

Kinetically Constrained Models

Juan P. Garrahan, Peter Sollich, and Cristina Toninelli

¹ School of Physics and Astronomy, University of Nottingham, Nottingham NG7 2RD, U.K. juan.garrahan@nottingham.ac.uk

² King's College London, Department of Mathematics, Strand, London WC2R 2LS, U.K. peter.sollich@kcl.ac.uk

³ LPMA – Univ.Paris VI-VII, CNRS UMR 7599, Case courrier 188, 4 Place Jussieu, 75252 Paris Cedex 05 cristina.toninelli@upmc.fr

Abstract. In this chapter we summarize recent developments in the study of kinetically constrained models (KCMs) as models for glass formers. After recalling the definition of the KCMs which we cover we study the possible occurrence of ergodicity breaking transitions and discuss in some detail how, before any such transition occurs, relaxation timescales depend on the relevant control parameter (density or temperature). Then we turn to the main issue: the prediction of KCMs for dynamical heterogeneities. We focus in particular on multipoint correlation functions and susceptibilities, and decoupling in the transport coefficients. Finally we discuss the recent view of KCMs as being at first order coexistence between an active and an inactive space-time phase.

1.1 Motivation

Kinetically constrained models (KCMs) are simple lattice models of glasses. They furnish a perspective on the glass transition problem which has its origin in the work of Glarum [1], Anderson [2] and coworkers [3], and Andersen and coworkers [4]. This perspective assumes that most of the interesting properties of glass forming systems are dynamical in origin, while thermodynamics plays a very limited role. KCMs tend to have simple and uninteresting thermodynamics, typically that of a non-interacting lattice gas. In contrast, they display rich dynamical behaviour as a consequence of kinetic constraints. This combination of simple thermodynamics and locally constrained dynamics is often assumed to be the result of coarse-graining of a dense molecular system [5]: dense fluids are structureless at distances beyond the molecular length, but interatomic forces at high densities are highly constraining, giving rise to local restrictions in the dynamics. As such KCMs are meant as models of glass forming systems at high densities or low temperatures, and aim to capture their dynamical behaviour for motion beyond the inter-molecular distance and for long times.

KCMs generally use local constraints. Nevertheless, as we discuss below, these give rise to collective dynamics due to a form of dynamical frustration: at low temperatures/high densities there is a conflict between the scarcity of excitations/vacancies and the need for them to “facilitate” local motion, leading to hierarchical and cooperative relaxation. Note that this frustration, in contrast with e.g. the random first-order approach [6], does not arise from quenched disorder, an ingredient which at any rate is not trivial to justify in models of real liquids. A further distinction from the latter approach is that KCMs offer a “non-topographic” [7] view of the glass transition problem. By this we mean that it is not a change in the topographic structure of the potential energy landscape (such as a transition between a saddle-dominated and a minima-dominated regime) that determines glassy relaxation, but a change in the degree of connectivity of the configuration space. Indeed, due to the presence of the constraints, the effective connectivity of this space depends on temperature/density, as we discuss in more detail in section 1.3. This picture provided by the KCM approach is appealing as an effective description of the physics for example in the case of hard spheres, where all allowed configurations have the same energy. For a more detailed discussion on the non-topographic KCM approach versus the energy landscape scenario (and additional references on the latter) we refer the reader to [7, 8]. Further information on disorder or landscape based approaches, along with other explanations of glassy behavior that invoke a thermodynamic transition, can be found in Chapter 1 of this book.

KCMs are simple enough to allow for detailed analysis. The current interest in KCMs originates in the fact that they exhibit explicit mechanisms for super-Arrhenius slowdown and stretched relaxation [9] as a consequence of local, disorder free interactions, and without the emergence of finite temperature singularities [10]. At the same time they provide a natural explanation [11] for the phenomenon of dynamical heterogeneity (for reviews see [12–14]). KCMs are to the constrained dynamics/facilitation perspective [3, 4] what the random energy model and the p -spin spin glass are to the random-first order transition approach [6]. Insights from the analysis of KCMs allow one to construct a comprehensive theoretical picture of the glass transition problem (see [15] for a review) which is quite distinct from other competing theories [6, 16–19]. Beyond their importance as models for describing glassy phenomenology, KCMs are of interest for the mathematical community. Indeed, even though they belong to the class of interacting particle systems with Glauber and Kawasaki dynamics, the rate at which elementary moves (birth/death or jump of particles) occur may degenerate to zero due to the presence of the kinetic constraints. This prevents the use of the standard probabilistic tools developed for such systems. Furthermore it gives rise to peculiar phenomena including the presence of several invariant measures, ergodicity breaking transitions [20], unusually long mixing times [21] and aging phenomena [22].

The most recent comprehensive review of KCMs is that of Ritort and Sollich [23] which covers the literature until roughly the end of 2001. Earlier

surveys can be found in Refs. [24–26]. The aim of this chapter is to briefly describe the developments in the last decade or so. It is organised as follows. In Section 1.2 we introduce the KCMs that we will cover, grouping them according to whether they are non-conservative (Glauber dynamics) or conservative (Kawasaki dynamics). Section 1.3 is devoted to a discussion of probably the most basic mathematical manifestation of glassiness in KCMs, i.e. whether they possess ergodicity breaking transitions. Beyond such a transition, the relaxation time to equilibrium is effectively infinite. Physically, of course, this may not be distinguishable from a finite but very long time, and so we consider in Section 1.4 how relaxation timescales in KCMs depend on the relevant control parameters (temperature/density). These times do indeed generically increase extremely fast, in particular in the so-called cooperative models. In Section 1.5, finally, we turn to the predictions of KCMs for dynamical heterogeneity in glasses. Here the models provide a appealingly intuitive picture that is directly based on considering the dynamics in real (rather than e.g. Fourier mode) space. We cover definition and predictions for multi-point correlation functions and susceptibilities, decoupling in the temperature or density-dependence of transport coefficients, and finally the view of KCMs as being at a first-order coexistence between two different dynamical (space-time) phases. We conclude in Section 1.6 with a summary and a critical discussion of the advantages and drawbacks of KCMs. We also provide a brief comparison with alternative approaches, and finally an outlook towards future work.

1.2 The models

The majority of KCMs are defined as stochastic lattice models with binary degrees of freedom, and it is on such models that we focus in this chapter. We do not attempt to give here an exhaustive survey of the wider range of models studied previously, but refer to this for the review listed in the introduction [23–26]. For our present purposes, then, a KCM has on each lattice site i an occupation variable $n_i \in \{0, 1\}$, and the collection of the n_i defines the overall configuration \mathbf{n} . *Conservative* KCMs are lattice gases where the n_i indicate the presence ($n_i = 1$) or absence ($n_i = 0$) of particles. The dynamics follows a continuous-time Markov process that consists of a sequence of particle jumps, subject to kinetic constraints as explained below. The total particle number $\sum_i n_i$ is conserved.

We will also discuss *non-conservative* KCMs. These are motivated by a conceptual coarse-graining to a length scale several times the particle diameter of the underlying physical glass that is being modelled. Each lattice site i then represents a small region of material containing at least a few particles, and one sets $n_i = 1$ (respectively $n_i = 0$) if the density in this region is above (respectively below) a certain threshold where local re-arrangements – of the particles inside the element – become possible. This representation no longer contains the precise values of all local densities, and consequently in non-

conservative KCMs moves that change $\sum_i n_i$ are allowed. The reader should be aware that in much of the literature a reverse convention for the two states of n_i is used, with $n_i = 1$ standing for high mobility, i.e. low density, and vice versa for $n_i = 0$. (These states are then often further identified as up- and down-spins, but we will avoid this terminology.) We opt for the version discussed above as it makes for a unified discussion of conservative and non-conservative KCMs: $n_i = 0$, an empty or low-density site, is mobile in both contexts and facilitates the dynamics on neighbouring sites. For brevity we use the terms “empty” and “low-density” interchangeably below, and similarly for “occupied” and “high-density”. When we talk about the density of the system, we correspondingly mean the fraction of sites with $n_i = 1$, both in conservative and non-conservative KCMs.

The key property of all KCMs is that in order to perform a move the configuration must satisfy a local constraint which usually corresponds to requiring a minimal number of empty sites in some appropriate neighbourhood. This represents the physical intuition that, when particles re-arrange in a glass, motion in any given region requires the presence of mobile regions around it [4, 27]. When the density of these mobile or facilitating sites decreases, the dynamics slows down. Another property shared by all models is that the rates satisfy detailed balance with respect to (w.r.t.) a Boltzmann distribution that factorizes over sites (mathematically, a Bernoulli product measure), so that there are no static interactions, beyond the effective hard core repulsion implemented by the restriction $n_i \in \{0, 1\}$. As explained in the introduction, this idealization means that KCMs can also be viewed as an attempt to find out how much of glass phenomenology can be explained on purely dynamical grounds, without recourse to e.g. static phase transitions.

KCMs as described above are plainly quite simplistic compared to more realistic interacting atomic or molecular systems. As we shall explain, however, for appropriate choices of the constraints they display a behavior which is in agreement with the broad phenomenology of glass forming liquids including super-Arrhenius slowing down of the dynamics, stretched exponential relaxation, dynamical heterogeneities, aging phenomena and ergodicity breaking transitions.

1.2.1 Facilitated spin models: FA, East and Spiral models

We next describe the non-conservative KCMs that we will consider. Because mobile ($n_i = 0$) sites facilitate motion on neighbouring sites, and because binary degrees of freedom with non-conservative dynamics can be interpreted as spins, such models are also called “facilitated spin models”. In these models, the rate for changing the state of site i is $f_i(\mathbf{n})[(1 - \rho)n_i + \rho(1 - n_i)]$, where f_i depends on the configuration \mathbf{n} in a finite neighbourhood of i , but *not* on n_i itself. It is then immediate to verify that detailed balance holds w.r.t. a Boltzmann distribution with energy function $-\sum_i n_i$ and inverse temperature β linked to the density by $\rho = 1/(1 + e^{-\beta})$: in equilibrium we have $n_i = 1$

with probability ρ and $n_i = 0$ with probability $1 - \rho$, independently at each site. In general f_i is chosen to be nonzero (allowed move) when an appropriate neighbourhood of i contains a minimal number of empty sites, otherwise it is zero (forbidden move). As the temperature $1/\beta$ decreases, the density ρ of occupied sites increases and the density $q = 1 - \rho = 1/(1 + e^\beta)$ of empty sites decreases: the dynamics must then slow down.

Non-conservative KCMs can be divided into two classes: non-cooperative and cooperative models. For the former it is possible to construct an allowed path – a sequence of configurations linked by transitions with nonzero rates – which completely empties *any* configuration provided that it contains somewhere an appropriate finite cluster of empty sites, the *mobile defect*. For cooperative models, no such finite mobile defects exist. As we will detail in Section 1.4, non-cooperative and cooperative models display respectively an Arrhenius and super-Arrhenius slowing down.

We now give the definitions of some specific models which cover these two different classes of behaviour. The first class of non-conservative KCMs was proposed by Fredrickson and Andersen (FA) [4], hence the name Fredrickson-Andersen models. The simplest of these models, which we write as FA-1 (“one-spin facilitated FA”), allows a change of state at site i only if at least one of the nearest neighbours is empty: $f_i(\mathbf{n}) = 1$ if $\sum_{j \sim i} (1 - n_j) > 0$, $f_i(\mathbf{n}) = 0$ otherwise, where the sum runs over the nearest neighbours j of site i . It is easy to check that the presence of a single empty site allows one to empty the whole lattice: the model is non-cooperative. For theoretical calculations it is often easier to define f_i somewhat differently, as $f_i(\mathbf{n}) = \sum_{j \sim i} (1 - n_j)$. This leaves the kinetic constraints enforced by vanishing rates as they are and just changes the rates for allowed moves when more than one empty neighbour is present; as such, it makes no qualitative difference to the behaviour of the model.

One can similarly define m -constrained FA models, FA- m , by setting $f_i(\mathbf{n}) = 1$ if at least m neighbouring sites of i are empty, and $f_i(\mathbf{n}) = 0$ otherwise. These models are normally considered on hyper-cubic lattices of dimension d , with $2 \leq m \leq d$ [4]. As can be checked directly, it is not possible in any of these models to devise a finite cluster of empty sites which always lets one empty the entire lattice. Consider for example the case $d = 2$, $m = 2$ in infinite volume and focus on a configuration which contains two adjacent infinite rows of occupied sites. Even if the rest of the lattice is completely empty, none of the sites in these two rows can ever change state because each has at most one empty neighbour. Thus it is not possible to devise a mobile defect which can unblock (empty) every configuration and the model is cooperative. The same conclusion applies on a finite lattice, e.g. with periodic boundary conditions. The restriction on m comes from the fact that, if $m > d$, it is possible to construct *finite* structures of occupied sites that are blocked, i.e. can never change state. They are therefore not suitable to describe the slow dynamics close to a glass or jamming transition because a finite fraction of the system is jammed at *any* density.

Another example of a cooperative model is the one-dimensional East model [28]. In this case the constraint requires for an allowed move that the left nearest neighbour be empty, that is $f_i = 1 - n_{i-1}$. Note that on a finite chain the presence of a single empty site at the left boundary lets one empty the entire chain; but this is not a mobile defect because it will work only when it is found in a specific position. In a chain that is full except for a finite cluster of empty sites in a generic location, the part of the chain to the left of the cluster cannot be emptied, and so the model is cooperative. One can show that due to the directed nature of the constraint the relaxation involves the cooperative rearrangements of increasingly large regions as q becomes small. This leads to super-Arrhenius behaviour of relaxation time scales [29,30]. (The East model shows up a slight shortcoming of our definition of co-operativity: with periodic boundary conditions, we would formally label this model non-cooperative because all sites can be emptied starting from a single empty site. Physically, the model is clearly co-operative independently of the boundary conditions.) We discuss below also higher-dimensional generalizations of the East model, e.g. in $d = 3$ the North-East-Front (NEF) model, where a local change of state is allowed if at least one of the neighbouring sites to the North or East or front is empty.

The models described above have been closely studied since they were proposed in the 80s [4] and 90s [28], with some work continuing into the present decade [23], particularly regarding the non-equilibrium behaviour (see Section 1.4.3) and detailed comparisons with glass phenomenology [5,31–35]. In parallel, progress has come from the development of new KCMs [5]. We discuss below one of these, which reproduces in particular the mixed character of the glass transition, where diverging lengthscales characteristic of second order transitions combine with a first order jump in the fraction of frozen degrees of freedom. This two-dimensional Spiral model [20,36,37] is defined as follows. Consider a square lattice and, for each site i , divide its first and second neighbours into the North-East (NE), South-West (SW), North-West (NW) and South-East (SE) pairs, as illustrated in Fig. 1.1. Then the constraint for a move to occur at i is that “both its NE *and/or* both its SW neighbours should be empty” *and* “both its SE *and/or* both its NW neighbours should be empty” (see Fig. 1.1); in other words, $f_i(\mathbf{n}) = 1$ if this condition is satisfied and $= 0$ otherwise. Stated in a more geometrical form, the constraint requires that all four sites in one of four contiguous sets of neighbours (NE and SE, or SE and SW, or SW and NW, or NW and NE) have to be empty. The motivation for the somewhat involved form of this constraint will become clear when we analyse the dynamical transition in the model below. We note that before the Spiral model a somewhat more involved choice of the constraints (the so called Knights model) was proposed [36,38] which leads to the same mixed first/second order transition. Other choices in the same spirit have also been investigated since, see for example [39].




Fig. 1.1. Site i and its NE, NW, SE and SW neighbours. The constraint is (is not) verified at the central site in case (a) (in case (b)).

1.2.2 Kinetically constrained lattice gases: KA and TLG models

Conservative KCMs are defined in a similar spirit to their non-conservative analogues, but with dynamics that conserves total particle number. A particle at i attempts at a fixed rate ($= 1$ without loss of generality) to jump to any empty nearest neighbour site j . The rate for this move is thus $n_i(1-n_j)f_{ij}(\mathbf{n})$. The factor $f_{ij}(\mathbf{n})$ again implements a kinetic constraint and is chosen not to depend on the configuration at sites i and j . The resulting dynamics preserves the number of particles, and so detailed balance on finite lattices is satisfied w.r.t. the Boltzmann distribution which is uniform on configurations with the appropriate fixed particle number. By linearly combining distributions with different particle numbers one can of course also create grand-canonical Boltzmann distributions where each site independently contains a particle with some probability ρ . Detailed balance is also satisfied w.r.t. these distributions, and they extend naturally to the limit of an infinite lattice.

Again, we can classify conservative KCMs into non-cooperative and cooperative models. For the former it is possible to construct a finite group of empty sites, the *macrovacancy*, such that for *any* configuration the macrovacancy can be moved all over the lattice and any jump of a particle to a neighbouring empty site can be performed when the particle is adjacent to the macrovacancy. Clearly the baseline lattice gas in which there are no kinetic constraints, which in $d = 1$ is the symmetric simple exclusion process (SSEP), is non-cooperative and the minimal macrovacancies are just isolated empty sites.

Chronologically the first conservative KCM is the Kob-Andersen model (KA) [27], which has cooperative dynamics. Here a particle can jump to a neighbouring site only if both in the initial and final position at least m of its nearest neighbouring sites are empty. The original KA model had $m = 3$, with particles on a cubic lattice, but one can similarly define KA- m models on hyper-cubic lattices of dimension d and for different values of the parameter m , with $2 \leq m \leq d$. The restrictions on m arise from the fact that $m = 1$ corresponds to an unconstrained lattice gas, while a model with $m > d$ has a finite fraction of frozen particles at any density (as for FA models in the same

parameter regime). Similarly to the case of FA- m models with $m \geq 2$, one can directly check that all KA- m models are cooperative. On other lattices, KA models can be non-cooperative. Consider for example the KA model with $m = 2$ on a triangular lattice. Here one can check that a “dimer” of two neighbouring empty sites forms the required macrovacancy, which can move across the lattice in a tumbling motion even when all other sites are occupied. Two other KCMs on triangular lattices were introduced in [40] and have more recently been analysed in [41]. For the (1)-TLG a particle can move from site i to a neighbouring site j if at least one of the two mutual neighbours of i and j is empty. For the (2)-TLG a particle can move from i to j if *both* the two mutual neighbour sites are empty. It is easy to verify that the (1)-TLG is non-cooperative, with – as for KA with $m = 2$ on the same lattice – a dimer of vacancies being a macrovacancy. The (2)-TLG, on the other hand, is cooperative: a chain of particles occupying an entire row of the lattice, for example, can never be destroyed, thus it is not possible to construct a macrovacancy.

1.3 Ergodicity breaking transitions

As explained above, the equilibrium distribution in KCMs is trivial because it factorizes over sites. In particular, then, no equilibrium phase transition can occur. On the other hand, the presence of constraints might induce transitions of purely dynamical type: detailed balance alone does not guarantee that the distribution over configurations will converge for large times to the Boltzmann equilibrium distribution. To see this, return to the example of the FA-2 model on a square lattice ($m = d = 2$), either finite with periodic boundary conditions or infinite, and a starting configuration $\mathbf{n}(0)$ which has two adjacent rows of occupied sites. Then since these sites are forever blocked, for any site i within the two rows, $\lim_{t \rightarrow \infty} \langle n_i \rangle_t = 1 \neq \langle n_i \rangle_{\text{eq}} = \rho$, where $\langle \cdot \rangle_t$ is the average over the stochastic dynamics which starts from $\mathbf{n}(0)$ and $\langle \cdot \rangle_{\text{eq}}$ is the equilibrium average.

Motivated by this, we can ask whether convergence to the equilibrium distribution is recovered at least in the thermodynamic limit: if we sample an initial configuration $\mathbf{n}(0)$ from the equilibrium distribution, then is the large time limit of the average of any function $g(\mathbf{n})$ equal (with probability one) to its equilibrium average, i.e. is $\lim_{t \rightarrow \infty} \langle g \rangle_t = \langle g \rangle_{\text{eq}}$? If this is the case we will say that the system is *ergodic*. The models we consider are all ergodic at sufficiently high density $q = 1 - \rho$ of facilitating (empty) sites. If they do become non-ergodic as q is reduced, we call q_c the critical density of empty sites at which this transition occurs. As will be discussed below, the models defined in the previous sections are ergodic at any positive $q = 1 - \rho$, so that $q_c = 0$, with the exception of the Spiral model and FA models on Bethe lattices which display ergodicity breaking transitions at $0 < q_c < 1$.

Owing to the factorized form of the equilibrium distribution, ergodicity corresponds to the fact that in the thermodynamic limit the configuration space is covered by a single irreducible set, i.e. a set of configurations which are connected to each other by allowed paths; for a proof see [21, Proposition 2.4] and [42, Section 2.3]. This in turn is equivalent to the requirement that for any site i (in non-conservative KCMs) or pair of nearest neighbour sites i, j (in conservative KCMs) there is an allowed path which transforms the configuration into one where the constraint is satisfied at i (i, j). In other words, the probability that any site i belongs to a cluster of forever blocked sites must vanish in the thermodynamic limit.

For *non-cooperative* KCMs one sees easily that ergodicity holds at any density $\rho < 1$, i.e. at any $q > 0$. In the non-conservative case, the probability of finding at least one mobile defect in an equilibrium configuration goes to one in the infinite volume limit, and then starting from this defect one can empty all sites. Thus, any configuration is connected to the “all empty” configuration with probability one. The situation is similar for non-cooperative but conservative KCMs.

The case of *cooperative* KCMs is more delicate and here an ergodicity breaking transition can occur. The non-conservative case is again simpler, so we begin with this. In order to determine the probability that a site belongs to a blocked cluster, consider the following deterministic procedure: iteratively empty all sites for which the constraint is verified until we reach either the completely empty configuration, or one in which there is a “backbone” of mutually blocked occupied sites. Then it is easy to verify that the sites that are blocked forever under the stochastic dynamics of the actual KCM are precisely those that are empty at the end of this deterministic procedure. In other words, the problem of the existence of an ergodicity breaking transition for cooperative non-conservative KCMs can be reformulated as a percolation transition for the final configuration of the above deterministic dynamics. For FA- m models this deterministic dynamics coincides with the well known algorithm of bootstrap percolation and the results in [43, 44] establish that $q_c = 0$ on hypercubic lattices in any dimension d and for any facilitating parameter $1 \leq m \leq d$ (while trivially $q_c = 1$ for $m > d$ since finite blocked structures can occur). This work disproved a long-standing conjecture from the original FA paper [4], namely that an ergodicity breaking transition would occur at some $q_c > 0$. Such a transition does take place, however, when one considers FA models on a Bethe lattice as will be explained in Section 1.3.1. An example of a finite-dimensional model displaying such a transition is provided by the Spiral model (see Section 1.3.2).

For cooperative conservative KCMs, the proof of ergodicity is more involved. For example, for the KA-2 model on a square lattice ($d = 2$), we have shown that the irreducible set of configurations that has unit probability in the thermodynamic limit is the one containing all configurations which can be connected by an allowed path to a configuration which has a frame of empty sites on the last shell before the boundary [45]. Establishing that this

set has unit probability in the thermodynamic limit is more complicated than in the corresponding non-conservative KCMs, i.e. FA models. This is because there is no deterministic bootstrap-like procedure which allows one to establish whether or not a configuration does or does not belong to the irreducible set. Nevertheless, a formal proof can be constructed [42,45], and demonstrates that the ergodicity breaking transition originally conjectured by Kob and Andersen [27] does not exist. Analogous arguments can be constructed for the other choices of m and d [45] to rule out the occurrence of a transition.

1.3.1 FA models on Bethe lattices

As a simple example of a model that does have an ergodicity breaking transition we consider next the FA- m model on a Bethe lattice, i.e. a random regular graph of connectivity $k + 1$. Exploiting the local tree-like structure of such a graph, it is easy to write a self-consistent equation for the probability P that a given site is occupied and blocked, conditional on the fact that its ancestor is occupied and blocked:

$$P = (1 - q) \sum_{i=0}^{m-1} \binom{k}{i} P^{k-i} (1 - P)^i. \quad (1.1)$$

The factor $1 - q$ is the probability for the given site to be occupied; the sum gives the probability that at most $m - 1$ of the descendants of the site are empty so that it is in fact blocked. The leading term on the r.h.s. for small P is $O(P^{k-m+1})$. For $m = k$, the largest value that does not produce finite blocked structures, this is linear and so one gets a continuous transition at $(1 - q_c)m = 1$. For all smaller m , the small P -increase is with a higher power of P and so a discontinuous transition results, with $P = 0$ for $q > q_c$ and $P = P_c + O(\sqrt{q_c - q})$ for $q < q_c$, where $P_c > 0$. The mechanism behind this combination of a discontinuous onset and a critical singularity has been analysed in detail [46, 47]. In particular, the singular square root behavior is due to the extreme fragility of the infinite spanning frozen or “jammed” cluster at the transition, and to the existence of a related length scale which diverges as a power law.

1.3.2 Spiral model

In this section we will explain the mechanism behind the ergodicity breaking transition of the Spiral model [20, 36, 37]. Consider the directed lattice that is obtained from the square lattice by putting two arrows from each site towards the neighbours in the NE pair, i.e. pointing North and North-East. The resulting lattice is a tilted and squeezed version of a two-dimensional oriented square lattice (see Figure 1.2). Therefore, if the density is larger than the critical density of oriented site percolation (directed percolation, DP), ρ_c^{DP} , there exists a cluster of occupied sites which spans the lattice following the direction

of the arrows. Consider now a site in the interior of this directed cluster (see Figure 1.2): by definition there is at least one occupied site in both its NE and SW neighbouring pairs, therefore the site is blocked with respect to the constraints of the Spiral model. Thus for $\rho > \rho_c^{\text{DP}}$ the system contains one or more blocked clusters and is therefore non-ergodic. This suggests that the ergodicity breaking transition occurs at $q_c = 1 - \rho_c^{\text{DP}}$, but the argument so far does not exclude a transition earlier, i.e. at a larger density q of empty sites. Indeed, since blocking can occur along either the NE-SW or the NW-SE direction (or both), the presence of a blocked cluster does not imply that there is a directed path through the lattice as in the DP argument. To establish that, nevertheless, blocked clusters do not occur for $q > 1 - \rho_c^{\text{DP}}$, one shows that empty regions of linear size much larger than the parallel length of DP, ξ_{\parallel} , act as “critical defects”: starting from any such empty region we can very likely empty the whole lattice. Because ξ_{\parallel} is finite, so is the size of these defects. In the thermodynamic limit at least one such defect will occur in the system, and so for $q > 1 - \rho_c^{\text{DP}}$ any configuration can be completely emptied: the system is ergodic. This confirms that indeed $q_c = 1 - \rho_c^{\text{DP}}$.




Fig. 1.2. The tilted and squeezed two dimensional oriented lattice obtained by drawing arrows from each site to its neighbours to the North and East. The site marked by the cross belongs to an oriented occupied cluster and is therefore blocked.

An important feature of the ergodicity breaking transition in the Spiral model is that it is discontinuous, in the sense that the fraction of blocked sites jumps to a nonzero value at $q = q_c$. This is relevant for the connection to real glasses, because it implies that two-point correlation and persistence

functions must display a plateau in their time-dependence when the transition is approached ($q > q_c$), as is observed experimentally. Earlier models with a transition at a finite q_c , e.g. a two-dimensional generalization of the East model [48], develop only a fractal cluster of blocked particles at the transition, which occupies a vanishing fraction of the system. The fact that in the Spiral model blocked clusters are compact at q_c follows by a direct construction of blocked structures [20, 36, 37] and is a consequence of the presence of the two transverse blocking directions in the constraints, and of the anisotropy of DP.

1.3.3 Summary: Presence/absence of ergodicity breaking transition

For the sake of clarity let us summarize the results of this section by listing the KCMs we have discussed according to whether or not they display an ergodicity breaking transition at a non-trivial (different from 0 or 1) critical defect density q_c .

The models that do display such a transition are the FA and KA models on Bethe lattices, the Spiral model and the North-East model. The latter was not mentioned above because it is not of direct interest for modelling the glass transition, due to the continuous character of its ergodicity breaking transition. It is a spin facilitated model on a square lattice with the constraint requiring both the North and East neighbour to be empty. It is easy to verify along the lines of the argument in the previous subsection that the North-East model as defined by this constraint displays a transition at the critical density of oriented percolation.

The models that do not display a transition are: the FA- m and KA- m models for any choice of m and on hypercubic lattices in any finite dimension d ; the East model; and the (1)-TLG and (2)-TLG models.

1.4 Bulk dynamics of KCMs

1.4.1 Glassy time scale divergences and static length scales

For KCMs to be useful as models of physical glasses they need to satisfy the basic requirement of dynamical slowing down as the glassy regime approached. In our case this corresponds to increasing density ρ , or decreasing density of empty sites q . We will therefore now give an overview of numerical and analytical results on how relaxation time scales τ diverge in KCMs when q approaches q_c . In concrete terms, we will take τ as the typical time in the relaxation of density-density correlation and persistence functions. In cooperative KCMs, τ turns out to be connected to statically defined blocking lengths, as we also discuss.

Starting with FA-1 models, relaxation occurs via the effective diffusion of empty sites: an empty site facilitates the emptying of a neighbour site,

with rate $q = 1 - \rho$, and the original site can then fill and does so with probability $1/2$ before the new empty site does. Empty sites can thus diffuse freely, with an effective diffusion constant $q/2$. This suggests that τ should grow as an inverse power of q . Because $q \simeq \exp(-\beta)$ at low temperatures $1/\beta$, this corresponds to Arrhenius scaling. Indeed, by an exact mapping to a diffusion-limited aggregation model one can derive [49] that $\tau \sim 1/q^z$ with $z = 3$ in $d = 1$ and $z = 2$ in $d \geq 2$. The $d = 1$ result is simple to understand: empty sites are typically a distance $1/q$ apart, and relaxation requires diffusion across this distance with the effective diffusion constant $q/2$, so $\tau \sim (1/q)^2/(q/2) \sim 1/q^3$.

For FA- m with $m > 1$, an isolated empty site is unable to move on its own and has to wait for a mobile defect, i.e. an appropriate region of empty sites, to move cooperatively into its neighbourhood and so facilitate its motion. As we will detail below, the typical number of moves involved in this cooperative process increases as $q \rightarrow 0$, thus a super-Arrhenius scaling of the relaxation time has to result. This has been confirmed by several numerical investigations which have proposed different forms for the density dependence of τ [50–52]. In order to better understand the cooperative mechanism let us consider for example the case $m = d = 2$. An $\ell \times \ell$ square of empty sites can be expanded by one lattice spacing in both directions, provided that at least one empty site is present on two adjacent sides of the square. This is very likely to be the case when $\ell \gg 1/q$. On the other hand, the probability of being able to empty all sites in a region of size $1/q$ without using external empty sites can be shown to be proportional to $\exp(-c/q)$ with c a constant of order one [43, 53]. This result is obtained by iterating the process outlined above, multiplying the relevant probabilities at each step: one can iteratively remove all the particles in any region starting from its interior provided there is a central 2×2 square of empty sites and one additional empty site somewhere on each of the four sides of each subsequent shell. In [54] it was conjectured that relaxation occurs via the diffusion of these critical defects leading to a relaxation time diverging as τ_D/ρ_D , with $\rho_D \simeq \exp(-c/q)$ the defect density and τ_D their diffusion time. The latter, as detailed in [45, Section 6], should be a sub-leading correction, giving $\tau \sim 1/\rho_D$ to leading order. In principle one might think that this is not the optimal relaxation mechanism and that other relaxation processes could be much faster, avoiding the super-Arrhenius scaling. However, the results in [43, 53] also give the typical size of the incipient blocked cluster to which any fixed site belongs and that has to be eroded via successive moves from its boundary before the site in question can be unblocked. This size, L_c , diverges as $\exp(c/2q)$ thus providing (due to finite speed of propagation) a lower bound on time scales of the form $\tau \geq \exp(c/2q)$ which confirms the super-Arrhenius scaling. Note that L_c^2 scales as the inverse of ρ_D , which (as the probability that a region of size $1/q \times 1/q$ is a critical defect) is essentially the number density of critical defects. Indeed, clusters of linear size larger than L_c are typically unblocked because they contain at least one critical defect from which relaxation can occur. In [44] the typical size of such incipient blocked

clusters was derived for generic m and d leading to τ growing at least as fast as $\exp^{o(m-1)}(c/q^{1/(d-m+1)})$. Here \exp^{os} is the exponential iterated s times, so that the divergence of the time scale is extremely rapid for any $m \geq 3$.

For KA models the basic relaxation mechanism is very similar [45], the main difference being that the critical defects are now regions in which empty sites and particles are arranged in such a way that one can find an allowed path – a sequence of allowed moves – to perform any nearest neighbour exchange. As shown in [45, 55], the properties as a function of q of these regions coincide with those for the FA models, leading to the same estimates for the scaling of the relaxation time.

Relaxation processes in the Spiral model proceed in a broadly analogous manner. But the size of the critical defects that can expand further diverges already at q_c , proportionally to the parallel length ξ_{\parallel} of DP clusters. The typical size of incipient blocked clusters grows much more quickly, as $L_c \simeq \exp[c/(q - q_c)^{\mu}]$ with $\mu = \nu_{\parallel}(1 - z)$, where $\nu \simeq 1.73$ is the critical exponent of ξ_{\parallel} and $z \simeq 0.64$ the exponent relating the parallel and transverse lengths of DP. As for FA models, since these clusters can be unblocked only from the boundary, τ should diverge at least as L_c .

Interestingly, in the other model we have considered that has an ergodicity breaking transition at nonzero q_c , the FA model on a Bethe lattice, simulation measurements of persistence and correlation functions [56] show that relaxation times grow only as power laws on approaching the transition, $\tau \propto (q - q_c)^{-\gamma}$ with $\gamma \simeq 2.9$. Qualitatively, thus, the transition for FA models on Bethe lattices has the characteristics of a mode-coupling theory (MCT) [16] arrest transition.

Finally we discuss the East model, where the origin of the super-Arrhenius time scale is possibly easiest to see. The basic relaxation mechanism can be understood in terms of the dynamics of domains of occupied sites separated by empty sites. In the limit of small q one can then argue [29] that the typical relaxation time should scale as the minimal time required for an empty site to facilitate the motion of the first empty site to its right, which is typically at distance $1/q$. The optimal path to create an empty site at a distance d involves an energy barrier of order $\log_2(d)$ (the minimum over all paths of the maximal number of empty sites which we encounter along the path). Setting $d = 1/q$ thus leads to the relaxation time estimate $\tau \sim q^{-\log_2(1/q)} \sim \exp[\beta^2 / \ln 2]$ [29]. The above argument does not take into account the behavior on scales smaller than the typical distance between empty sites. Incorporating this turns out to halve the coefficient in the exponent, to $\tau \sim \exp[\beta^2 / (2 \ln 2)]$ [21, 57].

We have discussed so far only the overall timescale for the decay of correlation or persistence functions in KCMs, but not the time dependence of this decay. For the cooperative models one generically finds stretched exponential forms, as seen experimentally, while non-cooperative models can exhibit power law tails reflecting the diffusive motion of the mobile defects. We refer to [23] for an overview and quote only the example of the East model, where one can show that the stretching becomes extremely strong at low tempera-

tures, with correlations and persistence both decaying as scaling functions of $(t/\tau)^{1/[\beta \ln 2]}$ [29, 30, 58].

1.4.2 Some rigorous results

In the recent years KCMs have been also analysed in the mathematical community. We summarize here those rigorous results which help to further understand the slow relaxation of KCMs, and have in some cases corrected conjectures resulting from numerical simulations or intuitive arguments.

The analysis of the large time behavior in the ergodic regime was started by [59] where for the East model it was established that $[1/\ln 2 - o(1)] \ln^2(1/q) \leq \ln(1/\text{gap}(q)) \leq [1/(2\ln 2) + o(1)] \ln^2(1/q)$ with $1/\text{gap}(q)$ the inverse of the spectral gap of the Liouvillian operator generating the dynamics. The latter, which in a finite system is just the inverse of the smallest nonzero eigenvalue of the transition matrix, represents the longest relaxation time for all one-time quantities and so can be identified with the relaxation time scale τ discussed above. The bounds of [59] quoted above then say that $\ln \tau$ scales for small q as $\ln^2(1/q) = \beta^2$, with a prefactor between $1/(2\ln 2)$ and $1/\ln 2$, the upper bound being the naive estimate of [29]. These bounds were sharpened in [21], establishing that it is in fact the lower bound that gives the correct asymptotics: written in terms of the gap, $\lim_{q \rightarrow 0} \ln(1/\text{gap})/\ln^2(1/q) = 1/(2\ln 2)$.

Positivity of the spectral gap guarantees in particular exponential convergence in the large time limit: for any $g(\mathbf{n})$ one has $\langle g(\mathbf{n}(t))g(\mathbf{n}(0)) \rangle - \langle g(\mathbf{n}(0)) \rangle^2 \leq \text{const} \times \exp(-2\text{gap} t)$, where $\langle \dots \rangle$ is the mean over the initial Boltzmann equilibrium at empty site density q and over the stochastic process governing the evolution in time. In [21] it was shown that positivity of the spectral gap also guarantees exponential convergence of the persistence function $P(t)$, which is the probability that a site does not change its state during a time interval of length t : $P(t) \leq \exp(-q\text{gap} t) + \exp(-(1-q)\text{gap} t)$. In the same paper, a multi-scale approach was developed which allows one to prove positivity of the spectral gap in the whole ergodic region $q > q_c$ for all the choices of constraints described in Section 1.2. With this technique one can also derive [21] the following (sometimes optimal) bounds when $q \downarrow q_c$. For FA-1 in $d = 1$, $\tau \propto 1/q^3$; in $d = 2$, $1/q^2 < \tau \leq \ln(1/q)/q^2$; and in $d \geq 3$, $1/q^{1+2/d} < \tau \leq 1/q^2$. These bounds are in agreement with the analytical results in [49], and confirm that the initial findings [60] in $d = 2$ and $d = 3$, based on a mapping to DP, were incorrect. For the cooperative KCMs FA-2 and FA-3, results in [21] show that $\exp(q^{-1}) \leq \tau \leq \exp(q^{-2})$ and $\exp[\exp(q^{-1})] \leq \tau \leq \exp[\exp(q^{-2})]$, respectively, thus establishing a super-Arrhenius scaling compatible with [50, 54].

For conservative KCMs with non-cooperative behaviour, the diffusive scaling $1/\text{gap} = O(L^2)$ in a volume of linear size L and the positivity of the self-diffusion coefficient at any density was established in [61]. Moreover the hydrodynamic limit was studied there for a special class of models leading

to a porous medium equation, namely a degenerate partial differential equation $\partial_t \rho = \nabla(D(\rho)\nabla\rho)$ with a diffusion coefficient $D(\rho)$ vanishing as a power law of $1 - \rho$ when $\rho \rightarrow 1$. The methods used to establish these results use as a key ingredient the existence of a path between configurations which allows any two particles to be exchanged, by exploiting the presence of mobile macrovancancies. This approach, then, cannot be extended to cooperative models. However, more recently a different technique has been devised [42] that proves also for cooperative models the diffusive scaling $1/\text{gap} = O(L^2)$ and establishes in $d = 2$ the diffusive decay in time $\sim 1/t$ of the density-density autocorrelation function. The self-diffusion coefficient for a specific cooperative model, namely the KA- m , was analysed in [62] where its positivity at any q was proved, modulo a conjecture on the behavior of a random walk in a random environment.

1.4.3 Non-equilibrium behaviour

Due to space constraints we focus in this paper almost exclusively on the equilibrium dynamics of KCMs. Non-equilibrium behaviour results if, for example, a non-conservative KCM is prepared in a configuration with a high density q of empty sites, and then q is reduced quickly. This corresponds to a sudden lowering of the temperature $1/\beta$ and so mimicks quench experiments in real glasses. If the final q is low enough, i.e. sufficiently close to q_c , the time scale τ for relaxation to the new equilibrium will be very long, and aging will occur: the properties of the system depend on the time elapsed since the quench. Aging can be monitored via two-time response and correlation functions, and one ask for example whether these are linked by a non-equilibrium fluctuation-dissipation theorem with some effective temperature [63]. Surprisingly, in some simple KCMs such as the East and FA-1 models, aging at low q is in fact easier to analyse than the dynamics at equilibrium [29, 64]. Work in this area up to late 2001 is summarized in [23], with more recent studies revealing a very rich phenomenology in the aging of KCMs [65–67]. There has also been progress in rigorous approaches, with e.g. Refs. [22, 68] establishing the non-equilibrium behavior of the East model derived in [29].

1.5 Dynamical heterogeneity and its consequences

The main success of KCMs has been the ability to account, at least qualitatively and sometimes also quantitatively, for many aspects of dynamical heterogeneity (DH) [12–14] observed in glass forming systems. DH emerges naturally in KCMs at high density or low temperature [5, 69], and this is the topic of this section. We focus on the simpler KCMs, such as the FA-1, East and TLG constrained lattice gases. For these models results are easier to obtain both analytically and computationally than for, say, the Spiral model or the FA- m model with $m > 1$; nevertheless they display all the important

physics. In particular, the dynamical scaling properties of the East model are close to those observed in realistic glass formers (see e.g. [15]).

The simplicity of KCMs allows one to make detailed predictions about the growth of dynamical correlation lengths and their scaling relations to growing relaxation times [10, 11, 33, 49, 69–71]. These dynamical lengthscales are an indication of spatial correlations that build up over time, so they are most easily extracted from multi-point correlation functions, which we discuss in Subsection 1.5.1. DH also tells us that the dynamics of glass formers is fluctuation dominated. A central consequence is transport decoupling, the breakdown in standard transport relations of liquid state theory that are obtained under the assumption of homogeneous dynamics, such as that of Stokes-Einstein relating self-diffusion rate to viscosity [72]. KCMs provide a direct explanation for decoupling based on the response of a moving molecule to the distribution of local relaxation timescales in the host fluid [31]. Such motion can be approximately quantified within a continuous-time random walk (CTRW) formalism [71]. Decoupling and CTRW are discussed in Subsection 1.5.2. Finally, the observed DH is only a mesoscopic phenomenon: dynamical lengthscales are always finite (and transient) at non-zero temperature or at less than maximal density, that is, for $q > q_c$ (note that all models discussed in this section have $q_c = 0$). Nevertheless, it is a precursor to a fully-fledged non-equilibrium, or “space-time”, phase transition [73–75], which we discuss briefly in Subsection 1.5.3.

1.5.1 Multipoint-correlations and susceptibilities

Figure 1.3 illustrates DH in KCMs. A convenient observable to quantify local relaxation is the persistence field, $p_i(t) = 0, 1$, where 0 indicates that site i has changed its state at least once up to time t , and 1 otherwise. The ensemble average, $P(t) \equiv \langle p_i(t) \rangle = \lim_{N \rightarrow \infty} N^{-1} \sum_i p_i(t)$, is the persistence function discussed above. Fig. 1.3(a) shows the persistence field $p_i(t_{1/2})$ in a three-dimensional version of the East model, or NEF (for North-East-Front model) [71], where $t_{1/2}$ is the time at which $P(t_{1/2}) = 1/2$, i.e., half the system has relaxed and half has not. We see from Fig. 1.3(a) that relaxation is heterogeneous: there is a clear spatial segregation of sites which have relaxed, $p_i = 0$ (colored black), from those that have not, $p_i = 1$ (colored white). Relaxation dynamics is spatially correlated. Fig. 1.3(a) shows the persistence field at two different values of q (two different temperatures). The spatial extension of dynamic correlations increases with decreasing temperature, and therefore with increasing relaxation time. Note that these dynamical correlations are unrelated to thermodynamic correlations since the equilibrium measure of the NEF is trivial at all temperatures. Fig. 1.3(a) shows similar persistence field plots for a density conserving model, the two-vacancy assisted (2)-TLG, but at a fixed high density ρ and for different observation times t . This figure shows that DH is a transient effect, with dynamical correlations becoming maximal at some intermediate time t^* (see below).




Fig. 1.3. Dynamical Heterogeneity in KCMs. (a) Spatial distribution of the local persistence in the NEF model, at time $t_{1/2}$ such that $P(t_{1/2}) = 1/2$ (i.e., 50 of sites, shown in black, have flipped by time $t_{1/2}$) at temperatures $T = 1.0$ (top) and $T = 0.15$ (bottom) for a system of size $N = 40^3$. Adapted from Ref. [71], with permission. (b) Same, but for (2)-TLG model, at fixed density $\rho = 0.77$, and varying observation times, $t = 10^3, 10^4, 10^5, 10^6$. From Ref. [69], with permission.

The extent of the dynamical correlations evident from Figure 1.3 can be quantified by means of multipoint functions [11, 33, 34, 76]. Consider for example the “four-point” structure factor,

$$S_4(\mathbf{k}, t) \equiv N^{-1} \mathcal{N}(t) \sum_{i,j} [\langle p_i(t) p_j(t) \rangle - P^2(t)] e^{i\mathbf{k} \cdot (\mathbf{r}_i - \mathbf{r}_j)} \quad (1.2)$$

This is the Fourier transform of the spatial correlation function of the persistence field. The factor $\mathcal{N}(t)$ is a convenient normalization factor. We adopt the choice $\mathcal{N}(t) = [P(t) - P^2(t)]^{-1}$, which make S_4 equal to unity if all the nonzero contributions come from “self” terms $i = j$. Fig. 1.4(a) shows S_4 for the NEF. This is the structure factor of the DH pictures of Fig. 1.3. It has the characteristic shape of that of a system correlated over finite distances. The zero wave vector limit of S_4 gives the “four-point” susceptibility, $\chi_4(t) = S_4(k \rightarrow 0, t)$, which provides an estimate of the correlation volume of DH. Fig. 1.4(b) shows $\chi_4(t)$ for the NEF model as a function of observation time t , at different temperatures. Several things are apparent. χ_4 is non-monotonic in time, indicating that DH is a transient phenomenon. It peaks at around the relaxation time of the persistence function. The peak value increases with decreasing temperature: dynamical correlations increase with increasing relaxation time.

Multipoint functions reveal the scaling properties of DH. The dynamic susceptibilities defined above have their peak, $\chi_4^* = \chi_4(t^*)$, at times t^* close to the relaxation times τ of the corresponding persistence functions. Scaling is controlled by the distance to the dynamical critical point at zero concentration

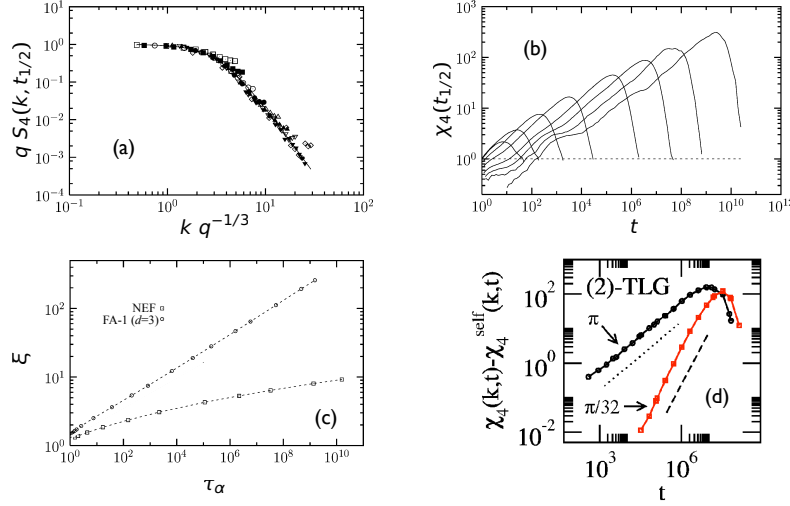


Fig. 1.4. Dynamical correlations and scaling. (a) DH structure factor, $S_4(k, t_{1/2})$ in the NEF model at various temperatures. The data collapses under the scaling $S_4 \rightarrow q S_4$ and $k \rightarrow k\xi$ with $\xi \sim q^{-1/3}$, suggesting the values $\gamma = 1$ and $\nu = 1/3$ for the scaling exponents (see text). (b) Time dependence of four point susceptibility $\chi_4(t)$ in the NEF. Dynamical fluctuations probed by this function are maximal at time t^* . The peak susceptibility $\chi_4^* = \chi_4(t^*)$ grows with decreasing temperature. The peak time is approximately the relaxation time of the system, $t^* \approx \tau$, as extracted from the persistence function. (c) Scaling of DH correlation lengths to relaxation times: FA-1 models follow a simple scaling law [49] ($\tau \sim \xi^4$ for the three-dimensional FA-1 shown in the figure), while for East-like models $\tau \sim \xi^{z(q)}$, where the dynamical exponent $z(q)$ increases with decreasing q (i.e. decreasing temperature) [9]. Adapted from Ref. [71], with permission. (d) The behaviour of four-point susceptibilities depend on the wavelength k of the observable which is used to define them. $\chi_4(k, t)$ is plotted with its self-part, i.e. the terms with $i = j$ in its definition, subtracted. From Ref. [34], with permission.

of excitations or vacancies, so we expect $\chi_4^* \sim q^{-\gamma}$ [33, 34]. Furthermore, we expect the four-point structure factor to behave as $S_4 \approx \chi_4^* f(k\xi)$, where f is a scaling function, and $\xi \sim q^{-\nu}$ a correlation length for DH at times $t = t^*$. These exponents – and indeed scaling forms – may vary from model to model. For the simpler KCMs they can be calculated analytically. Figure 1.4(a) shows the numerical collapse of S_4 in the NEF for $\gamma = 1$ and $\nu = 1/3$ [71]. For other models, such as the FA-1 these exponents can be calculated analytically [33]. For example, in $d = 3$ one finds $\tau \sim \xi^4$. This is most easily understood from the fact that the upper critical dimension of the model is $d_c = 2$, so exponents are d -independent above this [49]. But in $d = 2$, the characteristic lengthscale is the distance between vacancies of density q , $\xi \sim q^{-1/2}$, and the timescale is related to this by the vacancy diffusion constant $q/2$ (see Section 1.4.1), giving $\tau \sim \xi^2/(q/2) \sim \xi^4$.

Four-point and similar dynamical susceptibilities measure the dynamical fluctuations of global observables. These observables probe relaxation over a certain lengthscale, so it is important to note that the scaling properties of the corresponding susceptibilities depend on such lengthscales [34]. In the examples above the observable was the persistence function, which probes relaxation on a lengthscale of one lattice spacing. Let us now consider instead the following susceptibility for the case of conservative (conserved density) KCMs, $\chi_4(k, t) \equiv \lim_{k' \rightarrow 0} N^{-1} \sum_{ij} \langle \delta F_i(\mathbf{k}, t) \delta F_j(-\mathbf{k}, t) \rangle e^{i\mathbf{k}' \cdot (\mathbf{r}_i - \mathbf{r}_j)}$, where i, j label the N particles in the system, the position of the i th particle at time t is $\mathbf{r}_i(t)$, and $\delta F_i(\mathbf{k}, t) \equiv e^{i\mathbf{k} \cdot [\mathbf{r}_i(t) - \mathbf{r}_i(0)]} - \langle e^{i\mathbf{k} \cdot [\mathbf{r}_i(t) - \mathbf{r}_i(0)]} \rangle$. The wave vector \mathbf{k}' appearing here is the analogue of \mathbf{k} in $S_4(\mathbf{k}, t)$ above. A normalization factor $\mathcal{N}(t)$ could be included as before but this turns out to make no qualitative difference. This χ_4 measures the system to system fluctuations of the self-intermediate scattering function of wave vector $F(\mathbf{k}, t)$. It thus probes structural relaxation at lengthscales comparable to $2\pi/k$. Figure 1.4(d) shows how $\chi_4(k, t)$ changes in behavior as we go from large to small k in the (2)-TLG: as we probe larger lengthscales (smaller k) χ_4 peaks at later times, and the initial power law growth changes exponent [34]. This change in behavior is related to the non-Fickian to Fickian crossover of particle diffusion [32], which we discuss in the next subsection.

1.5.2 Transport decoupling

A central consequence of DH is transport decoupling. A prominent example is Stokes-Einstein breakdown [72]: in deeply supercooled liquids the rate for self-diffusion is orders of magnitude larger than what would be predicted from the Stokes-Einstein relation between self-diffusion constant and viscosity, $D_s \propto \eta^{-1}$. Similar transport relations of liquid state theory also break down near the glass transition [12, 78]. This is a consequence of the dynamical fluctuations associated with DH.

A major success of KCMs is the ability to rationalize this phenomenon. It does so in terms of the decoupling between the different fundamental timescales for local relaxation [31, 77]. How this comes about is illustrated in Fig. 1.5. Panels (a) and (b) show trajectories of a probe particle embedded in a KCM. This can be thought of as a molecule in a liquid that has been labelled for tracking while coarse-graining over the rest of the system, thus describing its effective dynamics via a KCM. The motion of the probe particle is determined by the underlying fluctuations of the host KCM to which it is coupled. A natural dynamical rule is that the probe can make a diffusive jump from site i to site j only if both sites i and j are excited, $n_i = n_j = 0$ [31]. Figs. 1.5(a) and (b) show the difference in the probe motion between high and low temperatures. At high T , Fig. 1.5(a), excitations in the KCM are plentiful and probe motion appears Brownian. At low T , Fig. 1.5(b), excitations are scarce, the dynamics of the KCM is heterogeneous, and probe motion is intermittent:

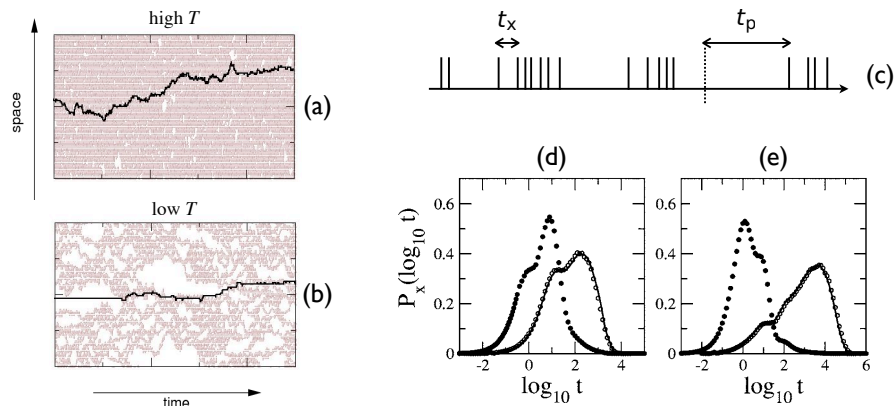


Fig. 1.5. Persistence/exchange decoupling. (a,b) Trajectory of a probe particle diffusing in a KCM. We show the case of the one-dimensional FA-1 for ease of illustration; similar behaviour is observed in other KCMs. The probe particle can make a diffusive step only if allowed by the local excitations in the KCM. (Sites with $n_i = 0$ are shown as grey in the background.) At high temperatures/low densities ($T = 3$) diffusion becomes Fickian after short times and distances. Plot (b) is on the same scale as (a), showing that at lower temperatures ($T = 0.8$) diffusion is intermittent and non-Fickian over much longer timescales and lengthscales. (c) Timeline of displacement events. The waiting time between events is termed local *exchange* time; it is the time measured to the next event with the knowledge of when the previous one took place. The waiting time until the next event from an arbitrarily chosen starting observation time is termed *persistence* time. When the time series of events in non-Poissonian, as a consequence of dynamical correlation in the KCMs, typical persistence and exchange times are different: exchange times are dominated by the clustering of events, while persistence times are determined by the long quiescent periods. (d,e) The decoupling between the distributions of exchange (earlier curves) and persistence (later curves) times becomes more pronounced the lower the temperature. Results are shown for the East model, with $T = 1$ in (d) and $T = 0.5$ in (e). Adapted from Refs. [31, 77], with permission.

the probe is immobile if immersed in an inactive space-time “bubble”; in order to move it has to wait for an excitation to come along.

There are two fundamental timescales that control this intermittent motion [31, 77]. The first one is the “persistence time”, t_p , that is, the time the probe needs to wait to start moving *for the first time*, given an arbitrary start time for observation. The second timescale is the (local) “exchange time”, t_x , the time *between moves*. Due to DH jump events are not Poissonian, but display “bunching”, and typical persistence times can become much larger than typical exchange times, see Fig. 1.5(c). This decoupling between persistence and exchange becomes more pronounced as temperatures is decreased, as illustrated in Figs. 1.5(d) and (e).

We can approximately quantify the motion of the probe particle [32] by means of a continuous-time random walk (CTRW) approach [79]. The probe makes random walk steps of unit size at random times determined by the fluctuations of the host KCM. Lets assume that this random clock ticks at times drawn from the exchange time distribution, $\phi(t_x)$, but which other than that are independently distributed. The probability for the probe to be at position r at time t , or van Hove function, is $G_s(r, t) = \sum_{m=0}^{\infty} \pi_m(t) \Gamma^{(m)}(r)$ (we assume $d = 1$ for simplicity, extension to higher dimensions is straightforward). Here $\pi_m(t)$ is the probability that the probe made m steps after time t , and $\Gamma^{(m)}(r)$ is the probability that a random walker is at a distance r after m steps. After Laplace transforming in time and Fourier transforming in space we obtain,

$$\hat{F}_s(k, \sigma) = \hat{P}(\sigma) + \cos(k) \frac{\hat{p}(\sigma)}{\sigma} \frac{1 - \hat{\phi}(\sigma)}{1 - \cos(k) \hat{\phi}(\sigma)}, \quad (1.3)$$

where $\hat{F}_s(k, \sigma)$ is the Laplace transform of the self-intermediate scattering function $F_s(k, t)$, which in turn is the Fourier transform of $G_s(r, t)$. Equation (1.3) is the Montroll-Weiss equation for the motion of the probe in the CTRW approximation. $\hat{\phi}(\sigma)$, $\hat{p}(\sigma)$ and $\hat{P}(\sigma)$ are the Laplace transforms of the exchange time distribution, the persistence time distribution, and the persistence function $P(t)$, respectively. The last two functions appear in Eq. (1.3) because the first step is determined by the persistence time: its distribution $p(t_p)$ is related to the persistence function $P(t) = \pi_0(t)$ by $P(t) = \int_t^{\infty} p(t_p) dt_p$. If the dynamics is stationary, then there is nothing special about time zero and $p(t_p) = \langle t_x \rangle^{-1} \int_{t_p}^{\infty} \phi(t') dt'$, corresponding to a uniform average over all earlier jump times; the normalization factor is the inverse of the average exchange time $\langle t_x \rangle = \int_0^{\infty} dt \phi(t) t$. From (1.3) we can define the wavelength dependent relaxation time: $\tau(k, T) = \lim_{\sigma \rightarrow 0} \hat{F}_s(k, \sigma)$. We obtain [32]

$$\tau(k) = \tau_p + \frac{\cos(k)}{1 - \cos(k)} \tau_x, \quad (1.4)$$

where τ_x and τ_p are the average exchange and persistence times, respectively. At low temperatures we have persistence/exchange decoupling, $\tau_p \gg \tau_x$ [77]. The structural or alpha relaxation time is often defined as $\tau_\alpha = \tau(k = \pi/a)$, where a is the lattice spacing which we set to unity. For these large wave vectors the first term of (1.4) dominates, and the structural relaxation time is set by the persistence time, $\tau_\alpha \approx \tau_p$. For small enough wave vectors the second term in (1.4) dominates, $\tau(k) \approx \tau_x/k^2$, and since the limit of $k \rightarrow 0$ defines the diffusion rate we find that $D \approx \tau_x^{-1}$. This explains Stokes-Einstein breakdown at low temperatures: $D\tau_\alpha \approx \tau_p/\tau_x \neq \text{const}$ [31].

The decoupling between persistence and exchange times is an effect of dynamical fluctuations. At low temperatures the self-diffusion constant seems to scale as a fractional power of the relaxation time $D \sim \tau_p^{-\delta}$, with $\delta < 1$ (the Stokes-Einstein relation is $\delta = 1$). This is the case for all KCMs with strong

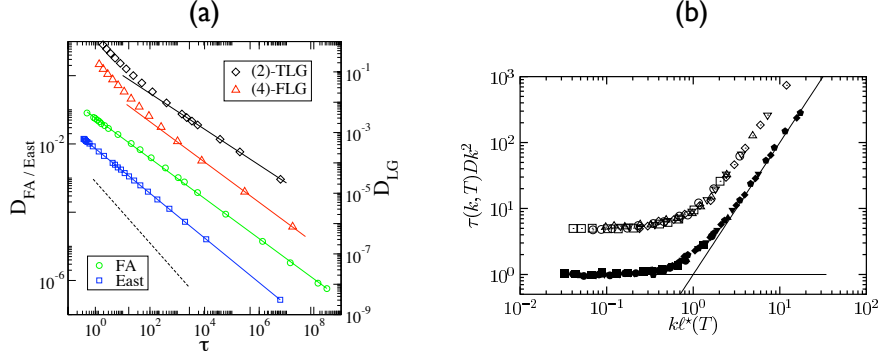


Fig. 1.6. Transport decoupling in KCMs. (a) Scaling of probe or self-diffusion constant with structural relaxation time in the one-dimensional FA-1 and East models (left scale) and two- and three- dimensional lattice gases (right scale). Solid lines indicate power law fits $D \sim \tau^{-\delta}$; the dashed line is the Stokes-Einstein relation $D \sim \tau^{-1}$. In all cases $\delta < 1$ at low enough q , indicating a breakdown of the Stokes-Einstein relation, as a consequence of dynamical fluctuations. (For the FA-1 model $\delta = 1$ for dimensions larger than its critical dimension $d_c = 2$ [49].) Adapted from Refs. [31, 69, 80], with permission. (b) Lengthscale dependent relaxation time $\tau(k)$ in the FA-1 and East models. At short lengthscales the relaxation time is k independent and determined by the persistence time, $\tau(k) \approx \tau_p$. At larger lengthscales it becomes diffusive, $\tau(k) \approx \tau_x/k^2$. This crossover is controlled by the Fickian lengthscale $l^* \sim \sqrt{\tau_p/\tau_x}$. The figure shows that $\tau(k)$ at different q collapse under $k \rightarrow k l^*$. From Ref. [32], with permission.

enough constraints, as shown in Fig. 1.6(a) for the FA-1 in $d = 1$, the East model and various constrained lattice gases [31, 69, 80]. The fractional exponent $\delta \approx 0.6 - 0.8$ is not distinct from that observed in experiments [72, 81].

Equation (1.4) indicates that there is a spectrum of timescales that interpolate between a k independent value τ_p at shorter lengthscales, to a diffusive timescale τ_x/k^2 at large lengthscales [32]. This crossover is shown in Fig. 1.6(b) for the FA-1 and East models. It is the crossover from non-Fickian diffusion at short lengthscales to eventual Fickian diffusion at long enough ones. The length l^* at which this crossover takes place, or “Fickian lengthscale”, is given by $l^* \propto \sqrt{\tau_p/\tau_x}$, and grows with decreasing temperature/increasing density. It is the distance a particle has to move before it forgets how long it took to make the first step. Fig. 1.6(b) shows how $\tau(k)$ at different temperatures collapse under $k \rightarrow l^* k$. The non-Fickian to Fickian crossover is also responsible for the wavelength dependence of four-point functions [34], Fig. 1.4(d).

The CTRW analysis can be extended to describe the effect of driving, for example by externally forcing the probes. The competition between timescales in this case leads to interesting non-linear response behaviour, such as non-monotonic differential mobility and giant diffusivity [82]. Furthermore, a study of (1.3) in the crossover regime between non-Fickian and Fickian explains [83]

the exponential tails observed [84] in van Hove functions at intermediate times. The waiting time distributions used in the CTRW analysis above are the ones that are obtained from the study of KCMs. The CTRW approach can also be used by assuming a different origin for the waiting time distributions, as for example in the analysis of metabasin transitions [85] in atomistic liquids.

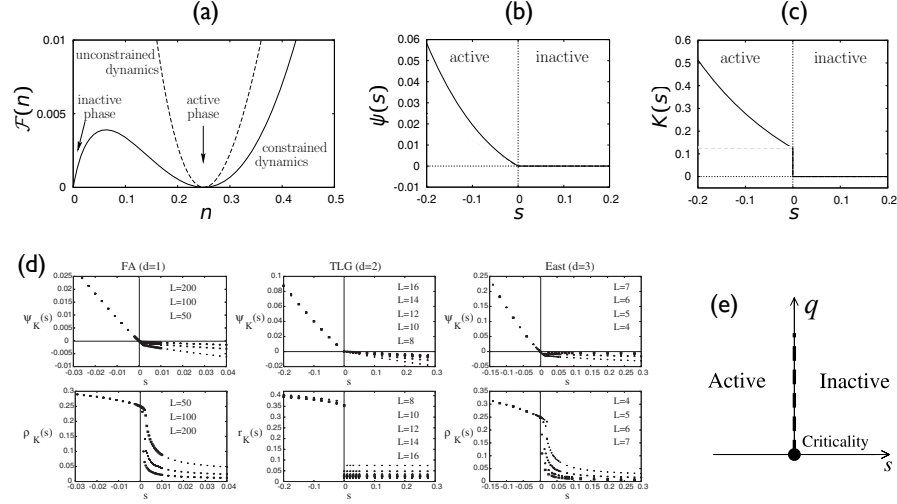


Fig. 1.7. Space-time phase transitions in KCMs. (a) Variational dynamical free-energy $\mathcal{F}(n)$ in terms of mean excitation density. Kinetic constraints give rise to a bistable $\mathcal{F}(n)$ (full line). In the absence of kinetic constraints the corresponding variational function is unistable (dashed line). (b) Mean-field estimate of large-deviation function $\psi(s)$. There is a singularity at $s = 0$, indicating a dynamical phase-transition. (c) Mean activity $K(s) \equiv -\psi'(s)$ as a function of s . This dynamical order parameter shows a discontinuous jump at $s = 0$: the transition is a first-order one between an active dynamical phase and an inactive dynamical phase. (d) The dynamical first-order scenario is also present in finite dimensions, as shown from numerical computation of $\psi(s)$ and $K(s)$ for various KCMs. (e) Dynamical phase diagram. The line $s = 0$ is one of first-order coexistence between the active dynamical phase ($s < 0$) and the inactive one ($s > 0$). It extends all the way along the q axis. The critical point at $q = 0$ controls the scaling behaviour discussed in Subsection 1.4. Adapted from Refs. [75, 86], with permission.

1.5.3 Space-time phase transitions

Thermodynamically KCMs are trivial, so all interesting behaviour is dynamical. Nevertheless, DH pictures such as those of Fig. 1.3 are suggestive of phase separation between two distinct phases. The phases of Fig. 1.3 are distinguished by their dynamics: dark regions are dynamically *active* while

light ones are dynamically *inactive*. Furthermore, the phase separation in the spatial projection of low temperature/high density equilibrium trajectories, such as those of Fig. 1.3, is only mesoscopic: when coarse-grained over large enough lengthscales, irrespective of time, the space projected trajectories are homogeneous. Here we show how these observations are directly related to a true non-equilibrium phase transition [75], which in contrast to thermodynamic transitions, occurs in ensembles of trajectories and is driven by non-equilibrium driving fields. This phase transitions can be studied by recourse to the large-deviation method [87, 88].

A convenient order parameter to discern active and inactive dynamics is the “dynamical activity” K [75], defined as the total number of configuration changes in a trajectory. In a non-conservative KCM it would amount to the total number of local changes from empty to occupied or vice versa, and in a lattice gas (conservative KCM) to the total number of particle displacements. The activity is extensive in space-time volume, i.e., typically $K = O(Nt)$, where N is the number of lattice sites and t the time extension of the trajectory. Each trajectory $x(t)$ in the ensemble of (equilibrium) trajectories of length t has a total activity, $\hat{K}[x(t)]$. The activity is thus distributed, $P_t(K) = \langle \delta(K - \hat{K}[x(t)]) \rangle$, where the average is over the set of equilibrium trajectories, $\{x(t)\}$. At long enough times this probability acquires a large-deviation form [87, 88]. $P_t(K) \approx e^{-t\varphi(K/t)}$. The function $\varphi(k)$ is called a large-deviation function, and it plays in this dynamical context the same role as, for example, the entropy density in the micro canonical ensemble of equilibrium statistical mechanics. Alternatively we can consider the generating function of K ,

$$Z_t(s) \equiv \sum_K e^{-sK} P_t(K) \approx e^{t\psi(s)}, \quad (1.5)$$

which also displays a large-deviation form. The large-deviation function $\psi(s)$ is akin to a free-energy density, and is related to $\varphi(K/t)$ by a Legendre transform, $\psi(s) = -\min_k [\varphi(k) + sk]$. Just like a free-energy in a thermodynamic problem, the function $\psi(s)$ carries the information of *dynamical* phase behaviour. Specifically, its singularities indicate *dynamical* phase transitions [75].

The calculation of the large-deviation function $\psi(s)$ is simplified greatly by the following observation [89]: if \mathbb{W} is the master operator that generates the stochastic dynamics, then $\psi(s)$ is the largest eigenvalue of a modified operator \mathbb{W}_s , where $\mathbb{W}_0 = \mathbb{W}$. This reduces the calculation from that of computing a “partition sum”, Eq. (1.5), to an eigenvalue problem. For example, for a spin facilitated model, $\mathbb{W}_s = \sum_i f_i(\mathbf{n}) \{e^{-s} [(1-\rho)\sigma_i^+ + \rho\sigma_i^-] - \rho n_i - (1-\rho)\}$, where $f_i(\mathbf{n})$ is the kinetic constraint on site i , and σ_i^\pm are the raising/lowering operators on site i . While it is not always possible to diagonalise such an operator analytically, bounds for its largest eigenvalue can be estimated variationally. For the FA- m this amounts to minimising a Landau free energy $\mathcal{F}(n) = -n^m (2e^{-s}\sqrt{qn} - q - n)$. The factor n^m comes from the kinetic constraint, and makes $\mathcal{F}(n)$ non-linear enough to allow for multiple minima, Fig.

1.7(a). The corresponding $\psi(s)$ has a singular structure, the first derivative being discontinuous at $s = 0$, Figs. 1.7(b) and (c). The meaning of this is the following. Trajectories are organized into two dynamical phases, an active one with $K > 0$ and an inactive one with $K = 0$. The field s determines the bias for or against activity. For $s < 0$ the active phase is the dominant one, while for $s > 0$ the inactive phase dominates. At $s = 0$ the probability of trajectories in either phase is equal in the $t \rightarrow \infty$ limit and we have dynamical first-order phase coexistence [75]. This situation occurs in all KCMs, Fig. 1.7(d). Interestingly, similar dynamical phase structure is observed in spin glass models [90, 91] and in atomistic liquids [92].

Actual dynamics takes place at $s = 0$. The results above show that this is the condition for dynamical coexistence *in the bulk*, i.e., infinitely far from boundaries active (ergodic) or inactive (non-ergodic) trajectories are equally likely. However, just like in the case of ordinary phase transitions, boundary fields can bias the bulk into one of the coexisting phases. In the case of dynamics, initial conditions play the role of a (time) boundary. In particular, almost all possible initial configurations chosen from the equilibrium static distribution at non-zero temperature or less than maximal density will select the active dynamical phase. In this case the inactive phase manifests via rare region effects, giving rise to DH only at the mesoscopic scale. An important open question is whether there are physical controllable fields that play the role of s in the analysis above.

1.6 Summary and outlook

In this chapter we have attempted to summarise recent developments in the study of KCMs as models of glass formers. In the long tradition of statistical mechanics KCMs provide simplified models that capture important ideas about the fundamental physics behind the phenomenology of glassy systems. Their simplicity allows for detailed study, which in turn gives rise to further physical insights into the glass transition problem. The central message from KCMs is that the complex and cooperative dynamics of glass forming systems can be achieved without recourse to complex thermodynamic behaviour: in KCMs thermodynamics plays essentially no role, and complex dynamics emerge from rather simple local kinetic rules. These rules are local and free of disorder, but nevertheless give rise to dynamical frustration.

The irrelevance of thermodynamics for glassy dynamics that the study of KCMs suggest contrasts sharply with approaches such as that of the random first-order transition theory [6, 93] where thermodynamics is essential. Whether thermodynamic aspects are relevant or not to glass transition phenomena is still a matter of debate, but to the extent that they are KCMs can say very little about them. This can either be seen as a flaw of the KCM based approach [94] or as an indication that these aspects are described by degrees of freedom that do not contribute too much to the long time dynam-

ics and have therefore been coarse-grained out [95]. Furthermore, by their coarse-grained and lattice based nature KCMs can in principle say very little about short distance/short time dynamics, such as beta-relaxation or anomalous vibrations. There is however evidence that these short scale phenomena are coupled to longer scale dynamic heterogeneity (see for example [96, 97]) so it may be possible to capture some of these effects with generalisations of KCMs [98, 99]. In any case, KCMs provide an explicitly real-space picture of glassy dynamics. Their main success has been the rationalization of dynamic heterogeneity. While DH can be analysed with other approaches, such as generalisations of mode-coupling theory [100], the immediacy of the results and explanations for DH related phenomena obtained from KCMs is remarkable.

While we know a lot about KCMs we still do not have a satisfactory understanding of how they emerge as an effective description from realistic systems. (This is also the case in alternative approaches, be it the random first-order transition [6], or frustrated limited domains [18], where the idealised models that display the proposed behaviour cannot be readily obtained from realistic liquid systems.) It is usually argued [5] that KCMs ought to emerge from some form of local coarse-graining of a microscopic system, but this procedure has not been shown to work just yet (except in highly simplified situations [101, 102]). Proving a direct connection between atomistic liquids and KCMs is arguably the central open problem in this field.

References

1. S H Glarum. Dielectric relaxation of polar liquids. *J. Chem. Phys.*, 33:1371–1375, 1960.
2. P W Anderson. Lectures on amorphous systems. In R Balian, R Maynard, and G Toulouse, editors, *Ill-Condensed Matter*, Les Houches Session XXXI. Elsevier Science, 1979.
3. R G Palmer, D L Stein, E Abrahams, and P W Anderson. Models of hierarchically constrained dynamics for glassy relaxation. *Phys. Rev. Lett.*, 53:958–961, 1984.
4. G Fredrickson and H Andersen. Kinetic Ising model of the glass transition. *Phys. Rev. Lett.*, pages 1244–1247, 1984.
5. J P Garrahan and D Chandler. Coarse-grained microscopic model of glass formers. *Proc. Natl. Acad. Sci.*, 100:9710–9714, 2003.
6. V Lubchenko and P G Wolynes. Theory of structural glasses and supercooled liquids. *Ann. Rev. Phys. Chem.*, 58:235–266, 2007.
7. L Berthier and J P Garrahan. Real space origin of temperature crossovers in supercooled liquids. *Phys. Rev. E*, 68:041201, 2003.
8. S Whitlam and J P Garrahan. Geometrical picture of dynamical facilitation. *J. Phys. Chem. B*, 108:6611–6615, 2004.
9. P Sollich and M R Evans. Glassy time-scale divergence and anomalous coarsening in a kinetically constrained spin chain. *Phys. Rev. Lett.*, 83:3238–3241, 1999.

10. C Toninelli, G Biroli, and D S Fisher. Spatial structures and dynamics of kinetically constrained models of glasses. *Phys. Rev. Lett.*, 92, 2004.
11. J P Garrahan and D Chandler. Geometrical explanation and scaling of dynamical heterogeneities in glass forming systems. *Phys. Rev. Lett.*, 89:035704, 2002.
12. M D Ediger. Spatially heterogeneous dynamics in supercooled liquids. *Ann. Rev. Phys. Chem.*, 51:99–128, 2000.
13. S C Glotzer. Spatially heterogeneous dynamics in liquids: insights from simulation. *J. Non-Cryst. Solids*, 274:342–355, 2000.
14. H C Andersen. Molecular dynamics studies of heterogeneous dynamics and dynamic crossover in supercooled atomic liquids. *Proc. Nat. Acad. Sci.*, 102:6686–6691, 2005.
15. D Chandler and J P Garrahan. *Ann. Rev. Phys. Chem.*, 61:191, 2010.
16. W Götze and L Sjögren. Relaxation processes in supercooled liquids. *Rep. Prog. Phys.*, 55:241–376, 1992.
17. P G Debenedetti and F H Stillinger. Supercooled liquids and the glass transition. *Nature*, 410:259–267, 2001.
18. S A Kivelson and G Tarjus. In search of a theory of supercooled liquids. *Nature Mater.*, 7:831–833, 2008.
19. A Cavagna. Supercooled liquids for pedestrians. *Phys. Rep.*, 476:51–124, 2009.
20. C Toninelli and G Biroli. *J. Stat. Phys.*, 130:83–112, 2008.
21. N Cancrini, F Martinelli, C Roberto, and C Toninelli. Kinetically constrained spin models. *Probab. Theory Related Fields*, 140:459–504, 2008.
22. A Faggionato, F Martinelli, C Roberto, and C Toninelli. In preparation.
23. F Ritort and P Sollich. Glassy dynamics of kinetically constrained models. *Adv. Phys.*, 52:219–342, 2003.
24. J Jäckle. Models of the glass-transition. *Rep. Prog. Phys.*, 49:171–231, 1986.
25. G H Fredrickson. Recent developments in dynamical theories of the liquid-glass transition. *Ann. Rev. Phys. Chem.*, 39:149–180, 1988.
26. R G Palmer. Models for slow relaxation in glassy systems. In H Takayama, editor, *Cooperative Dynamics in Complex Physical Systems*, volume 43 of *Springer Series in Synergetics*, pages 118–127. Springer, Heidelberg, 1989.
27. 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*, 48:4359–4363, 1993.
28. S Eisinger and J Jäckle. A hierarchically constrained kinetic Ising model. *Z. Phys. B*, 84:115–124, 1991.
29. M R Evans and P Sollich. Glassy time-scale divergence and anomalous coarsening in a kinetically constrained spin chain. *Phys. Rev. Lett.*, 83:3238–3241, 1999.
30. P Sollich and M R Evans. Glassy dynamics in the asymmetrically constrained kinetic Ising chain. *Phys. Rev. E*, 68:031504, 2003.
31. YJ Jung, J P Garrahan, and D Chandler. Excitation lines and the breakdown of Stokes-Einstein relations in supercooled liquids. *Phys. Rev. E*, 69:061205, 2004.
32. L Berthier, D Chandler, and J P Garrahan. Length scale for the onset of Fickian diffusion in supercooled liquids. *Europhys. Lett.*, 69:320–326, 2005.
33. 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*, 71:041505, 2005.

34. 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*, 74:051501, 2006.
35. Y S Elmatad, D Chandler, and J P Garrahan. Corresponding states of structural glass formers. *J. Phys. Chem. B*, 113:5563–5567, 2009.
36. C Toninelli, G Biroli, and D S Fisher. *Phys. Rev. Lett.*, 98:129602, 2007.
37. C Toninelli and G Biroli. *Eur. Phys. J. B*, 130, 2008.
38. C Toninelli, G Biroli, and D S Fisher. Jamming percolation and glass transitions in lattice models. *Phys. Rev. Lett.*, 96:035702, 2006.
39. M Jeng and J M Schwarz. On the study of force balance percolation. Preprint arXiv:0806.1552.
40. J Jäckle and A Krönig. A kinetic lattice-gas model for the triangular lattice with strong dynamic correlations. 1: Self-diffusion. *J. Phys. Condens. Matter*, 6:7633–7653, 1994.
41. L O Hedges and J P Garrahan. Dynamic propensity in a kinetically constrained lattice gas. *J. Phys. Condens. Matter*, 19:205124, 2007.
42. N Cancrini, F Martinelli, C Roberto, and C Toninelli. Kinetically constrained lattice gases. 2010.
43. M Aizenman and J L Lebowitz. Metastability effects in bootstrap percolation. *J. Phys. A*, 21:3801–3813, 1988.
44. R H Schonmann. On the behavior of some cellular automata related to bootstrap percolation. *Ann. Probab.*, 20:174–193, 1992.
45. C Toninelli, G Biroli, and D S Fisher. Cooperative behavior of kinetically constrained lattice gas models of glassy dynamics. *J. Stat. Phys.*, 120:167–238, 2005.
46. A V Goltsev S N Dorogovtsev and J F F Mendes. *Phys. Rev. Lett.*, 96:040601, 2006.
47. A J Liu J M Schwarz and L Q Chayes. *Europhys. Lett.*, 73:560, 2006.
48. J Reiter, F Mauch, and J Jäckle. Blocking transitions in lattice spin models with directed kinetic constraints. *Physica A*, 184:458–476, 1992.
49. R L Jack, P Mayer, and P Sollich. Mappings between reaction-diffusion and kinetically constrained systems: $A + A \leftrightarrow A$ and the Fredrickson-Andersen model have upper critical dimension $d_c = 2$. *J. Stat. Mech. Theor. Exp.*, P03006, 2006.
50. P Harrowell S Butler. *J. Chem. Phys.*, 85:4466, 1991.
51. G H Fredrickson and S A Brawer. *J. Chem. Phys.*, 84:3351, 1986.
52. I S Graham, L Piché, and M Grant. Model for dynamics of structural glasses. *Phys. Rev. E*, 55:2132–2144, 1997.
53. A E Holroyd. Sharp metastability threshold for two-dimensional bootstrap percolation. *Probab. Theory Related Fields*, 125:195–224, 2003.
54. J Reiter. Statics and dynamics of the 2-spin-facilitated kinetic Ising-model. *J. Chem. Phys.*, 95:544–554, 1991.
55. C Toninelli, G Biroli, and D S Fisher. Spatial structures and dynamics of kinetically constrained models for glasses. *Phys. Rev. Lett.*, 92:185504, 2004.
56. M Sellitto, G Biroli, and C Toninelli. *Europhys. Lett.*, 69:496, 2005.
57. N Cancrini, F Martinelli, C Roberto, and C Toninelli. Relaxation times of kinetically constrained spin models with glassy dynamics. *J. Stat. Mech. Theor. Exp.*, L03001, 2007.
58. A Buhot and J P Garrahan. Crossover from fragile to strong glassy behavior in kinetically constrained systems. *Phys. Rev. E*, 64:021505, 2001.

59. D Aldous and P Diaconis. The asymmetric one-dimensional constrained Ising model: rigorous results. *J. Stat. Phys.*, 107:945–975, 2002.
60. L Berthier, J P Garrahan, and S Whitlam. Renormalization group study of a kinetically constrained model for strong glasses. *Phys. Rev. E*, 71:026128–026142, 2005.
61. L Bertini and C Toninelli. Exclusion processes with degenerate rates: convergence to equilibrium and tagged particle. *J. Stat. Phys.*, 117:549–580, 2004.
62. C Toninelli and G Biroli. Dynamical arrest, tracer diffusion and kinetically constrained lattice gases. *J. Stat. Phys.*, 117:27–54, 2004.
63. A Crisanti and F Ritort. Violation of the fluctuation-dissipation theorem in glassy systems: basic notions and the numerical evidence. *J. Phys. A*, 36:R181–R290, 2003.
64. Peter Mayer and Peter Sollich. Ageing in one-dimensional coagulation-diffusion processes and the Fredrickson-Andersen model. *J. Phys. A*, 40:5823–5856, 2007.
65. P Mayer, S Leonard, L Berthier, J P Garrahan, and P Sollich. Activated aging dynamics and negative fluctuation-dissipation ratios. *Phys. Rev. Lett.*, 96:030602, 2006.
66. S Léonard, P Mayer, P Sollich, L Berthier, and J P Garrahan. Non-equilibrium dynamics of spin facilitated glass models. *J. Stat. Mech. Theor. Exp.*, P07017, 2007.
67. F Corberi and L F Cugliandolo. Out-of-equilibrium dynamics of the spiral model. *J. Stat. Mech. Theor. Exp.*, P09015, 2009.
68. N Cancrini, F Martinelli, R H Schonmann, and C Toninelli. Facilitated oriented spin models: some non-equilibrium results. 2010.
69. 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*, 72:041106, 2005.
70. S Whitlam, L Berthier, and J P Garrahan. Dynamic criticality in glass-forming liquids. *Phys. Rev. Lett.*, 92:185705, 2004.
71. L Berthier and J P Garrahan. Numerical study of a fragile three-dimensional kinetically constrained model. *J. Phys. Chem. B*, 109:3578–3585, 2005.
72. 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.*, 90:015901, 2003.
73. M Merolle, J P Garrahan, and D Chandler. Space-time thermodynamics of the glass transition. *Proc. Nat. Acad. Sci.*, 102:10837–10840, 2005.
74. R L Jack, J P Garrahan, and D Chandler. Space-time thermodynamics and subsystem observables in a kinetically constrained model of glassy materials. *J. Chem. Phys.*, 125:184509, 2006.
75. J P Garrahan, R L Jack, V Lecomte, E Pitard, K van Duijvendijk, and F van Wijland. Dynamical first-order phase transition in kinetically constrained models of glasses. *Phys. Rev. Lett.*, 98:195702, 2007.
76. N Lacevic, F W Starr, T B Schroder, and S C Glotzer. Spatially heterogeneous dynamics investigated via a time-dependent four-point density correlation function. *J. Chem. Phys.*, 119:7372–7387, 2003.
77. Y Jung, J P Garrahan, and D Chandler. Dynamical exchanges in facilitated models of supercooled liquids. *J. Chem. Phys.*, 123:084509, 2005.
78. I Chang and H Sillescu. Heterogeneity at the glass transition: Translational and rotational self-diffusion. *J. Phys. Chem. B*, 101:8794–8801, 1997.

79. E W Montroll and G H Weiss. Random walks on lattices. 2. *J. Math. Phys.*, 6:167, 1965.
80. D J Ashton. PhD thesis, University of Nottingham, 2009.
81. S F Swallen, K Traynor, R J McMahon, M D Ediger, and T E Mates. Self-diffusion of supercooled tris-naphthylbenzene. *J. Phys. Chem. B*, 113:4600–4608, 2009.
82. R L Jack, D Kelsey, J P Garrahan, and D Chandler. Negative differential mobility of weakly driven particles in models of glass formers. *Phys. Rev. E*, 78:011506, 2008.
83. P Chaudhuri, L Berthier, and W Kob. Universal nature of particle displacements close to glass and jamming transitions. *Phys. Rev. Lett.*, 99, 2007.
84. D A Stariolo and G Fabricius. Fickian crossover and length scales from two point functions in supercooled liquids. *J. Chem. Phys.*, 125, 2006.
85. A Heuer. Exploring the potential energy landscape of glass-forming systems: from inherent structures via metabasins to macroscopic transport. *J. Phys. Condens. Matter*, 20:373101, 2008.
86. J P Garrahan, R L Jack, V Lecomte, E Pitard, K van Duijvendijk, and F van Wijland. First-order dynamical phase transition in models of glasses: an approach based on ensembles of histories. *J. Phys. A*, 42:075007, 2009.
87. V Lecomte, C Appert-Rolland, and F van Wijland. Thermodynamic formalism for systems with Markov dynamics. *J. Stat. Phys.*, 127:51–106, 2007.
88. H Touchette. The large deviation approach to statistical mechanics. *Phys. Rep.*, 478:1–69, 2009.
89. J L Lebowitz and H Spohn. A gallavotti-cohen-type symmetry in the large deviation functional for stochastic dynamics. *J. Stat. Phys.*, 95:333–365, 1999.
90. K van Duijvendijk, R L Jack, and F van Wijland. Second-order dynamic transition in a $p = 2$ spin-glass model. *Phys. Rev. E*, 81:011110, 2010.
91. R L Jack and J P Garrahan. Metastable states and space-time phase transitions in a spin-glass model. *Phys. Rev. E*, 81:011111, 2010.
92. L O Hedges, R L Jack, J P Garrahan, and D Chandler. Dynamic order-disorder in atomistic models of structural glass formers. *Science*, 323:1309–1313, 2009.
93. M Mezard and G Parisi. Statistical physics of structural glasses. *J. Phys. Condens. Matter*, 12:6655–6673, 2000.
94. G Biroli, J P Bouchaud, and G Tarjus. Are defect models consistent with the entropy and specific heat of glass formers? *J. Chem. Phys.*, 123, 2005.
95. D Chandler and J P Garrahan. Thermodynamics of coarse-grained models of supercooled liquids. *J. Chem. Phys.*, 123:044511, 2005.
96. A. Widmer-Cooper, P. Harrowell, and H. Fynewever. How reproducible are dynamic heterogeneities in a supercooled liquid? *Phys. Rev. Lett.*, 93:135701, 2004.
97. C Brito and M Wyart. Heterogeneous dynamics, marginal stability and soft modes in hard sphere glasses. *J. Stat. Mech. Theor. Exp.*, L08003, 2007.
98. A J Moreno and J Colmenero. Anomalous relaxation in binary mixtures: a dynamic facilitation picture. *J. Phys. Condens. Matter*, 19:205144, 2007.
99. D. Ashton and J.P. Garrahan. Relationship between vibrations and dynamical heterogeneity in a model glass former: Extended soft modes but local relaxation. *Eur. Phys. J. E*, 30:303–307, 2009.
100. G Biroli, J P Bouchaud, K Miyazaki, and D R Reichman. Inhomogeneous mode-coupling theory and growing dynamic length in supercooled liquids. *Phys. Rev. Lett.*, 97:195701, 2006.

101. J P Garrahan and M E J Newman. Glassiness and constrained dynamics of a short-range nondisordered spin model. *Phys. Rev. E*, 62:7670–7678, 2000.
102. J P Garrahan, A Stannard, M O Blunt, and P H Beton. Molecular random tilings as glasses. *Proc. Nat. Acad. Sci.*, 106:15209–15213, 2009.