to obtain the probability that the system will absorb to a particular state irrespective of when it exits,

$$\vec{p}_{\text{abs.}}^T = \vec{v}_I^T \mathbf{NR}. \tag{2.14}$$

The fundamental matrix can also be used to determine the average time to leave the transient subspace

$$\langle \tau \rangle = \vec{v}_I^T \mathbf{N} \vec{e}. \tag{2.15}$$

Once our system is in the initial state we can generate an exit time by solving the inequality

$$\vec{v}_I^T \mathbf{T}^m \vec{e} < \overline{r} \le \vec{v}_I^T \mathbf{T}^{m-1} \vec{e}, \tag{2.16}$$

where \overline{r} is a random number between 0 and 1. Next, we use a second random number to choose an absorption state from the distribution in equation (2.12) and then we can update the system appropriately. The new state will become the initial state in another absorbing Markov chain, and so on. A successful MCAMC algorithm will choose transient states such that the system tends to move between them many times before exiting.

In this framework the CT algorithm is an MCAMC algorithm with a single transient state s = 1. Therefore MCAMC can be thought of as a generalisation of the rejection free techniques that fast forwards past not only failed moves, but futile moves as well.

2.4.2 Application to the East Model

When applied to magnetic reversal in the Ising model the MCAMC algorithm has a well defined initial transient state - all spins pointing in the same direction. The next state that goes into the transient subspace is that of a single spin reversing. This could be any one of the spins so this state groups together all of the N equivalent states. To leave the transient subspace requires two spins to reverse and these states form the absorbing subspace.

As already discussed, the East model is ideally suited to the MCAMC method because it also falls into deep energy traps from which it is difficult to escape. Unlike magnetic reversal there is no single initial transient state, instead we have to define a class of trapping states that will change as the dynamics evolve. The initial state we are interested in is one where excitations are isolated in space

 $\cdots 100 \cdots 100 \cdots 100 \cdots$

When in this state the only choice is to excite a spin and pay the accompanying energy penalty. Consequently it is highly likely that the raised spin is immediately relaxed, returning to the previous state. It is important to note that in this isolated state all excitations must be separated by a minimum of two unexcited spins. When separated by only a single spin there are two possible outcomes from the creation of an excitation, i.e. one can create a double state 110 or a triplet 111. This produces an unnecessary complication since the algorithm can no longer be classified by two simple transient states.

By analogy to the Ising model one may define two transient states for the system, the isolated configuration described above and states in which a single excitation pair exists. However, it is clear that neither of the transient states identified for the East model are unique. It is possible to construct numerous configurations which satisfy the above criterion, in essence we have identified two classes of transient state. The absorbing states consist of all configurations attainable by the excitation of two spins, either forming two isolated doubles or a triplet state,

 $\cdots 110 \cdots 110 \cdots 100 \cdots$

 $\cdots 111 \cdots 100 \cdots 100 \cdots$

For the East model it is possible to classify each lattice site according to its local neighbourhood. Taking a site along with its nearest and next-nearest neighbour to the right, each site can be classed according to a binary labelling scheme, i.e. $100 \equiv 4, 110 \equiv 5$, etc., where the number of sites in each class is, N_4 , N_5 , etc. Using this notation we define the entry condition for the algorithm with s = 2 transient states as the point at which the number of sites in class 4 equals the total number of excitations present in the lattice, M, i.e. $N_4 = M$.

Before constructing the transient and recursive matrices it is necessary to determine the probabilities for all possible transitions between the different states. The transient and recursive states may be labelled as follows,

$\cdots 100 \cdots 100 \cdots$	v_1
$\cdots 110 \cdots 100 \cdots$	v_2
$\cdots 110 \cdots 110 \cdots$	u_1
$\cdots 111 \cdots 100 \cdots$	u_2

with the following transition probabilities

$$P(v_1 \rightarrow v_2) = \frac{\epsilon N_4}{N},$$

$$P(v_2 \rightarrow v_1) = \frac{1}{N},$$

$$P(v_1 \rightarrow u_1) = 0,$$

$$P(v_1 \rightarrow u_2) = 0,$$

$$P(v_2 \rightarrow u_1) = \frac{\epsilon (N_4 - 1)}{N}$$

$$P(v_2 \to u_2) = \frac{\epsilon}{N},$$

where N is the system size.

These transition probabilities are then used to build the transient and recursive matrices for the system

$$\mathbf{T} = \begin{pmatrix} 1-x & x \\ y & 1-x-y \end{pmatrix}, \qquad (2.17)$$

$$\mathbf{R} = \begin{pmatrix} 0 & 0 \\ x - \epsilon y & \epsilon y \end{pmatrix}, \qquad (2.18)$$

where $x = \frac{\epsilon N_4}{N}$ and $y = \frac{1}{N}$.

The absorption probabilities for the u_1 and u_2 states can be found by taking the fundamental matrix, **N**, and solving equation (2.14) giving

$$P(u_1) = 1 - \frac{1}{N_4}, \tag{2.19}$$

$$P(u_2) = \frac{1}{N_4}, (2.20)$$

where we have used an initial state vector $\vec{v}_I^T = (1 \ 0)$.

To determine the exit time one must choose a random number and iteratively solve the inequality given in equation (2.16). One then proceeds to choose an exit state from the distribution formed by the exit probabilities, equations (2.19) and (2.20). It is clear that matrices \mathbf{T} and \mathbf{R} are characterised by the variable N_4 and as such both the probability distribution for the absorption states and the exit time are governed by the entry state, each state having its own unique solution.

This s = 2 construction provides an algorithm that improves on standard continuous time, s = 1, by a factor proportional to e^{β}/N_4 . This improvement in computational speed is offset by the algorithmic complexity required to formulate the s = 2 model.

2.4.3 Approximations for the Update Time

Computationally, the most expensive part of the algorithm as described above is the procedure used to determine the time to exit from the transient state. To perform the calculation exactly involves diagonalising the **T** matrix and iteratively solving the inequality using the halving method [32] or something similar. There are, however, a number of approximations that we can employ to get around this. The exact form for $\vec{v}_I^T \mathbf{T}^m \vec{e}$ for the s = 2 case is

$$\vec{v}_I^T \mathbf{T}^m \vec{e} = \frac{1}{2} \left[\lambda_2^m + \lambda_1^m - (\lambda_2^m + \lambda_1^m) \left(\frac{1}{4z} + z \right) \right], \qquad (2.21)$$

where $z = \frac{1}{2}\sqrt{1 + 4\epsilon N_4}$ and λ_1 , λ_2 are the eigenvalues of **T**. Both eigenvalues are quite close to (and less than) 1. However, in the limit of large m, we have that $(\lambda_2/\lambda_1)^m \ll 1$, allowing us to simplify equation (2.21). If we drop the restriction that m must be discrete then we can write (2.16) as an equality,

$$m \approx \log\left(\frac{2\overline{r}}{1+z+1/4z}\right) / \log(\lambda_1),$$
 (2.22)

where again \overline{r} is a random number between 0 and 1. Both z and λ_1 depend on N_4 and can be stored in a lookup table. The approximation works best when m is large, so for low temperatures (T < 1) where the time steps are larger it works very well. At higher temperatures one must be careful using this approach.

Another possibility is to free oneself from the requirement to pick the update time from a distribution and use instead the average. This does mark a departure from the exact Monte Carlo algorithm, but in most cases it turns out to be a reasonable simplification (it is analogous to the approximation made when going from the n-fold algorithm [29] to the CT one [32]). If we take the average time, then we can use equation (2.15) which requires calculation of the fundamental matrix, \mathbf{N} , either analytically or numerically. For the East model, \mathbf{N} only depends on the number of excitations M, so the time updates can be stored in a lookup table allowing for a significant increase in speed.

To check the validity of using the average value for time updates instead of picking them from a distribution, we can use the result

$$\langle \tau^2 \rangle = \vec{v}_I^T (2\mathbf{N}^2 - \mathbf{N})\vec{e}, \qquad (2.23)$$

with (2.15) to calculate the mean square fluctuations. This shows that for lower temperatures the error on any given measurement is $\sim \langle \tau \rangle$. Whilst this seems large it is important to remember that we are always looking at logarithmic time and on this axis the error is less significant. Also there are many iterations between sampling points and the measurements are averaged over many runs which will help to reduce any discrepancy. All the simulations for this paper were performed using the average time update.

2.4.4 Generalisation to Any Dimension

The method described in the previous section can easily be extended allowing one to construct generalised transient and recursive matrices for the East model in any spatial dimension d. Considering the transient states for the system it is clear that the s = 2 algorithm is triggered when all excitations within the lattice are isolated by a region of space which encompasses all moves attainable by two successive spin flips. The d = 2 analog of the "100" above is:

$$\begin{array}{rcrcrcr}
0 \\
100 \equiv & 0 & 0 \\
& 1 & 0 & 0
\end{array}$$

where no triangles may overlap if the algorithm is to trigger. As for the d = 1 case, the **T** and **R** matrices are obtained by evaluating the probabilities for all transitions between

the transient and absorbing states. This analysis yields matrices of the following form

$$\mathbf{T} = \begin{pmatrix} 1-xd & xd \\ y & 1-y-xd - (d-1)\epsilon y \end{pmatrix}$$
(2.24)

$$\mathbf{R} = \begin{pmatrix} 0 & 0\\ xd - \epsilon y & \epsilon yd \end{pmatrix}$$
(2.25)

It is easy to solve for the exit probabilities to each of the absorbing states, where N_4 again indicates the number of excitations in the isolated state,

$$P(u_1) = 1 - \frac{d}{d - 1 + N_4 d}$$
(2.26)

$$P(u_2) = \frac{d}{d - 1 + N_4 d}.$$
 (2.27)

The average lifetime to exit from the transient subspace may be obtained from (2.15),

$$\langle \tau \rangle = \frac{e^{2\beta} + (2N_4d + d - 1)e^{\beta}}{N_4d(N_4d + d - 1)}.$$
(2.28)

For systems in which β is high and N_4 is large compared to d, one may approximate the exit time to

$$\langle \tau \rangle \approx \frac{e^{2\beta}}{(N_4 d)^2},$$
(2.29)

hence time steps are a factor of d^2 smaller than those of the East model in d = 1.

2.4.5 FA-East Crossover model

The MCAMC algorithm described above can also be applied to the FA model [22,33] and, more generally, to the model that interpolates between the FA and East models [34], which serves as a simple model for fragile-to-strong transitions. This model is characterised by the rates

$$11 \xrightarrow{b} 01, \ 01 \xrightarrow{b\epsilon} 11, \ 11 \xrightarrow{(1-b)} 10, \ 10 \xrightarrow{(1-b)\epsilon} 11.$$
 (2.30)

The limit $b \to 0$ corresponds to the East model, and the limit $b \to 1/2$, to the FA model. For intermediate values of b the model displays a crossover between East-like dynamics at higher temperature, to FA dynamics at low temperature.

The MCAMC algorithm is applied in much the same way as in the East model case, except that now we have to allow for the possibility of movement to the *west*. In the simplest version, the transient states are the same as for the East model, and the absorption states are increased to include any move to the left. The result is an s = 2, r = 5 absorbing Markov chain, with the following transient states

$$0100\cdots 0100\cdots v_1,$$

 $0110\cdots 0100\cdots v_2,$

and absorbing states

The two states in u_4 are the same as far as this algorithm is concerned. Caution would be required if they were to be used as transient states. The transient and recursive matrices are as follows

$$\mathbf{T} = \begin{pmatrix} 1 - xd & xda \\ ya & 1 - y - xd - a(d-1)\epsilon y \end{pmatrix}$$
(2.31)

$$\mathbf{R} = \begin{pmatrix} 0 & 0 & bxd & 0 & 0\\ a(xd - \epsilon y) & a\epsilon dy & 0 & bxd & by \end{pmatrix}$$
(2.32)

where $a \equiv 1 - b$, and x and y are the same as for the East model case.

From here on the procedure is exactly the same as with the East model except that there are more absorbing states to choose from. One may use Eq. (2.14) to obtain values for the absorption probabilities. Caution is required as the approximation breaks down in the regime of high temperature and high symmetry $(b \rightarrow 1/2)$. It should be noted that, while less striking, even in the FA limit of b = 1/2 this algorithm outperforms standard CT.

2.4.6 Higher order MCAMC

The entirely isolated state is problematic in terms of the dynamics of the East model. In order to relax the isolated excitations must propagate in the lattice until they encounter another excitation along the direction of facilitation. Movement of this nature is promoted by the occurrence of branching events,

$$100 \rightarrow \cdots \rightarrow 111 \rightarrow 101$$

The "triplet" absorption state, u_2 , is the rate limiting step for branching events, and hence the propagation of excitations in the lattice. However, from the absorption probabilities, Eqs. (2.19) and (2.20), we see that compared to the u_1 state the u_2 exit state is suppressed by a factor of $\frac{1}{N_4-1}$. Hence, the formation of triplets is unlikely, particularly when the system size is large.

To overcome this problem it is possible to extend the MCAMC algorithm to include the u_1 state as a transient state of the system. There are now three transient and three absorbing states, one absorbing state corresponds to the u_2 state of the s = 2 algorithm, the remainder corresponding to configurations attainable from the new transient state. In one dimension one can represent the states as follows,

$$\cdots 100 \cdots 100 \cdots 100 \cdots v_1, \\ \cdots 110 \cdots 100 \cdots 100 \cdots v_2, \\ \cdots 110 \cdots 110 \cdots 100 \cdots v_3, \\ \cdots 110 \cdots 110 \cdots 110 \cdots u_1, \\ \cdots 111 \cdots 100 \cdots 100 \cdots u_2, \\ \cdots 111 \cdots 110 \cdots 100 \cdots u_3.$$

Once again the transient and recursive matrices may be constructed by considering all possible transitions between the states.

$$\mathbf{T} = \begin{pmatrix} 1-x & x & 0 \\ y & 1-x-y & x-\epsilon y \\ 0 & 2y & 1-2y-x \end{pmatrix}$$
(2.33)
$$\mathbf{R} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \epsilon y & 0 \\ x-2\epsilon y & 0 & 2\epsilon y \end{pmatrix}.$$
(2.34)

The s = 2 transient matrix, Eq. (2.17), is now a submatrix of **T**; the addition of an extra transient state has appended one extra row and column to the matrix, the rest of the structure remaining intact.

Unlike the case of s = 2, it is not so simple to generalise the s = 3 matrices for any dimension. This arises from the non-equivalence of the v_3 state in dimensions d > 1, i.e. in two dimensions

0				0				0		
1	0		\neq	1	0		• • •	1	0	
1	1	0		1	0	0		1	0	0

When considered as absorption states the two configurations above may be treated identically since the probability of exiting to each state is the same. However, as transient states each configuration has different exit probabilities and as such must be treated independently. In essence, one requires an s = 4 algorithm to provide the equivalent result in dimension two and above.

Returning to the d = 1 example, we find that the u_2 is now the most likely absorption state. This is because all other exit states require the excitation of an additional spin, i.e., they are suppressed by a factor of e^{β} . Solving for the average lifetime gives

$$\langle \tau \rangle \approx \frac{e^{2\beta}}{N_4}.$$

Hence, s = 3 improves on s = 2 by a factor of N_4 . While this may seem a modest en-

hancement in performance, note that the extra algorithmic complexity required to develop s = 3 is negligible. Having made a working s = 2 algorithm one may essentially use s = 3 for free.

As CT enables one to obtain an e^{β} speed increase over traditional MC, s = 2 enables one to achieve a further improvement of e^{β} over CT. In effect, s = 2 enables one to bypass all e^{β} processes (i.e. those that involve the excitation of a single spin) by insisting that two successive spins are excited. The double and triplet states of the s = 2 model are examples of $e^{2\beta}$ processes. In order to construct an algorithm with a further e^{β} speed gain requires one to identify all of the $e^{2\beta}$ arrangements and include them as transient states. This means that the absorption states now correspond to all configurations attainable from the transient states which result in the simultaneous excitation of three spins. This analysis leads to an s = 7 algorithm consisting of seven transient and seven absorbing states.

In one dimension, s = 7 may be triggered when all excitations within the lattice are separated by at least three unexcited spins, i.e. $1000 \cdots 1000$. To maximise performance it is useful to use a hybrid algorithm consisting of s = 1, 3 and 7 components with each sub-algorithm activated by its own triggering condition.

In order to improve algorithmic efficiency it is convenient to compute absorption probabilities using Eq. (2.14) rather than the exact form of Eq. (2.12). Unlike the case of s = 2, where it may be shown that two expressions are identical, for higher order algorithms the solution of Eq. (2.14) only provides an approximation. In general one must employ caution when using this approach. For both s = 3 and s = 7 it has been shown that the approximation is good for all regimes in which the algorithms are effective, the approximation breaking down at higher temperatures.

2.4.7 Speed Tests

In the low T limit, the average exit time from an s = 2 MCAMC iteration for the East model is approximately $e^{2\beta}/(N_4d)^2$. The corresponding average time step for standard CT is e^{β}/N_4d . The s = 2 time step becomes larger by a factor of e^{β}/N_4d . It gets a speed-up from e^{β} , and a slowdown from N_4 , as the more excitations that are present upon entering the algorithm the quicker it exits, and from d, as the higher the dimension the more facilitated sites are available. At low temperature, however, $N_4 \approx N\epsilon$, so for fixed system size N the speed-up factor of s = 2 MCAMC with respect to CT grows as $e^{2\beta}$.

In figures 2.3 and 2.4 we show speed tests comparing the performance of the MCAMC algorithms to standard MC and CT on East model simulations. Fig. 2.3 shows the temperature dependence of the CPU time required for generating an equilibrium trajectory of total Monte Carlo time $10^7 \times e^{2\beta}$ in an East model of $N = 10^5$ sites. In an s = 1 CT algorithm the average CPU time for such a simulation is independent of T. Fig. 2.3 shows that at very high temperatures standard MC is the fastest method, but as T is lowered CT soon outperforms it. At lower temperatures s = 2 MCAMC becomes more efficient than CT by a approximately a factor of $e^{2\beta}$. As the temperature is dropped further, s = 7 MCAMC provides a further improvement of approximately $e^{2\beta}$, and so on.

As discussed above, the efficiency of MCAMC depends on the system size. In addition to a reduced time step this also determines the probability of encountering the isolated



Figure 2.3: Temperature dependence of CPU time for equilibrium East model trajectories of total Monte Carlo time $t = 10^7 \times e^{2\beta}$ and system size $N = 10^5$, for MC, CT, and MCAMC algorithms. The straight lines indicate the approximate speed-up of the MCAMC simulations. CPU time shown relative to the average time needed when using CT dynamics.



Figure 2.4: System size dependence of CPU time (relative to that for CT) for equilibrium East model trajectories of MC time $t = 3 \times 10^{12}/N$ at T = 0.2.



Figure 2.5: Persistence time τ_{α} as a function of inverse temperature $\beta = 1/T$ in the East model in dimensions d = 1 - 5, 10, 13. The lines through the data points are quadratic fits, $\log \tau_{\alpha} = a_0 + a_1\beta + a_2\beta^2$, with a_i fitting parameters. The fit suggests that $a_1 \sim 1$ in general. The bottom-right panel shows a_2 as a function of d. This coefficient seems to go as $a_2 \approx b/d$, with the constant $b \approx 0.8$ (shown as a full line). This value is between $(\ln 2)^{-1}$ (dotted line) and $(2 \ln 2)^{-1}$ (dashed line).

entry state for the s > 1 algorithms. Fig. 2.4 shows the CPU time, now for different system sizes, at fixed temperature and total MC time $t = 3 \times 10^{12}/N$. Again, the CPU time for such a simulation using CT is constant. As expected, Fig. 2.4 shows that as Nbecomes larger the MCAMC algorithms are less and less effective; beyond $N\epsilon^2 \approx O(1)$ the CT scheme works better. This means that in order to maximise the MCAMC efficiency one needs to simulate the smallest possible system sizes. This is limited by the need to be compatible with bulk behaviour, which in the case of facilitated models requires that the system in average contains a sufficient number of excitations, i.e., $N\epsilon$ cannot be too small.

2.4.8 Example of results

In this section we present an example of numerical results obtained with the MCAMC.

A useful correlation function to study the relaxation of facilitated models is the persistence function P(t), e.g. [34–36], which gives the probability that a site has not changed its state up to time t. In terms of the local persistence field $p_i(t) = 0, 1$, where 1 indicates that site i has not flipped up to that time, and 0 that it has flipped at least once, the persistence function reads $P(t) = N^{-1} \sum_i p_i(t)$. In contrast to standard MC or CT simulations, the MCAMC algorithm could run into problems when trying to measure persistence. By construction, it misses some of the events that could occur whilst in the transient sub-



Figure 2.6: Concentration of excitations c(t) as a function of scaled time $T \ln t$, in the d = 1 East model after a quench from infinite temperature, from simulations with s = 7 MCAMC.

space, for example, from the isolated state many spins could flip up and then flip back down before finally exiting to an absorbing state. At low temperatures in equilibrium, however, the contribution of these events is negligible, as the vast majority of changes to the global persistence function is from existing excitations spreading out into unmoved territory, and this is captured by the MCAMC algorithm. In fact, the only sites one needs to be concerned with are those immediately next to the initial excitations, which are very few in equilibrium at low T.

Figure 2.5 shows the equilibrium persistence time [35,37], τ_{α} , of the East model in various dimensions d, calculated using the MCAMC algorithm with s = 2. For all dimensions studied we find that τ_{α} is a super-Arrhenius function of T. This seems to indicate that the East model is a fragile in all dimensions. Given that any simple mean-field estimate of the relaxation in this model would give Arrhenius behaviour, the above result would suggest that the East model has no upper critical dimension to its dynamics [38]. The data is compatible with $\log \tau_{\alpha} = a_0 + a_1\beta + a_2\beta^2$, as expected if relaxation processes in the East model in any d are quasi one-dimensional [38]. The coefficient a_2 of the quadratic fits is compatible with $a_2 \approx b/d$ [38], with the constant b obeying $(\ln 2)^{-1} \ge b \ge (2 \ln)^{-1}$, reminiscent of the rigorous d = 1 result of Ref. [39]. Within these limits b appears much closer to $(2 \ln 2)^{-1}$ which is consistent with new results from Ref. [40]. Note that the timescales reached with the MCAMC in Fig. 2.5 are between three and five orders of magnitude longer than in previous studies [36].

The MCAMC proves also useful when simulating out-of-equilibrium dynamics. Consider the aging of the East model following a quench from infinite temperature. As the system relaxes towards its equilibrium the dynamics proceeds by stages characterised by the distance between isolated excitations [41]. These domains grow as $\overline{d} \sim t^{T \ln 2}$. Consequently, the isolated transient state also plays an important role in such out-of-equilibrium dynamics of the East model, and the MCAMC algorithm is also applicable in this regime. Figure 2.6 shows the aging of the concentration of excitations, c(t), with time after a quench to low temperatures in the East model, using the s = 7 MCAMC algorithm.

In these aging simulations the nature of the speed-up due to the MCAMC becomes evident. Each stage of the dynamics is associated with an isolated domain of the form $10 \cdots 0$. The k-th stage corresponds to domains of typical length $l \sim 2^k$, and a corresponding energy barrier of size k to further relaxation [41]. In essence, at each successive plateau of c(t) one requires an algorithm that produces time steps comparable to the activation time, $e^{k\beta}$. An s = 2 MCAMC enables one to push simulations one plateau further than CT (s = 1), the s = 7 algorithm helps overcome the next energy barrier, and so on.

2.5 Discussion

In this chapter we have shown that it is possible to dramatically improve the efficiency of numerical simulations of KCMs by using the MCAMC method. Even the simplest s = 2 algorithm can improve simulation times at low temperature by a factor of $e^{2\beta}$ over the n-fold or continuous time MC. By increasing the number of transient states s even larger computational gains can be achieved. There is always a competition between the complexity of the algorithm (the book keeping costs) and the extra efficiency that brings.

So far this competition has been the main stumbling block for applying MCAMC to other models. The *f*-spin facilitated FA model with f > 1 can be an incredibly slow to simulate model, even using CT [22, 27, 42]. Several features of these systems make the application of MCAMC less straightforward: since their kinetic constraints depend on more than one site, i.e. facilitation by two or more excitations in the FA models or two or more vacancies in the lattice gases, for generic entry states the tree of possible transient states is much larger than for, say, East models. This means that the computational cost of the necessary bookkeeping will be much higher (bookkeeping could be simplified by reducing the possible entry states, at the expense of triggering less frequently the MCAMC). This problem is compounded by the fact that f > 2 FA models are very slow even at moderate temperatures, so that the potential exponential in β gains from excitation rates are very modest, and may not even be enough to offset the bookkeeping cost.

In the next chapter we will be studying another class of KCM, the constrained lattice gases [43,44], where barriers are purely entropic. Here we do not get any speed gain from overcoming energy barriers which makes applying an MCAMC an even bigger challenge. Given that at high densities these systems are much slower than the East models a clever MCAMC algorithm which overcomes these hurdles could prove extremely useful.

Chapter 3

A KCM in Three Dimensions: The FCC Constrained Lattice Gas

The facilitated spin models discussed in chapter 2 are very useful for showing how glassy dynamics can arise from the application of local kinetic constraints. The FA model constraints are physically reasonable, although it should be noted that, at this time, there is no process for deriving any particular choice of constraint from microscopic models [45,46]. Perhaps more important is how, on slightly changing the constraint (to the East model), the mode of relaxation drastically changes.

Another branch of kinetically constrained model (KCM) are the constrained lattice gases [22]. In contrast to the FSMs these models retain the notion of particles. Restricting these particles onto a lattice significantly reduces the number of degrees of freedom and thus the complexity. Static interactions between the particles do not exist so that the thermodynamic properties are trivial, in keeping with other KCMs. Kinetic constraints in lattice gas models tend to have a more natural origin. For example, simple exclusion prevents two particles occupying the same site. Beyond this type of constraint there are a number of different models to choose from; Kob and Andersen (KA) introduced a model where a particle could only move if its number of nearest neighbours, before and after a move, were below a threshold [43]. This extra constraint dramatically slows down the dynamics of the system.

In this chapter we review the constrained triangular lattice gas (TLG). Introduced by Jäckle and Krönig [44], it is very similar to the KA model relying on purely geometrical constraints to achieve glassiness. The advantage over the KA model is that the constraints are explicitly steric. The TLG is a two-dimensional model and detailed studies have shown that it displays glassy behaviour such as slowing down and Stokes-Einstein breakdown. Later in this thesis we will be using a version of the TLG to study vibrations in glass formers; when it comes to diffusion or vibrations the dimensionality of the model is very important and this is the motivation to introduce a three dimensional version of the TLG using an FCC lattice. Jäckle and Krönig suggested this in the original TLG paper but



Figure 3.1: An allowed move in the two-vacancy assisted TLG, the (2)-TLG.

up to now there are no detailed studies of the model. After discussing both models new results for the FCC lattice gas (FLG) will be presented alongside reproduced results for the TLG. Both models are used extensively in the remainder of this thesis so this chapter also serves as a prerequisite for chapters 4 and 5.

3.1 The constrained triangular lattice gas

The triangular lattice gas consists of particles on a two-dimensional triangular lattice with diameter no greater than the lattice spacing [44]. Static interactions are absent so the equilibrium configuration is for all sites to be occupied with equal probability determined by the density, ρ . The dynamics proceed by choosing a site at random, if there is a particle there then one of its six neighbouring sites is chosen. If this site is vacant then the particle is moved, otherwise nothing happens. One unit of time is defined as N attempted moves, where N is the number of sites. In this version of the model moves are only dependent on the states of the two chosen sites and no surrounding sites - this is the zero-vacancy assisted TLG, or (0)-TLG.

By only considering the states before and after a move the (0)-TLG takes no account of how neighbouring particles might block any given transition. If the particles are hard discs with a diameter greater than $\frac{\sqrt{3}}{2}a$, where *a* is the lattice spacing, then such moves could involve the overlap of particles. To account for a steric interaction one can introduce an extra kinetic constraint by demanding that one of the mutual neighbours of the target and destination site must also be vacant. The extra vacancy creates enough space for the particle to be moved continuously between the two sites provided that it does not have to follow a straight path. A stricter requirement still is to require both mutual neighbours to be vacant which would allow the straight path to be followed. This variant is the (2)-TLG as shown in Fig. 3.1.

It is only once these kinetic constraints are added that we begin to see some very interesting dynamical behaviour [47] with many of the features discussed in chapter 1 being reproduced. The most important of which is a slowing of the dynamics as density is increased. The dynamics of the (1)-TLG are relatively fast because pairs of vacancies are able to diffuse about the system without much long range cooperation. The (2)-TLG does



Figure 3.2: An Angell plot for the (1)-TLG (circles) (2)-TLG (squares) similar to that of Ref. [47]. Temperature is defined by $1/T \equiv -\ln(1-\rho)$ and T_g is defined as the temperature when τ_{α} reaches 6×10^6 Monte Carlo sweeps, this occurs at $\rho_g = 0.9995$ for the (1)-TLG and $\rho_g = 0.8$ for the (2)-TLG. The fitting forms are $\ln \tau_{\alpha} \propto \Delta/T$ (Arrhenius with $\Delta \approx 2.3$) and $\ln \tau_{\alpha} \propto \exp(a/T)$ respectively.

not have such an easily identifiable excitation and as such the dynamics are considerably slower and more complicated with timescales that increase dramatically with increasing density.

To relate the lattice gas models, that do not have temperature or energy, to the facilitated spin models we follow the scheme from Ref. [47] that the concentration of vacancies, $1 - \rho$, should be linked to the concentration of excitations in the FSMs, $c \approx e^{-\beta}$. This gives us an effective inverse temperature of

$$\frac{1}{T} \equiv \beta \equiv -\ln(1-\rho) \tag{3.1}$$

In this scheme the timescales of the (1)-TLG scale with temperature in an Arrhenius 'strong' manner similar to the FA model and the (2)-TLG scales in a super-Arrhenius 'fragile' manner. In the same way that for the FA and East models a small change in kinetic constraint can dramatically change the behaviour, so too for the TLG models we see very different dynamics. Fig. 3.2 shows the strong and fragile behaviour with an Angell plot using the α -relaxation time, τ_{α} for both models.

Apart from a dramatic slowing down with increasing density, we also see other key features of glass formers, most notably dynamical heterogeneity. The (2)-TLG in particular can have large variations in the structural relaxation times of different parts of the liquid [48]. This is in turn responsible for the breakdown of the Stokes-Einstein relation that relates structural relaxation to diffusion [49]. The appeal of the TLG is that these effects arise from very tangible constraints.
3.2 The FCC lattice gas

3.2.1 Definition of the model

In two dimensions the triangular lattice is the most efficient way to pack circular discs to maximise the density. In three dimensions the FCC lattice and the HCP lattice are just as efficient at packing but the FCC lattice has greater rotational symmetry so this is the favoured lattice for our purposes. Fig. 3.3 shows how the FCC lattice can be thought of as successive planes of triangular lattices that repeat every third layer (so called ABCABC formation). In a coordinate frame where one of the triangular lattice layers falls in the x-y plane the unit vectors that we use are

$$\vec{n}_1 = (1 \ 0 \ 0), \ \vec{n}_2 = \left(\frac{1}{2} \ \frac{\sqrt{3}}{2} \ 0\right), \ \vec{n}_3 = \left(0 \ \frac{1}{\sqrt{3}} \ \sqrt{\frac{2}{3}}\right)$$
 (3.2)

and the 12 nearest neighbours are then made from combinations $\pm \vec{n}_1$, $\pm \vec{n}_2$, $\pm (\vec{n}_1 - \vec{n}_2)$, $\pm \vec{n}_3$, $\pm (\vec{n}_3 - \vec{n}_2)$, $\pm (\vec{n}_3 - \vec{n}_2 + \vec{n}_1)$.

Any two neighbouring sites share four mutual neighbours and the equivalent steric constraint of the (2)-TLG requires that all four of these neighbours are vacant, we refer to this as the (4)-FLG. One can see how this constraint works in Fig. 3.3. Another way to establish the constraint is to consider the dot product between any two neighbours, $\vec{n}_i \cdot \vec{n}_j$. If this dot product is positive then a movement towards *i* would also involve, at least initially, moving towards *j*. With hard spheres that are tightly packed this would not be possible. All blocked moves in the (4)-FLG turn out to have a dot product of $\vec{n}_i \cdot \vec{n}_j = 1/2$, this is a nice feature of the symmetry that no site is more important in blocking a move than any other and there is no preferred plane. Softer versions of the model could require n < 4 mutual neighbours to be vacant and in general are reffered to as (n)-FLGs.

The dynamics of the FLG then proceeds in exactly the same manner as the TLG. Once again there are no static interactions so the equilibrium configuration is for a fixed number of particles to be placed at random. The numerical results in this chapter were carried out using periodic boundary conditions with a box size, L = 50, containing 50^3 lattice sites. The dynamics were simulated using a rejection free Monte Carlo algorithm (see chapter 2) that chose time updates from a distribution to preserve short time dynamics as well as efficiently probe long times. Distances are quoted in units of the lattice spacing and time in Monte Carlo sweeps.

3.2.2 Backbones

Part of the reason that the (2)-TLG is so slow is that large structures can exist that are essentially frozen until they can be unlocked by a long series of coordinated moves. In a finite system these structures can wrap around the periodic boundary conditions and form an unmovable backbone. The simplest such structure is an unbroken straight line of particles. Jäckle and Krönig showed that in the limit of infinite system size these backbones do not grow fast enough to span the system and there is thus no dynamic phase transition at $\rho_c < 1$ [44,50]. For practical purposes, when running numerical simulations, it is not possible to go above a density of $\rho \gtrsim 0.81$ as backbones become larger than any



Figure 3.3: The nearest neighbours of a target particle (red) on an FCC lattice. (a) shows the TLG plane with the nearest neighbours from the adjacent layers (shifted vertically for clarity). (b), from a different angle, shows the four particles that would block a movement into a vacancy out of the page.



Figure 3.4: The size of backbones, L, at density, ρ , in the (4)-FLG. Points represent the density at which the fraction of systems that were jammed, f_J , crosses 0.5. Error bars indicate where $0.45 \leq f_J < 0.55$. Inset is the same plotted against L^{-1} , if there exists a dynamic phase transition then this line will touch zero at some density $\rho_c < 1$.

reasonable simulation box.

In three dimensions the (4)-FLG also has such backbone structures. Similar results for the three dimensional KA model suggest that will be no dynamic transition for $\rho_c < 1$ [50]. For finite numerical simulations there is certainly a limit on the system size for any given density. To make sure that we always operate well above this size it is important to know the size of backbones as a function of density. To do this we took random configurations at variable box lengths, L, and tested for backbones using the following procedure: all particles that are able to move are removed from the system. If in doing this any new particles become unblocked then they too are removed. This is repeated recursively until no more particles can be removed (there is a backbone) or there are none left (no backbone). This is done for many systems and the fraction of jammed to not-jammed systems, f_J , is measured.

For a fixed density, ρ , the fraction of systems with a backbone goes through a sharp transition between 0 and 1 at a particular length, L. The crossover point as a function of density is plotted in Fig. 3.4 and serves as a guide as to how big a simulation box one needs to use. It is not possible to extrapolate beyond the data to prove there is no dynamic transition. However, the highest density that was possible to be simulated was $\rho = 0.69$ and this is limited by the slow dynamics, not the box size.



Figure 3.5: The self-intermediate scattering function, $F_s(k, t)$, at wave vector $k = \pi$ and the mean squared displacement $\langle \Delta \vec{r_i}(t)^2 \rangle$. (a) and (c) on the left are for the (4)-FLG at densities $\rho = 0.05 - 0.65$, 0.68 in steps of 0.05 and (b) and (d) on the right are for the (2)-TLG with $\rho = 0.05 - 0.8$ again in steps of 0.05.

3.2.3 Two-point correlation functions

A benefit of the lattice gas models is the ability to make microscopic measurements relating to the displacement of particles that are easily comparable to experiments and more realistic simulations. The most common measure of structural relaxation is the self-intermediate scattering function, $F_s(\vec{k}, t)$, defined as

$$F_s(\vec{k}, t) = \frac{1}{N} \sum_i e^{i\vec{k} \cdot (\vec{r}_i(t) - \vec{r}_i(0))}$$
(3.3)

where \vec{k} is the wave vector and \vec{r}_i is the position of particle, *i*, at time *t*. The wave vector, \vec{k} , sets the lengthscale on which the particles have to move before the correlation falls to zero. The α -relaxation time, τ_{α} , is usually defined as the time for $F_s(\vec{k}_0, t)$ to fall to e^{-1} at a wave vector, \vec{k}_0 , that corresponds to the size of the particles. For the lattice gases \vec{k}_0 falls along the reciprocal lattice vectors with magnitude $k_0 = \pi$.

The self-intermediate scattering function for the (4)-FLG (Fig. 3.5(a)) is qualitatively similar to the (2)-TLG and other real glassy systems [6, 51]. At high densities it becomes increasingly stretched and for late times can be fitted with a stretched exponential $F_s(\vec{k}_0, t) \approx \exp[-(t/\tau_\alpha)^\beta]$, where here β is a density dependent constant. Measurements



Figure 3.6: α -relaxation time, τ_{α} , against density ρ . The fitted line is of the form $\tau_{\alpha} \propto \exp[\exp[\exp[-a\ln(1-\rho)]]]$. Inset shows the relaxation time as a function of effective inverse temperature, $\beta \equiv -\ln(1-\rho)$. For the (2)-TLG (squares) the fit is a double exponential and for the (4)-FLG (circles) the fit is a triple exponential.

of the structural relaxation time, τ_{α} , reveal that the (4)-FLG is super-Arrhenius with timescales growing even faster than the (2)-TLG, making it even more fragile (Fig. 3.6).

Naive comparison to the KA model in multiple dimensions suggests that the timescales might scale as a triple exponential, $\tau_{\alpha} \propto \exp[\exp[\exp[-a \ln(1-\rho)]]]$ [50]. This fit does appear to work well for the available data (solid lines in Fig. 3.6) although higher density results (involving very large timescales) would be necessary to be sure. The (2)-TLG, for comparison, is fitted well with a double exponential.

Another revealing two-point correlation function is the mean-squared-displacement, $\langle \Delta \vec{r_i}(t)^2 \rangle$, where $\Delta \vec{r_i}(t) = \vec{r_i}(t) - \vec{r_i}(0)$. Once again the (4)-FLG is in qualitative agreement with the (2)-TLG showing diffusive behaviour at low densities and then showing crossover behaviour at higher densities (Fig. 3.5 bottom). At short times the dynamics are dominated by the fast particles moving short distances and later we enter a subdiffusive regime where they become trapped inside dynamical heterogeneities. Finally at long times, comfortably beyond τ_{α} , system-wide diffusive motion is recovered. The long time displacement can be used to extract the self-diffusion constant

$$D_s \equiv \lim_{t \to \infty} \frac{\langle \Delta \vec{r_i}(t)^2 \rangle}{4t} \tag{3.4}$$

and Fig. 3.7 shows this plotted against effective temperature for the (4)-FLG and the (2)-TLG. The self-diffusion constant is a characteristic timescale of the liquid and, like τ_{α} , rapidly changes with increasing density. As we will see in the next section the two do not scale with density in quite the same manner showing that, for a supercooled liquid,



Figure 3.7: Self-diffusion constant, D_s , against effective inverse temperature, β .

motion cannot be explained simply in terms of one timescale.

3.2.4 Stokes-Einstein breakdown

Liquids above their melting point have homogeneous dynamics and the structural relaxation time, τ_{α} , also sets the length of time needed for a diffusive step to be uncorrelated with any previous steps. This assertion is the basis of the Stokes-Einstein (SE) relation that says that the structural relaxation time and the diffusion constant are inversely related, $D_s \propto \tau_{\alpha}^{-1}$.¹

As we have already seen, both diffusion and structural relaxation dramatically slow down in the supercooled regime. However, the SE relation fails on approach to the glass transition with diffusion appearing to be relatively enhanced [53–55]. This phenomena is a direct consequence of dynamical heterogeneity. Particles in fast parts of the liquid are able to translate further than the slow regions and this has a much greater effect on the mean squared displacement than it does on τ_{α} (which is insensitive to distances further than the particle spacing). Both the 2D and 3D models reproduce SE breakdown at different crossover densities as shown in Fig 3.8. The product $D_s \tau_{\alpha}$ is more or less constant for low densities but at higher densities increases by two orders of magnitude.

Experimentally it is observed that when the SE relation breaks down it crosses over into a fractional SE relation, $D_s \propto \tau_{\alpha}^{-\xi}$, where ξ is a constant, $\xi < 1$ [53]. This is also observed in the lattice gas models and Fig. 3.9 shows that the (2)-TLG and the (4)-FLG crossover to exponents of $\xi \approx 0.59$ and $\xi \approx 0.665$ respectively. That the exponent, ξ ,

¹The Stokes-Einstein relation is normally expressed as $D\eta = \text{const}$ where η is the reduced shear viscosity $\eta \equiv \eta_s/T$, η_s being the shear viscosity. Experimental evidence suggests that $\tau_{\alpha} \propto \eta$ which gives our definition for the SE relation [52, 53].



Figure 3.8: Breakdown of the Stokes-Einstein relation, $D_s \tau_{\alpha} = \text{const}$, at high densities in both the (2)-TLG and the (4)-FLG.



Figure 3.9: Fractional Stokes-Einstein relationship between the self-diffusion constant, D_s , and the structural relaxation time, τ_{α} . The (2)-TLG (squares shifted by factor of 10 for clarity) has an exponent of -0.59 and the (4)-FLG (circles) has an exponent of -0.665. Marked points represent possible crossover densities for the two models.

is larger for the three dimensional model is to be expected as the increased connectivity of an extra dimension will dampen fluctuations. That it is not larger still is also quite interesting and is the result of the highly hierarchical nature of the relaxation. When the kinetic constraint is softened the decoupling of diffusion and structural relaxation is much weaker. The (1)-TLG has $\xi = 0.88$ [47] and results (not shown here) for the softer FLG models give a smaller variation of $\xi \approx 0.96$, 0.9 and 0.87 for the (1),(2) and (3)-FLG respectively. The extra constraint of the (4)-FLG therefore makes a big difference to modes of relaxation. This behaviour is also seen in the spin facilitated models of Ref. [56]. Exponents for real liquids fall somewhere in between the (4) and (3)-FLGs at around $\xi \approx 0.77$ [53, 55].

Spin facilitated models provide a useful insight to the decoupling of timescales by adding probe particles that interact with the mobility field [57]. In this simplified picture the decoupling of diffusion and structural relaxation is equivalent to the decoupling of persistence times (the time it takes to move for the first time) and exchange times (the times between subsequent moves). With the lattice gases it can be confirmed that persistence and structural relaxation times are equivalent. The two-point correlation, p(t), which measures the fraction of particles that have not yet moved, is the same as $F_s(q = \pi, t)$ for later times. The exchange times on the other hand do not simply relate to diffusion due to the effect of rattling.

3.2.5 Persistence and exchange times

So far we have only talking about consequences of dynamical heterogeneity. Stokes-Einstein breakdown and stretched exponentials can be seen in experiments and explanations invoking heterogeneous dynamics are compelling. To gain more confidence that this is the right explanation we turn to the kinetically constrained models that allow us to directly observe broadly distributed dynamics [13]. For this we use the distribution of persistence times, $\Pi(\log(t))$. From a given initial configuration a particle's persistence time, t_p , is the time taken to make its first move. The distribution of these times, shown in Fig.3.10, shows us the range of time scales that different parts of the system are taking to rearrange. After moving for the first time the exchange times, t_x , give the times between all subsequent moves. The distribution of exchange times, $X(\log(t))$, therefore gives an indication of the amount of activity at a given time scale.

At very low densities the two distributions are more or less identical [56] implying that all steps are independent, a statement that is in the spirit of Stokes-Einstein. As the density increases persistence and exchange decouple with the former becoming broadly distributed over six orders of magnitude in time for the (4)-FLG (see Fig. 3.10). By the time we reach the peak in the exchange time distribution most of the particles are still jammed and will remain so for a long time. The exchange dynamics are therefore dominated by a relatively small number of fast particles. The distributions for the (2)-TLG look qualitatively the same although it should be noted that in Ref. [47] Pan *et al.* use the *site* persistence, which monitors events at lattice sites, and here we use the *particle* persistence, this accounts for a slightly different shape in the distributions.

The crossover from simple dynamics to glassy dynamics is a fairly smooth process over



Figure 3.10: Persistence (a) and exchange time distributions (b) for the (4)-FLG at densities of $\rho = 0.05, 0.15, 0.25, 0.35, 0.45, 0.5, 0.55, 0.6, 0.65$ and 0.68. The thicker brown line for $\rho = 0.5$ indicates a possible crossover density.

a range of densities so it is difficult to pinpoint a well defined onset-density. The plots of Fig. 3.9 and 3.10 do, however, provide some visualisation of the process. In Fig. 3.9 the points that correspond to densities of $\rho = 0.45$ and $\rho = 0.5$ in the (4)-FLG do not quite fall on any of the fitted lines for SE or fractional-SE behaviour. Around these densities the persistence time distributions are beginning to show a shoulder developing and most notably the exchange time distribution at $\rho = 0.5$ has the broadest distribution of any density (Fig. 3.10). This would indicate that at $\rho = 0.5$ we are starting to see a little of all dynamical behaviours, heterogeneous and homogeneous, but neither is yet dominating.

It is even more difficult to pick out a crossover density for the (2)-TLG. Fig. 3.9 would seem to suggest it is somewhere in the region of $\rho = 0.6$, but it is a much smoother transition. The relevance of these densities will become apparent later on in chapter 5 when we consider rigidity transitions in the same models.

3.3 Conclusion

Kinetically constrained models reproduce all the essential features of a glass former with the fewest possible ingredients. In this chapter and the previous chapter we've reviewed two of the most important classes of KCM that are able to drastically change the dynamical behaviour with small changes to these ingredients.

The triangular lattice gas of Jäckle and Krönig is a particularly neat model that contains very intuitive geometrical constraints. While the three dimensional version has been proposed, numerical results for the dynamics have never been performed. In this chapter we have confirmed that the dynamics are qualitatively the same as for two dimensions. The (4)-FLG is more fragile than the (2)-TLG and glassiness sets in at a lower density. At high densities the (4)-FLG is very dynamically heterogeneous and displays fractional Stokes-Einstein behaviour $D_s \propto \tau_{\alpha}^{-0.665}$.

Understanding the dynamics of the (4)-FLG in detail is useful because in chapter 5 we will be using it alongside the (2)-TLG to study vibrations in glassy systems. Here having a three dimensional model for comparison is important because for vibrations dimensionality can be very important. For example in two dimensions one sees divergences in the Debye-Waller factor that can be problematic. Unlike other three dimensional models, such as the KA model, the explicit steric constraints of the (4)-FLG provide a natural way to build vibrations into the model.

Chapter 4

Fast Degrees of Freedom

The usual procedure in constructing a kinetically constrained model is to either ignore, or effectively integrate out, the fast degrees of freedom in such a way that they leave their mark on the slow degrees of freedom only through the kinetic constraints. This is a bold move but done in the right way has proved very successful at describing the structural relaxation of a broad range of glass formers. It does of course completely miss all of the short time dynamical behaviour that, as we have seen in chapter 1, can be rich in its own right.

Natural questions that arise include: how do the dynamical heterogeneities, which are seen in the long time dynamics, relate to the faster processes? Are these faster processes heterogeneous as well, and if so how does this relate to the structure of the liquid? The aim of this chapter is to reintroduce fast degrees of freedom to the KCMs in an attempt to recover some of the early time dynamics that are usually ignored. In the spirit of the simple model we will be attempting to add as little as possible to do this. First we review what behaviour the model is aiming to produce and then we discuss different schemes for including fast degrees of freedom. In section 4.2 a model that adds coupled Ising variables to the facilitated spin models (FSMs) is proposed. To check how the more complicated interactions affect the underlying FSM the thermodynamics are solved in one dimension. Section 4.3 extends the model to two dimensions and also onto the triangular lattice gas (TLG). Here the thermodynamics resemble closely the well studied Blume-Capel model. While the static properties cannot be solved exactly we can employ numerical cluster techniques to study the phase behaviour. Finally scrutiny is turned to the dynamics and specifically the interplay between the different timescales and the correlations that arise.

4.1 Theories of fast dynamics

For a simple liquid at high temperature, the number of distinct time scales is restricted to a ballistic regime at very short times, where particles move freely between collisions, and a diffusive regime at longer times. On a log-log plot of mean squared displacement (MSD), $\langle |\mathbf{r}(t) - \mathbf{r}(0)|^2 \rangle$, against time (as in Fig. 4.1(a) at T = 5.0 [6]), this shows up as two straight lines corresponding to $\langle |\mathbf{r}(t) - \mathbf{r}(0)|^2 \rangle \propto t^2$ at early times, and later crossing



Figure 4.1: (a) Mean squared displacement of A-particles in a binary Lennard-Jones mixture (taken from Kob [6]), from high temperature regime at T = 5.0 to a glassy regime at T = 0.466. Units are defined relative to the interaction energy of the particles. (b) A single particle trajectory from a dense colloidal suspension, tracked using conformal microscopy (taken from Weeks et al. [9]) over 100 minutes, showing the caging effect.

over to $\langle |\mathbf{r}(t) - \mathbf{r}(0)|^2 \rangle \propto t$ when the dynamics become diffusive [58]. In a real liquid this crossover is usually complete by the picosecond time scale.

With supercooled liquids, instead of crossing over to diffusion, the particles become trapped in the cages of their nearest neighbours - who are in turn trapped in their own cages. On the MSD plot this shows up as a plateau covering many decades of time where the particles make very little progress away from their starting position. After a long time 'rattling' in these cages the particles eventually begin to find a way out and structural relaxation (the α -relaxation) can occur. A dramatic illustration of this process is shown in Fig. 4.1(b) for a tagged particle in a dense colloidal mixture.

Finally, on time scales even longer than the α -relaxation (due to Stokes-Einstein breakdown - see chapter 3), diffusion is once again recovered. In liquids close to the glass transition the separation of time scales can be as high as 14 orders of magnitude - much of this covered by the so-called fast processes.

4.1.1 The β -relaxation

The β -relaxation is a fairly broad term that describes the dynamics of a glass former on time scales where the overall structure cannot appreciably change but the motion of particles has become collective and highly non-trivial. On a two-point correlation plot, such as Fig. 1.4, the β -regime is usually taken to cover from the fall to the plateau (early- β) all the way to the beginning of the α -relaxation.

There are not many theories that cover the β -relaxation and most of what we know comes from experiments or computer simulations. What predictions there are mostly come from mode coupling theory (MCT). MCT makes approximations to the equations of motion to construct a 'first principles' theory [59]. Its most famous prediction is that there exists a critical temperature, T_c , where the system becomes non-ergodic and timescales diverge. The existence of a T_c has not been verified experimentally and this is the most famous failing of MCT. Above the critical temperature MCT works quite well and makes some useful predictions for the short-time dynamics. For a general two-point correlation, C(t), MCT predicts the two-step relaxion to a plateau [14], and close to T_c it says that on approach to the plateau it takes the form

$$C(t) = f + A_1 t^{-a} (4.1)$$

where f is sometimes referred to as the Debye-Waller factor and A_1 is an arbitrary constant. Experimental data seems to support this relation [15], as does computer simulation [6]. Whether or not the MCT fit is the correct one, it is clear from the experiments that the initial relaxation takes a slow form characteristic of collective behaviour - it is not simply cage rattling in the sense of independent motion within a static cage.

The β -relaxation described by MCT and observed in scattering experiments is not the only fast process seen in glassy materials. Dielectric loss experiments show that, as well as a strong peak at frequencies corresponding to the α -relaxation, there is another peak corresponding to shorter time scales [60, 61]. This secondary process is distinct from the β -relaxation described above and is often referred to as the slow β -relaxation or the Johari-Goldstein process [2]. The time scale associated with this process scales with temperature in an Arrhenius manner in contrast to the fast β -process that has a weak temperature dependence.

4.1.2 Correlated dynamics

As mentioned at the beginning of the chapter, Widmer-Cooper and Harrowell have attempted to look at the spatial correlation between the fast β -relation and the slow α relaxation in binary Lennard-Jones mixtures. Initially they were looking for a way of predicting the dynamical heterogeneity from the static structure of the liquid. To make a deterministic link between the dynamics and the structure they introduced the idea of dynamic propensity [62]. For a given starting configuration all trajectories that pass through are sampled by re-running the simulation with different initial velocities. In this way one can obtain the average dynamic that goes with a particular configuration.

It proved very difficult to come up with a structural measure that correlated to the long time dynamics. The local free volume and the local potential energy were both poor predictors. What did work, however, was the fast dynamics. They found that the parts of the liquid that moved the most on a short time scale, the loosest regions, also moved the most on longer time scales [63]. Whatever information that was contained in the structure was clearly accessible at both short and long times.

For some KCMs it is already possible to make predictions on the long time dynamics by searching the structure. Hedges and Garrahan showed that extended clusters of particles connected to vacancies were the main excitations in the (2)-TLG [48]. If we introduce fast variables we would like to know if they are sensitive to such excitations.

4.1.3 Schemes for fast KCMs

After the coarse graining procedure that goes with creating a FSM each cell on the lattice contains many hidden internal degrees of freedom. These extra states can generally be ignored when studying the dynamics although it is sometimes necessary to address them when considering the thermodynamics [64]. In order to recover short time dynamics it will be necessary to add more states to the cell that somehow reflects the fast character of the liquid. Before proposing any such model we first state a few requirements that we would like the model to satisfy:

- 1. By adding fast degrees of freedom we do not want to change the character of the underlying KCM. Any variable that is added in should be able to be coarse grained back out.
- 2. The thermodynamics should remain controlled. Most importantly we would like to avoid adding static correlations between slow variables.
- 3. There must be some coupling between fast and slow degrees of freedom.

There are many ways to create a model that satisfy these requirements. In the next section we will study one such model that achieves this using a physically reasonable basis.

4.2 One dimension

Our model is an extension of the FA model [22,27,57,65] to include fast degrees of freedom. The FA model consists of spins, n_i , on a lattice that can either be mobile, $n_i = 1$, or immobile, $n_i = 0$. There are no static interactions between cells and being mobile comes with an energy cost, hn_i . Transitions between these states must be facilitated by f neighbouring mobile sites such that, summing over nearest neighbours, $\sum_j n_j \ge f$. In one dimension f can only be 1, in higher dimensions the FA model is taken to mean f = 1unless specified otherwise.

Added to the FA model is a fast variable, $s_i \in \{-1, 1\}$, that is unconstrained by the mobility field but interacts with it via the Hamiltonian

$$H = h \sum_{i} n_{i} - J \sum_{i} (1 - n_{i})(1 - n_{i+1}) s_{i} s_{i+1}$$
(4.2)

The dynamics of the model proceed by attempting to change the mobility, n, subject to the usual kinetic constraint, or changing the *s*-field with equal rate, where the only constraint is obeying detailed balance using the Metropolis rule [25]. The rates are therefore

$$n_{i} = 0 \quad \stackrel{C_{i} \min\{1, e^{-\beta \Delta E}\}}{\longleftrightarrow} \quad n_{i} = 1$$

$$s_{i} = -1 \quad \stackrel{\min\{1, e^{-\beta \Delta E}\}}{\longleftrightarrow} \quad s_{i} = +1$$

$$(4.3)$$

where $C_i = 1$ if the kinetic constraint is satisfied and $C_i = 0$ otherwise.

The physical motivation for this model can be seen from the effect of the mobility field on the fast variables. One can define a local coupling constant, $J_{ij} = J(1 - n_i)(1 - n_j)$, such that $J_{ij} = J$ if both sites are immobile and $J_{ij} = 0$ if either site is mobile. The mobility field disrupts the *s*-field and an *s*-spin inside a mobile cell can change state at no energy cost. In one dimension a mobility excitation kills long range order in the *s*-field by disconnecting regions.



Figure 4.2: (left) $\langle n \rangle$ as a function of coupling, J, with h = 1. Inverse temperature, $\beta = 1/T$, ranges from 1-3. (right) $\Gamma_{n,s}(R)$ at $\beta = 2$, h = 1 and coupling J = 0.5.

4.2.1 Thermodynamics

Before we study the dynamical properties of this model it is essential to understand the thermodynamics. By adding an interaction term between the cells we have potentially added correlations that could significantly change the model. A feature of KCMs that we would like to preserve is simple static properties. In one dimension we can solve these exactly using transfer matrices [66]. The detailed solution is included in appendix A and so here we will just present the results.

The transfer matrix associated with the Hamiltonian in Eq. 4.2 is a symmetric 4×4 matrix with non-degenerate eigenvalues, λ_i , and eigenvectors, $|i\rangle$. To calculate the effect of the *s*-field on the concentration of mobility excitations (the *n*-field) we use

$$\langle n \rangle = \langle 0 | \hat{n} | 0 \rangle \tag{4.4}$$

where $|0\rangle$ is the eigenvector corresponding to the largest eigenvalue, λ_0 , of the transfer matrix and \hat{n} is the operator that projects the value of n_i associated to a site (see appendix A for its definition). The result is a little cumbersome algebraically but is plotted in Fig. 4.2.

As one might expect the s-field has a large effect on the concentration. In the FA model limit, $J \to 0$, the concentration is given by $\langle n \rangle = (1 + \exp[\beta h])^{-1}$ At low temperatures, where most of the s-spins will be aligned, an excitation will cost energy h + 2J due to the broken bonds. If one takes the away the cost of excitations, $h \to 0$, then the cost of mobility excitations comes purely from disrupting the s-field. This is an interesting limit of the model that will be discussed later.

An important feature of the FA model is that excitations are not statically correlated; if we do introduce such correlations then we would expect the dynamics to significantly change. Preferably we'd like to avoid this. We can measure the spatial corrlations between a general variable σ over a distance, R, with the quantity $\Gamma_{\sigma}(R) = \langle \sigma_0 \sigma_R \rangle - \langle \sigma_0 \rangle \langle \sigma_R \rangle$. We find that while correlations in the *s*-field are given by

$$\Gamma_s(R) = f(\beta h, \beta J) \left(\frac{\lambda_1}{\lambda_0}\right)^R \tag{4.5}$$



Figure 4.3: The specific heat capacity, $c_v(T)$ (left) and the entropy, S(T) (right), for the ns-model in one dimension all with h = 1. Dashed curves are the pure FA model. At values of J > 0.5 the fluctuations are dominated by domain walls and for J < 0.5 there are two separate features, one for mobility and one at lower temperature for the fast *s*-field.

where λ_1 is the second biggest eigenvalue, correlations between in the mobility field go to the next biggest eigenvalue

$$\Gamma_n(R) = g(\beta h, \beta J) \left(\frac{\lambda_2}{\lambda_0}\right)^R \tag{4.6}$$

as plotted in figure (4.2) for $\beta = 2$ and J = 1/2. Because we are using the second biggest eigenvalue the correlation very quickly decays for increasing R. This is encouraging because it means that even at R = 1 the correlations in the mobility field are small. With a suitable rescaling to account for the lower concentration of mobility excitations, it is reasonable to expect that the dynamics of the mobility field should closely resemble those of the FA model.

Finally, we can also look at some of the bulk thermodynamic properties of the nsmodel. It has been shown that the FA model, and a version that mixes with the East model (see chapter 2), can be used to fit the viscosity data for a range of glass formers while also being able to account for the difference in heat capacity measurements related to fragility [38]. For this to be valid it is necessary to consider extra states within the FA cell that do not affect the dynamics [64, 67]. Without these extra states the immobile regions cannot contribute to the heat capacity and the model fails.

The ns-model takes one step towards this with the s-field providing two extra states. At very low temperatures, when the concentration of mobility excitations is low, we can still get a large contribution from fluctuations in s. This is shown in Fig. 4.3, where for small values of the coupling, J, we get an extra peak in heat capacity at temperatures where the model is dynamically slow. The same story can be told with the entropy (Fig. 4.3), with s-field fluctuations keeping it bouyant at low temperatures. For larger values of J the extra peak in heat capacity moves to higher temperatures and eventually removes the peak due to h completely.

4.2.2 Dynamics

With the static properties under control it only remains to see what happens dynamically. For these numerical results we have used Monte Carlo simulations taking the rates from Eq. 4.3. The usual FA model constraints are applied to changes in the *n*-field, and the *s*-field is left unconstrained. Moves are attempted at the same rate with both fields although it is possible to introduce a relative factor in the rates to account for the different physical nature of each field. To generate equilibrium configurations one can again use transfer matrix techniques; how this is applied to configurations is described in appendix A.

To study the relaxation behaviour we used two different two-point time correlation functions. For the s-field this was the autocorrelation function, $C_s(t)$, defined as

$$C_s(t) = \frac{\langle s_i(0)s_i(t)\rangle - \langle s_i(0)\rangle\langle s_i(t)\rangle}{1 - \langle s_i(0)\rangle\langle s_i(t)\rangle}$$
(4.7)

where the averages are over all sites, *i*. $C_s(t)$ goes from $C_s(0) = 1$ at the chosen starting point to $C_s(t \to \infty) = 0$ at long times. For the mobility field, *n*, we used the persistence function. For a given site its persistence, $P_i(t)$, is defined to be 1 if it has never moved and 0 for all times after the first move. The persistence function, P(t), is the system average and also goes between P(0) = 1 and $P(t \to \infty) = 0$. These correlations are plotted in Fig. 4.4 for a fixed value of J = 0.1 and h = 1. At this value of coupling we see a very large separation of timescales. The *s*-field is very quick to relax compared to the mobility field and the temperature dependence of the relaxion time is much weaker.

Fig. 4.5 shows the J-dependence of the correlation functions at a fixed temperature, T = 1/3, and again h = 1. We see that the coupling has a strong influence on both correlations. This is not to be unexpected as we already know that it strongly affects the concentration of mobility defects. In the pure FA model the persistence time scales with the concentration of defects, $c = \langle n \rangle$, as $\tau_p \sim c^{-3}$ [57]. This relationship appears to carry over to the ns-model as shown in Fig. 4.5 where we have scaled time in the persistence function by a factor of c^3 . The data collapse works particularly well for smaller values of J; at larger values of J we start to see a small contribution from the static correlations in the *n*-field. It is remarkable that despite the large effect adding a coupling term has on the thermodynamics, the fundamental processes of the FA model are unchanged. This shows that by considering the concentration of excitations, and not the particular values of h or J that achieved it, we can coarse grain back to the FA model as required in section 4.1.3.

At low temperatures the relaxation of the s-field is dominated by the diffusion of domain walls. This gives a similar scaling argument for the relaxation times of the FA model except the diffusion constant is of order unity. If we define the concentration of domain walls as, $w = \frac{1}{2}(1 - \langle s_i s_{i+1} \rangle)$, then we expect the relaxation time to scale as $\tau_s \sim w^{-2}$. Again Fig. 4.5 shows that this scaling works fairly well over the a small range of Js. Given these scaling relationships one can be quite confident that the ns-model still contains the physics of the FA model provided the right observables are considered.

If the decay of the spin autocorrelation function is to represent the early decay of a correlator in a real glass then the functional form of $C_s(t)$ at late times tells us how the



Figure 4.4: Spin autocorrelation function, $C_s(t)$, and the persistence function, P(t), for the ns-model in 1d with h = 1, J = 0.1 and temperatures, $\beta = 1/T = 1$ -5, with temperature decreasing left to right.



Figure 4.5: Same as Fig. 4.4 but with fixed temperature, $\beta = 1/T = 3$, and variable J = 0.1-0.5 in equal steps from left to right. The right hand figure shows the same plot scaled as a function of concentration of excitations, c, and domain wall density (see text).

system relaxes to the β -plateau. Although fits are not presented here, the tails of the various $C_s(t)$ plots are difficult to distinguish from exponential functions with the data available - particularly for small J. While the relaxation of the spins is certainly not trivial, in one dimension it does not show any of the more collective behaviour discussed in section 4.1.

4.2.3 Spatial correlation

As well as the average relaxation of the system we are also interested in the spatial distribution of that relaxation. We know that the FA model is dynamically heterogenous and this will clearly continue for the slow part of the ns-model. What happens with the *s*-field is not quite so clear, although at low temperatures we can make some observations: Any spin that starts in a mobile cell will relax on a timescale of order 1. A mobility excitation disconnects different parts of the chain, so a passing defect will break up any correlations either side of it. Lastly, in a mostly ordered chain it costs energy 4J to create a domain wall. Next to a defect this only costs 2J; the defects can therefore be seen as sources of domain walls.

An impression of the spatial form of the dynamics is given by a space-time plot as shown in Fig 4.6 (cf. Ref. [68]). This shows a single trajectory for 500 sites and parameters h = 1, J = 1.7 and $\beta = 1.2$, chosen to show as many of the dynamical features in the available space. One can see the mobility excitations (shown in black) separating spin domains and also domain walls originating from them in many cases.

To better visualise how features in the static structure affect dynamics further down the line we turn to the idea of dynamic propensity, introduced by Widmer-Cooper and Harrowell [62]. For any time dependent quantity we can define its propensity as the average over all possible trajectories from a fixed starting configuration. These trajectories belong to the isoconfigurational ensemble and for any particular starting point the quantity's isoconfigurational average is a deterministic function in time. Using this with our two-point correlations we can define the isoconfigurational persistence, $[P_i(t)]_{\rm IC}$, and the isoconfigurational spin autocorrelation function, $[s_i(0)s_i(t)]_{\rm IC}$, where $[\cdot]$, implies an average over trajectories.

Fig. 4.6 shows these quantities for the same starting configuration as the single trajectory. This time an average over 100 trajectories is shown. A site fades from blue to white as its average spin correlation decays and from white to black as its average persistence decays. The strongest features in the spin dynamics are the domain walls that exist at the start. Secondary to this one can see a white haze drifting out from the mobility excitations. This is the result of the increased chance of forming a domain wall next to these sites.

It appears that there might be some spatial correlation between the relaxation dynamics of the spins and the mobility but to be sure it would be better to try and use the propensity fields to make a quantitative measure. For this we define a two-time correlation

$$C(t_1, t_2) = \langle [s_i(0)s_i(t_1)]_{\rm IC}[P_i(t_2)]_{\rm IC} \rangle - \langle s_i(0)s_i(t_1) \rangle \langle P_i(t_2) \rangle$$
(4.8)



Figure 4.6: (Left) A single space-time trajectory for 500 sites in the ns-model over 500 Monte Carlo sweeps with h = 1, $\beta = 1.2$ and J = 1.7. Red and blue indicate up and down spins $(s = \pm 1)$ whereas black indicates the position of mobility excitations (n = 1). (Right) From the same starting configuration the isoconfigurational ensemble, over 100 trajectories, of spin autocorrelation (blue decaying to white) and mobility persistence (white decaying to black).



Figure 4.7: Correlation, $C(t_1, t_2)$, between fast and slow dynamics for coupling J = 0.3 and inverse temperature, $\beta = 3$. The correlation peaks at the relaxation times of the spins and mobility respectively.

where $\langle \cdot \rangle$ implies an average over particles. $C(t_1, t_2)$ is plotted in Fig. 4.7 for a system with inverse temperature, $\beta = 3$ and relatively weak coupling, J = 0.3. There is a single peak in the correlation close to the fast relaxation time, τ_s , along t_1 and the slow relaxation time, τ_p , along t_2 . This is perhaps surprising given how strong the influence of the preexisting domain walls appears to be in Fig. 4.6, however, once the connected part of the correlation is removed there is still a clean signal from the interaction with the mobility.

4.2.4 Summary

The one dimensional ns-model is a useful model because it extends the FA model in a non-trivial way to include a fast β -relaxation without breaking the spirit of the original model. The role of mobility excitations in breaking up the fast ordering is physically intuitive and leads to interesting correlations between the two dynamics. It also offers an alternative view of the FA model where the energy penalty for mobility excitations comes solely from breaking up long range order.

A possible downside to the model is that, even though the fast variables are coupled, in one dimension we do not see the sort of collective behaviour that is observed in real glasses. Indeed it is hard to imagine any model in one dimension that would be able to support such a property. In the next section we will take the model to two dimensions where the spins are likely to be much more correlated.

4.3 Two dimensions

The two dimensional version of our model uses a square lattice and the interactions are characterised by the Hamiltonian

$$H = h \sum_{i} n_{i} - J \sum_{\langle ij \rangle} (1 - n_{i})(1 - n_{j})s_{i}s_{j}$$
(4.9)

where the second sum is over nearest neighbour pairs. Excitations still perform the role of breaking up order in the s-field although now they are not direct sources of domain walls. Without any disorder at all $(h \to \infty)$ we recover a standard spin-1/2 Ising model which has a phase transition at $T_c = 2J/\ln(1 + \sqrt{2})$. The prescence of a phase transition means that the thermodynamics are going to be much more complicated than in one dimension.

It is also possible to migrate the ns-model to the triangular lattice gas (TLG) (see chapter 3 for an introduction). In this version of the model spins reside on the particles and the defects come from unoccupied lattice sites. The crucial difference is that the number of defects is now fixed by the density, ρ , instead of a chemical potential. If the particle field $\phi(\mathbf{r})$ is 1 where there is a particle and 0 for the vacancies then the Hamiltonian is nearly identical to Eq. 4.9

$$H = -J \sum_{\langle ij \rangle} \phi(\mathbf{r}_i) \phi(\mathbf{r}_j) s_i s_j \tag{4.10}$$

except this time the sum is over nearest neighbours on a triangular lattice. Before going

into the detail of these models we will review some closely related models that have already received a great deal of attention.

4.3.1 Diluted Ising model

If we make the approximation that the mobility field, n, is a quenched random variable that is fixed at all times then the effect of an excitation, $n_i = 1$, is to remove the *s*-spin that is on that site (at least from our consideration). If we let $n_i = 0$ with independent probability, p, then the model reduces to a random site-diluted Ising model with occupancy p. This is a well studied model introduced by Griffiths [69] and it is known to have a second order phase transition at a temperature, $T_c(p)$, that is equivalent to the pure case at p = 1and decreases as p is reduced and falls to zero when $p \to 1/2$ at the percolation threshold.

Between the pure transition temperature, sometimes referred to as the Griffiths temperature $T_G = T_c(p = 1)$ and the new critical temperature, $T_c(p) < T < T_G$, lies the Griffiths phase. In this temperature range the spins are beginning to line up but the ordering is broken up just enough by the dilution to prevent the correlation length diverging. On crossing T_G there is no jump in the order parameter, $\langle s_i \rangle$, although the dynamics do begin to slow down.

Inside the Griffiths phase the dynamics are even slower and are dominated by rare regions of local order [70, 71]. Considering a two-point correlation function, $C(t) = \langle s_i(0)s_i(t) \rangle$, Bray showed that in the long time limit this can be written in terms of the relaxation of regions of size, L, and local dilution, p', such that

$$C(t) = \sum_{L,p'} P(L,p') \exp\left(-\frac{t}{\tau(L,p')}\right)$$
(4.11)

where P(L, p') is the probability of a site belonging to such a region and $\tau(L, p')$ is the subsequent relaxation time. Bray finds that this sum is dominated by terms with p' = 1, that is to say regions with unusually low dilution. The probability of belonging to such a cluster of volume L^d , to leading order, is

$$P(L,1) \sim p^{L^d} = \exp(-cL^d)$$
 (4.12)

and the relaxation time for such a cluster is obtained by considering the free energy required to create an interface of length, L, giving

$$\tau(L,1) \sim \tau_p \exp(\sigma L^{d-1}/T) \tag{4.13}$$

where σ is the surface tension and τ_p is the relaxation time in the pure system below T_c . By putting Eqs. 4.12 and 4.13 into Eq. 4.11 one gets

$$C(t) \sim \sum_{L} \exp(-cL^d - (t/\tau_p)e^{-\sigma L^{d-1}})$$
 (4.14)

In the limit of large t this can be solved using the saddle point approximation by finding the value of L that gives the biggest term in the sum; this being the only one left as $t \to \infty$. Once this is done we arrive at a scaling form for the dynamics inside the Griffiths phase as

$$\ln C(t) \sim -A(\ln t)^{d/(d-1)} \tag{4.15}$$

At $T \to T_c(p)$ this breaks down as clusters of all dilutions begin to become important and the dynamics become critical.

In the quenched disorder version of the ns-model we therefore find that the dynamics do indeed slow down over a range of temperatures that are not immediately near to a phase transition. The dynamics at longer times in the Griffiths phase are dominated by clusters with low dilution, or in our case a low concentration of mobility. It is therefore promising that we will see the same behaviour in the unquenched version because, due to the kinetic constraint, it will take a long time for regions that have low mobility to acquire it from another part of the system.

4.3.2 Blume-Capel model

The results from the diluted Ising model are quite promising for our fast dynamics. Obviously in our model the mobility field, n, is not a quenched random variable but an annealed one. This will undoubtedly change the static properties, the question is how much?

If we make a change of variable $\sigma_i = (1 - n_i)s_i$ then the Hamiltonian can be rewritten from Eq. 4.9 as

$$H = h \sum_{i} (1 - \sigma_i^2) - J \sum_{\langle ij \rangle} \sigma_i \sigma_j$$
(4.16)

with $\sigma_i \in \{0, 0, \pm 1\}$. Apart from the degeneracy of the $\sigma_i = 0$ state this is Hamiltonian for the Blume-Capel model [72, 73] in zero magnetic field, with a spin-1 Ising variable. A similar model by Blume, Emery and Griffiths has been used to model the phase behaviour of He³-He⁴ mixtures [74]. The phase behaviour is well known and preserves many of the features of the diluted Ising model.

The extra degrees of freedom allow for a richer phase structure. For positive values of h it looks similar to the diluted case with defects lowering the second order transition temperature. For negative values of h we see the appearance of a first order phase transition where the spins become separated into a liquid-vapour-like coexistence region. For the right values of h and T the first and second order transition lines come together at a tri-critical point [74].

Energetically our two models look very similar to the Blume-Capel model [72,73]. As we have already mentioned the ns-model has a slightly more degenerate $\sigma = 0$ state and the TLG version has a fixed density. In the next section we will map out the phase structure and thermodynamic properties for both models before we proceed to the dynamics.

4.3.3 Thermodynamics

The partition functions from the Hamiltonians in Eq. 4.9 and Eq. 4.10 cannot be solved exactly so we must resort to numerical methods to map out the phase structure. For the most part a standard Metropolis Monte Carlo algorithm is sufficient, although around



Figure 4.8: Phase diagrams for the 2D ns-model (a) and the TLG with Ising spins (b). Solid lines with circles indicate the lines of second order transitions (lines in (b) are a guide to the eye). The dashed line with squares indicates the line of first order transitions. In both figures the dotted line shows pure Ising T_c on a square lattice (a) and a triangular lattice (b). The system sizes were of length L = 100 for the ns-model, where finite size effects give only a small correction, and L = 500 for the TLG.



Figure 4.9: Example configurations for the TLG at density $\rho = 0.8$ for a system in the Griffiths-like region , the ordered phase and demixed phase. Red and blue represent up and down spins.

the second order phase transitions the s-spin dynamics become sufficiently slow to require something a little more sophisticated.

To get around the problem of critical slowing down we employed a Wolff cluster algorithm (see appendix B) in combination with a single flip Metropolis algorithm. The standard equilibration procedure was to run the Wolff algorithm for a time to equilbrate the s-spins on a given realisation of the n-field (or particle configuration) and then subsequently run the Metropolis algorithm on all variables to shake up the disorder. By alternating between the two the system equilibrates very quickly.

To calculate the position of the second order phase transition, $T_c(h, J)$ (or $T_c(\rho, J)$ for the TLG), it is possible to use an invaded cluster algorithm to search for the temperature where critical clusters begin to span the system (see appendix B for details). To find this temperature with annealed disorder it is once again necessary to run in combination with a Metropolis algorithm to make sure we are averaging correctly. This procedure works better with the ns-model than the TLG. If one is not careful with the TLG the first order transitions can pull the algorithm away from the critical point greatly reducing its efficiency.

The resulting phase diagrams are plotted in Fig. 4.8. Both models have a line of second order phase transitions with a Griffiths-like phase sitting between the new transition temperature and the pure T_c . The ns-model with positive h has no first order transition because for low temperatures it is energetically favourable to take $\langle n \rangle \to 0$ rather than to separate. At low temperatures in negative h the second order line will eventually terminate at the tri-critical point meeting the first order line. The physical meaning of negative h is that it is now energetically favourable to be in the mobile state (although breaking the order in the *s*-field does complicate this somewhat). This is a bit counter intuitive to the original idea of the FA model and so, preferably, we'd like to avoid this situation. For this reason we do not include the first order transition line as it is far away from our region of interest.

With the particle number conserved, the TLG phase diagram looks slightly different. We still have a second order transition between ordered and disordered phases, but at temperatures a bit lower the correlations between vacancies begin to grow and a first order transition into a demixed phase occurs. Example configurations in the different phases are shown in Fig. 4.9. The first order transition lines in Fig. 4.8 were calculated



Figure 4.10: Particle correlation function, $\Gamma(R)$, for the coupled-TLG at fixed density, $\rho = 0.7$,(a) and fixed distance, R = 1, on the (b).

by observing the sudden growth of the particle correlation length. As the density is lowered towards the percolation threshold long range order can no longer be supported and so the second order line hits the first order line at the tri-critical point around $\rho \approx 0.55$ and $T \approx 1$.

Growing correlations between the particles, or mobility excitations, will have significant effects on the dynamics. It is important to be able to quantify this at different parts of the phase diagram. For the TLG we define the average over all pairs of lattice sites,

$$\langle \phi(0)\phi(R)\rangle = \sum_{i,j} \phi(\mathbf{r}_i)\phi(\mathbf{r}_j)\delta(\mathbf{r}_i - \mathbf{r}_j - \mathbf{R})$$
(4.17)

and then use this to calculate the correlation function

$$\Gamma(R) = \frac{\langle \phi(0)\phi(R) \rangle - \rho^2}{\rho(1-\rho)}$$
(4.18)

that goes between $\Gamma(0) = 1$ and $\Gamma(\infty) = 0$.

 $\Gamma(R)$ is plotted in Fig. 4.10 at density $\rho = 0.7$ for temperatures corresponding to the



Figure 4.11: Specific heat capacity for ns-model in 2D with $N = 10^4$ sites. In units with h = 1 curves are plotted for J = 0 (equivalent to the pure FA model), J = 0.05, 0.1, 0.2 and 0.3.

Griffiths-like region (T = 3), the ordered magnetic phase (T = 1.5) and the fully separated phase (T = 0.5). The intermediate temperatures of T = 1 and T = 2 lie close to the first and second order phase transitions respectively (see Fig. 4.8). The correlation length only gets larger than 1 on approach to the first order transition. Even in the ordered phase particle correlations are not very long ranged, and for higher temperatures the particles are hardly correlated at all. The ns-model (that doesn't have a first order transition for h > 0) has much weaker correlations in the *n* field. It is only at values of J > h that any significant effect is seen.

In the previous section we showed that the heat capacity for the FA model can be boosted by contributions from extra states within the cell. In two dimensions these contributions can become very large indeed due to the prescence of the phase transitions. Fig. 4.11 shows the specific heat capacity for the ns-model at various different couplings, J. For very small J the FA result can still be seen, but as J is increased it quickly becomes eclipsed by the spin contribution.

Without a coupling term the TLG is purely entropic and as such does not have a heat capacity. When interactions are turned on we get a heat capacity profile like that of Fig. 4.12. The most notable feature is that, even with a simulation box containing 10⁶ lattice sites, the divergence in heat capacity at the second order line is very difficult to see - especially at low densities. If one looks only at the contribution of the magnetisation (see inset of Fig. 4.12) then the phase transition can be seen to be very sharp. The particle field acts to dampen energy fluctuations caused by fluctuations in the magnetisation making it difficult to see the transition point in a finite system. No such problem exists at the first order line where the heat capacity rapidly diverges.



Figure 4.12: Specific heat capacity the for TLG with Ising spins on a lattice of $N = 10^6$ sites. Inset shows the fluctuations in the average magnetisation, $S = \sum_i s_i$, which diverge at the second order phase transition. Energy fluctuations are significantly larger near the first order transition.

With a more complicated phase structure the models would seem to be poor candidates for an extended KCM. When one looks closer, however, in the majority of regimes the conditions are actually quite good. Moving to two dimensions has a large effect on the fast variables while, on the whole, leaving the original models intact. Only at low temperatures in the TLG (or strong coupling in the ns-model) do we see static correlations in the slow variables, and even around the ordering transition these do not grow significantly. The prescence of the phase transition itself does mark a departure from the simple idea of a KCM and we will discuss this later.

4.3.4 Dynamics

Before we try to explicitly solve the dynamics of the two dimensional models we can make some comments from the static properties. The region of most interest in the phase diagram (Fig. 4.8) is going to be the Griffiths-like region below the pure transition temperature but above the actual ordering transition. Here we expect to see interesting spin dynamics, similar to the diluted Ising model, without causing any long range correlations in the slow variables. Below the phase transition the spin dynamics will be that of full magnetic reversal, a process that is much slower than we want.

At this point we begin to see the differences between the ns-model and the TLG. The FA model in two dimensions does not show slow dynamics unless the concentration of mobility excitations, $\langle n \rangle$, is very small (~ 10⁻²). If we want to disrupt the *s*-field then we need a much larger fraction of defects to have an appreciable effect. The (2)-TLG, on the



Figure 4.13: Fast and slow relaxation in the (2)-TLG with Ising spins. The left most curves are the spin autocorrelation functions, $C_s(t)$, and the right most curves are the particle persistence functions, P(t). For each density four different temperatures are shown: $\beta = 1/T = 0.2$ (black), 0.3 (red), 0.35 (green) and 0.4 (blue). The dashed lines indicate the persistence of the pure (2)-TLG for comparison. Note that at low density (a)the persistence functions are slower with decreasing temperature whereas at high density the trend is initially reversed ((c) and (d)).

other hand, is known to be very slow even at densities as low as $\rho = 0.7$ (chapter 3), it is therefore possible to have a high concentration of defects while retaining a slow structural relaxation. For this reason this section only considers the dynamics of the (2)-TLG with coupled spins.

The dynamics of the coupled (2)-TLG follow much the same procedure as the ns-model in section 4.2.2. We attempt to change spins or move particles with equal probability and accept or reject moves according to the Metropolis algorithm, $P(\text{Accept}) = \min\{1 : \exp{-\beta\Delta E}\}$, to maintain detailed balance. Particle moves are subject to the (2)-TLG kinetic constraint whereby both mutual neighbours of the target and destination site must vacant (see chapter 3 for this in more detail). Recall that spins are attached to *particles* and not the lattice. When a particle moves it takes its spin with it.

As in section 4.2.2 we use the spin-autocorrelation function, $C_s(t)$, from Eq. 4.7 for the spin dynamics and the persistence function, P(t), for the particle dynamics. The results for these two-point correlation functions are plotted in Fig. 4.13 for a range of densities and temperatures. Below the pure critical temperature $(T_c^{\text{pure}} \approx 3.6J)$ on a triangular lattice) we see a slow down in the spin dynamics with non-exponential relaxation even quite far from the real phase transition temperature $T_c(\rho)$. This appears to confirm that the dynamics are the same as the diluted Ising model although we don't have enough



Figure 4.14: Self intermediate scattering function, $F_s(k,t)$ at wave vector $k = \pi$ along the x axis showing a two-step relaxation. All curves have temperature, T = 3.33, and densities are $\rho = 0.6, 0.65, 0.7, 0.75, 0.8$ running left to right

statistics to verify the long time scaling relation. Closer to $T_c(\rho)$ we begin to see critical slowing down with $C_s(t) \sim t^{-\gamma}$ as one would expect.

The computational load of running simulations that include rapidly fluctuating spins with slow moving particles is very high and as such the coupled-TLG cannot run to the timescales available without the spins making the long time particle dynamics difficult to reach. The uncoupled curves are included in Fig. 4.13 as dashed lines as a reference.

For the most part the coupled-TLG persistence functions stay quite close to the uncoupled case. At low densities the effect of the coupling is to create effective energy barriers for particle moves thus slowing down structural relaxation. At higher densities these barriers become negligable compared to the kinetic constraint and by $\rho = 0.75$ all the persistence curves come together.

Past $\rho = 0.75$ more peculiar behaviour sets in. At density $\rho = 0.77$ we can see that the low temperature curves initially relax *faster* than the uncoupled TLG, while coming back together at later times. This can only be caused by static correlations in the particle field. As shown in Fig. 4.10, at high densities these correlations are not much bigger than anywhere else, but with the dynamics becoming so drastically slow for $\rho > 0.75$ the system is much more sensitive to small changes. It is important to note, however, that the effect is still quite small compared to the α -relaxation in general and the character of the structural relaxation is largely preserved.

To give a flavour of how these two relaxation functions could add up to give the two-step behaviour that we see in real supercooled liquids one can consider the self intermediate scattering function, $F_s(\mathbf{k}, t) = \langle e^{\mathbf{k} \cdot \Delta \mathbf{r}(t)} \rangle$ where the particle position is offset by an amount, $\delta x = \pm \epsilon/2$, depending on its spin state. In Fig. 4.14 this is plotted for $\epsilon = 0.3$ for a range of densities. By $\rho = 0.8$ we begin to see a definite plateau. In assigning a spatial value to the fast variables we start to see something that looks very much like the β -relaxation in glasses. While this is only an example of how one could combine the variables it is useful for comparison with the real liquids (see Fig. 1.4) and allows for a physical interpretation of the spin autocorrelation function as the early decay of the scattering function.

4.3.5 Dynamic Heterogeneity

In this final section we would like to turn our attention to the spatial structure of the dynamics. The dynamics of the TLG are known to be very heterogenous in space, but what about the spin dynamics? In the diluted Ising model the tail of the distribution is dominated by rare high density regions [69] which, in itself, is a form of dynamic heterogeneity. The natural question is therefore: what is the spatial correlation between the fast spin dynamics and the slow particle dynamics?

We can get a picture of the two spatial distributions by considering the isoconfigurational autocorrelation function, $[C_s(t_1)]_{IC}$, for the spins and the isoconfigurational persistence, $[P(t_2)]_{IC}$, for the particles. For each particle, and its resident spin, these quantities are averaged over 100 trajectories and projected onto the initial configuration - as shown in Fig. 4.15. The two distributions are qualitatively similar. To better quantify this we use the correlation function, $C(t_1, t_2)$ from equation 4.8, that averages the correlation between a particle and its spin at different times. The result is as we might expect (see Fig. 4.16); there is a positive correlation between the two propensity fields that peaks at a time, t_1 , corresponding to the characteristic relaxation time of the spins, and at time, t_2 , a bit before the α -relaxation time.

The shape of $C(t_1, t_2)$ is more or less the same for all densities, becoming more stretched as the temperature is dropped or the density raised in a way similar to the two-point correlations. We'd also like to know how the magnitude of the correlation depends on these parameters. The bottom of Fig. 4.16 shows the cross sections through the peak along t_1 for different temperatures. The magnitude of the peak grows as temperature is lowered reaching a maximum at the phase transition. At low temperatures, with a growing correlation length, the spins become sensitive to more of their surrounding structure. Beyond the spin-ordering phase transition the dynamics become too slow for this analysis.

Also shown in Fig. 4.16 is the correlation for a version of the model where the spin dynamics and the particle dynamics are run separately. In this version both dynamics use a shared initial configuration but then the spins are evolved with the particle field frozen in place and the particle dynamics are run without considering the spins at all. The strength of the correlation is not much different between the two models.

This is a useful result; it tells us that the correlation between the dynamics does not rely on the back reaction of the fast variables on the slow variables. Particles are moving sufficiently slowly such that spins are behaving as though there is a background of quenched disorder. Likewise for the particles, it is only the initial configuration, coupled through the kinetic constraint, that affects their relaxation. The spin field has time to completely rearrange itself many times over before most particles even make their first move. If we attempted to move particles less often, which would be physically reasonable, then we would expect the quenched disorder approximation to be improved further. This is an idea that will be explored further in the next chapter.



Figure 4.15: Top shows a typical starting configuration of particles and spins (red and blue) for the (2)-TLG. The lower pictures show the same configuration but this time particles are faded out according to the on site isoconfigurational autocorrelation, $[C_s(t_1)]_{\rm IC}$, (left) and persistence, $[P(t_2)]_{\rm IC}$. The smaller the persistence (autocorrelation) the more faded the particle is drawn, the most mobile regions appear as white. Times correspond to roughly the fast and slow relaxation times respectively, circles are for orientation. Density $\rho = 0.75$ and T = 3.



Figure 4.16: (Top) The connected correlation $C(t_1, t_2)$ between the isoconfigurational spin autocorrelation function at time, t_1 , and the particle persistence at time, t_2 . This plot is for a system of $N = 10^4$ lattice sites with $\rho = 0.7$ and temperature T = 3.33. Bottom figure shows a slice along t_1 through the peak value for different inverse temperature, $\beta = 1/T$, for the full model with annealed disorder and also a version where the spin dynamics are run with a quenched particle field and the particle dynamics are run without the spins at all.

4.4 Concluding Remarks

In this chapter we have shown that it is possible to extend two different KCMs to include fast degrees of freedom in a way that doesn't change the underlying character of the original models. As can be shown exactly in one dimension, and numerically in two, significant spatial correlations can be built into the fast variables while only introducing extremely short range correlations to the original slow variables. These extra variables are also candidates to correct for the apparently small heat capacities of facilitated spin models in the glassy regime.

In two dimensions the static correlations between the fast variables can become strong enough to undergo a phase transition into an ordered state. While the phase transition is always present we find that the disorder from the underlying models can be enough to disrupt its onset over a large range of temperatures. The dynamics of the fast variables within this range display very interesting collective behaviour with a slow relaxation that is reminiscent of the early β regime of real glasses.

We find that dynamic heterogeneity, the key property of the KCMs, is not only preserved but is also replicated in the fast dynamics. In a result similar to Widmer-Cooper and Harrowell's [63] we find that regions that are slow to relax on the short time scales are also slow to relax on long time scales as well. This correlation is transmitted through the disordered structure that both dynamics live with. We can even remove the coupling completely and find that, provided we use the same disorder, the dynamics will still be well correlated.

Chapter 5

Vibrations

One of the most intriguing aspects of glassformers is the wide range of time scales on which interesting behaviour occurs. Taken from the perspective of supercooled liquids we see increasingly heterogenous dynamics alongside a rapidly increasing time scale for structural relaxtion. From the perspective of amorphous solids we see anomalous vibrational behaviour on a short time scale, giving rise to the 'Boson Peak', an excess of vibrational modes in the THz range [8]. An important question is whether these two characteristic aspects of the dynamics are related, and in particular whether there is a common structural origin for both of them.

In a series of recent studies of systems of hard spheres [75], hard discs [76], and soft discs [77] it has been shown that the spatial localisation of the anomalous low frequency modes give a good indication of the spatial distribution of structural relaxation at much longer times. It has been argued [75, 76] that this correspondence is causal [77], in that it is these soft vibrational modes which provide the underlying structural mechanism to long time relaxation.

In this final chapter we address this problem by studying spatial correlation between structural relaxation and vibrational modes in suitably generalized kinetically constrained models of glasses [22]. In the previous chapter we found that it is possible to introduce fast degrees of freedom in such a way that, while sensitive to the background disorder of the KCM, the back reaction does not unduely affect the long time dynamics. Here we go a step further and introduce a vibrational model, built on the structure of two KCMs, but studied in isolation.

Before introducing our model we begin by reviewing some vibrational properties of glasses, especially in the anomalous low frequency regime, and the relationship to marginally rigid structures. In section 5.2 we return to the constrained lattice gases and then in section 5.3 we show how to generalise these models to build an elastic network on which to study vibrations. We find that defects in the elastic network lead to the appearance of anomolous low frequency modes and that these modes are spatially correlated to the long time dynamics (section 5.4). We end by comparing the vibrations of the elastic network to its response to a small deformation.

The work in this chapter has been submitted to Phys. Rev. Lett. and is available in e-print form as *Relationship between vibrations and dynamical heterogeneity in a model* glass former: extended soft modes but local relaxation, Douglas J. Ashton and Juan P. Garrahan, arXiv:0808.2412.

5.1 Vibrations and rigidity

At a long enough length scale, corresponding to a low enough frequency, the vibrations of almost any material will be those of a continuous elastic medium [78]. The dispersion relation, $k = \omega/c$, c being the speed of sound in the medium, gives rise to a vibrational density of states (DoS) of the form $D(\omega) \propto \omega^{d-1}$. This relation applies at low frequency and continues until the frequency is high enough (length scale small enough) to be sensitive to the microscopic structure. The crossover frequency, ω_D , is often referred to as the Debye frequency.

From experiments using neutron or x-ray scattering [21, 79, 80] it is observed that in glasses, compared to what one would expect for a continuous elastic material, there is an excess of states at frequencies below the Debye frequency. This excess in the THz range is known as the Boson peak (see Fig. 1.5 for examples) and is seen in a wide variety of amorphous materials [81] including proteins [82]. The structural origin of the excess is not entirely agreed upon and there are a number of proposed explanations [17,83]. It appears that the peak is more pronounced in strong glass formers than fragile ones [84] and it can drop in frequency on heating [21].

Wyart and co-workers have developed a theory whereby the Boson peak is a natural consequence of a weakly connected, marginally rigid solid [19,20]. As a material begins to lose its rigidity one starts to see the growth of so called 'floppy modes'. A floppy mode is a normal mode that is disconnected from the rest of the system, or in other words it is a zero-frequency mode of motion without a restoring force. In this picture the anomolous low frequency modes are modes that are on the way to becoming floppy but not yet fully disconnected.

A large scale example of a floppy mode is in the classic engineers problem of building a strong structure. A square arrangement of supporting beams will buckle if force is applied along the top (a floppy mode), a triangular structure, on the other hand, is strong in all directions; it is rigid. Maxwell pointed out that the condition for being a rigid structure is that the number of degrees of freedom, N_d , is exceeded by the number of constraints, N_c , on those degrees of freedom. Unlike the triangular structure, the square has degrees of freedom that are not matched by a constraint and as such is mechanically unstable. The same criteria applies at the microscopic level for materials. Here constraints come in the form of contact forces between neighbouring particles, either from mutual exclusion or friction.

Closely related to the theme of rigidity is that of jamming in granular media. Granular media exist essentially at zero temperature. At a sufficient density, the jamming threshold, ϕ_c , a system will go from a flowing medium to a jammed solid [85]. At this point the material gains a finite pressure and shear modulus. If the pressure is gradually released then the average number of contacts begins to fall away. Accompanying this is the growth of an excess of low frequency vibrational modes similar to the Boson peak. At the point


Figure 5.1: Propensity maps for a (2)-TLG at density $\rho = 0.75$ (second to fifth panels) for a typical configuration (shown on the leftmost panel) averaged over 100 trajectories. Black circles indicate the average distance travelled by each particle after time, from left to right, $t = 5 \times 10^1$, 5×10^2 , 5×10^3 , 5×10^4 . At this density, $\tau_{\alpha} \approx 10^4$.

where the pressure is just zero, the isostatic point, the excess moves to zero frequency and the solid breaks up [86]. In glasses, where temperature is not zero, we do not have such an idealised system. However, useful analogies can be drawn.

5.2 Dynamics of the constrained lattice gases

The models we study are a generalization of constrained lattice gases [44], as studied in detail in chapter 3. We use the non-interacting version of these models where hard-core particles occupy the vertices of a lattice, with single occupancy per site, and with no static interactions between them. For the dynamics, particles try to hop to neighboring sites, but the local hopping rates depend on the occupancy of surrounding sites, so as to mimic steric interactions. Recall that these models have the trivial thermodynamics of a non-interacting lattice gas, but their dynamics at high densities displays many of the features of glass formers, such as non-exponential relaxation [43, 50], dynamic heterogeneity [47], transport decoupling [47], and aging [87] (see Ref. [22] for a review).

In particular, we focus on the two-vacancy assisted lattice gas model on the triangular lattice [44], or (2)-TLG, where the kinetic constraint is explicitly due to steric restrictions: a particle can hop to an empty nearest neighbour site only if the common two neighbouring sites are also empty. We also consider its three dimensional variant on an FCC lattice, the (4)-FLG, where the constraint is that the four common neighbours of the sites undergoing the transition are empty (see chapter 3). Dynamics are measured numerically using a continuous time Monte Carlo algorithm (Ref. [32] and chapter 2).

Figure 5.1 shows propensity maps for the (2)-TLG, which illustrate the spatial localization of relaxational dynamics in the model. The leftmost panel is a typical configuration of the model at a density, $\rho = 0.75$ in this case, for which relaxation is highly non-exponential and heterogeneous [47]. Panels two to five show the average particle displacement at increasing times in the so-called iso-configurational ensemble [62], i.e., for all trajectories starting from the same initial configuration (the one in the leftmost panel). For a detailed study of propensity in the (2)-TLG model see Ref. [48].

5.3 Central Force Network

In order to study vibrations together with dynamical heterogeneity we generalize the models above by adding harmonic interactions between neighboring particles. That is, any two occupied nearest neighbour sites interact through a linear spring of force constant k and of rest length equal to the lattice spacing. The vibrational Hamiltonian then reads,

$$H_{\rm vib} = \sum_{\langle ij\rangle} \frac{1}{2} n_i n_j k \left[(\delta \vec{r_i} - \delta \vec{r_j}) \cdot \hat{r}_{ij}^0 \right]^2, \qquad (5.1)$$

where $\langle ij \rangle$ means that the sum is over nearest neighbor pairs, $n_i = 0, 1$ indicates whether lattice site *i* is empty or occupied, $\delta \vec{r_i}$ is the displacement of the particle whose equilibrium position is site *i*, and \hat{r}_{ij}^0 is the unit vector between sites *i* and *j*. Each configuration thus gives rise to a disordered elastic network due to the presence of vacancies in the particle configuration. This elastic problem is well known, as it corresponds to the "central force" problem of rigidity percolation [88,89]. In particular, Eq. (5.1) for the (2)-TLG and (4)-FLG models correspond to the site diluted triangular and FCC central force networks studied in Ref. [89–91]

Expanding (5.1) up to second order in $\delta \vec{r_i}$ defines the dynamical matrix

$$H_{\rm vib} = \sum_{ij} \delta \vec{r}_i \mathcal{M}_{ij} \delta \vec{r}_j \tag{5.2}$$

which we diagonalise directly to obtain the normal modes. In dimension, d, the dynamical matrix always has d modes with eigenvalue, $\omega^2 = 0$, corresponding to translations - beyond this any zeroes indicate the presence of floppy modes. The constraint counting method mentioned in section 5.1 can give us an effective medium estimate for the number of floppy modes, N_f . A lattice with N sites and where each site is occupied with probability ρ will have $N_d = \rho N d$ degrees of freedom. With z nearest neighbours there will be on average $N_c = \frac{1}{2}\rho^2 N z$ constraints. The number of floppy modes is the difference between these quantities, $N_f = N_d - N_c$. The point at which $N_f = 0$, given here by $\rho_c = 2d/z$, is a critical point separating the rigid and floppy networks [92].

To calculate N_f exactly for any given system one can diagonalise the dynamical matrix, \mathcal{M} , and count the number of zero frequency modes or make use of localised constraint counting techniques [90, 93]. Above the critical density there should be an insignificant number of zero frequency modes whereas for densities below ρ_c the number becomes extensive in the system size (see Fig. 5.2). Fluctuations around the critical point have the effect of shifting the critical density away from the effective medium prediction but the result is always close [90].

Figure 5.3 shows the average density of states (DoS), $D(\omega)$, of vibrational modes for elastic networks corresponding to the (2)-TLG model and (4)-FLG model at various den-



Figure 5.2: The fraction of zero frequency (floppy) modes, N_f/Nd , for the (2)-TLG and the (4)-FLG. Below the critical density, $\rho_c \approx 0.7$ (TLG), $\rho_c \approx 0.5$ (FLG), the number of zero frequency modes becomes extensive in system size. Both curves use systems with approximately 4000 modes.



Figure 5.3: Density of states for the central force networks of the (2)-TLG model (left panel) and the (4)-FLG model (right panel), at various densities ρ above the isostatic density ρ_c . The curves at $\rho = 1$ are exact. The curves for all other densities were obtained by numerical diagonalization of the dynamic matrix of systems with 4000 normal modes.



Figure 5.4: Top left: The particle configuration of Fig. 5.1 is the basis of an elastic network. Top right is the lowest frequency normal mode $\vec{e}_{i\omega}$ at $\omega \sim 10^{-2}$, bottom right is a high frequency, $\omega \approx 2.4$, and bottom left is a middle range frequency $\omega \approx 1$.

sities. For densities lower than that of the full lattice, $\rho < 1$, the DoS presents an excess of low frequency modes. Debye scaling, $D(\omega) \propto \omega^{d-1}$, is recovered at low enough frequencies. The presence of excess modes becomes more pronounced as the density decreases. At the critical density, ρ_c , the excess tail extends all the way to $\omega = 0$. This is the isostatic point at which the system is marginally rigid [88,89,91]. We find that for the (2)-TLG $\rho_c \approx 0.7$ and for the (4)-FLG $\rho_c \approx 0.5$, which agrees with the results of Ref. [89] for the triangular and FCC lattices, respectively.

Moving from frequencies to the modes themselves; for all densities below that of the full lattice, $\rho < 1$, the vibrational modes are spatially distributed in a non-trivial manner. Figure 5.4 shows a selection corresponding to low ($\omega \sim 10^{-2}$), medium ($\omega \sim 1$) and high frequency ($\omega = 2.4$) for the (2)-TLG configuration of Fig. 5.1. Arrows show the 2-dimensional vector, $\vec{e}_{i\omega}$, that is the component of the eigenvector of frequency ω corresponding to particle *i*. The low frequency modes take on a characteristic swirl-like formation whereas the high frequencies become tightly localised. Figure 5.7 illustrates this elastic heterogeneity with the spatial weight, $\nu_i(\omega)$, of modes of low frequency (Fig.5.7, top-right) and high frequency (Fig.5.7, bottom-right), with $\nu_i(\omega) \equiv \vec{e}_{i\omega} \cdot \vec{e}_{i\omega}$. This will be



Figure 5.5: Inverse participation ratio, P^{-1} , for the TLG. Close to the isostatic point only high frequency modes are localised. As density increases some of the low frequency modes become quasi-localised.

discussed in more detail below.

One way to quantify the localisation is to measure the inverse participation ratio

$$P^{-1} = \sum_{i} \nu_i^2(\omega) = \sum_{i} |\vec{e}_{i\omega} \cdot \vec{e}_{i\omega}|^2$$
(5.3)

The eigenvectors are normalised such that $\sum_i \nu_i(\omega) = 1$ for all ω , so a mode localised entirely on one site would have $P^{-1} = 1$ and an homogenous mode would have $P^{-1} = N^{-1}$. Figure 5.5 suggests that near the rigidity threshold the highest frequency modes are strongly localised in space compared to all other frequencies. At higher densities all the modes become more extended except for some of the low frequency modes that start to become quasi-localised.

More detailed information on the localisation can be obtained using the method of Ref. [94], where a mode localisation length is extracted by tracking the change in its eigenvalue due to asymmetric perturbations of the dynamical matrix [95]. The eigenvalues of the symmetric dynamical matrix are all real. An asymmetry term is added to the matrix such that

$$\mathcal{M}_{ij} \to e^{\dot{h} \cdot \hat{r}_{ij}^0} \mathcal{M}_{ij} \tag{5.4}$$

where \vec{h} is the asymmetry parameter. It is shown in detail in Ref. [94] that provided the localisation length of a mode is small compared to the asymmetry, $\xi_{\omega} < 1/h$, then the eigenvalue, ω^2 , will be unaffected. Starting with a small value for h, it can gradually be turned up until the eigenvalue collides with another one nearby and moves into the complex plane. By monitoring the value for h where this occurs we can extract a localisation length,



Figure 5.6: Localization length of the vibrational modes of the (2)-TLG, calculated using the technique or Ref. [94]. The dashed line indicates the box length, L = 50. As the density decreases towards the isostatic point more modes become localized; for densities below the isostatic one the system breaks up into disconnected elastic components.

 $\xi_{\omega} \approx 1/h.$

Figure 5.6 (top) shows the change in the localisation length computed in this way for different densities in the (2)-TLG. In the full lattice limit, $\rho = 1$, there are no vibrational anomalies and all the modes are extended. As soon as a small density of vacancies is present some of the modes become localised, as illustrated in Fig. 5.6 for $\rho = 0.95$. More modes become localized as the density of particles decreases. Near to the isostatic point most modes have localization lengths that are smaller than our system size, see top right panel of Fig.5.6. Beyond the isostatic density the system becomes elastically unstable.

5.4 Comparison of the Two Models

The propensity dynamics of the (2)-TLG and the elastic properties of the site diluted triangular lattice (and their three dimensional counterparts) are well studied, but separate problems. Their common connection is through the lattice structure and the directional element of the interactions. Both models have spatially non-trivial dynamics as shown in Fig. 5.1 for the (2)-TLG and in Fig. 5.4 for the CFN. The spatial weight of the modes of the CFN, where the elastic heterogeneity becomes most apparent, are shown in Fig. 5.7 overlayed on the propensity map from Fig. 5.1. The picture suggests a close correlation between areas of high propensity for relaxational motion and the location of soft modes. Fig. 5.7 (bottom-right) shows a similar spatial correlation between regions of low propensity (shown as the inverse of Fig. 5.7, top-left) and the location high frequency modes.



Figure 5.7: Clockwise from top left: (a) Propensity map of the (2)-TLG at density $\rho = 0.75$ averaged over 100 trajectories for the same configuration as Fig. 5.1. Black circles indicate the isoconfigurational average distance travelled after time $t = 5 \times 10^3$. (b) The average participation, $\nu_i(\omega)$ of the lowest frequency modes, $\omega < 0.1$, and (c) the highest frequency modes, $\omega > 2.4$, of the central force network. (d) The connected correlation, $C(t,\omega)$, between isoconfigurational persistence, $p_i^{\rm IC}(t)$, and the participation, $\nu_i(\omega)$, as defined in the text. From bottom to top the frequencies are $\omega = 0.02, 0.1, 1, 2, 2.3, 2.4$.

In order to quantify this spatial correlation we define the following cross correlation between local propensity for motion and vibrations:

$$C(t,\omega) = \langle p_i^{\rm IC}(t)\nu_i(\omega)\rangle - \langle p_i^{\rm IC}(t)\rangle.$$
(5.5)

Here $p_i(t)$ denotes the persistence function [47] of particle *i*, i.e. $p_i(t) = 1$ if particle *i* has not moved up to time *t*, and $p_i(t) = 0$ otherwise. $p_i^{IC}(t)$ is the iso-configurational average of $p_i(t)$, i.e. the average of the persistence field over all trajectories that start from a given configuration [48]. The average $\langle \cdot \rangle$ in Eq. (5.5) is over all equilibrium configurations at a given density. Fig. 5.7 (bottom-left) shows $C(t,\omega)$ as a function of time *t* for various vibrational frequencies for the (2)-TLG at density $\rho = 0.75$. For high frequencies the correlation is positive, indicating that particles which are more persistent than average also participate in high frequency vibrational modes. For low frequencies the correlation is negative, indicating that fast relaxing particles (low p_i^{IC}) are also those which participate in soft vibrations. In both cases the correlation is non-monotonic in time, peaking at times around τ_{α} .

Figure 5.8 shows $C(t, \omega)$ over the full range of frequencies for the (2)-TLG at density $\rho = 0.75$ and for the (4)-FLG at density $\rho = 0.6$. In both plots the correlation is strongest at high frequencies, however, this is not always true. Figure 5.9 shows how the spatial correlation between vibrations and dynamic heterogeneity depends on density. It plots the peak value of $C(t, \omega)$ for vibrational modes of the lowest and highest non-trivial frequencies accessible in the simulations. The fast-relaxation/soft-mode correlation increases with density for all densities larger than the isostatic one, $\rho > \rho_c$. In contrast, the slow-relaxation/high-frequency correlation decreases with increasing density. These trends are similar in dimension two, (2)-TLG, and dimension three, (4)-FLG. Below the isostatic point the correlation between slow regions and low frequency modes changes sign as the vibrational spectrum becomes plagued by zero modes.

An alternative way to visualise the relationship between fast (slow) dynamics and low (high) frequency vibrations is in the structure of the modes themselves. The eigenvectors are normalised so that $\sum_i \nu_i(\omega) = 1$ and the symmetry of the dynamical matrix means that the same is also true in the frequency basis, $\sum_{\omega} \nu_i(\omega) = 1$. This has the consequence that if a particle features strongly in a mode at one frequency it must play a smaller role at other frequencies. When this is combined with information on the particle dynamics we see the emergence of an ordered structure in the modes. Figure 5.10 shows a density map of the eigenvectors for the (2)-TLG at $\rho = 0.75$ where the particle basis has been sorted by the average mean square displacement after $t = 10^3$ time steps. In agreement with the results above we see the variation from a flat distribution at the extremes of frequency and propensity.

Similar results to ours have been found in atomistic simulations of glassy systems. Studies on hard sphere systems have shown that sudden movements, collective jumps referred to as avalanches, often follow the direction of the low frequency normal modes [75]. On a longer time scale, a study on a binary mixture of soft discs [77] found that parts of the liquid that were able to make large structural rearrangements correlated well to areas where the anomalous modes were strong. It has been suggested that the link is a causal



Figure 5.8: Site correlation, $C(t, \omega)$, between the isoconfigurational persistence, $p_i^{\rm IC}(t)$, of particle *i* and the participation, $\nu_i(\omega)$ of the same particle in the normal modes at $\rho = 0.75$ for the (2)-TLG (top) and $\rho = 0.6$ for the (4)-FLG (bottom). Particles that tend not to move until late times are correlated with high frequency modes with a peak around τ_{α} . Particles that tend to move early on are correlated with low frequency modes with a broad trough just before τ_{α} .



Figure 5.9: Peak in the absolute value of the correlation $C(t, \omega)$ between propensity for relaxational motion and vibrational modes, in the (2)-TLG and (4)-FLG, as a function of density; cf. Fig.5.8. The correlation between fast regions/high-frequency modes decreases with increasing density. The correlation between slow regions/low-frequency modes increases with increasing density.



Figure 5.10: Coarse grained density map of the matrix of eigenvectors, $|\vec{e}_{i\omega}|^2$, for the TLG at density $\rho = 0.75$. The particle basis, *i*, has been sorted by the isoconfigurational mean square displacement, $\langle \Delta \vec{r}_i^2 \rangle_{\rm IC}$ at $t = 10^3$. White regions indicate a higher than average participation.

one, whereby the anomalous modes are the channels through which structural relaxation can occur. In our model, however, relaxation is via the propagation of localized clusters of vacancies [47, 48, 50], and does not emanate from the soft vibrational modes. In fact, it is precisely the presence of these localized defects in the elastic network that gives rise to the anomalous soft modes. (5.8).

5.5 Nonaffine Displacement

In this final section we take another perspective on the elastic heterogeneity already observed in the normal mode structure by considering the non-affine response of our system to an external strain. At a large enough length scale a material can be viewed as a continuous elastic material with a constant elastic-modulous. The response to an applied stress on the boundary will be a uniform, affine strain [78,96]. At the level of particles, where local defects disturb the elastic-modulous, the individual response will in general be non-affine.

We follow the technique of Ref. [97] where instead of applying a stress to the edges of our system we maintain the periodic boundary conditions and directly apply an affine strain so that all particle positions are transformed, $\vec{r}_{i0} \rightarrow \vec{r}'_{i0}$. In this case we perform a stretch in the *x*-direction of magnitude, $\epsilon \ll 1$, with an equal compression along the *y*-axis to maintain a constant volume. After the stretch we minimize the energy using conjugate gradient descent. Because of the periodic boundary conditions, once the affine displacement has been made the system will not be able to relax back to its original configuration. The final positions of the particles can be expressed as the sum of an affine part, \vec{r}'_{i0} and a nonaffine part, \vec{u}_i ,

$$\vec{r}_i' = \vec{r}_{i0}' + \vec{u}_i \tag{5.6}$$

where the overall affinity ensures that $\sum_{i} \vec{u}_{i} = 0$.

For high densities the conjugate gradient algorithm is very efficient at finding the minimum energy configuration. As density is lowered towards the rigidity threshold the time to converge starts to increase rapidly. This is due to the formation of weakly connected clusters that can make large displacements at a small energy cost. These clusters are the beginnings of the floppy modes that are able to operate almost independently from the rest of the system. At the isostatic point gradient descent cannot converge. To prevent these divergences we have added a small confining harmonic potential of force constant Ω at each site to Eq. (5.1) so that

$$H = H_{\rm vib} + \sum_{i} \frac{1}{2} \Omega^2 |\vec{u}_i|^2$$
 (5.7)

This extra term is diagonal, and so the normal modes remain unchanged but with a shift in frequency. An example of the kind of field, \vec{u}_i , obtained using this method is shown in Fig. 5.12 for a (2)-TLG at density $\rho = 0.75$ and confining potential, $\Omega^2 = 10^{-3}$. Here we have used a stretching parameter, $\epsilon = 10^{-4}$, although it should be noted that once the non-affine displacement field has been rescaled the particular choice for ϵ is not important provided it is small. The fields produced contain vortex patterns, reminiscent of the low



Figure 5.11: Average non-affine response, $\langle (\vec{u}_i/\epsilon)^2 \rangle^{1/2}$, as a function of density in the (2)-TLG. The different curves are for different strengths Ω of the confining potential. In the limit of $\Omega \to 0$, the non-affine response appears to diverge at ρ_c . Inset: Projection of the non-affine displacement into the vibrational eigenmodes, at $\rho = 0.75$ for $\Omega = 10^{-5}$.

frequency modes of Fig. 5.4, the shape of which are dependent on the confining potential used.

Figure 5.11 shows the average magnitude of the non-affine response, $\langle (\vec{u}_i/\epsilon)^2 \rangle^{1/2}$, as a function of the density in the (2)-TLG model for confinement potentials $\Omega^2 = 10^{-3}$, 10^{-4} and 10^{-5} . Close to the isostatic point the average non-affine deformation increases very rapidly, and would appear to diverge when $\Omega \to 0$. The inset to Fig. 5.11 has the projection of the non-affine deformation on the normal modes, $\langle u|\omega\rangle \equiv \sum_i \vec{u}_i \cdot \vec{e}_{i\omega}$, as a function of frequency of the modes, for the (2)-TLG at density $\rho = 0.75$ and the smallest confinement, $\Omega^2 = 10^{-5}$. As is suggested by Fig. 5.12 this shows that non-affinity is carried preferably by softer modes.



Figure 5.12: The nonaffine displacement field after an affine stretch along the *x*-axis for the (2)-TLG at density $\rho = 0.75$ and confinement potential $\Omega^2 = 10^{-3}$.

5.6 Conclusion

In this chapter we have studied the correlation between dynamic heterogeneity and anomalous vibrations in a two and a three dimensional constrained lattice gas model of glasses. The structural relaxation of these models at high density is similar to that of glass formers, displaying non-exponential relaxation and dynamic heterogeneity [47, 48]. Their vibrational properties are those of well-studied random networks [89], and mimic characteristic aspects of the anomalous vibrations of glasses: excess low frequency modes, non-affinity and elastic heterogeneity, all related to the presence of an isostatic point. We have found that the location of anomalous vibrational modes correlates to dynamic heterogeneity of structural relaxation, as is observed in atomistic systems [75–77]. In our case, however, structural relaxation, and therefore dynamic heterogeneity, originate in localized vacancies [48, 50], and not in the extended structures that the soft-modes span. In fact, vacancies act as quenched localized defects for the vibrations, cf. Eq. (5.1), giving rise to the anomalous elastic behaviour observed. We have thus shown through these simple examples that a correlation between soft modes and propensity does not imply a causal relation for relaxation mechanisms. A similar situation may hold in atomistic models as well.

Chapter 6

Conclusion

The aim of this thesis was to create new models, building on the kinetically constrained models, that can describe both the fast and slow dynamical regimes in glass forming systems. KCMs have been very successful in describing slow processes but up to now have not been able to comment on fast processes such as the β -relaxation. Here we have taken significant steps in this direction.

As well as extending the KCMs we have also studied the properties of existing models. In chapter 2 we exploited the simple nature of the facilitated spin models to introduce a new algorithm that can speed up numerical simulation by several orders of magnitude. This allowed us to test theoretical predictions in regimes previously unreachable with existing techniques. In chapter 3 we obtained new results for a three dimensional constrained lattice gas on an FCC lattice. The reason for studying this model was that the kinetic constraints arise from explicit steric interactions between particles. This is an important detail as it provides a natural way to include vibrations in chapter 5.

In chapter 4 we added fast degrees of freedom to several KCMs. These degrees of freedom were coupled Ising spins that interacted with the underlying KCM in such a way that excitations in the KCM caused disorder in the fast variables. We showed that it is possible to reproduce much of the fast dynamical behaviour seen in real glasses without significantly changing slow dynamics. The new variables were also able to account for some thermodynamic features of glasses that are not described by standard KCMs.

In order to study vibrations we added harmonic interactions into two constrained lattice gas models. We found that the resulting elastic network contains features similar to those found in glasses such as the Boson peak. The low frequency vibrational modes that form the Boson peak also correlated with regions of high propensity for motion. This is something that has only recently been observed in simulations of real glasses and it is remarkable that such a simplified system can reproduce this behaviour.

Over the course of this thesis we have shown that it is possible to consider both fast and slow dynamics using simple models. By adding the minimum number of fast degrees of freedom we have been able to undo some of the coarse graining that goes into constructing a KCM giving us extra confidence that KCMs are good models for glass formers. Future work that can improve our understanding of how dynamics cross from fast to slow would be very worth while.

Appendix A

Transfer Matrices

In this appendix we will go through, in detail, the exact solution to the thermodynamics of the ns-model in one dimension. This is done using the transfer matrix approach; for a clear introduction see Ref. [66]. We will first go through some general results for transfer matrices using our model as an example. Finally the model-specific results will be calculated.

The basic idea of transfer matrices is to show that the sum over all possible configurations that is in the partition function is equivalent to taking the matrix product of a series of matrices. For the ns-model in one dimension the Hamiltonian can be written

$$H = h \sum_{i} n_{i} - J \sum_{i} (1 - n_{i})(1 - n_{i+1}) s_{i} s_{i+1}$$
(A.1)

The partition function, $Z = \sum e^{-\beta H(n,s)}$, summed over all possible configurations, can be written in the more symmetrical form

$$Z = \sum_{\{n,s\}} \prod_{i=0}^{N-1} \exp\left[-\beta \left(\frac{h(n_i + n_{i+1})}{2} - J(1 - n_i)(1 - n_{i+1})s_i s_{i+1}\right)\right]$$
(A.2)

where β is inverse temperature, $\beta \equiv 1/T$, and the sum inside the Hamiltonian has been turned into a product outside the exponential. Given that n_i and s_i are both binary variables we define a variable, $\sigma_i = f(n_i, s_i)$, that can take an index between 0 and 3 to represent the entire state of site *i*. The partition function can then be rewritten as $Z = \sum_{\{\sigma\}} \prod_i T(\sigma_i, \sigma_{i+1})$, where *T* takes the summand of Eq.A.2. To get a better feel for this one can write the product and the sum explicitly

$$Z = \sum_{\sigma_0=0}^{3} \sum_{\sigma_1=0}^{3} \cdots \sum_{\sigma_{N-1}=0}^{3} T(\sigma_0, \sigma_1) T(\sigma_1, \sigma_2) \cdots T(\sigma_{N-1}, \sigma_0)$$
(A.3)

and it is now clear that the summations are equivalent to a matrix product over N identical transfer matrices, **T**, that contains all 16 possible terms between site i and i + 1. The final summation over σ_0 , thanks to the periodic boundaries, takes the trace. This gives a

greatly compacted partition function as

$$Z = \text{Tr}(\mathbf{T}^N) = \langle \sigma_0 | T^N | \sigma_0 \rangle = \sum_{k=0}^3 \lambda_k^N$$
(A.4)

where λ_k is the k^{th} eigenvalue of **T**. From here on we will use Dirac notation where $|\sigma\rangle$ corresponds to the spin states (with a suffix, $|\sigma_x\rangle$, this refers to a specific spin x) and $|i\rangle$ corresponds to the eigenvectors of T.

A.1 General results

A.1.1 Bulk properties

If we write the eigenvalues in order of size, $\lambda_0 > \lambda_1 > \lambda_2 \cdots$, and factorise the largest out, we can rewrite the sum as

$$Z = \lambda_0^N \left[1 + \sum_{k=1}^3 \frac{\lambda_k^N}{\lambda_0^N} \right] \tag{A.5}$$

and taking the thermodynamic limit, $N \to \infty$, we can recover the free energy as

$$F = -\frac{N}{\beta} \ln \lambda_0 \tag{A.6}$$

All of the bulk thermodynamic quantities come directly from the free energy. For some microscopic quantities we will need more information than just the largest eigenvalue as we will see in a moment. Of greatest interest to us are the average site energy

$$E = -\frac{\partial \ln \lambda_0}{\partial \beta} \tag{A.7}$$

specific heat capacity

$$c_v = \beta^2 \frac{\partial^2 \ln \lambda_0}{\partial \beta^2} \tag{A.8}$$

and the specific entropy

$$S = \ln \lambda_0 - \beta \frac{\partial \ln \lambda_0}{\partial \beta} \tag{A.9}$$

A.1.2 Magnetisation

The magnetisation for a general variable, σ , requires a little extra work. We start with the weighted sum

$$\langle \sigma_0 \rangle = \frac{\sum_{\{\sigma\}} \sigma_0 e^{-\beta H(\sigma)}}{Z} \tag{A.10}$$

which in transfer matrix terms becomes

$$\langle \sigma_0 \rangle = \frac{\sum_{\sigma_0} \sigma_0 \langle \sigma_0 | T^N | \sigma_0 \rangle}{\lambda_0^N} \tag{A.11}$$

If we write σ_0 as an operator such that $\sigma_0 = \langle \sigma_0 | \hat{\sigma}_0 | \sigma_0 \rangle$ and switch into the diagonal basis, then this becomes

$$\langle \sigma_0 \rangle = \frac{\sum_i \langle i | \hat{\sigma}_0 \lambda_i^N | i \rangle}{\lambda_0^N} \tag{A.12}$$

where $|i\rangle$ refers to the i^{th} eigenvector. In the thermodynamic limit only the i = 0 term survives and so the general result is

$$\langle \sigma_0 \rangle = \langle 0 | \hat{\sigma_0} | 0 \rangle \tag{A.13}$$

To get the magnetisation we simply project the $\hat{\sigma}_0$ operator into the diagonal basis and read off the 0,0 element.

A.1.3 Spatial Correlation

Here we are interested in the quantity $\langle \sigma_0 \sigma_R \rangle$ as a measure of the spatial correlation over a distance R. In the transfer matrix notation this is most easily written down by breaking up the product

$$\langle \sigma_0 \sigma_R \rangle = \sum_{\{\sigma\}} \sigma_0 \langle \sigma_0 | T^R | \sigma_R \rangle \sigma_R \langle \sigma_R | T^{N-R} | \sigma_0 \rangle / \lambda_0^N \tag{A.14}$$

and then jumping straight into the diagonal basis

$$\langle \sigma_0 \sigma_R \rangle = \sum_i \langle 0 | \hat{\sigma}_0 | i \rangle \left(\frac{\lambda_i}{\lambda_0} \right)^R \langle i | \hat{\sigma}_R | 0 \rangle \tag{A.15}$$

The i = 0 term gives the separate magnetisations. We are usually interested in the function, $\Gamma_{\sigma}(R) = \langle \sigma_0 \sigma_R \rangle - \langle \sigma_0 \rangle \langle \sigma_R \rangle$, which is given by excluding the first term

$$\Gamma_{\sigma}(R) = \sum_{i \neq 0} \langle 0 | \sigma_0 | i \rangle \left(\frac{\lambda_i}{\lambda_0} \right)^R \langle i | \sigma_R | 0 \rangle \tag{A.16}$$

and so for spatial correlations we need the smaller eigenvalues as well.

A.2 Results for the ns-model

To calculate the specific results for our system we need the explicit form for \mathbf{T} which is given by

$$\mathbf{T} = \begin{pmatrix} z^{-1} & z^{-1} & z^{-1/2} & z^{-1/2} \\ z^{-1} & z^{-1} & z^{-1/2} & z^{-1/2} \\ z^{-1/2} & z^{-1/2} & x & x^{-1} \\ z^{-1/2} & z^{-1/2} & x^{-1} & x \end{pmatrix}$$
(A.17)

where $z \equiv e^{\beta h}$ and $x \equiv e^{\beta J}$. The top left quadrant is where both n_i and $n_{i+1} = 1$. Only in the bottom right quadrant where they are both equal to 0 do we see the terms related to s_i and s_{i+1} . For a given n_i and s_i the index of the matrix is $2n_i + (s_i + 1)/2$. The eigenvalues, in order largest to smallest, are

$$\lambda_{0} = \frac{zx^{2} + 2x + z + \sqrt{(z + x(xz + 2))^{2} - 8(x - 1)^{2}xz}}{2xz}$$

$$\lambda_{1} = x - x^{-1}$$

$$\lambda_{2} = \frac{zx^{2} + 2x + z - \sqrt{(z + x(xz + 2))^{2} - 8(x - 1)^{2}xz}}{2xz}$$

$$\lambda_{3} = 0$$
(A.18)

and the corresponding eigenvectors are then

$$\vec{u}_{0}^{T} = (v + w, v + w, 1, 1)$$

$$\vec{u}_{1}^{T} = (0, 0, -1, 1)$$

$$\vec{u}_{2}^{T} = (v - w, v - w, 1, 1)$$

$$\vec{u}_{3}^{T} = (-1, 1, 0, 0)$$

(A.19)

where $v \equiv \frac{2x - (x^2 + 1)z}{4x\sqrt{z}}$ and $w \equiv \frac{\sqrt{(z + x(xz + 2))^2 - 8(x - 1)^2xz}}{4x\sqrt{z}}$. The diagonalising matrix $\mathbf{D} = (\vec{u}_0, \vec{u}_1, \vec{u}_2, \vec{u}_3)$ has a simple inverse which makes it easy to transform between the spin basis and the diagonal basis. From here it is straight forwarded to compute the thermodynamic quantities that we are interested in, starting with magnetisations.

A.2.1 Magnetisation

The FA model has a magnetisation that is $\langle n \rangle = (1+z)^{-1}$. Now that we have interactions between the cells we might expect something very different. In the spin basis the \hat{n} operator is defined by the matrix $\mathbf{n}_{00} = \mathbf{n}_{11} = 1$ and all other elements are 0. Using the standard result from Eq. A.13 this gives the concentration of mobility excitations as

$$\langle n \rangle = \langle 0 | \hat{n} | 0 \rangle = \frac{v+w}{2w}$$
 (A.20)

This is plotted in chapter 4 LINK. Expanding to lowest order for $\beta J \ll 1$ we can see the departure from the FA model result.

$$\langle n \rangle \approx \frac{1}{1+z} - \frac{z^2}{(1+z)^3} (\beta J)^2$$
 (A.21)

The concentration of mobile cells is therefore influenced quite strongly by the s-field. Whether or not this changes the underlying physics of the FA model depends more on whether there are any long range static correlations in space. Due to symmetry there is no average magnetisation in the s-field and so $\langle 0|s|0\rangle = \langle s\rangle = 0$.

A.2.2 Spatial Correlations

Again, here we can directly apply the results from general transfer matrix theory. For correlations in the FA mobility field we can reuse the operator from the previous section. When we put this into Eq. A.16 we get

$$\Gamma_{\sigma}(R) = \sum_{i \neq 0} \langle 0|\hat{n}|i\rangle \left(\frac{\lambda_i}{\lambda_0}\right)^R \langle i|\hat{n}|0\rangle \tag{A.22}$$

The sum is usually cutoff at i = 1 as this is by far the most dominant term. However, for the \hat{n} operator $\langle 0|n|1 \rangle = 0$ and so we rely on the next biggest eigenvector, λ_2 , to get

$$\Gamma_n(R) = \frac{1}{4} \left(1 - \frac{v^2}{w^2} \right) \left(\frac{\lambda_2}{\lambda_0} \right)^R \tag{A.23}$$

which decays quite rapidly with increasing R because for most systems $\frac{\lambda_2}{\lambda_0} \ll 1$

To calculate the correlations for the *s*-field we define the *s*-matrix in the spin basis as $s_{ij} = \delta_{ij}(-1)^j$ and putting this into Eq. A.16 gives

$$\Gamma_s(R) = \frac{w - v}{2w} \left(\frac{\lambda_1}{\lambda_0}\right)^R \tag{A.24}$$

For R = 0 the term involving $\lambda_3 = 0$ should be included as it does not vanish there. This retrieves $\Gamma_s(0) = 1$ as required. This function also gives the density of domain walls, $\langle s_i s_{i+1} \rangle$, for R = 1.

Finally, a possible quantity of interest is the static correlation between a domain wall and an n excitation. This is obtained from

$$\langle n_0 s_R s_{R+1} \rangle = \sum_{j,k} \langle 0|n|j \rangle \left(\frac{\lambda_j}{\lambda_0}\right)^R \langle j|s|k \rangle \frac{\lambda_k}{\lambda_0} \langle k|s|0 \rangle \tag{A.25}$$

The only R dependent term that is nonzero in this sum depends on $(\lambda_2/\lambda_0)^R$ which, as we have seen, quickly decays to zero.

A.3 Equilibrium configurations

The transfer matrix technique can also be used to generate equilibrium configurations. This is done by reducing the terms in the partition function every time we make a choice for a particular site. For a general one-dimensional system, with no other lattice sites fixed, the probability distribution for the first site comes from the diagonal elements of \mathbf{T}^N in the spin basis

$$P(\sigma_0) = \frac{\langle \sigma_0 | T^N | \sigma_0 \rangle}{\lambda_0^N} \tag{A.26}$$

A choice of σ_0 is taken from this distribution and then fixed. For the next site σ_0 is now treated as a constant that cannot be summed over

$$P(\sigma_1|\sigma_0) = \frac{\langle \sigma_0|T|\sigma_1\rangle\langle \sigma_1|T^{N-1}|\sigma_0\rangle}{\langle \sigma_0|T^N|\sigma_0\rangle}$$
(A.27)

and in turn for the next spin

$$P(\sigma_2|\sigma_0,\sigma_1) = \frac{\langle \sigma_0|T|\sigma_1\rangle\langle \sigma_1|T|\sigma_2\rangle\langle \sigma_2|T^{N-1}|\sigma_0\rangle}{\langle \sigma_0|T|\sigma_1\rangle\langle \sigma_1|T^{N-1}|\sigma_0\rangle}$$
(A.28)

The pattern continues until all the spins have been fixed. The trickiest part of this method is dealing with the matrix $\langle \sigma_x | T^{N-x} | \sigma_0 \rangle$. In practice it can be calculated cheaply using the approximation, $(\lambda_1/\lambda_0)^{N-x} \approx 0$, provided that N-x is reasonably large. For the last few spins, where N-x is small, the matrix can be calculated by direct multiplication. In this way equilibrium configurations can be generated quickly without the need for slowly converging techniques such as Monte Carlo.

Appendix B

Cluster Algorithms

In this appendix we will outline the use of cluster-flip Monte Carlo algorithms to study the Ising model near criticality. As with all specialist algorithms these techniques make use of the underlying physics near a critical point to improve efficiency. Before we go into cluster algorithms it will be a good idea to remind ourselves of some of the basic concepts. For a good review on Monte Carlo techniques, including those presented here, see Ref. [25].

B.1 Detailed Balance

If we are resorting to numerical methods to solve a problem then it is a fair bet that the space of all possible configurations is very large. Even a modest simulation of an Ising model on a 50×50 lattice has roughly 10^{75} possible states. It is therefore necessary to make sure that we mostly sample statistically likely states. This is called importance sampling and the best way to do it is to visit states with a probability proportional to their Boltzmann weight. The best way to achieve this is through detailed balance.

Detailed balance says that, if we are sampling correctly in equilibrium, there should be no net current through any given state. A way of guarenteeing this is to set

$$p_{\mu}P(\mu \to \nu) = p_{\nu}P(\nu \to \mu) \tag{B.1}$$

where $P(\mu \rightarrow \nu)$ is the probability of moving to state ν given we are in state μ . The stationary probabilities are taken from the Boltzmann distribution such that

$$\frac{P(\mu \to \nu)}{P(\nu \to \mu)} = e^{-\beta(E_{\nu} - E_{\mu})} \tag{B.2}$$

and so provided we keep faithful to these transition rates we will visit all states with the desired probability and sample the system in a manner that is faithful to the Boltzmann distribution.

B.2 The Metropolis algorithm

Given that it is only the ratio of transition rates that are restricted the Metropolis algorithm seeks to maximise the number of transitions by making each term as large as possible. Within Metropolis we attempt to change the system from state μ to state ν and we accept the move with rate

$$P_{\text{accept}} = \min\{1, e^{-\beta(E_{\nu} - E_{\mu})}\}$$
(B.3)

satisfying detailed balance. The algorithm works best when it is accepting lots of moves, this means only attempting small changes in the configuration. Because it uses small steps Metropolis can be used to give information about the dynamics a system as well, for example single spin flips for an Ising model. This is something we make a lot of use of in this thesis.

Provided the dynamics of the system are relatively fast and the free energy landscape does not have too many traps then the Metropolis algorithm is very efficient. Around a critical point, where lengthscales are diverging and relaxation times growing, it begins to run into problems. To make progress under these conditions it is necessary to attempt much bigger changes.

B.3 Wolff algorithm

Considering an Ising model around the critical temperature, we know that the dynamics slow down as $\tau \sim \xi^z$, where ξ is the correlation length and z the dynamical critical exponent. Near to T_c this will be limited by the size of our system but none-the-less the timescales are still very long due to large clusters of correlated spins developing that are very difficult to flip over one at a time. The Wolff algorithm attempts to get around this by flipping a large number of spins all at once. If chosen carefully it is even possible to do this with an acceptance rate of unity.

The algorithm starts by choosing a spin at random and using this as the starting point of our cluster. If we were to flip this spin we would break some bonds, say m bonds, with neighbouring aligned spins, and create n new bonds with anti-aligned spins. Satisfying detailed balance for this move implies that

$$\frac{P(\mu \to \nu)}{P(\nu \to \mu)} = e^{-\beta(E_{\nu} - E_{\mu})} = e^{-2\beta J(m-n)}$$
(B.4)

Instead of breaking a bond with an aligned neigbour we can try and add it to the cluster with a probability P_{add} . If successful then upon flipping, the energy between these spins will not change. For each added spin we can try to add any of its aligned neighbours that have become adjacent to the cluster, repeating recursively until eventually there are no new bonds to try. If we were to flip at this point then we would again be breaking a number of bonds, m, and making a number of new bonds, n. The probability of selecting such a state is $g(\mu \to \nu) = (1 - P_{add})^m$, and similarly, the probability of selecting the reverse move is $g(\nu \to \mu) = (1 - P_{add})^n$.

We can write the transition probability as a product of selection and acceptance probabilities

$$P(\mu \to \nu) = g(\mu \to \nu)A(\mu \to \nu) \tag{B.5}$$

such that the detailed balance equation becomes

$$\frac{g(\mu \to \nu)A(\mu \to \nu)}{g(\nu \to \mu)A(\nu \to \mu)} = (1 - P_{\text{add}})^{m-n} \frac{A(\mu \to \nu)}{A(\nu \to \mu)} = e^{-2\beta J(m-n)}$$
(B.6)

The clever bit is then to notice that if we set

$$P_{\rm add} = 1 - e^{-2\beta J} \tag{B.7}$$

then we are simply left with $A(\mu \to \nu) = A(\nu \to \mu)$ and we are free to set both to 1. Therefore the Wolff algorithm therefore consists of recursively adding aligned spins to the cluster with probability P_{add} and then when this is finished flipping all the spins at once with unit probability.

The marvelous thing about the Wolff algorithm is that the sizes of the clusters that it flips grow inline with the growing correlation length making it immune to the critical slowing down that plagued the Metropolis algorithm. At $T = T_c$ it will generate clusters of all sizes and below T_c the clusters will typically span the system (thus exploring both the up and down magnetised states equally). In fact one way of knowing that we have reached T_c could be to say that it corresponds to the value of P_{add} that just creates a spanning cluster. This is exactly what the invaded cluster algorithm does.

B.4 Invaded cluster algorithm

If it is not the thermodynamic properties of the Ising model near T_c that we are interested in, but the value of T_c itself, then it is possible to turn the cluster algorithm on its head using the idea of the percolating cluster. If instead of knowing a temperature we happen to know the value of P_{add} then we can invert the relationship in Eq. B.7 to get

$$T = -2J/\ln(1 - P_{\rm add})$$
 (B.8)

The critical temperature, for a given system size, is marked by the point where one of the clusters *just* spans the system. At this percolation threshold the system will have clusters of all sizes, hence the scale free relaxation dynamics. If we have a means of calculating the required value of $P_{\rm add}$ to achieve percolation, then we obtain a value for T_c . The invaded cluster algorithm provides a way to do just this.

Before going through how the algorithm works we must first introduce the cluster algorithm by Swendsen and Wang (SW). The physics of this algorithm works in exactly the same way as the Wolff algorithm but instead of growing one cluster at a time and flipping with unit probability, the SW algorithm creates all clusters in one go. It does this by considering every bond between aligned spins and either keeping it with probability $P_{\rm add}$ (the same as $P_{\rm add}$ as with Wolff), or breaking it. This procedure puts every spin into a cluster, possibly on its own, possibly very large, and each cluster is flipped with probability 1/2. The SW algorithm is more or less equivalent to the Wolff algorithm, Wolff is slightly easier to implement so it is usually the favoured choice.

The invaded cluster algorithm mimics the SW algorithm by adding bonds and creating clusters. However, in this case we are trying to calculate a temperature so we do not know the value of P_{add} . To find it we simply add bonds at random into the system and monitor the clusters that are forming. After a while we will have added enough bonds that one of the clusters will wrap around the system (assuming periodic boundaries) and form a spanning cluster. As soon as this is achieved we stop adding bonds and calculate P_{add} from $P_{\text{add}} = N_{\text{added}}/N$, where N_{added} is the total number of bonds we added and N is the number of bonds.

What we have effectively calculated, from Eq. B.8, is the temperature that is required to form a spanning cluster for a particular configuration of spins. The next step is to flip the clusters with probability 1/2 creating a new configuration, characteristic of this temperature, and then start again. We will now see how this pushes the system towards the critical point.

Starting with a zero temperature configuration (all spins aligned) on a square lattice the value of P_{add} that creates a percolating cluster is known to be 1/2. This is the smallest number we can get and as such corresponds to a high temperature, T = 2.9J the configuration after flipping the clusters will be quite disordered. On the next iteration we have the opposite; it will be much more difficult to create a percolating cluster (because there are fewer available bonds) and we are likely to come up with a value for P_{add} that is much higher (and so T much lower). The net effect is for the algorithm to drive itself towards T_c and fluctuate around it. After waiting a short while for this to equilibrate we can then calculate the average temperature and estimate T_c to a high accuracy.

Bibliography

- P. G. Debenedetti and F. H. Stillinger. "Supercooled liquids and the glass transition." Nature, 410(6825), 259 (2001). ISSN 0028-0836.
- [2] M. D. Ediger, C. A. Angell and S. R. Nagel. "Supercooled liquids and glasses." J. Phys. Chem., 100, 13200 (1996).
- [3] C. A. Angell. "Formation of glasses from liquids and biopolymers." Science, 267(5206), 1924 (1995).
- [4] S. F. Swallen, K. L. Kearns, M. K. Mapes, Y. S. Kim, R. J. McMahon, M. D. Ediger, T. Wu, L. Yu and S. Satija. "Organic glasses with exceptional thermodynamic and kinetic stability." Science, **315**(5810), 353 (2007).
- [5] I. E. T. Iben, D. Braunstein, W. Doster, H. Frauenfelder, M. K. Hong, J. B. Johnson, S. Luck, P. Ormos, A. Schulte, P. J. Steinbach, A. H. Xie and R. D. Young. "Glassy behavior of a protein." Phys. Rev. Lett., 62(16), 1916 (1989).
- [6] W. Kob. "Computer simulations of supercooled liquids and glasses." J. Phys. Conds. Matt., 11(10), R85 (1999). ISSN 0953-8984.
- [7] L. Santen and W. Krauth. "Absence of thermodynamic phase transition in a model glass former." Nature, 405(6786), 550 (2000). ISSN 0028-0836.
- [8] C. A. Angell, K. L. Ngai, G. B. McKenna, P. F. McMillan and S. W. Martin. "Relaxation in glassforming liquids and amorphous solids." J. Appl. Phys., 88(6), 3113 (2000).
- [9] E. R. Weeks, J. C. Crocker, A. C. Levitt, A. Schofield and D. A. Weitz. "Threedimensional direct imaging of structural relaxation near the colloidal glass transition." Science, 287(5453), 627 (2000).
- [10] M. D. Ediger. "Spatially heterogeneous dynamics in supercooled liquids." Annu. Rev. Phys. Chem., 51, 99 (2000).
- [11] D. N. Perera and P. Harrowell. "Relaxation dynamics and their spatial distribution in a two-dimensional glass-forming mixture." J. Chem. Phys., 111(12), 5441 (1999).
- [12] W. Gotze and L. Sjogren. "Relaxation processes in supercooled liquids." Rep. Prog. Phys., 55, 241 (1992).

- [13] L. O. Hedges, L. Maibaum, D. Chandler and J. P. Garrahan. "Decoupling of exchange and persistence times in atomistic models of glass formers." J. Chem. Phys., 127(21), 211101 (2007).
- [14] D. R. Reichman and P. Charbonneau. "Mode-coupling theory." J. Stat. Mech., 2005(05), P05013 (2005). ISSN 1742-5468.
- [15] W. Gotze. "Recent tests of the mode-coupling theory for glassy dynamics." J. Phys. Conds. Matt., 11(10A), A1 (1999). ISSN 0953-8984.
- [16] W. Schirmacher, G. Diezemann and C. Ganter. "Harmonic vibrational excitations in disordered solids and the iboson peakl." Phys. Rev. Lett., 81(1), 136 (1998).
- [17] V. L. Gurevich, D. A. Parshin and H. R. Schober. "Anharmonicity, vibrational instability, and the boson peak in glasses." Phys. Rev. B, 67(9), 094203 (2003).
- [18] V. Gurarie and A. Altland. "Phonons in random elastic media and the boson peak." Phys. Rev. Lett., 94(24), 245502 (2005).
- [19] M. Wyart, L. E. Silbert, S. R. Nagel and T. A. Witten. "Effects of compression on the vibrational modes of marginally jammed solids." Phys. Rev. E, 72(5), 051306 (2005).
- [20] N. Xu, M. Wyart, A. J. Liu and S. R. Nagel. "Excess vibrational modes and the boson peak in model glasses." Phys. Rev. Lett., 98(17), 175502 (2007).
- [21] A. I. Chumakov, I. Sergueev, U. van Bürck, W. Schirmacher, T. Asthalter, R. Rüffer, O. Leupold and W. Petry. "Collective nature of the boson peak and universal transboson dynamics of glasses." Phys. Rev. Lett., 92(24), 245508 (2004). Boson peak experimental figures.
- [22] F. Ritort and P. Sollich. "Glassy dynamics of kinetically constrained models." Adv. Phys., 52(4), 219 (2003). ISSN 0001-8732.
- [23] D. J. Ashton, T. C. Jarrett and N. F. Johnson. "Effect of congestion costs on shortest paths through complex networks." Phys. Rev. Lett., 94(5), 058701 (2005).
- [24] T. Gleim, W. Kob and K. Binder. "How does the relaxation of a supercooled liquid depend on its microscopic dynamics?" Phys. Rev. Lett., 81(20), 4404 (1998).
- [25] M. Newman and G. Barkema. Monte Carlo Methods in Statistical Physics. Clarendon Press (1999).
- [26] L. Berthier and W. Kob. "The monte carlo dynamics of a binary lennard-jones glassforming mixture." J. Phys. Conds. Matt., 19(20), 205130 (2007). ISSN 0953-8984.
- [27] G. H. Fredrickson and H. C. Andersen. "Kinetic ising model of the glass transition." Phys. Rev. Lett., 53(13), 1244 (1984).
- [28] S. Eisinger and J. Jäckle. "Analytical approximations for the hierarchically constrained kinetic ising chain." J. Stat. Phys., 73(3), 643 (1993).

- [29] A. B. Bortz, M. H. Kalos and J. L. Lebowitz. "A new algorithm for monte carlo simulation of ising spin systems." J. Comput. Phys., 17, 10 (1975).
- [30] M. A. Novotny. "Monte carlo algorithms with absorbing markov chains: Fast local algorithms for slow dynamics." Phys. Rev. Lett., 74(1), 1 (1995).
- [31] M. A. Novotny. A tutorial on advanced dynamic Monte Carlo methods for systems with discrete state spaces. Singapore: World-Scientific (2001).
- [32] G. T. Barkema and M. E. J. Newman. "New monte carlo algorithms for classical spin systems." cond-mat, page 9703179 (1997).
- [33] G. H. Fredrickson and H. C. Andersen. "Facilitated kinetic ising models and the glass transition." J. Chem. Phys., 83(11), 5822 (1985).
- [34] A. Buhot and J. P. Garrahan. "Crossover from fragile to strong glassy behavior in kinetically constrained systems." Phys. Rev. E, 64(2), 021505 (2001).
- [35] J. P. Garrahan. "Glassiness through the emergence of effective dynamical constraints in interacting systems." J. Phys. Conds. Matt., 14(7), 1571 (2002). ISSN 0953-8984.
- [36] L. Berthier and J. Garrahan. "Numerical study of a fragile three-dimensional kinetically constrained model." J. Phys. Chem. B, 109(8), 3578 (2005). ISSN 1520-6106.
- [37] L. Berthier and J. P. Garrahan. "Real space origin of temperature crossovers in supercooled liquids." Phys. Rev. E, 68(4), 041201 (2003).
- [38] J. P. Garrahan and D. Chandler. "Coarse-grained microscopic model of glass formers." PNAS, 100(17), 9710 (2003).
- [39] D. Aldous and P. Diaconis. "The asymmetric one-dimensional constrained ising model: Rigorous results." Journal of Statistical Physics, 107(5), 945 (2002).
- [40] N. Cancrini, F. Martinelli, C. Roberto and C. Toninelli. "Facilitated spin models: recent and new results." cond-mat, page 0712.1934 (2007).
- [41] P. Sollich and M. R. Evans. "Glassy time-scale divergence and anomalous coarsening in a kinetically constrained spin chain." Phys. Rev. Lett., 83(16), 3238 (1999).
- [42] M. Sellitto, G. Biroli and C. Toninelli. "Facilitated spin models on bethe lattice: Bootstrap percolation, mode-coupling transition and glassy dynamics." Europhys. Lett., 69(4), 496 (2005). ISSN 0295-5075.
- [43] W. Kob and H. C. Andersen. "Kinetic lattice-gas model of cage effects in high-density liquids and a test of mode-coupling theory of the ideal-glass transition." Phys. Rev. E, 48(6), 4364 (1993).
- [44] J. Jäckle and A. Krönig. "A kinetic lattice-gas model for the triangular lattice with strong dynamic correlations. i. self-diffusion." J. Phys. Conds. Matt., 6(38), 7633 (1994). ISSN 0953-8984.

- [45] M. Vogel and S. C. Glotzer. "Spatially heterogeneous dynamics and dynamic facilitation in a model of viscous silica." Phys. Rev. Lett., 92(25), 255901 (2004).
- [46] M. T. Downton and M. P. Kennett. "Connecting microscopic simulations with kinetically constrained models of glasses." Phys. Rev. E, 76(3), 031502 (2007).
- [47] A. C. Pan, J. P. Garrahan and D. Chandler. "Heterogeneity and growing length scales in the dynamics of kinetically constrained lattice gases in two dimensions." Phys. Rev. E, 72(4), 041106 (2005).
- [48] L. O. Hedges and J. P. Garrahan. "Dynamic propensity in a kinetically constrained lattice gas." J. Phys. Conds. Matt., 19(20), 205124 (2007). ISSN 0953-8984.
- [49] I. Chang and H. Sillescu. "Heterogeneity at the glass transition: Translational and rotational self-diffusion." J. Phys. Chem. B, 101(43), 8794 (1997). ISSN 1520-6106.
- [50] C. Toninelli, G. Biroli and D. S. Fisher. "Spatial structures and dynamics of kinetically constrained models of glasses." Phys. Rev. Lett., 92(18), 185504 (2004).
- [51] D. Chandler, J. P. Garrahan, R. L. Jack, L. Maibaum and A. C. Pan. "Lengthscale dependence of dynamic four-point susceptibilities in glass formers." Phys. Rev. E, 74(5), 051501 (2006).
- [52] J. L. Barrat, J. N. Roux and J. P. Hansen. "Diffusion, viscosity and structural slowing down in soft sphere alloys near the kinetic glass transition." Chemical Physics, 149(1-2), 197 (1990).
- [53] S. F. Swallen, P. A. Bonvallet, R. J. McMahon and M. D. Ediger. "Self-diffusion of tris-naphthylbenzene near the glass transition temperature." Phys. Rev. Lett., 90(1), 015901 (2003).
- [54] S. R. Becker, P. H. Poole and F. W. Starr. "Fractional stokes-einstein and debyestokes-einstein relations in a network-forming liquid." Physical Review Letters, 97(5), 055901 (2006).
- [55] F. Fernandez-Alonso, F. J. Bermejo, S. E. McLain, J. F. C. Turner, J. J. Molaison and K. W. Herwig. "Observation of fractional stokes-einstein behavior in the simplest hydrogen-bonded liquid." Physical Review Letters, 98(7), 077801 (2007).
- [56] Y. Jung, J. P. Garrahan and D. Chandler. "Dynamical exchanges in facilitated models of supercooled liquids." The Journal of Chemical Physics, 123(8), 084509 (2005).
- [57] Y. Jung, J. P. Garrahan and D. Chandler. "Excitation lines and the breakdown of stokes-einstein relations in supercooled liquids." Phys. Rev. E, 69(6), 061205 (2004).
- [58] C. Toninelli, M. Wyart, L. Berthier, G. Biroli and J.-P. Bouchaud. "Dynamical susceptibility of glass formers: Contrasting the predictions of theoretical scenarios." Phys. Rev. E, **71**(4), 041505 (2005).
- [59] K. Binder and W. Kob. Glassy Materials and Disordered Solids. World Scientific (2005). ISBN 981-256-510-8.

- [60] G. P. Johari and M. Goldstein. "Viscous liquids and the glass transition. iii. secondary relaxations in aliphatic alcohols and other nonrigid molecules." The Journal of Chemical Physics, 55(9), 4245 (1971).
- [61] L. Wu and S. R. Nagel. "Secondary relaxation in o-terphenyl glass." Phys. Rev. B, 46(17), 11198 (1992).
- [62] A. Widmer-Cooper and P. Harrowell. "On the relationship between structure and dynamics in a supercooled liquid." J. Phys. Conds. Matt., 17(49), S4025 (2005). ISSN 0953-8984.
- [63] A. Widmer-Cooper and P. Harrowell. "Predicting the long-time dynamic heterogeneity in a supercooled liquid on the basis of short-time heterogeneities." Phys. Rev. Lett., 96(18), 185701 (2006).
- [64] D. Chandler and J. P. Garrahan. "Thermodynamics of coarse-grained models of supercooled liquids." The Journal of Chemical Physics, 123(4), 044511 (2005).
- [65] 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 (2007).
- [66] J. M. Yeomans. Statistical Mechanics of Phase Transitions. OUP (1992). ISBN-13: 978-0-19-851730-6.
- [67] G. Biroli, J.-P. Bouchaud and G. Tarjus. "Are defect models consistent with the entropy and specific heat of glass formers?" The Journal of Chemical Physics, 123(4), 044510 (2005).
- [68] M. Merolle, J. P. Garrahan and D. Chandler. "Space time thermodynamics of the glass transition." Proceedings of the National Academy of Sciences of the United States of America, 102(31), 10837 (2005).
- [69] R. B. Griffiths. "Nonanalytic behavior above the critical point in a random ising ferromagnet." Phys. Rev. Lett., 23(1), 17 (1969).
- [70] A. J. Bray. "Dynamics of dilute magnets above tc." Phys. Rev. Lett., 60(8), 720 (1988).
- [71] S. Jain. "Anomalously slow relaxation in the diluted ising model below the percolation threshold." Physica A, 218(3-4), 279 (1995).
- [72] M. Blume. "Theory of the first-order magnetic phase change in uo2." Phys. Rev., 141(2), 517 (1966).
- [73] H. Capel. "On the possibility of first-order phase transitions in ising systems of triplet ions with zero-field splitting." Physica, 32(5), 966 (1966).
- [74] M. Blume, V. J. Emery and R. B. Griffiths. "Ising model for the lambda transition and phase separation in he³-he⁴ mixtures." Phys. Rev. A, 4(3), 1071 (1971).

- [75] C. Brito and M. Wyart. "Heterogeneous dynamics, marginal stability and soft modes in hard sphere glasses." J. Stat. Mech., 2007(08), L08003 (2007). ISSN 1742-5468.
- [76] C. Brito and M. Wyart. "Normal modes analysis of the microscopic dynamics in hard discs." cond-mat, page 0804.2439 (2008).
- [77] A. Widmer-Cooper, H. Perry, P. Harrowell and D. R. Reichman. "Irreversible reorganization in a supercooled liquid originates from localized soft modes." Nat Phys, advanced online publication, (2008). ISSN 1745-2481.
- [78] L. D. Landau and E. Lifshitz. *Theory of Elasticity, Third Edition*. Butterworth-Heinemann (1986).
- [79] Philips. Amorphous Solids: Low Temperature Properties. Springer-Verlag, Berlin (1986).
- [80] J. Wuttke, W. Petry, G. Coddens and F. Fujara. "Fast dynamics of glass-forming glycerol." Phys. Rev. E, 52(4), 4026 (1995).
- [81] D. Engberg, A. Wischnewski, U. Buchenau, L. Börjesson, A. J. Dianoux, A. P. Sokolov and L. M. Torell. "Origin of the boson peak in a network glass b2o3." Phys. Rev. B, 59(6), 4053 (1999).
- [82] H. Leyser, W. Doster and M. Diehl. "Far-infrared emission by boson peak vibrations in a globular protein." Phys. Rev. Lett., 82(14), 2987 (1999).
- [83] G. Parisi. "Euclidean random matrices, the glass transition and the boson peak." The European Physical Journal E - Soft Matter, **9**(3), 213 (2002).
- [84] V. N. Novikov, Y. Ding and A. P. Sokolov. "Correlation of fragility of supercooled liquids with elastic properties of glasses." Phys. Rev. E, 71(6), 061501 (2005).
- [85] C. S. O'Hern, L. E. Silbert, A. J. Liu and S. R. Nagel. "Jamming at zero temperature and zero applied stress: The epitome of disorder." Phys. Rev. E, 68(1), 011306 (2003).
- [86] L. E. Silbert, A. J. Liu and S. R. Nagel. "Vibrations and diverging length scales near the unjamming transition." Phys. Rev. Lett., 95(9), 098301 (2005).
- [87] J. Kurchan, L. Peliti and M. Sellitto. "Aging in lattice-gas models with constrained dynamics." EPL (Europhysics Letters), 39(4), 365 (1997). ISSN 0295-5075.
- [88] E. J. Garboczi and M. F. Thorpe. "Density of states for random-central-force elastic networks." Phys. Rev. B, 32(7), 4513 (1985).
- [89] M. F. Thorpe and E. J. Garboczi. "Site percolation on central-force elastic networks." Phys. Rev. B, 35(16), 8579 (1987).
- [90] M. V. Chubynsky and M. F. Thorpe. "Algorithms for three-dimensional rigidity analysis and a first-order percolation transition." Phys. Rev. E, **76**(4), 041135 (2007).
- [91] S. Feng and P. N. Sen. "Percolation on elastic networks: New exponent and threshold." Phys. Rev. Lett., 52(3), 216 (1984).

- [92] S. Feng, M. F. Thorpe and E. Garboczi. "Effective-medium theory of percolation on central-force elastic networks." Phys. Rev. B, 31(1), 276 (1985).
- [93] D. J. Jacobs and M. F. Thorpe. "Generic rigidity percolation: The pebble game." Phys. Rev. Lett., 75(22), 4051 (1995).
- [94] Z. Zeravcic, W. van Saarloos and D. R. Nelson. "Localization behavior of vibrational modes in granular packings." cond-mat.soft, page arXiv:0802.3440v1 (2008). To be published in EPL.
- [95] N. Hatano and D. R. Nelson. "Localization transitions in non-hermitian quantum mechanics." Phys. Rev. Lett., 77(3), 570 (1996).
- [96] B. A. DiDonna and T. C. Lubensky. "Nonaffine correlations in random elastic media." Phys. Rev. E, 72(6), 066619 (2005).
- [97] F. Leonforte, R. Boissière, A. Tanguy, J. P. Wittmer and J.-L. Barrat. "Continuum limit of amorphous elastic bodies. iii. three-dimensional systems." Phys. Rev. B, 72(22), 224206 (2005).