Ivailo Hartarsky and Cristina Toninelli

Kinetically constrained models

February 15, 2025

To our friend, mentor and collaborator Fabio Martinelli,
without whom this book would have ended here.

Contents

Preface

This manuscript focuses on Kinetically Constrained Models (KCM), a topic which lies at the intersection between probability and statistical mechanics. KCM are a class of Markov processes. They belong to the larger class of interacting particle systems with stochastic dynamics on discrete lattices. KCM were introduced in the physics literature in the 1980's to model the liquid-glass transition, a longstanding open problem in condensed matter physics. The key feature of KCM is that the update at a given lattice site can occur only if the configuration verifies a kinetic constraint requiring that there are no particles in a suitable neighbourhood. Extensive numerical simulations indicate that KCM display a remarkable behavior typical of glassy systems. Therefore, they have been the subject of several investigations in the last forty years with the aim of providing a deeper understanding of the liquid-glass transition and of more general jamming transitions.

Mathematically, KCM pose very challenging and interesting problems. In fact, the presence of the constraints induces non-attractiveness, the occurrence of several invariant measures, and the failure of many powerful tools to analyze relaxation to equilibrium (coercive inequalities, coupling, censoring...). Remarkably, the degeneracy of the rates caused by the constraints is not a mere technical obstacle which prevents using the classic tools. Indeed, the behavior of KCM is qualitatively different from that of interacting particle systems without constraints. Peculiar features include anomalously long mixing times, aging, singularities in the dynamical large deviation function, dynamical heterogeneities, and atypical ergodicity breaking transitions corresponding to the emergence of a large variety of amorphous structures. All in all, we can definitely say that KCM open a new chapter in the well established field of interacting particle systems.

Major progress has been made in the last twenty years towards a full and rigorous understanding of the large time behavior of KCM at stationarity. We present these results, illustrating both the high level ideas and some novel technical tools that have been devised to deal with the presence of constraints and with the lack of attractiveness. On the way, we unveil some remarkable connections of KCM with other mathematical subjects, in particular with bootstrap percolation cellular automata. We also present a choice of open problems concerning particularly the out of equi-

librium dynamics. Indeed, despite some achievements, robust tools to analyse KCM in this regime are still lacking and several beautiful questions remain open, even for simple choices of constraints.

This book aims at being accessible to both mathematicians and physicists. Hopefully it will be a useful tool to reinforce the bridge between the two communities which, in our opinion, have still much to learn from each other on KCM and glassy dynamics.

Outline

The content of the manuscript is as follows.

- In Chapter 1 we provide the physics background and motivation for studying KCM.
- In Chapter 2 we formally introduce KCM along with the relevant notation and key quantities of interest. It may be viewed as defining the scope of the manuscript.
- In Chapter 3 we discuss deterministic monotone cellular automata known as bootstrap percolation and their fundamental relation to KCM.
- In Chapter 4 we explore KCM in one dimension and introduce some basic tools, notably the bisection-constrained method. We focus particularly on the Fredrickson–Andersen 1-spin facilitated model (FA-1f) and on the East model, which not only serve as a warm-up for more advanced models, but also as a tool for their study.
- In Chapter 5 we consider the Fredrickson–Andersen 2-spin facilitated model in 2 dimensions. We develop progressively more sophisticated tools for its study, culminating with determining its sharp asymptotic behaviour at low temperature. These tools, which are flexible enough to be generalised to treat other models, include a robust long range Poincaré inequality and a very flexible multi-scale renormalisation tool, the Matryoska dolls.
- In Chapter 6 we examine the universality theory for KCM in one and two dimensions. It further elaborates our techniques and establishes a detailed map of the domain.
- In Chapter 7 we turn our attention to results on KCM out of equilibrium. Convergence to equilibrium and mixing times are investigated, using a set of tools completely separate from previous chapters.
- In Chapter 8 we briefly mention several settings, other than the one of Chapter 2, in which KCM have been studied. We also mention some closely related models, and provide more detailed references for the interested reader.

In view of Figure 0.1, Chapters 1 and 8 can be regarded as optional general knowledge. Chapters 2 and 3 are indispensable core material. A graduate course on the subject could cover these two chapters and a selection of Chapters 4 and/or 7, which both introduce a large variety of techniques in an accessible setting. The remaining Chapters 5 and 6 are intended for a more expert audience, particularly for newcomers to the field, who have already covered the basics, but need some background and intuition before delving into the details of specific papers.
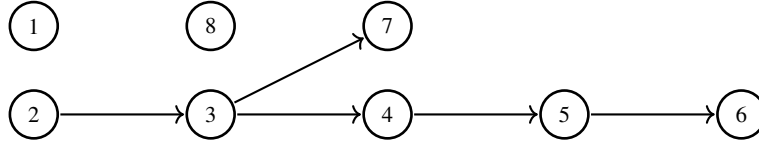
Fig. 0.1: Diagram of the dependencies between different chapters.

The more basic material (Chapters 1-4) is covered including full proofs or detailed sketches and featuring exercises to help assimilating the content. Subsequent chapters are less detailed and often refer to original papers for technical details. We apologise for the inevitable inaccuracies due to favouring simplicity over technical completeness. Indeed, we aim at highlighting the heuristic ideas and the guiding lines behind each method and result.

We have tried to keep the presentation as self-contained as possible, but there are some prerequisites. We do not assume, but hope for some familiarity with the basics of standard textbooks in the field such as [144] on Markov chains and [145, 146] on interacting particle systems, while some of the contents of an undergraduate course in probability theory, as covered for instance in [72, 77, 88, 142], will be used without notice. While it is possible to only refer to these books as needed, it may be a good investment to first acquire some superficial experience with their content, which is excellent to have in any case.

Acknowledgements

Chapter 1
Background from physics

In this chapter we discuss the physics background. We recall the basic phenomenology of the liquid-glass transition and more general jamming transitions, and explain the role of Kinetically Constrained Models (KCM) as models for glassy dynamics.

1.1 The liquid-glass transition

From the point of view of statistical physics, a key motivation behind the study of KCM is the effort to understand the *liquid-glass transition*. Glass is widely present in our daily life: it is a very versatile material, easily produced and manipulated on an industrial scale by cooling different liquid mixtures (e.g. silica, sodium carbonate and calcium oxide). And yet a microscopic understanding of this state of matter (which, according to archaeological findings in Egypt and Eastern Mesopotamia, people have been manufacturing since 3000 B.C.) and of how glass forms is still an open challenge for condensed matter physicists (see e.g. the review [8]). In 1995, Nobel prize Philip W. Anderson [7] wrote: *"The deepest and most interesting unsolved problem in solid state theory is probably the theory of the nature of glass and the glass transition."* He added, *"This could be the next breakthrough in the coming decade."* Thirty years later, physicists still disagree about the nature of glass and on how it forms.

At the heart of this puzzle lies the intriguing fact that the glasses display properties of both solids and liquids. In fact, we could either say "glass is an extremely viscous liquid that does not flow" or "glass is an unstructured, amorphous solid". Indeed, despite its macroscopic rigidity, the microscopic structure of a glass has the same disordered arrangement of molecules as a liquid. In other words, based on a single snapshot, liquid and glass are essentially indistinguishable.

This lack of order might seem in contrast with the thermodynamics paradigm that, when a liquid is cooled below its melting temperature, an ordered structure should form and the liquid should become a crystal. The secret is to perform the cooling sufficiently fast. This way the nucleation of the crystal is prevented and

the liquid enters a long-lived metastable state, the *super-cooled liquid* phase. Very roughly speaking, the liquid-crystal transition is avoided because molecules do not have enough time to reorganise and form the ordered crystal structure. The molecules move slower and slower forming a thick syrup and eventually they get trapped in the structureless glass state. Since the nucleation time of the stable crystal structure is out of reach for any reasonable experiment, though the state of super-cooled liquid is not thermodynamically stable, for all practical purposes it can be considered as an equilibrium system. In particular, one can define a relaxation time (and measure it experimentally via viscosity) and establish fluctuation-dissipation relations connecting response to an external driving force and correlations functions. Essentially, we can forget about the crystal and just focus on the super-cooled phase.

A common feature of super-cooled liquids around the glass transition is the sharp slowdown of dynamics. As shown in Figure 1.1, viscosity can increase by over 14 orders of magnitude upon a small decrease in temperature. It also highlights the fact that super-cooled liquids can be classified into two groups: *strong* and *fragile*. If we let η be the viscosity, T be the temperature and we define the activation energy as $E := T \log \eta$, strong liquids are characterised by a temperature-independent E, while for fragile liquids E increases as T decreases. The corresponding scaling forms for η are called *Arrhenius* and *super-Arrhenius*, respectively.

This dramatic growth of time scales[1] is related to the fact that when the temperature is lowered, the density increases: molecules tend to obstruct each other, blocked structures may arise, and motion becomes very cooperative[2]. A key experimental observation of the cooperative motion is the fact that when a glass cools, the molecules do not slow down uniformly. There is indeed a clear coexistence of fast and slow regions. This phenomenon is called *dynamical heterogeneity*: some regions of the liquid jam, while in other regions molecules continue to move around [21, 78]. Thus, even if a change of structure does not occur when the glass is formed, an underlying dynamical phase transition separating slow and fast trajectories seems to occur (see Figure 1.2). An indirect experimental evidence of dynamical heterogeneity is the decoupling of self-diffusion coefficient D_s and viscosity η: super-cooled liquids violate the phenomenological Stokes–Einstein relation, $D_s \eta / T = $ constant, which holds in homogeneous liquids. Though both D_s^{-1} and η increase when the temperature is lowered, the former does not increase as fast as the latter. This leads to $D_s \eta / T$ increasing by 2-3 orders of magnitude when temperature is decreased towards the glass transition temperature (see e.g. [149]). The reason why this decoupling of self-diffusion and viscosity is interpreted as a hallmark of dynamical heterogeneity is that the diffusion of the tracer particle (see Section 8.6 for a more formal description) should be dominated by the fastest regions and the structural relaxation (measured by the viscosity) by the slowest regions.

---

[1] For example, for fragile glass-forming liquids time scales at the melting temperature are typically of the order of the picosecond (which is also roughly the time scale of molecular motion), and are of the order of 100 seconds when temperature equals 2/3 of the melting temperature.

[2] That is, a growing number of particles need to move in a coordinated way in order for relaxation to occur.

Despite a great deal of experimental and theoretical investigation, a complete understanding of this behaviour and of other peculiar phenomena occurring in the vicinity of the glass transition (aging, hysteresis, rejuvenation, anomalous transport phenomena . . .) is still far out of reach. None of the numerous theories covers all the above phenomenology and a common consensus around "the" theory of the glass transition is still lacking in the physics community (see [8] for a review of various theories). A central theoretical difficulty is the fact that from the point of view of critical phenomena the situation is very peculiar: the liquid/glass transition displays a *mixed character*. Indeed, diverging time and length scales (typical of second order phase transitions), are accompanied by a discontinuous order parameter (typical of first order transitions). The jump of the order parameter corresponds to the discontinuous emergence of an amorphous density profile. Furthermore, both from the experimental and the theoretical point of view, the degeneracy of ground states complicates the problem. Thus, everybody agrees at least on one point: this is certainly not a standard type of ergodicity breaking transition!

The active research on the glass transition is also enhanced by the fact that a dynamical arrest towards an amorphous state displaying similar properties occurs in a large variety of systems upon tuning a proper external parameter (see e.g. [8, Sec. IV]). These phenomena, that are generically dubbed *jamming transitions*, occur for several materials: grains in powders (granular media), emulsions, foams, colloidal suspensions, polymers, plastics, ceramics, etc. Last but not least, understanding the glass transition would probably yield to novel theoretical and numerical tools that could be useful in other fields of science handling systems displaying a non trivial global collective behavior. This is the vast realm of systems that goes under the general name of *complex systems*.
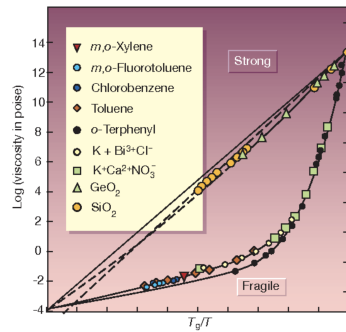
Fig. 1.1 Logarithm of the viscosity of several glass-forming liquids plotted against the inverse temperature. Here temperature is rescaled by the empirical glass transition temperature, T_g, defined as the value at which the viscosity equals $10^{12} Pa \cdot s$. Reprinted from [69] with permission. Copyright © 2001, Springer Nature Limited.
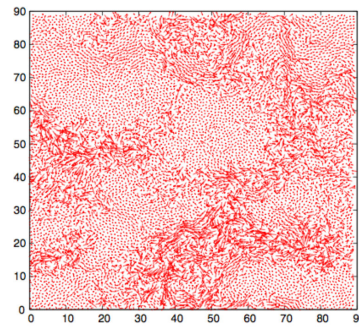
Fig. 1.2 Spatial map of single particle displacements in a molecular dynamics simulation of a super-cooled liquid in two dimensions. Arrows show the displacement of each particle in a trajectory of length comparable to the relaxation time. Reprinted from [20] with permission. Copyright © 2011 American Physical Society.

1.2 KCM as models for glassy dynamics

Kinetically constrained models (KCM) are toy models for the liquid-glass and jamming transitions. They rely on the idea that these are dynamical phenomena in which static interactions play a minor role. The kinetic constraints are therefore devised to mimic the mechanism of local caging which slows the dynamics down at low temperature or high density. Originally motivated by free volume theories [101], KCM have been promoted in the last decades as a paradigm model for the so called *dynamical facilitation theory* for the glass transition (see e.g. [27, 53]). Indeed, despite their simplicity and their trivial statics, KCM display many key dynamical features of real materials that undergo glass or jamming transitions: anomalous ergodicity breaking transitions, percolation of blocked structures, dynamical arrest, non-trivial spatio-temporal fluctuations, dynamical heterogeneities, aging. . . Furthermore, depending on the choice of the constraints, they feature either a super-Arrhenius or an Arrhenius behavior for the relaxation time, thus we recover both the behavior of fragile and strong supercooled liquids.

On the other hand, a major criticism to KCM is that a convincing derivation of these toy models via a coarse graining from realistic molecular models of liquids is missing (though some attempts in this direction have been recently performed both in experiments of granular glasses [49, 50] and in numerical simulations of super-cooled liquids [170]). In particular, it is not clear how one should identify at the molecular level the facilitating (empty) sites. We refer the reader to [8, 27, 99, 178] for further comments on successes and limitations of KCM as toy models for real glass forming liquids as well as for the illustration of alternative theories. What we can confidently say, adopting a sentence from [8], is that KCM have been influential and very instructive in order to develop a theoretical understanding of glassy phenomena.

Regarding derivations of KCM, we should also mention a different class of models which have been introduced to prove that that kinetic constraints can emerge spontaneously from static interactions, the so-called plaquette models. These are spin models with the usual Glauber dynamics reversible w.r.t. a Gibbs measure corresponding to a particular Hamiltonian $\mathcal{H}$. For certain choices of $\mathcal{H}$, the relaxation at low temperature is dominated by the motion of "excited" plaquettes. These excitations act as a source of mobility since the energetic barrier to flip a spin is smaller in their vicinity. Thus, at low temperature, their dynamics can be somehow mapped to a KCM. We will return to plaquette models in Section 8.4.

We also mention that in recent years there has been an increasing amount of work on the quantum versions of KCM that have been proposed in the study of Rydberg atoms [201] and as models for quantum many-body localization [95]. For example, in [171], the quantum version of the East KCM has been analyzed and the occurrence of a first-order quantum transition at which the ground state becomes exponentially localized (with a consequent slowdown of dynamics) has been detected.

Chapter 2
Models

In this chapter we provide a general definition of KCM and associated quantities of interest. We introduce the choices of constraints corresponding to some of the most studied models, including East, Friedrickson-Andersen j-spin facilitated (FA-jf), Duarte, North-East (NE) and Spiral models.

2.1 Notation

We denote the sets of nonnegative integers, integers and reals by $\mathbb{N}$, $\mathbb{Z}$ and $\mathbb{R}$ respectively. The models that we consider in this manuscript are most often defined on the infinite integer lattice $\mathbb{Z}^d$ or a subset thereof. We denote by $\mathbf{x} = (x_1, \dots x_d)$ the *sites* of $\mathbb{Z}^d$, by $\mathbf{e}_1, \dots, \mathbf{e}_d$ its *canonical basis vectors* and by $\Omega = \{0, 1\}^{\mathbb{Z}^d}$ the *configuration space*. Elements of Ω are called *configurations* and denoted by Greek letters $\sigma, \eta, \omega, \dots$. For a configuration $\omega \in \Omega$ and a site $\mathbf{x} \in \mathbb{Z}^d$, $\omega_{\mathbf{x}}$ denotes the *occupancy variable* (the value) of ω at $\mathbf{x}$. When $\omega_{\mathbf{x}} = 1$, we say that site $\mathbf{x}$ is *occupied*. When $\omega_{\mathbf{x}} = 0$ we say that $\mathbf{x}$ is *empty*. For $\omega \in \Omega$, we write $|\omega| = |\{\mathbf{x} \in \mathbb{Z}^d : \omega_{\mathbf{x}} = 0\}|$ for the number of empty sites in ω.

KCM are endowed with a parameter $q \in (0, 1]$ called *vacancy density* and corresponding to the inverse temperature β (in the right units) via the relation

$$q = 1/(1 + e^{\beta}). \tag{2.1}$$

In particular, the limit $q \to 0$ corresponds to the zero temperature limit, $\beta \to \infty$. Given $q \in [0, 1]$, we denote the product Bernoulli$(1 - q)$ measure on Ω, under which a site is empty with probability q, by μ_q or simply μ, when q is clear from the context.

The mean with respect to a measure ν on Ω of a function $f : \Omega \to \mathbb{R}$ is denoted by $\nu(f)$, while its variance is denoted by $\text{Var}_\nu(f)$ or simply $\text{Var}(f)$ when $\nu = \mu$. Functions f we are interested in only take a finite number of values, so no integrability issues arise.

We sometimes work on a subset $\Lambda \subset \mathbb{Z}^d$ of the lattice. Correspondingly, we set $\Omega_\Lambda = \{0,1\}^\Lambda$. The *restriction* of a configuration $\omega \in \Omega$ to Λ is denoted by $\omega_\Lambda \in \Omega_\Lambda$. For any measure ν that is the product of a measure on Ω_Λ and Ω_{Λ^c}, we similarly denote by ν_Λ the restriction to Λ. When $\Lambda = \{\mathbf{x}\}$ is a singleton, we simply write $\nu_{\mathbf{x}}$ for ν_Λ. We write Var_Λ for $\mathrm{Var}_{\mu_\Lambda}$, that is the variance with respect to the occupation variables in Λ. Given disjoint sets $\Lambda_1, \Lambda_2 \subset \mathbb{Z}^d$, and $\omega^{(1)} \in \Omega_{\Lambda_1}$ and $\omega^{(2)} \in \Omega_{\Lambda_2}$, we write $\omega^{(1)} \cdot \omega^{(2)} \in \Omega_{\Lambda_1 \cup \Lambda_2}$ for the configuration such that

$$(\omega^{(1)} \cdot \omega^{(2)})_{\mathbf{x}} = \begin{cases} \omega^{(1)}_{\mathbf{x}} & \mathbf{x} \in \Lambda_1, \\ \omega^{(2)}_{\mathbf{x}} & \mathbf{x} \in \Lambda_2. \end{cases} \tag{2.2}$$

We denote the fully occupied (resp. empty) configurations by $\mathbf{1}$ (resp. $\mathbf{0}$) and omit the domain Λ in $\mathbf{1}_\Lambda$ (resp. $\mathbf{0}_\Lambda$), when it is clear from the context.

2.2 Update families

KCM are Glauber type Markov processes on Ω (or Ω_Λ) reversible w.r.t. μ (μ_Λ, respectively). We give the general definition introduced in [44] to cover all the models studied in physics. Each KCM is characterised by its *update family*, namely a finite collection $\mathcal{U} = \{U_1, \ldots, U_m\}$ of finite subsets of $\mathbb{Z}^d \setminus \{0\}$ called *update rules*. Given a vertex $\mathbf{x} \in \mathbb{Z}^d$, we will say that the *constraint* at $\mathbf{x}$ is satisfied by the configuration $\omega \in \Omega$ if at least one of the update rule translated at $\mathbf{x}$ is completely empty, namely,

$$c_{\mathbf{x}}(\omega) = c_{\mathbf{x}}^{\mathcal{U}}(\omega) = \begin{cases} 1 & \exists U \in \mathcal{U}, \forall \mathbf{u} \in U, \omega_{\mathbf{x}+\mathbf{u}} = 0, \\ 0 & \text{otherwise.} \end{cases} \tag{2.3}$$

Observe that constraints $c_{\mathbf{x}}$ are non-increasing w.r.t. the product partial order on configurations in Ω given by $\omega \leqslant \omega'$ if $\omega(\mathbf{y}) \leqslant \omega'(\mathbf{y})$ for all $\mathbf{y} \in \mathbb{Z}^d$. Inversely, any non-increasing function $c_0 : \Omega \to \{0,1\}$ depending only on the restriction of the configuration to a finite subset of $\mathbb{Z}^d \setminus \{0\}$ can be written in the form $c_0^{\mathcal{U}}$ for some update family $\mathcal{U}$. Correspondingly, there is a natural partial order on update families: $\mathcal{U}_1 \leqslant \mathcal{U}_2$, if $c_0^{\mathcal{U}_1}(\omega) \leqslant c_0^{\mathcal{U}_2}(\omega)$ for every $\omega \in \Omega$.[1] The maximal and minimal update families correspond to $\mathcal{U} = \varnothing$ (i.e. $c_0^{\mathcal{U}} \equiv 0$) and $\mathcal{U} = \{\varnothing\}$ (i.e. $c_0^{\mathcal{U}} \equiv 1$). Since the goal of KCM is to explore the effect of constraints on the dynamics, we discard these cases, by systematically assuming update families and update rules to be nonempty. Moreover, we usually drop $\mathcal{U}$ from the notation, as $\mathcal{U}$ is fixed or arbitrary.

We now present some of the most commonly considered update families (see Figure 2.1 for their two-dimensional representation). While at this point they may

---

[1] Note that since we only care about the update family through the constraint it induces, we identify update families yielding the same constraints, e.g. $\{\{1\}\}$ and $\{\{1\}, \{1,2\}\}$ in one dimension.
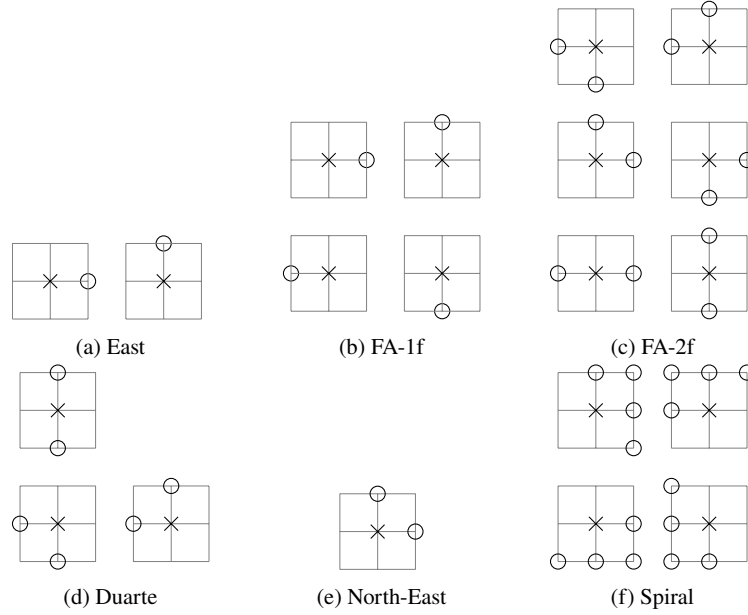
Fig. 2.1: The two-dimensional update families introduced in Section 2.2.

seem arbitrary, in Chapter 6, we will see that they are representatives of different universality classes displaying very different behaviour.

- The **East** [9, 131] update family is $\mathcal{U} = \{\{\mathbf{e}_1\}, \ldots, \{\mathbf{e}_d\}\}$. That is, the East constraint at $\mathbf{x} \in \mathbb{Z}^d$ is satisfied, if $\mathbf{x}$ has an empty neighbour in some of the positive coordinate directions.
- The **Frederickson–Andersen j-spin facilitated (FA-jf or j-neighbour)** [89, 90] update family is $\mathcal{U} = \{U \subset \{\mathbf{e}_1, \ldots, \mathbf{e}_d, -\mathbf{e}_1, \ldots, -\mathbf{e}_d\} : |U| = j\}$, where $j \in \{1, \ldots, 2d\}$ is a parameter. That is, the FA-jf constraint at $\mathbf{x} \in \mathbb{Z}^d$ is satisfied, if $\mathbf{x}$ has at least j empty neighbours.
- The **Duarte** [71] update family is $\mathcal{U} = \{\{-\mathbf{e}_1, \mathbf{e}_2\}, \{-\mathbf{e}_1, -\mathbf{e}_2\}, \{-\mathbf{e}_2, \mathbf{e}_2\}\}$ in two dimensions. That is, the Duarte constraint at $\mathbf{x} \in \mathbb{Z}^2$ is satisfied, if $\mathbf{x}$ has at least two empty neighbours other than $\mathbf{x} + \mathbf{e}_1$.
- The **North-East (NE)** [177] update family is $\mathcal{U} = \{\{\mathbf{e}_1, \ldots, \mathbf{e}_d\}\}$ for $d \geq 2$ (for $d = 1$ we get the East model). That is, the constraint at $\mathbf{x} \in \mathbb{Z}^d$ is satisfied, if all the neighbours of $\mathbf{x}$ in the positive coordinate directions are empty.
- The **Spiral** [196] update family is $\mathcal{U} = \{U_1, U_2, U_3, U_4\}$ in two dimensions, where $U_1 = \{(1, -1), (1, 0), (1, 1), (0, 1)\}$ and U_2, U_3, U_4 are obtained by rotating U_1 by $\pi/2$, π and $3\pi/2$ around the origin respectively.

2.3 The Markov process

Let us first informally describe KCM on $\mathbb{Z}^d$ via its so-called *graphical representation*. Each vertex $\mathbf{x} \in \mathbb{Z}^d$ is equipped with a unit intensity Poisson process, whose atoms $t_{\mathbf{x},k}$ for $k \in \mathbb{N}$ are the *clock rings*. We are further given independent Bernoulli random variables $s_{\mathbf{x},k}$ with parameter $1 - q$ called *coin tosses*. At the clock ring $t_{\mathbf{x},k}$, if the current configuration ω satisfies the constraint at $\mathbf{x}$, we set the occupation variable $\omega_{\mathbf{x}}$ to $s_{\mathbf{x},k}$. Such updates are called *legal*. If, on the contrary, the constraint is not satisfied, the configuration remains unchanged at time $t_{\mathbf{x},k}$.

More formally, the Markov process can be constructed via its self-adjoint Markov *semigroup* $P_t := e^{t\mathcal{L}}$ on $\mathrm{L}^2(\mu)$, where the *generator* $\mathcal{L}$ is a non-negative self-adjoint operator with domain $\mathrm{Dom}(\mathcal{L})$ that can be constructed in a standard way (see e.g. [146, Sections I.3, IV.4]) starting from its action on *local functions* (i.e. functions depending on the occupancy variables on a finite number sites):

$$\mathcal{L}f = \sum_{\mathbf{x} \in \mathbb{Z}^d} c_{\mathbf{x}} \cdot (\mu_{\mathbf{x}}(f) - f) \,. \tag{2.4}$$

Spelling out the definition of $\mu_{\mathbf{x}}$, the generator can be equivalently rewritten as

$$\mathcal{L}f(\omega) = \sum_{\mathbf{x} \in \mathbb{Z}^d} c_{\mathbf{x}}(\omega) \left((1 - \omega_{\mathbf{x}})(1 - q) + \omega_{\mathbf{x}}q\right) (f(\omega^{\mathbf{x}}) - f(\omega)) \tag{2.5}$$

with $\omega^{\mathbf{x}}$ the configuration obtained from ω by flipping its value at $\mathbf{x}$, i.e.

$$\omega_{\mathbf{y}}^{\mathbf{x}} = \begin{cases} \omega_{\mathbf{y}} & \mathbf{y} \neq \mathbf{x}, \\ 1 - \omega_{\mathbf{x}} & \mathbf{y} = \mathbf{x}. \end{cases} \tag{2.6}$$

We further introduce the Dirichlet form $\mathcal{D} : \mathrm{Dom}(\mathcal{L}) \to \mathbb{R}$ defined as

$$\mathcal{D}(f) = -\mu(f \cdot \mathcal{L}f) = \sum_{\mathbf{x} \in \mathbb{Z}^d} \mu \left(c_{\mathbf{x}} \mathrm{Var}_{\mathbf{x}}(f)\right). \tag{2.7}$$

Using the formulation (2.5) it is not hard to verify that

$$\mathcal{D}(f) = \frac{1}{2} \int \mu(\mathrm{d}\omega) \sum_{\mathbf{x} \in \mathbb{Z}^d} c_{\mathbf{x}}(\omega) \left((1 - \omega_{\mathbf{x}})(1 - q) + \omega_{\mathbf{x}}q\right) (f(\omega^{\mathbf{x}}) - f(\omega))^2 \,.$$

When the initial distribution at time $t = 0$ is ν, the law and expectation of the KCM process on the Skorokhod space $D([0, \infty), \Omega)$ of càdlàg functions are denoted by $\mathbb{P}_\nu$ and $\mathbb{E}_\nu$ respectively (see [26, Chapter III] for background). If ν is concentrated over a single configuration, $\nu = \delta_\sigma$, we write simply $\mathbb{P}_\sigma$ and $\mathbb{E}_\sigma$ for $\mathbb{P}_\nu$ and $\mathbb{E}_\nu$, while if $\nu = \mu$, we simply write $\mathbb{P}$ and $\mathbb{E}$. We use $\omega(t)$ to denote the state of the KCM at time $t \geqslant 0$.

We next discuss an important property of KCM—reversibility (see [146, Section II.5] for background). It is not hard to verify that, since the constraint $c_{\mathbf{x}}(\omega)$ does

not depend on $\omega_{\mathbf{x}}$, the dynamics satisfies detailed balance w.r.t. the product measure μ. Therefore, μ is reversible (i.e. $\mu(f \cdot P_t g) = \mu(g \cdot P_t f)$ for all $f, g \in L^2(\mu)$ and $t \geqslant 0$) and therefore it is an invariant measure for the process (i.e. $\mu P_t = \mu$ for all $t \geqslant 0$). However, μ is *not* the unique invariant measure—e.g. the Dirac measure on the fully occupied configuration is clearly invariant in view of (2.3) and (2.4). We nevertheless refer to the KCM with initial condition μ as the *stationary process*.

2.4 Boundary conditions

KCM can also be defined on finite or infinite subsets $\Lambda \subset \mathbb{Z}^d$ (we write $\Lambda \Subset \mathbb{Z}^d$ when we assume that Λ is finite). In this case the most natural choice is to imagine that the configuration is defined also outside Λ, where it is frozen and equal to some reference configuration $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$, the *boundary condition*. Then, for $\mathbf{x} \in \Lambda$, $\omega \in \Omega_\Lambda$, the constraint is defined as

$$c_{\mathbf{x}}^\sigma(\omega) = c_{\mathbf{x}}(\omega \cdot \sigma) \tag{2.8}$$

(recall (2.2) and (2.3)). We denote by $\mathcal{L}_\sigma$ and $\mathcal{D}_\sigma$ the generator and Dirichlet form of this process on Ω_Λ that are obtained by restricting the sums in (2.4) and (2.7) to sites in Λ and substituting $c_{\mathbf{x}}$ with $c_{\mathbf{x}}^\sigma$. We similarly denote by $\mathbb{P}_\nu^\sigma$ and $\mathbb{E}_\nu^\sigma$ ($\mathbb{P}_\zeta^\sigma$ and $\mathbb{E}_\zeta^\sigma$, if $\nu = \delta_\zeta$, and $\mathbb{P}^\sigma$ and $\mathbb{E}^\sigma$, if $\nu = \mu$) the law and expectation of the process with initial distribution ν and by $\omega^\sigma(t)$ the process at time t. Note that μ_Λ is reversible for this process.

2.5 Characteristic times and critical parameters

Having defined KCM, we next discuss the kinds of questions and quantities that we seek to tackle for them. Let us start with three natural observables. The *emptying*, *occupation* and *persistence* times of a KCM $(\omega(t))_{t \geqslant 0}$ are given by

$$\tau_0 = \inf\{t \geqslant 0 : \omega_0(t) = 0\}, \quad \tau_1 = \inf\{t \geqslant 0 : \omega_0(t) = 1\}, \quad \tau_\vee = \max(\tau_0, \tau_1) \tag{2.9}$$

respectively. Turning to more analytic quantities, the *relaxation time* (also known as the inverse of the *spectral gap* of $\mathcal{L}$) is defined as

$$T_{\mathrm{rel}} = \frac{1}{\mathrm{gap}}, \qquad\qquad \mathrm{gap} = \inf_{\substack{f \in \mathrm{Dom}(\mathcal{L}) \\ \mathrm{Var}(f) \neq 0}} \frac{\mathcal{D}(f)}{\mathrm{Var}(f)} \tag{2.10}$$

where $\mathcal{D}$ is the Dirichlet form (2.7). This definition is equivalent to saying that T_{rel} is the smallest constant $C \geqslant 0$ such that the *Poincaré inequality*

$$\mathrm{Var}(f) \leqslant C \mathcal{D}(f) \tag{2.11}$$

is satisfied for any $f \in \mathrm{Dom}(\mathcal{L})$. A finite relaxation time is equivalent to the fact that the measure μ is mixing for the semigroup P_t with exponentially decaying correlations (see e.g. [179]), namely for all $f \in \mathrm{L}^2(\mu)$ it holds that[2]

$$\mathrm{Var}\,(P_t f) \leqslant e^{-2t/T_{\mathrm{rel}}} \,\mathrm{Var}(f). \tag{2.12}$$

Thus, the relaxation time controls the decay of correlations in the stationary process. Of course, the above time scales have no reason to be finite, so it is natural to consider the corresponding critical parameters. We define the *ergodicity* and *exponential decay* critical parameters

$$q_{\mathrm{c}} = q_{\mathrm{c}}(\mathcal{U}) = \inf\left\{q > 0 : \mathbb{P}(\tau_0 < \infty) = 1\right\}, \tag{2.13}$$

$$\tilde{q}_{\mathrm{c}} = \tilde{q}_{\mathrm{c}}(\mathcal{U}) = \inf\left\{q > 0 : T_{\mathrm{rel}} < \infty\right\}. \tag{2.14}$$

Our main goals are to determine (or characterise) q_{c} and $\tilde{q}_{\mathrm{c}}$ and to study the asymptotics of τ_0 and T_{rel} for the stationary process as $q \to q_{\mathrm{c}}+$, as well as the behaviour of the KCM out of equilibrium (with initial condition different from μ) for any $q > 0$.

Once a Poincaré inequality (2.11) is established it is natural to ask whether it would be possible to prove also a stronger coercive inequality for the generator. In particular one can investigate whether a logarithmic or modified logarithmic Sobolev inequality holds (see [179] for background). These correspond, respectively, to the existence of a finite constant C_{LS} and C_{MLS} such that for any nonnegative $f \in \mathrm{Dom}(\mathcal{L})$ we have

$$\mathrm{Ent}(f) \leqslant C_{\mathrm{LS}} \mathcal{E}\left(\sqrt{f}, \sqrt{f}\right), \tag{2.15}$$

$$\mathrm{Ent}(f) \leqslant C_{\mathrm{MLS}} \mathcal{E}(f, \log f), \tag{2.16}$$

where for any two functions f, g we set

$$\mathrm{Ent}(f) = \mu\left(f \log(f/\mu(f))\right), \qquad \mathcal{E}(f, g) = -\mu(f \mathcal{L} g).$$

The (stronger) inequality (2.15) is known to be equivalent to the hypercontractivity property of the semigroup P_t [70]. Instead, (2.16) is equivalent to exponential decay of the relative entropy for the generator P_t (see [35]), namely for each probability measure ν on Ω it holds that[3]

$$H(\nu P_t || \mu) \leqslant e^{-t/C_{\mathrm{MLS}}} H(\nu || \mu) \tag{2.17}$$

where for any two measures ν_1, ν_2 on Ω we denote by $H(\nu_1 || \nu_2)$ the relative entropy (or Kullback–Leibler divergence) of ν_1 w.r.t. ν_2.

---

[2] Indeed, by reversibility, $\frac{\mathrm{d}}{\mathrm{d}t} \mathrm{Var}(P_t f) = -2\mathcal{D}(P_t f)$.

[3] Indeed, if we let f_t be the relative density of νP_t w.r.t. μ it holds that $\frac{\mathrm{d}}{\mathrm{d}t} H(\nu P_t || \mu) = \frac{\mathrm{d}}{\mathrm{d}t} \mathrm{Ent}(f_t) = -\mathcal{E}(f_t, \log f_t)$.

The definitions in (2.9), (2.10), (2.15), (2.16) naturally extend to the finite volume setting. We denote the relaxation time (logarithmic or modified logarithmic Sobolev constants) for the KCM with boundary condition σ by T_{rel}^{σ} (resp. C_{MLS}^{σ} and C_{MLS}^{σ}). When $\sigma = \mathbf{0}_{\mathbb{Z}^d \setminus \Lambda}$ for $\Lambda \subset \mathbb{Z}^d$, we simplify the notation to T_{rel}^{Λ} (resp. C_{MLS}^{Λ} and C_{MLS}^{Λ}).

We finally define another natural time scale known as mixing time of the KCM on $\Lambda \subset \mathbb{Z}^d$ with boundary condition $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ (see [144] for background). The total variation distance d_{TV} between two measures μ_1, μ_2 on Ω_{Λ} is given by

$$d_{\text{TV}}(\mu_1, \mu_2) = \sup_{A \in \mathcal{F}} |\mu_1(A) - \mu_2(A)|, \tag{2.18}$$

where $\mathcal{F}$ denotes the Borel σ-field generated by the open sets of Ω_{Λ}. For $\varepsilon \in (0, 1)$, the mixing time is

$$t_{\text{mix}}^{\sigma}(\varepsilon) = \inf \left\{ t > 0, \ \max_{\omega \in \Omega_{\Lambda}} d_{\text{TV}}(\delta_{\omega} P_t, \mu) \leqslant \varepsilon \right\}. \tag{2.19}$$

If $\sigma = \mathbf{0}_{\mathbb{Z}^d \setminus \Lambda}$, we simplify the notation to $t_{\text{mix}}^{\sigma} = t_{\text{mix}}^{\Lambda}$. The mixing time has a natural probabilistic interpretation: it is the time when it is impossible to distinguish (e.g. by a statistical test) the law of the process at that time from the equilibrium measure, regardless of the initial state.

Chapter 3
From bootstrap percolation to kinetically constrained models

In this chapter we introduce bootstrap percolation cellular automata, which are instrumental for the study of KCM. We first state some relevant known results for these automata. We then show that several fundamental properties of KCM, including ergodicity, mixing and exponential relaxation, can be directly related to their bootstrap percolation counterparts.

3.1 Bootstrap percolation

Bootstrap percolation (BP) is a family of monotone cellular automata. They may be viewed as the discrete time synchronous monotone analogue of KCM. Specific instances of BP have been studied since the 1970s [54, 138, 176], but it is convenient to directly introduce them in greater generality as considered in [40, 106, 181]. Like KCM, BP is defined by an update family $\mathcal{U}$ (recall Section 2.2). Given $\omega \in \Omega$, we define $\mathcal{B}_{\mathcal{U}}(\omega) \in \Omega$ by[1]

$$(\mathcal{B}_{\mathcal{U}}(\omega))_{\mathbf{x}} = \begin{cases} 0 & \omega_{\mathbf{x}} = 0 \text{ or } \exists U \in \mathcal{U}, \forall \mathbf{u} \in U, \omega_{\mathbf{x}+\mathbf{u}} = 0, \\ 1 & \text{otherwise} \end{cases} \tag{3.1}$$

for all $\mathbf{x} \in \mathbb{Z}^d$. In words, in one discrete time step we empty all sites for which the constraint (recall (2.3)) is satisfied. In BP, empty sites remain empty. In view of this monotonicity, it is natural to define the *closure*

$$[\omega] = [\omega]_{\mathcal{U}} = \inf_{t \in \mathbb{N}} \mathcal{B}_{\mathcal{U}}^{\circ t}(\omega) = \lim_{t \to \infty} \mathcal{B}_{\mathcal{U}}^{\circ t}(\omega) \in \Omega, \tag{3.2}$$

where $\mathcal{B}_{\mathcal{U}}^{\circ t}$ denotes the t-fold iteration of $\mathcal{B}_{\mathcal{U}}$ and the infimum and limit are taken with respect to the product partial order and product topology respectively. That is, $[\omega]$ is the configuration obtained upon iterating the bootstrap percolation map of

---

[1] Note that, in most of BP literature, the roles of the 0 and 1 states are reversed.

(3.1). Similarly to (2.9), we define the *BP emptying time*

$$\tau_0^{\mathrm{BP}} = \inf \left\{ t \in \mathbb{N} : \left(\mathscr{B}_{\mathcal{U}}^{\circ t}(\omega) \right)_0 = 0 \right\} \in \mathbb{N} \cup \{\infty\}. \tag{3.3}$$

On a domain $\Lambda \subset \mathbb{Z}^d$ with boundary condition $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$, we set

$$\left(\mathscr{B}_{\mathcal{U}}^{\sigma}(\omega) \right)_{\mathbf{x}} = \begin{cases} 0 & \omega_{\mathbf{x}} = 0 \text{ or } \exists U \in \mathcal{U}, \forall \mathbf{u} \in U, (\sigma \cdot \omega)_{\mathbf{x}+\mathbf{u}} = 0, \\ 1 & \text{otherwise} \end{cases} \tag{3.4}$$

for $\omega \in \Omega_{\Lambda}$ and $\mathbf{x} \in \Lambda$. We further define $[\omega]^{\sigma} = \lim_{t \to \infty} (\mathscr{B}_{\mathcal{U}}^{\sigma})^{\circ t}(\omega)$ pointwise.

So far BP is completely deterministic. We next introduce randomness by considering an initial condition ω distributed according to the product Bernoulli measure μ_q with parameter (density of initially occupied sites) $1 - q \in [0, 1]$ (recall Section 2.1). Following (2.13) and (2.14), we define the *emptying* and *exponential decay* critical thresholds

$$q_{\mathrm{c}}^{\mathrm{BP}} = \inf \left\{ q > 0 : \mu_q \left(\tau_0^{\mathrm{BP}} < \infty \right) = 1 \right\}, \tag{3.5}$$

$$\tilde{q}_{\mathrm{c}}^{\mathrm{BP}} = \inf \left\{ q > 0 : \liminf_{t \to \infty} \frac{-\log \mu_q(\tau_0^{\mathrm{BP}} > t)}{t} > 0 \right\}. \tag{3.6}$$

Note that by ergodicity of the product measure (with respect to translations, see e.g. [136] for background), we have $\mu_q(\tau_0^{\mathrm{BP}} < \infty) = 1$ if and only if $\mu_q([\omega] = \mathbf{0}_{\mathbb{Z}^d}) = 1$. As for KCM, one is primarily interested in determining these thresholds and the asymptotics of τ_0^{BP} as $q \to q_{\mathrm{c}}^{\mathrm{BP}}+$.

In the present text our focus is on KCM, so we take BP results for granted. Let us therefore gather a few facts about the BP models corresponding to the update families introduced in Section 2.2.

Theorem 3.1 (BP background)

- *For East BP in dimension $d \geqslant 1$, $q_{\mathrm{c}}^{\mathrm{BP}} = \tilde{q}_{\mathrm{c}}^{\mathrm{BP}} = 0$ and $q^{1/d} \tau_0^{\mathrm{BP}}$ converges to a Weibull distribution, as $q \to 0$: $\mu(q^{1/d} \tau_0^{\mathrm{BP}} \geqslant t) \to e^{-t^d/d!}$ for $t \geqslant 0$.*
- *For 1-neighbour BP in dimension $d \geqslant 1$, $q_{\mathrm{c}}^{\mathrm{BP}} = \tilde{q}_{\mathrm{c}}^{\mathrm{BP}} = 0$ and $q^{1/d} \tau_0^{\mathrm{BP}}$ converges to a Weibull distribution, as $q \to 0$: $\mu(q^{1/d} \tau_0^{\mathrm{BP}} \geqslant t) \to e^{-(2t)^d/d!}$ for $t \geqslant 0$.*
- *For j-neighbour BP in dimension d with $d \geqslant j \geqslant 2$ we have $q_{\mathrm{c}}^{\mathrm{BP}} = \tilde{q}_{\mathrm{c}}^{\mathrm{BP}} = 0$ and $q^{1/(d-j+1)} \log^{\circ(j-1)} \tau_0^{\mathrm{BP}}$ converges in probability to a constant[2] $\lambda(d, j) > 0$, as $q \to 0$.*
- *For 2-neighbour BP in $d = 2$ dimensions, there exist positive constants[3] λ, λ_2 such that, as $q \to 0$, we have*

$$\mu_q \left(\left| \log \tau_0^{\mathrm{BP}} - \frac{\lambda}{q} + \frac{\lambda_2}{\sqrt{q}} \right| \leqslant \frac{\log^2(1/q)}{\sqrt[3]{q}} \right) \to 1. \tag{3.7}$$

---

[2] See [18, (1)-(3)] for an explicit expression of $\lambda(d, j)$.

[3] We have $\lambda = \pi^2/18$ and $\lambda_2 \approx 7.0545$, see [123, Section A.1.2] for an explicit expression.

- *For Duarte BP we have $q_c^{BP} = \tilde{q}_c^{BP} = 0$ and $q \log \tau_0^{BP}/\log^2(1/q)$ converges in probability to a positive constant, as $q \to 0$.*
- *For North-East BP in $d \geqslant 2$ dimensions we have $q_c^{BP} = \tilde{q}_c^{BP} = 1 - p_c^{OP,d}$, where $p_c^{OP,d} \in (0,1)$ is the critical probability of oriented site percolation in d dimensions (see e.g. [75, 121, 145, 146] for background) and $\mu_{q_c^{BP}}(\tau_0^{BP} = \infty) = 0$.*
- *For Spiral BP we have $q_c^{BP} = \tilde{q}_c^{BP} = 1 - p_c^{OP,2}$ and $\mu_{q_c^{BP}}(\tau_0^{BP} = \infty) > 0$.*
- *For j-neighbour BP in dimension d with $j > d \geqslant 1$ we have $q_c^{BP} = \tilde{q}_c^{BP} = 1$.*

Proof For East BP, we only sketch the argument and leave the details as an exercise to the reader. First, we verify that for any $t \in \mathbb{N}$, we have $\mu_q(\tau_0^{BP} > t) = (1-q)^{N_d(t)}$, where $N_d(t) = \{(x_1, \ldots, x_d) \in \mathbb{N}^d : \sum_{i=1}^d x_i \leqslant t\}$. One can check that $N_d(t) = \binom{t+d}{d} = t^d/d! + O_d(t^{d-1})$, where the implicit constant may depend on the dimension d. In particular, for any $q > 0$, $\mu_q(\tau_0^{BP} = \infty) = 0$ and $\mu_q(\tau_0^{BP} > t)$ decays at least exponentially, so $q_c^{BP} = \tilde{q}_c^{BP} = 0$. Moreover, the convergence in distribution follows from the asymptotics of $N_d(t)$.

The proof for 1-neighbour BP is analogous, replacing $N_d(t)$ by the volume of the discrete ℓ^1 ball of radius t, whose cardinal is asymptotically equivalent to $2^d N_d(t)$.

For j-neighbour BP with general $2 \leqslant j \leqslant d$, the asymptotics of $\log^{\circ(j-1)} \tau_0^{BP}$ is due to Balogh, Bollobás, Duminil-Copin and Morris [17] (the case $j = 2$ was established by Holroyd [126, 127]), while the identification of the critical value is due to Schonmann [182, Theorem 3.1].

The result for 2-neighbour BP is due to Hartarsky and Teixeira [123].

The quantitative result for Duarte BP is due to Bollobás, Duminil-Copin, Morris and Smith [38], while the qualitative one is due to Schonmann [181, 182].

The result for North-East BP follows from the fact that $\tau_0^{BP} = \infty$ if and only if there is an infinite oriented path of occupied sites from the origin (see [181]), together with classical results on oriented site percolation [1, 23, 163].

The result for Spiral BP is due to Toninelli and Biroli [196] (also see [112, 133, 195, 198, 199]).

For j-neighbour BP with $j > d \geqslant 1$ it suffices to observe that for $q < 1$, we have

$$\mu_q\left(\tau_0^{BP} = \infty\right) \geqslant \mu_q\left(A_0 \cap \{0,1\}^d = \varnothing\right) = (1-q)^{2^d} > 0.$$

This concludes the proof for all models.   □

The reader may have noticed the following pattern [112, Conjecture 8.1] (also see [182]) in Theorem 3.1.

Conjecture 3.2 (Sharp phase transition) For any update family $\mathcal{U}$ in any dimension it holds that $q_c^{BP} = \tilde{q}_c^{BP}$.

This is an important open problem in BP theory, which has so far been resolved for update families contained in an open half-space with the origin on its boundary [114, Theorem 1.6], as well as those with $q_c^{BP} = 0$. For the latter assertion, note that [15] proves a stretched exponential bound on the tail of τ_0^{BP}, but a standard renormalisation argument [182] can be used to recover an exponential decay.

Concerning the continuity of the phase transition, there is not even a guess what the answer should be, leaving the following problem wide open.

Problem 3.3 Determine which update families satisfy $\mu_{q_c^{\mathrm{BP}}}(\tau_0^{\mathrm{BP}} = \infty) = 0$ like North-East BP and unlike Spiral BP.

3.2 Legal paths

We next introduce the notion of legal path that will be instrumental in several proofs. In words, a legal path is a sequence of configurations differing by a legal update (recall Section 2.3).

Definition 3.4 (Legal paths) Given a domain $\Lambda \subset \mathbb{Z}^d$, boundary condition $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ and two configurations $\omega, \omega' \in \Omega_\Lambda$, a *legal path from ω to ω' in Λ* is a finite sequence $(\omega^{(i)})_{i=0}^n$ of configurations in Ω_Λ such that $\omega^{(0)} = \omega$, $\omega^{(n)} = \omega'$ and for each $i \in \{1, \ldots, n\}$ it holds that either $\omega^{(i)} = \omega^{(i-1)}$ or there exists $\mathbf{x}^{(i)} \in \Lambda$ such that $\omega^{(i)} = (\omega^{(i-1)})^{\mathbf{x}^{(i)}}$ (recall (2.6)) and $c_{\mathbf{x}^{(i)}}^{\sigma}(\omega^{(i)}) = 1$ (recall (2.8)). The *length* of the legal path $(\omega^{(i)})_{i=0}^n$ is n. Notice that if $(\omega^{(i)})_{i=0}^n$ is a legal path from ω to ω' in Λ, then its *inverse* $(\omega^{(n-i)})_{i=0}^n$ is a legal path from ω' to ω in Λ.

It turns out that BP provides a simple way to know when a legal path exists and to construct it. We start by observing that the BP closure is invariant along legal paths.

Lemma 3.5 (Invariance of closure) *Let $\omega \in \Omega$ and $\mathbf{x} \in \mathbb{Z}^d$ be such that $c_{\mathbf{x}}(\omega) = 1$. Then $[\omega] = [\omega^{\mathbf{x}}]$ and for any $\Lambda \subset \mathbb{Z}^d$ such that $\mathbf{x} \in \Lambda$ we have $[\omega_\Lambda]^{\omega_{\mathbb{Z}^d \setminus \Lambda}} = [\omega_\Lambda^{\mathbf{x}}]^{\omega_{\mathbb{Z}^d \setminus \Lambda}}$. Consequently, for $\omega', \omega'' \in \Omega_\Lambda$ connected by a legal path in Λ with boundary condition $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ we have $[\omega' \cdot \sigma] = [\omega'' \cdot \sigma]$ and $[\omega']^{\sigma} = [\omega'']^{\sigma}$.*

Proof Assume without loss of generality that $\omega_{\mathbf{x}} = 0$. By (2.3) and (3.1), we have that $(\mathscr{B}_{\mathcal{U}}(\omega^{\mathbf{x}}))_{\mathbf{x}} = 0 = \omega_{\mathbf{x}}$, so $\mathscr{B}_{\mathcal{U}}(\omega^{\mathbf{x}}) \leqslant \omega \leqslant \omega^{\mathbf{x}}$. By (3.2), since $[\cdot]$ is non-increasing, this yields $[\omega] = [\omega^{\mathbf{x}}]$. The proof of $[\omega_\Lambda]^{\omega_{\mathbb{Z}^d \setminus \Lambda}} = [\omega_\Lambda^{\mathbf{x}}]^{\omega_{\mathbb{Z}^d \setminus \Lambda}}$ is analogous. The statement on legal paths follows by induction on the length.   □

Lemma 3.6 *Let $\Lambda \subset \mathbb{Z}^d$, $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ and $\omega \in \Omega_\Lambda$. Let $\mathbf{x} \in \Lambda$ be such that $\omega_{\mathbf{x}} = 1$. There exists a legal path from ω to $\omega^{\mathbf{x}}$ if and only if $[\omega]_{\mathbf{x}}^{\sigma} = 0$. Moreover, if it exists, the legal path can be chosen with length at most $2|\Lambda|$.*

Proof The only if direction was proved in Lemma 3.5. Assume that $[\omega]_{\mathbf{x}}^{\sigma} = 0$. Then there exists a finite set $\Lambda' \subset \Lambda$ such that $[\omega_{\Lambda'}]_{\mathbf{x}}^{\sigma \cdot \omega_{\Lambda \setminus \Lambda'}} = 0$. Up to replacing Λ by Λ', we may assume that Λ is finite. We construct a legal path in Λ of length at most $|\Lambda|$ from ω to $[\omega]^{\sigma}$. To achieve this, we empty an arbitrary occupied site whose constraint with boundary condition σ is satisfied. Such a vertex always exist until we reach $[\omega]^{\sigma}$. We similarly obtain a legal path from $\omega^{\mathbf{x}}$ to $[\omega^{\mathbf{x}}]^{\sigma} = [\omega]^{\sigma}$ and conclude by concatenating its inverse with the legal path from ω.   □

Deducing or proving the following corollary is left as an exercise to the reader.

Corollary 3.7 (Legal paths and closure) *Let $\Lambda \subset \mathbb{Z}^d$, $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ and $\omega, \omega' \in \Omega_\Lambda$ be such that $\sum_{\mathbf{x} \in \Lambda} |\omega_\mathbf{x} - \omega'_\mathbf{x}| < \infty$. Then there exists a legal path from ω to ω' in Λ with boundary condition σ, if and only if $[\omega]^\sigma = [\omega']^\sigma$ (recall (3.4)).*

Definition 3.8 (Ergodic boundary condition) Let $\Lambda \Subset \mathbb{Z}^d$. We say that a boundary condition $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ is *ergodic*, if $[\mathbf{1}_\Lambda]^\sigma = \mathbf{0}_\Lambda$. By Corollary 3.7, this is equivalent to $\mathcal{L}^\sigma$ defining an ergodic process on Ω_Λ.

3.3 Ergodicity

The simple deterministic statements of Section 3.2 entail the following result of Cancrini, Martinelli, Roberto and Toninelli [44, Proposition 2.4], whose fundamental importance for KCM is apparent in view of Theorem 3.1.

Theorem 3.9 (Ergodicity) *For any update family $\mathcal{U}$ and $q \in (0, 1)$, the following are equivalent*

   (i) $\mathbb{P}_{\mu_q}(\tau_\vee < \infty) = 1$,
   (ii) $\mathbb{P}_{\mu_q}(\tau_0 < \infty) = 1$,
   (iii) $\mu_q(\tau_0^{\mathrm{BP}} < \infty) = 1$,
   (iv) 0 is a simple eigenvalue of $\mathcal{L}$ on $\mathrm{L}^2(\mu_q)$, that is, the dynamics is ergodic,
   (v) for all $f \in \mathrm{L}^2(\mu_q)$ we have $\lim_{t \to \infty} P_t f = \mu_q(f)$, that is, the dynamics is mixing.

In particular, $q_c = q_c^{\mathrm{BP}}$ (recall (2.13) and (3.5)).

Proof **(i) implies (ii).** This follows directly from the definition (2.9).
(ii) implies (iii). Consider the function $f : \Omega \to \{0, 1\} : \omega \mapsto [\omega]_0$ in $\mathrm{L}^2(\mu_q)$. By Lemma 3.5, we have that $\mathcal{L}_n f = 0$, where $\mathcal{L}_n$ is the KCM generator on Ω defined by restricting the sum in (2.4) to $\Lambda_n = \{-n, \dots, n\}^d$. Moreover, clearly $\mathcal{L}_n g \to \mathcal{L}g$ for any local function g, so $\mathcal{L}_n f \to \mathcal{L}f$ (see [146, Corollary I.3.14][4]). Thus, $\mathcal{L}f = 0$, so that for any $t \geqslant 0$, $\mathbb{P}_{\mu_q}$-a.s. $f(\omega(t)) = f(\omega(0))$. Consequently,

$$\mathbb{P}_{\mu_q}(\forall t \in \mathbb{Q} \cap [0, \infty), f(\omega(t)) = f(\omega(0))) = 1. \tag{3.8}$$

Assume that $\mathbb{P}_{\mu_q}(\tau_0 < \infty) = 1$. Then $\omega_0(t) = 0$ for all $t \in [\tau_0, \tau_0 + \varepsilon]$ for some random $\varepsilon > 0$. Thus, (3.8) gives $\mathbb{P}_{\mu_q}(f(\omega(0)) = f(\omega(\tau_0)) = 0) = 1$, which concludes the proof, since $\mu_q(\tau_0^{\mathrm{BP}} < \infty) = 1 - \mu_q(f) = 1$.
(iii) implies (iv). Fix $f \in \mathrm{L}^2(\mu_q)$ such that $\mathcal{L}f = 0$. Then, recalling (2.7), we have $\sum_{\mathbf{x} \in \mathbb{Z}^d} \mu_q(c_\mathbf{x} \mathrm{Var}_\mathbf{x}(f)) = \mathcal{D}(f) = 0$, so each of the (non-negative) summands is 0. We seek to prove that $\mathrm{Var}(f) = 0$ and our starting point is the unconstrained Poincaré

---

[4] As explained in [146, Section IV.4], the corollary and other results apply in $\mathrm{L}^2(\mu_q)$ instead of the space of continuous functions for the product topology.

inequality[5]

$$\mathrm{Var}(f) \leqslant \sum_{\mathbf{x} \in \mathbb{Z}^d} \mu(\mathrm{Var}_{\mathbf{x}}(f)). \tag{3.9}$$

Assume further that (iii) holds and fix some $\mathbf{x} \in \mathbb{Z}^d$. Then

$$\mu\left(\mathrm{Var}_{\mathbf{x}}(f)\right) \leqslant \mu\left(\omega_0 \left(f\left(\omega^{\mathbf{x}}\right) - f(\omega)\right)^2\right) = \sum_{n=1}^{\infty} \mu\left(\mathbb{1}_{\mathcal{E}_n \setminus \mathcal{E}_{n-1}}(\omega)\left(f\left(\omega^{\mathbf{x}}\right) - f(\omega)\right)^2\right)$$

where $\mathcal{E}_n = \{\omega \in \Omega : [\omega_{\Lambda_n}]_{\mathbf{x}}^{\mathbf{1}_{\mathbb{Z}^d \setminus \Lambda_n}} = 0\}$ and $\Lambda_n = \mathbf{x} + \{-n, \dots, n\}^d$. But by Lemma 3.6 for any $\omega \in \mathcal{E}_n \setminus \mathcal{E}_0$ there exists a legal path $(\omega^{(i)})_{i=0}^{2|\Lambda_n|}$ from ω to $\omega^{\mathbf{x}}$ in Λ_n. Writing $f(\omega^{\mathbf{x}}) - f(\omega)$ telescopically along this legal path and using the Cauchy–Schwarz inequality yields

$$\mathbb{1}_{\mathcal{E}_n \setminus \mathcal{E}_{n-1}}(\omega)\left(f\left(\omega^{\mathbf{x}}\right) - f(\omega)\right)^2 \leqslant 2|\Lambda_n| \sum_{i=1}^{2|\Lambda_n|} c_{\mathbf{x}^{(i)}}\left(\omega^{(i)}\right)\left(f\left(\omega^{(i)}\right) - f\left(\omega^{(i-1)}\right)\right)^2,$$

where $\mathbf{x}^{(i)}$ is the location of the legal update from $\omega^{(i-1)}$ to $\omega^{(i)}$. In order to recover $\mu(c_{\mathbf{x}^{(i)}} \mathrm{Var}_{\mathbf{x}^{(i)}}(f))$ from the μ-average of the last summand, we only need to perform a change of measure and observe that $\mu(\omega)/\mu(\omega^{(i)}) \leqslant (q(1-q))^{-|\Lambda_n|}$. Putting everything together, we obtain

$$\mu(\mathrm{Var}_{\mathbf{x}}(f)) \leqslant \sum_{n=1}^{\infty} 4|\Lambda_n|^2 (q(1-q))^{-|\Lambda_n|-1} \sum_{\mathbf{y} \in \Lambda_n} \mu\left(c_{\mathbf{y}} \mathrm{Var}_{\mathbf{y}}(f)\right) = 0.$$

Recalling (3.9), this gives that f is μ-a.s. constant as desired.
(iv) implies (i). This follows from the ergodic theorem (see e.g. [25]) applied to the function $\omega \mapsto \omega_0$.
(iv) is equivalent to (v). This is [146, Theorem IV.4.13] (based on operator spectral theory). $\square$

Before moving on, let us comment on the proof, which showcases two important ideas. Firstly, in order to empty the origin, we need to be able to do so in BP. Secondly, if BP is able to empty some site, we can turn that into a Poincaré inequality for the corresponding KCM. Building on these two insights, one can go surprisingly far. The first one is useful for obtaining lower bounds on τ_0, while the second one enables upper bounds. Finally, let us mention that the proof of the implication from (iii) to (iv) is our first encounter with the canonical path technique for Markov chains originating in [141, 187] (also see e.g. [144, Section 13.5]).

---

[5] In other words, the spectral gap of the generator of the product KCM corresponding to $\mathcal{U} = \{\varnothing\}$ is 1, since it is the tensor product of irreducible 2-state Markov processes with total rate 1. This classical fact also follows e.g. by taking $\mathcal{X} = \mathbb{X}_1$ in Lemma 4.7.

3.4 Exponential decay

In Section 3.3 we saw that ergodicity and mixing of KCM are equivalent to BP a.s. emptying $\mathbb{Z}^d$. All these results are purely qualitative and provide no quantitative control whatsoever, e.g. on the tails of emptying times. Our next task is to transfer the tail behaviour of BP to KCM, by proving the following result adapted from [44, 45, 112, 161]. Once again, its interest is made clear by Theorem 3.1 (also recall Conjecture 3.2 and Theorem 3.9).

Theorem 3.10 (Exponential decay) *For any update family $\mathcal{U}$ and $q \in (0, 1)$, the following are equivalent.*

(i) $\liminf_{t \to \infty} - \log \mathbb{P}_{\mu_q}(\tau_\vee > t)/t > 0$,
(ii) $\liminf_{t \to \infty} - \log \mathbb{P}_{\mu_q}(\tau_0 > t)/t > 0$,
(iii) $\liminf_{t \to \infty} - \log \mu_q(\tau_0^{\mathrm{BP}} > t)/t > 0$,
(iv) $T_{\mathrm{rel}} < \infty$ *(recall (2.10)).*

In particular, $\tilde{q}_{\mathrm{c}} = \tilde{q}_{\mathrm{c}}^{\mathrm{BP}}$ (recall (2.14) and (3.6)).

Proof (Sketch) **(i) implies (ii).** This follows directly from the definition (2.9).
(ii) implies (iii). We claim that there exists a constant $\delta > 0$ (depending on $\mathcal{U}$ and q) such that if we run BP and KCM from the same initial configuration $\omega \in \Omega$,

$$\mathbb{P}_\omega \left(\tau_0 \leqslant \delta \tau_0^{\mathrm{BP}} \right) \leqslant e^{-\tau_0^{\mathrm{BP}}(\omega)}. \tag{3.10}$$

The idea is that, in order to empty the origin, legal updates have to occur in the right order along some path of length $\tau_0^{\mathrm{BP}}(\omega)$, starting at the origin. Consecutive vertices of the path are allowed to be at distance at most $\max_{U \in \mathcal{U}, \mathbf{u} \in U} \|\mathbf{u}\|$. However, the number of such paths is at most $e^{C\tau_0^{\mathrm{BP}}(\omega)}$ for some $C > 0$. Moreover, the probability that the sum of N exponential random variables of mean 1 is smaller than δN is at most e^{-2CN}, choosing δ sufficiently small depending on C. We conclude the proof of the claim by a union bound on the possible paths. See [161, Lemma 4.3] for more details.
(iii) implies (iv). For simplicity of the presentation, we focus on the two-dimensional case and assume that $\{(2, 1), (1, 2)\} \in \mathcal{U}$. The proof proceeds in three steps. First, we prove that for $q = 1 - \varepsilon_0$ sufficiently close to 1, $T_{\mathrm{rel}} < \infty$ for the update family $\mathcal{U}' = \{U'\} = \{\{(1, 0), (1, 1), (0, 1)\}\}$. Second, we perform a renormalisation,[6] by tessellating space into large square boxes, which are deemed good if the $\mathcal{U}$-KCM restricted to the box is able to empty most of the bottom and left boundaries of the box in the current configuration. Third, we show how to completely empty a given (possibly non-good) box, assuming that the three neighbouring boxes corresponding to $\mathcal{U}'$ are good. We postpone the discussion of the first step to Section 5.1, where the bisection technique is presented (for the full details, refer to [44, Section 4]).

We turn to the second step. Fix $\varepsilon > 0$ small enough depending on ε_0 and then take $n \in \mathbb{N}$ large enough. Let $\Lambda = \{0, \dots, n - 1\}^2$ be the renormalisation box. Then

---

[6] In physics, one would rather speak of coarse-graining, but we adopt the mathematical jargon.

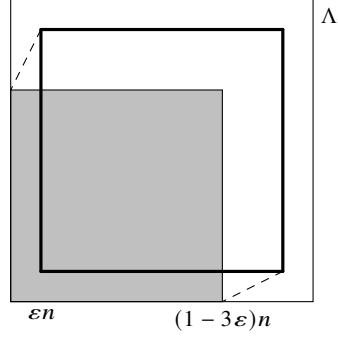εn $(1 - 3\varepsilon)n$

Fig. 3.1 A good box in the renormalisation of the proof of Theorem 3.10. The thick box is emptied in (3.11). This allows emptying the shaded one in (3.12). The dashed lines with slopes $1/2$ and 2 indicate how empty sites propagate via the rule $\{(2,1),(1,2)\}$.
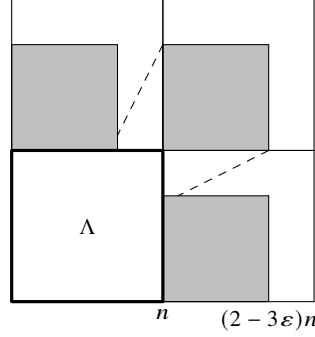

n $(2 - 3\varepsilon)n$

Fig. 3.2 If the three shaded boxes are empty, we are able to empty the thick box Λ, as indicated by the dashed lines with slopes $1/2$ and 2, thanks to the update rule $\{(2,1),(1,2)\}$.

the exponential decay provided by (iii) and a union bound give

$$\mu_q \left([\omega_\Lambda]_{[\varepsilon n,(1-\varepsilon)n)^2}^{\mathbf{1}_{\mathbb{Z}^2 \setminus \Lambda}} = \mathbf{0}_{[\varepsilon n,(1-\varepsilon)n)^2} \right) \geqslant 1 - \varepsilon_0. \qquad (3.11)$$

In words, it is likely that BP in Λ empties all of Λ except a thin frame (see Figure 3.1), in which case we say that Λ is *good*. Thus, (3.11) states that the probability that a box is good is at least $1 - \varepsilon_0$. However, since $\{(2,1),(1,2)\} \in \mathcal{U}$, we have

$$\left[\mathbf{0}_{[\varepsilon n,(1-\varepsilon)n)^2} \cdot \mathbf{1}_{\Lambda \setminus [\varepsilon n,(1-\varepsilon)n)^2} \right]_{[0,\dots,(1-3\varepsilon)n)^2}^{\mathbf{1}_{\mathbb{Z}^2 \setminus \Lambda}} = \mathbf{0}_{[0,\dots,(1-3\varepsilon)n)^2}, \qquad (3.12)$$

that is, BP empties all but a thin strip along the top and right boundaries of Λ, as desired (see Figure 3.1). See [112, Section 7.4] for more details on the second step.

We move on to the third step. Again using $\{(2,1),(1,2)\} \in \mathcal{U}$, we have

$$\left[\mathbf{0}_{[0,(1-3\varepsilon)n)+nU'} \cdot \mathbf{1}_{[0,2n)^2 \setminus ([0,\dots,(1-3\varepsilon)n)+nU')} \right]_{\Lambda}^{[0,2n)^2} = \mathbf{0}_{\Lambda} \qquad (3.13)$$

(see Figure 3.2). Consequently, if the three boxes of the form $\Lambda + n\mathbf{u}$ with $\mathbf{u} \in U'$ are good, then Λ can be emptied by BP in the union of these four boxes. Then Lemma 3.6 provides a legal path in $[0,2n)^2$ from $\omega_{[0,2n)^2}$ to $\omega_{[0,2n)^2}^{\mathbf{x}}$ for any $\mathbf{x} \in \Lambda$. Using this path and the canonical path approach as in the proof of Theorem 3.9, this yields that for some $\gamma < \infty$ depending on n and q and any local function f,

$$\mu \left(\mathbb{1}_{\forall \mathbf{u} \in U', \Lambda+n\mathbf{u} \text{ is good}} \operatorname{Var}_\Lambda(f) \right) \leqslant \sum_{\mathbf{x} \in \Lambda} \mu \left(\mathbb{1}_{\forall \mathbf{u} \in U', \Lambda+n\mathbf{u} \text{ is good}} \operatorname{Var}_{\mathbf{x}}(f) \right)$$

$$\leqslant \gamma \sum_{\mathbf{x} \in [0,2n)^2} \mu(c_{\mathbf{x}} \operatorname{Var}_{\mathbf{x}}(f)).$$

Finally, it remains to see that for some $\gamma' < \infty$ and any local function f,

$$\text{Var}(f) \leqslant \gamma' \sum_{\mathbf{x} \in \mathbb{Z}^2} \mu \left(\mathbb{1}_{\forall \mathbf{u} \in U', n\mathbf{x}+\Lambda+n\mathbf{u} \text{ is good}} \, \text{Var}_{n\mathbf{x}+\Lambda}(f) \right).$$

Recalling (2.3) and (2.11), we recognise exactly the Poincaré inequality for the $\mathcal{U}'$-KCM except that each site (box) now has more than two states and the parameter q is replaced by $\mu_q(\Lambda \text{ is good}) \geqslant 1 - \varepsilon_0$. We will frequently encounter such general state-space KCM arising from renormalisation procedures (see Section 4.4) and will see that usually the corresponding Poincaré inequalities are proved just like the ones for ordinary KCM. Since $T_{\text{rel}} < \infty$ for the $\mathcal{U}'$-KCM at $q \geqslant 1 - \varepsilon_0$, by the first step, this completes the proof. See [44, Section 5] for more details.

(iv) implies (i). Define

$$\lambda_0 = \inf \left\{ \mathcal{D}(f) : f \in L^2(\mu), \mu\left(f^2\right) = 1, \forall \omega \in \Omega, \omega_0 = 0 \Rightarrow f(\omega) = 0 \right\}.$$

Observe that any f as above satisfies $\text{Var}(f) \geqslant q$, since $\mu^2(f) = (1-q)^2 \mu^2(f|\omega_0 = 1) \leqslant (1-q)^2 \mu(f^2|\omega_0 = 1) = (1-q)$. Recalling (2.10), this implies $\lambda_0 \geqslant q/T_{\text{rel}}$. Finally, a general result on hitting times for Markov chains [5, Proposition 3.21] yields

$$\mathbb{P}(\tau_0 > t) \leqslant e^{-t\lambda_0} \leqslant e^{-qt/T_{\text{rel}}} \tag{3.14}$$

for any $t \geqslant 0$. This proves (ii) and one can recover (i), by proceeding analogously for τ_1 (recall (2.9)). See [44, Theorem 3.6] for a different proof. $\qquad \square$

The next result [44, Lemma 2.11, Proposition 2.13] reduces the relaxation time of a KCM in infinite volume to the one in finite volume (recall Section 2.4).

Proposition 3.11 (Finite volume reduction) *Let $\mathcal{U}$ be an update family and $q \in (0, 1)$. Given $\Lambda \Subset \mathbb{Z}^d$, recall that T_{rel}^Λ denotes the relaxation time of the finite-volume generator $\mathcal{L}_{\mathbf{0}_{\mathbb{Z}^d \setminus \Lambda}}$. Then*

$$T_{\text{rel}} = \lim_{\Lambda \Subset \mathbb{Z}^d, \Lambda \to \mathbb{Z}^d} T_{\text{rel}}^\Lambda.$$

The same holds for C_{LS} and C_{MLS}.

Proof We prove the chain of inequalities

$$T_{\text{rel}} \leqslant \sup_{\Lambda \Subset \mathbb{Z}^d} T_{\text{rel}}^\Lambda \leqslant \lim_{\Lambda \Subset \mathbb{Z}^d, \Lambda \to \mathbb{Z}^d} T_{\text{rel}}^\Lambda \leqslant T_{\text{rel}}. \tag{3.15}$$

For the first one, it suffices to prove the Poincaré inequality (2.11) with $C = \sup_{\Lambda \Subset \mathbb{Z}^d} T_{\text{rel}}^\Lambda$ for any local function f. Since f is local, for Λ large enough we have $\text{Var}(f) = \text{Var}_\Lambda(f)$ and $\mathcal{D}(f) = \mathcal{D}_{\mathbf{0}_{\mathbb{Z}^d \setminus \Lambda}}(f)$. So we are done by (2.10).

For the second and third inequalities and the existence of the limit in (3.15), it suffices to show that for $V \subset \Lambda \subset \mathbb{Z}^d$ we have

$$T_{\text{rel}}^V \leqslant T_{\text{rel}}^\Lambda. \tag{3.16}$$

Consider $f : \Omega_V \to \mathbb{R}$ and extend f to $\tilde{f} : \Omega_\Lambda \to \mathbb{R}$ by $\tilde{f}(\omega) = f(\omega_V)$. Clearly, $\mathrm{Var}(f) = \mathrm{Var}(\tilde{f})$, since $\mu_\Lambda = \mu_V \otimes \mu_{\Lambda \setminus V}$, so it remains to check that

$$\mathcal{D}_{\mathbf{0}_{\mathbb{Z}^d \setminus \Lambda}}(\tilde{f}) \leqslant \mathcal{D}_{\mathbf{0}_{\mathbb{Z}^d \setminus V}}(f).$$

But, recalling (2.7), this follows from $c_{\mathbf{x}}^{\mathbf{0}_{\mathbb{Z}^d \setminus \Lambda}} \leqslant c_{\mathbf{x}}^{\mathbf{0}_{\mathbb{Z}^d \setminus V}}$. The latter is true, using (2.8) and the fact that $c_{\mathbf{x}}$ is non-increasing.

The proof for the (modified) logarithmic Sobolev constant is identical.   □

3.5 Stronger functional inequalities

We next consider the time scales corresponding to mixing, modified logarithmic Sobolev and logarithmic Sobolev inequalities (recall Section 2.5). We start with general bounds on the mixing time.

Proposition 3.12 (Basic mixing time bounds) *For any update family $\mathcal{U}$ and $q > 0$ there exists $C > 0$ such that the following holds. If $q > \tilde{q}_c$, then for all $\varepsilon \in (0, 1)$ and $\Lambda \subset \mathbb{Z}^d$,*

$$t_{\mathrm{mix}}^{\mathbf{0}_{\mathbb{Z}^d \setminus \Lambda}}(\varepsilon) \leqslant C(\log(1/\varepsilon) + |\Lambda|). \tag{3.17}$$

For all $\varepsilon \in (0, 1)$ and all n large enough, for any $\sigma \in \Omega_{\mathbb{Z}^d \setminus \{1, \dots, n\}^d}$,

$$t_{\mathrm{mix}}^{\sigma}(\varepsilon) \geqslant n/C. \tag{3.18}$$

Proof For (3.17), we use a general inequality for reversible Markov chains (see e.g. [144, Theorem 20.6])

$$t_{\mathrm{mix}}^{\mathbf{0}_{\mathbb{Z}^d \setminus \Lambda}}(\varepsilon) \leqslant \log\left(\frac{1}{\varepsilon \mu_*}\right) T_{\mathrm{rel}}^{\Lambda},$$

with $\mu_* = \min_{\omega \in \Omega_\Lambda} \mu(\omega) = \min(q, 1-q)^{|\Lambda|}$. By (2.14), for $q > \tilde{q}_c$, we have $T_{\mathrm{rel}} < \infty$. Moreover, by (3.16), $T_{\mathrm{rel}} \geqslant T_{\mathrm{rel}}^{\mathbf{0}_{\mathbb{Z}^d \setminus \Lambda}}$, concluding the proof of (3.17).

For (3.18), we choose $\omega = \mathbf{1}_\Lambda$ in (2.19) with $\Lambda = \{1, \dots, n\}^d$. We denote the $\mathcal{U}$-KCM with parameter q, boundary condition $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ and initial condition ω by $(\eta(t))_{t \geqslant 0}$. Fix $C > 0$ large enough and set $\Lambda' = \{\lceil n/\sqrt{C} \rceil, \dots, n - \lceil n/\sqrt{C} \rceil\}^d$. It suffices to show that for any $q \in [0, 1]$

$$\lim_{n \to \infty} \mathbb{P}\left(\eta_{\Lambda'}(n/C) = \mathbf{1}_{\Lambda'}\right) = 1. \tag{3.19}$$

Let us denote $\|\mathcal{U}\| = \max\{\|u\| : \exists U \in \mathcal{U}, u \in U\}$. A *path from $\mathbf{x} \in \Lambda'$ to $\partial\Lambda$* is a finite sequence of sites $(\mathbf{x}_i)_{i=1}^N$ in Λ with $\mathbf{x}_1 = \mathbf{x}$, $\mathbf{x}_N$ at distance at most $\|\mathcal{U}\|$ from $\mathbb{Z}^d \setminus \Lambda$, and, for each $i \in \{2, \dots, N\}$, $\|\mathbf{x}_i - \mathbf{x}_{i-1}\| \leqslant \|\mathcal{U}\|$. Assume that $\eta_t(\mathbf{x}) = 0$ for some $t > 0$. Then we can construct a path from $\mathbf{x}$ to $\partial\Lambda$ inductively as follows, assuming that $\mathbf{x}_i$ is chosen (by definition, $\mathbf{x}_1 = \mathbf{x}$). If $d(\mathbf{x}_i, \mathbb{Z}^d \setminus \Lambda) \leqslant \|\mathcal{U}\|$, we set $N = i$ and stop. Otherwise, consider the first time $t' \leqslant t$ such that $\eta_{t'}(\mathbf{x}_i) = 0$ and

let $\mathbf{x}_{i+1}$ be arbitrarily chosen in $\mathbf{x} + \bigcup_{U \in \mathcal{U}} U$ so that $\eta_{\mathbf{x}_{i+1}}(t'-) = 0$. We call this a *decreasing path of* $\mathbf{x}$, since the emptying times of its sites are decreasing. For a given path from $\mathbf{x}$ to $\partial\Lambda$ of length k, the probability that it is decreasing is at most the probability that the sum of k i.i.d. standard exponential random variables E_i is at most t. Indeed, this follows from the graphical construction of Section 2.3. Moreover, $k \geqslant (n/\sqrt{C})/\|\mathcal{U}\|$, so, setting $t = n/C$, for n large enough we obtain

$$
\begin{aligned}
\mathbb{P}(\eta_t(\mathbf{x}) = 0) &\leqslant \sum_{k \geqslant \frac{n}{\|\mathcal{U}\|\sqrt{C}}} |\{(\mathbf{x}_i)_i \text{ path from } \mathbf{x} \text{ to } \partial\Lambda \text{ of length } k\}| \cdot \mathbb{P}\left(\sum_{l=1}^{k} E_l \leqslant t\right) \\
&\leqslant \sum_{k \geqslant \frac{n}{\|\mathcal{U}\|\sqrt{C}}} (2\|\mathcal{U}\| + 1)^{kd} \cdot \mathbb{P}\left(\sum_{l=1}^{k} E_l \leqslant \frac{k}{C^{1/3}}\right) \\
&\leqslant \sum_{k \geqslant \frac{n}{\|\mathcal{U}\|\sqrt{C}}} (2\|\mathcal{U}\| + 1)^{kd} \cdot \frac{(\mathbb{E}(e^{-C^{1/3}E_1}))^k}{e^{-k}} \leqslant e^{-n/(\sqrt{C}\|\mathcal{U}\|)},
\end{aligned}
$$

where we used the exponential Markov inequality in the last line and took into account that C can be chosen large enough depending on $\|\mathcal{U}\|$ and d. The desired (3.19) then follows by a union bound over $\mathbf{x} \in \Lambda'$. The argument used in the proof (3.18) is commonly referred to as *finite speed of propagation* and applies equally well to any finite-range interacting particle system [145, Section I.1]. In fact, this is a way to show that the graphical construction of the KCM produces a well-defined Markov process. $\qquad\square$

Corollary 3.13 (Infinite logarithmic Sobolev constants) *For any update family* $\mathcal{U}$ *and* $q \in (0, 1)$, *we have*

$$
C_{\mathrm{LS}} = C_{\mathrm{MLS}} = \infty. \tag{3.20}
$$

Furthermore, there exists $C = C(q) > 0$ *such that the following holds*

(i) Fix n *and let* $\Lambda = \{-n, \ldots, n\}^d$. *For any* $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ *and for* n *large enough, we have*

$$
C_{\mathrm{LS}}^\sigma \geqslant C_{\mathrm{MLS}}^\sigma \geqslant n/C. \tag{3.21}
$$

(ii) If $q > \tilde{q}_c$, *then, for any* $\Lambda \Subset \mathbb{Z}^d$, *we have*

$$
C_{\mathrm{LS}}^\Lambda \leqslant C|\Lambda|. \tag{3.22}
$$

Proof The first inequality in (3.21) is a general and classical result (see [70, Lemma 2.7]). To prove the second inequality we set

$$
f(\omega) = (1 - q)^{-|\Lambda|} \mathbb{1}_{\omega_\Lambda = \mathbf{1}_\Lambda},
$$

so that $\mu(f) = 1$. Let μ^f be the probability measure with density f w.r.t. μ. Then, by (3.19), for a properly chosen $C > 0$, we have

$$
\lim_{n \to \infty} \mathbb{E}_{\mu^f}(\eta_0(Cn)) = 1. \tag{3.23}
$$

Combining (2.17), (3.23) and Pinsker's inequality implies the second inequality of (3.21).

Inequality (3.22) follows from Theorem 3.10(iv) together with (3.15) and the general and classical bound (see [179] Corollary 2.2.10)

$$C_{\mathrm{LS}}^{\sigma} \leqslant \frac{\log((1/\mu^*) - 1)}{1 - 2\mu^*} T_{\mathrm{rel}}^{\sigma} \tag{3.24}$$

where $\mu^* := \min_{\omega \in \Omega_\Lambda} \mu(\omega)$. Finally, (3.20) follows from (3.21) and Proposition 3.11. □

3.6 Conclusion

In Theorems 3.9 and 3.10 we completely reduced the critical values q_c and $\tilde{q}_c$ to their BP counterparts. As we saw in Theorem 3.1 and will see more generally in Chapter 6, modulo Conjecture 3.2, this gives complete information about q_c and $\tilde{q}_c$. Our next goal is to find the asymptotics of τ_0 as $q \to q_c+$, as in the BP Theorem 3.1. In reality, even in the simplest BP models with $q_c \in (0, 1)$, such as North-East BP, the asymptotics of τ_0^{BP} as $q \to q_c+$ and of $\mu_{q_c}(\tau_0^{\mathrm{BP}} > t)$ as $t \to \infty$ remain inaccessible, even though it is classically conjectured that they should be governed by certain critical exponents. Therefore, in the remainder of the manuscript we mostly focus on update families for which $q_c = 0$.

Before we direct our efforts to seeking asymptotics as $q \to 0$, let us summarise the techniques we learned in this chapter, as several of them will be of further use.
BP lower bound. In (3.10) we have seen that BP provides a natural lower bound for KCM time scales. Let us also record the general quantitative bound obtained by combining (3.10) and (3.14):

$$\delta \mu_q \left(\tau_0^{\mathrm{BP}} \right) \leqslant \mathbb{E}_{\mu_q}(\tau_0) \leqslant \frac{T_{\mathrm{rel}}}{q} \tag{3.25}$$

for some $\delta > 0$ depending only on $\mathcal{U}$ (not necessarily equal to the one in (3.10)).
Test functions. Thanks to the variational definition (2.10) of T_{rel}, one can obtain lower bounds by plugging well-chosen functions. We used this in the proof of the implication (ii) to (iii) of Theorem 3.9. The test function we used here simply reflected BP, but in the next chapter we will need a more subtle choice.
Canonical paths. This method for proving upper bounds on T_{rel} was used in the implication (iii) to (iv) of Theorem 3.9. In this instance, the canonical paths were simply granted by BP, but in the next chapter, paths will incorporate more refined heuristics on the KCM dynamics.
Renormalisation. In the proof of the implication (iii) to (iv) we used the idea of renormalisation. We regarded large boxes as single sites in an auxiliary (generalised) KCM dynamics. This splits the problem of proving upper bounds on T_{rel} in two. First, we prove an upper bound on the relaxation time of the auxiliary dynamics.

Second, we show (e.g. by canonical paths) how to "locally" reconstruct the original Dirichlet form from the auxiliary one, taking advantage of the auxiliary constraint. We will make extensive use of this technique in the next chapters.

Finite speed of propagation. This observation allows us to show that with very high probability no information about the state of the process at a given place or its boundary condition can travel faster than linearly. The proof of the lower bound of Proposition 3.12 was an application of this fact.

Stronger functional inequalities. In Corollary 3.13, we established that, contrary to Poincaré inequalities, logarithmic and modified logarithmic Sobolev inequalities in infinite volume are *not* a suitable tool for studying KCM. Nonetheless, such techniques can be of use on suitably chosen finite volumes.

Chapter 4
One-dimensional models

In this chapter we investigate one-dimensional KCM. Most notably, these include the FA-1f and East models. We present the techniques used to determine the scaling of their characteristic times as $q \to 0$. We familiarise ourselves with the use of test functions, non BP-based canonical paths, combinatorial bottlenecks and bisection in the simplest possible setting. We then move on to FA-2f and general KCM, still in one dimension. These one-dimensional models are not only interesting in their own right, but will also serve as tools for the study of higher-dimensional KCM via renormalisation.

In one dimension there are three nearest neighbour KCM corresponding to the update families $\{\{-1\}, \{1\}\}$, $\{\{1\}\}$ and $\{\{-1, 1\}\}$ (excluding the trivial cases $\{\varnothing\}$ and $\varnothing$, as usual). We recognise the FA-1f, East and FA-2f models respectively. The latter is not very interesting from our viewpoint, but we consider it for completeness. FA-1f and East on the other hand are not only of interest themselves, but also provide fundamental building blocks for renormalisation arguments for more complex models (as in the proof of Theorem 3.10).

4.1 FA-1f

In this section, we consider $\mathcal{U} = \{\{-1\}, \{1\}\}$. From Theorem 3.1, we have $q_c = 0$. We therefore look for asymptotics as $q \to 0$, which are provided by the work of Cancrini, Martinelli, Roberto and Toninelli [44] and Shapira [184] (and its arXiv version).

Theorem 4.1 (FA-1f asymptotics) *For FA-1f in $d = 1$ dimension there exists $C > 0$ such that for q small enough*

$$1/C \leqslant q^3 T_{\text{rel}} \leqslant C, \qquad 1/C \leqslant q^3 \mathbb{E}_{\mu_q}[\tau_0] \leqslant C, \qquad \mathbb{P}_\mu(\tau_0 > t) \leqslant e^{-Cq^3 t}.$$

Remark 4.2 (Arrhenius law) Note that if we rewrite Theorem 4.1 in terms of the inverse temperature, using (2.1), the above scaling corresponds to an Arrhenius divergence, $\tau_0, T_{\text{rel}} \approx \exp(c\beta)$, as for strong glass forming liquids (see Figure 1.1).

Before embarking on the proof, let us explain the heuristics behind Theorem 4.1. For q small empty sites are typically isolated (and therefore cannot be immediately removed). However, if $x \in \mathbb{Z}$ is empty, at rate q we can empty $x - 1$ or $x + 1$. Suppose the first event occurs. At this point the constraint at x is also satisfied and, in a time of order one, we will (with equal probability) either occupy $x - 1$ or x. In the latter case, the net result is that we have "moved" the empty site from x to $x - 1$. So, we intuitively expect empty sites to behave like random walks moving at rate q until they meet. When they meet, they typically coalesce. This explains the scaling $1/q^3$: it is the time required to overcome the typical distance $\ell = 1/q$ between two consecutive empty sites (inverse rate times distance squared, that is $1/q \times 1/q^2$).

Proof (Theorem 4.1) In order to show that $T_{\text{rel}} \geqslant q^{-3}/C$, we recall (2.10). Define the test function $f(\omega) = \min\{k \geqslant 1 : \omega_k \omega_{-k+1} = 0\}$, that is, the distance from $1/2$ to the nearest empty site rounded up. As in the proof of Theorem 3.1, one can check that $\mu(f(\omega) = k + 1) = (2q - q^2)(1 - q)^{2k}$ for $k \geqslant 0$. This geometric random variable has $\text{Var}(f) = (1 - q)^2/(2q - q^2)^2$. On the other hand, by (2.7),

$$\mathcal{D}(f) = 2 \sum_{x \geqslant 1} \mu(c_x \, \text{Var}_x(f)) = 2q(1-q) \sum_{x \geqslant 1} \mu\left((1 - \omega_{x+1}) \prod_{y=-x+1}^{x-1} \omega_y \right)$$

$$= 2q^2 \sum_{x \geqslant 1} (1-q)^{2x} = \frac{q(1-q)^2}{1 - q/2}.$$

Hence, by (2.10),

$$T_{\text{rel}} \geqslant \frac{\text{Var}(f)}{\mathcal{D}(f)} = \frac{1}{4q^3(1 - q/2)} \geqslant \frac{1}{4q^3}.$$

The inequality $\mathbb{E}_\mu[\tau_0] \geqslant 1/(Cq^3)$ is proved in a similar way, but using a more subtle variant of (2.10) regarding hitting times. See [184, Section 4.2] and [183].

The proof of $T_{\text{rel}} \leqslant C/q^3$ proceeds similarly to the implication from (iii) to (iv) in Theorem 3.10. However, we use canonical paths reflecting the heuristic mechanism discussed above rather than the brutal legal paths provided by bootstrap percolation in Lemma 3.6. The first step is proving a Poincaré inequality for the generalised East model, which will be discussed in Section 4.4. It implies that there exists a constant $C < \infty$ independent of q such that for any local function f,

$$\text{Var}(f) \leqslant C \sum_{x \in \mathbb{Z}} \mu(\tilde{c}_x \, \text{Var}_{\Lambda_x}(f)), \tag{4.1}$$

where $\Lambda_x = \{x\lceil 1/q \rceil, \dots, (x + 1)\lceil 1/q \rceil - 1\}$ and $\tilde{c}_x = 1 - \prod_{y \in \Lambda_{x+1}} \omega_y$. This is precisely the unidimensional analogue of the renormalisation of Figure 3.2.

In view of (4.1), we seek to bound $\mu(\tilde{c}_x \operatorname{Var}_{\Lambda_x}(f))$ with terms $\mu(c_y \operatorname{Var}_y(f))$ of the Dirichlet form from (2.7). To do so, for each $x \in \mathbb{Z}$, $\omega \in \Omega$ such that $\omega_{\Lambda_{x+1}} \neq \mathbf{1}$ and $y \in \Lambda_x$ with $\omega_y = 1$, we define a legal path $(\omega^{(i)})_{i=0}^l$ from ω to ω^y as follows. Set $\xi = \min\{z > y : \omega_z = 0\}$, $l = 2(\xi - y - 1) + 1$ and define $\omega^{(0)} = \omega$, $\omega^{(l)} = \omega^y$ and

$$\omega^{(2i-1)} = \omega^{\xi-i}, \qquad\qquad \omega^{(2i)} = (\omega^{\xi-i})^{\xi-i-1} \qquad (4.2)$$

for $i \in \{1, \dots, \xi - y - 1\}$ ($\omega^{(0)} = \omega$ and $\omega^{(l)} = \omega^y$ by definition). In words, this is the legal path sending an empty interval of length oscillating between one and two sites from ξ to y. Let $x^{(i)}$ be the site such that $\omega^{(i)} = (\omega^{(i-1)})^{x^{(i)}}$. Observe that for odd $i \in \{3, \dots, l\}$ we have $\omega_{x^{(i)}}^{(i)} = 1 = \omega_{x^{(i+1)}}^{(i)}$, so it is convenient to set $j_i = 2\lceil i/2 \rceil - 1$ (that is, the odd number in $\{i, i-1\}$) for $i \in \{2, \dots, l\}$ and $j_1 = 0$.

Thus, by the Cauchy–Schwarz inequality, any ω such that $\tilde{c}_x(\omega) = 1$ satisfies

$$q\omega_y \left(f(\omega^y) - f(\omega)\right)^2 \leqslant ql \sum_{i=1}^l c_{x^{(i)}}\left(\omega^{(i)}\right) \left(f\left(\omega^{(i)}\right) - f\left(\omega^{(i-1)}\right)\right)^2,$$

$$= ql \sum_{i=1}^l \omega_{x^{(i)}}^{(j_i)} c_{x^{(i)}}\left(\omega^{(j_i)}\right) \left(f\left(\omega^{(j_i)}\right) - f\left(\left(\omega^{(j_i)}\right)^{x^{(i)}}\right)\right)^2.$$

Next note that, given y, $\omega^{(j_i)}$ and whether $i = 1$ or not, we can recover ω and that $\mu(\omega)/\mu(\omega^{(j_i)}) \leqslant (1 - q)/q$ (equality holds except for $i = 1$). Integrating the last display over ω, we obtain

$$\mu(\tilde{c}_x \operatorname{Var}_y(f)) \leqslant 4\lceil 1/q \rceil \sum_{z \in \Lambda_x \cup \Lambda_{x+1}} \mu\left(\tilde{c}_x(\omega)c_z(\omega)\omega_z \left(f(\omega) - f(\omega^z)\right)^2\right),$$

since $l \leqslant 2\lceil 1/q \rceil$. Summing the last result over $y \in \Lambda_x$ and then $x \in \mathbb{Z}$, we obtain

$$\sum_{x \in \mathbb{Z}} \mu(\tilde{c}_x \operatorname{Var}_{\Lambda_x}(f)) \leqslant \sum_{x \in \mathbb{Z}} \mu\left(\tilde{c}_x \sum_{y \in \Lambda_x} \operatorname{Var}_y(f)\right) \leqslant 8\lceil 1/q \rceil^3 \mathcal{D}(f),$$

recalling (2.7) and (3.9). Plugging this into (4.1) and recalling (2.10) concludes the proof.

The proof that $\mathbb{P}_\mu(\tau_0 > t) \leqslant e^{-Cq^3 t}$ follows similar but more delicate lines (see [184, Section 4.1] for the details). Finally the upper bound on $\mathbb{E}_\mu[\tau_0]$ follows directly from the last inequality. $\qquad\qquad\square$

We note that higher dimensional analogues of Theorem 4.1 are known and the scaling is $\log(1/q)/q^2$ in $d = 2$ and q^{-2} in $d \geqslant 3$ [44, 184] with the exception of the following conjecture, whose lower bound remains open.

Conjecture 4.3 (Relaxation time in two dimensions) For FA-1f in $d = 2$ dimensions, there exists $C > 0$ such that for q small enough $1/C \leqslant q^2 T_{\text{rel}}/\log(1/q) \leqslant C$.

4.2 East

Recall that the East model in one dimension corresponds to the update family $\mathcal{U} = \{\{1\}\}$. It is, in a sense, the simplest non-trivial KCM and lies at the base of the theory. Indeed, the very first rigorous results on KCM were proved for the East model around the turn of the century [4, 62]. Furthermore, the East model appears also in other contexts including the study of random walks on the group of upper triangular matrices [93, 172].

Once again, by Theorem 3.1, we have $q_c = 0$, so we are interested in asymptotics as $q \to 0$. We address upper and lower bounds separately, showcasing two important techniques. Together they give the following result of Aldous and Diaconis, and Cancrini, Martinelli, Roberto and Toninelli [4, 44].

Theorem 4.4 (East asymptotics) *For the East KCM in $d = 1$ dimension we have*

$$\lim_{q \to 0} \frac{\log T_{\mathrm{rel}}}{(\log(1/q))^2} = \frac{1}{2 \log 2}, \qquad \lim_{q \to 0} \mathbb{P}_\mu \left(\left| \frac{2 \log 2 \cdot \log \tau_0}{(\log(1/q))^2} - 1 \right| < \varepsilon \right) = 1$$

for any $\varepsilon > 0$.

We refer the reader to [55] for finer results including the scaling of T_{rel}^Λ as a function of q and $|\Lambda|$ and for the equivalence, up to a length scale $|\Lambda| = O(1/q)$, of the relaxation and mixing time of the East model; to [87] for a survey dedicated to this model; to [57] for higher dimensions; to [64] for a multicolour version.

Remark 4.5 (Super-Arrhenius law) If we rewrite Theorem 4.4 in terms of temperature using (2.1), we get $\tau_0, T_{\mathrm{rel}} \approx \exp(\beta^2/(2 \log 2))$. This super-Arrhenius divergence is reminiscent of the scaling for fragile super-cooled liquids (see Figure 1.1), which is an important reason for interest in this model among physicists. Also note that this scaling diverges much faster than the emptying time for the corresponding BP (recall Theorem 3.1) or FA-1f (recall Theorem 4.1 and Remark 4.2).

4.2.1 Lower bound: combinatorial bottleneck

The lower bound is proved via a *combinatorial bottleneck*. In rough terms, the strategy is as follows. We consider the stationary KCM started at a typical configuration under μ_q with q small. We identify a set of configurations around the origin which (deterministically) cannot be avoided if the origin is to be infected. For instance, this could be having an atypically large number of empty sites in the vicinity of the origin, but will usually be more subtle. We then seek to evaluate the probability and the number (entropy) of these *bottleneck configurations*. Finally, we use stationarity and a union bound on time to show that if these configurations are unlikely and there are few of them, thus one needs to wait a lot of time to observe any of them close to the origin. The hard part of such arguments is identifying the correct bottleneck.

4.2.1.1 Combinatorics for the East model

The key ingredient to understand the behaviour of the one-dimensional East model is the following combinatorial result of Chung, Diaconis and Graham [62].

Proposition 4.6 (Combinatorial bottleneck for East) *Consider the East model on* $\Lambda = \{\ldots, -2, -1\} \subset \mathbb{Z}$ *with boundary condition* $\mathbf{0}_{\{0,1,\ldots\}}$*. For* $n \geqslant 0$*, let* $V(n)$ *be the set of all configurations that the process can reach from* $\mathbf{1}_\Lambda$ *via a legal path (recall Definition 3.4) in which all configurations contain at most n empty sites. For* $k \in \{0, \ldots, n\}$*, let* $V(n, k) = \{\omega \in V(n) : \sum_{x \in \Lambda}(1 - \omega_x) = k\}$ *be the configurations in* $V(n)$ *with k empty sites. Finally, let* $\ell(n)$ *be the largest distance of an empty site (in* Λ*) from 0 in* $V(n)$*, that is*

$$\ell(n) = \sup\left\{-y : y \in \Lambda, \exists \omega \in V(n), \omega_y = 0\right\}$$

with the convention $\sup \varnothing = 0$*. For* $n \geqslant 0$*,*

$$\ell(n) = 2^n - 1, \tag{4.3}$$

$$|V(n, n)| \leqslant n! 2^{\binom{n}{2}}. \tag{4.4}$$

The first part, (4.3), was already observed in [189, 190]. The second statement, (4.4), is not very far from being tight. Indeed, [62] proved that $c^n \leqslant |V(n)|/(2^{\binom{n}{2}} n!) \leqslant C^n$ with $c \approx 0.67$ the largest root of $384x^3 - 336x^2 + 54x - 1$ and $C = 1/\log 4 \approx 0.72$. Proving Proposition 4.6 is an excellent exercise, which we invite the reader to do before moving on. We provide a full proof, as it is very instructive of the mechanism governing the relaxation of the East model (see Figure 4.1 for an illustration).
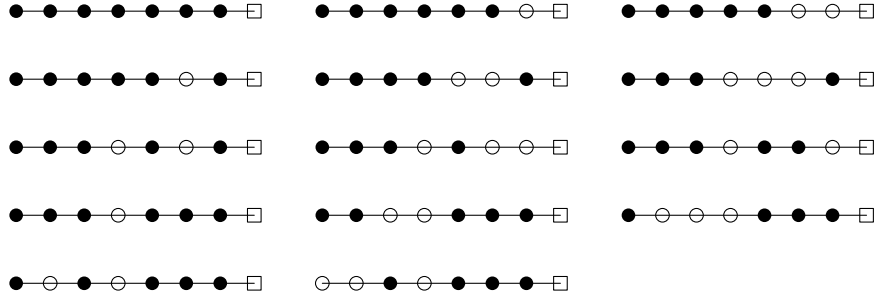

Fig. 4.1: A legal path on $\{\ldots, -2 - 1\}$ with at most 3 simultaneous empty sites starting from a completely occupied configuration and creating a vacancy at $-(2^3-1)$. Successive steps of the path should be read from left to right and top to bottom. The empty square at site 0 stands for the empty boundary condition.

Proof (Proposition 4.6) We prove (4.3) by induction on n. The statement is trivial for $n = 0$. Fix $n \geqslant 1$ and assume that for all $m < n$ we have $\ell(m) = 2^m - 1$. Given

$\omega \in V(n) \setminus \{\mathbf{1}_\Lambda\}$, let $k \leqslant n$, $x_k < \cdots < x_1 < x_0 = 0$ and $X = \{x_1, \ldots, x_k\}$ be such that $\omega = \mathbf{0}_X \cdot \mathbf{1}_{\Lambda \setminus X}$.

Claim There exists $i \in \{1, \ldots, k\}$ such that $x_i - x_{i-1} \leqslant 2^{n-k}$. $\qquad\qquad$ □

Proof Assume the contrary. It is important to recall that the inverse of a legal path is legal, so there exists a legal path from ω to $\mathbf{1}$ via configurations with at most n empty sites. If $k = n$, this immediately gives that there exists $i \in \{1, \ldots, k\}$ with $x_i = x_{i-1} - 1$, since otherwise no legal move can remove empty sites and there are already n of them.

Assume $k < n$ and let $\omega^{(j)}$ be a legal path from ω. We prove by a further induction on $j \geqslant 0$ that $\omega_X^{(j)} = \mathbf{0}_X$. The statement is trivial for $j = 0$. If it is true for all $t < j \geqslant 1$, then we can decompose the dynamics into the intervals $\{x_i + 1, \ldots, x_{i-1} - 1\}$ for $i \in \{1, \ldots, k\}$. Each such interval starts with $\mathbf{1}$ initial condition and has $\mathbf{0}$ boundary condition. Thus, by the first induction hypothesis for $m = n - k$, we necessarily have $\omega_{x_i+1}^{(j-1)} = 1$. Recalling that $\mathcal{U} = \{\{1\}\}$ and (2.3), this implies that $\omega_{x_i}^{(j)} = \omega_{x_i}^{(j-1)} = 0$, completing the second induction. Since there is a legal path from ω to $\mathbf{1}$, but the empty sites of ω cannot be removed, we obtain the desired contradiction proving the claim. $\qquad\qquad$ □

Returning to the first induction, let $i \in \{1, \ldots, k\}$ such that $x_i - x_{i-1} \leqslant 2^{n-k}$. By the induction hypothesis for $m = n - k$, we can find a legal path γ in $\Lambda_i = \{x_i + 1, \ldots, x_{i-1} - 1\}$ with boundary condition $\omega_{\Lambda \setminus \Lambda_i} \cdot \mathbf{0}_{\mathbb{Z} \setminus \Lambda}$ from $\omega_{\Lambda_i} = \mathbf{1}_{\Lambda_i}$ to a configuration $\omega' \in \Omega_{\Lambda_i}$ with $\omega'_{x_i+1} = 0$ and γ features at most $n - k$ empty sites in Λ_i simultaneously. To see this, consider a legal path placing an empty site at $-l(n-k) \leqslant x_i + 1 - x_{i-1}$, truncate it at the first step when an empty site is placed at $x_i + 1 - x_{i-1}$ and shift this path by x_{i-1}.

We can now form the path from ω to ω^{x_i} by performing γ, then occupying x_i and then performing the inverse of γ (which is still legal, because only the boundary condition at x_{i-1} is used due to the orientation of the East update family). By construction this path never creates more than $k + (n - k) = n$ empty sites in Λ, so $\omega^{x_i} \in V(n, k - 1)$. Let $\omega^{x_i} = \mathbf{0}_{X'} \cdot \mathbf{1}_{\Lambda \setminus X'}$ with $X' = \{x'_1, \ldots, x'_{k-1}\}$ and $0 > x'_1 > \cdots > x'_{k-1}$. Then

$$-x_k = \sum_{j=1}^{k} (x_{j-1} - x_j) \leqslant 2^{n-k} + \sum_{j=1}^{k-1} (x'_{j-1} - x'_j) = 2^{n-k} - x'_{k-1}.$$

Iterating this inequality, we obtain

$$-x_k \leqslant \sum_{j=n-k}^{n} 2^j \leqslant \sum_{j=0}^{n-1} 2^j = 2^n - 1,$$

so $l(n) \leqslant 2^n - 1$ as desired. To see that this is an equality it suffices to follow the equalities above, which naturally leads to the path depicted in Figure 4.1.

Finally, (4.4) also follows easily from the above. Namely, each configuration $\omega \in V(n, k)$ can be encoded by the index $i \in \{1, \ldots, k\}$, the distance $x_{i-1} - x_i \in$

$\{1, \dots, 2^{n-k}\}$ and the configuration $\omega^{x_i} \in V(n, k-1)$. Iterating this encoding gives

$$|V(n, n)| \leqslant \prod_{k=1}^{n} (k2^{n-k}) = n!2^{\binom{n}{2}},$$

which concludes the proof. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

4.2.1.2 From the combinatorial result to the emptying time

We now deduce the lower bound of Theorem 4.4 from Proposition 4.6. This was done in somewhat different ways in [4, 48], but we rather present a proof along the lines of [116], which is more adapted to generalisations. Let $n = \lfloor (\log(1/q) - \log\log(1/q))/\log 2 \rfloor$ and $\Lambda_n = \{0, \dots, 2^n - 1\}$. In view of Proposition 4.6, we identify configurations in $V(n)$ and $V(n, n)$ with their restriction to Λ_n. Let $\mathcal{A} = \{\omega \in \Omega : \omega_{\Lambda_n} = \mathbf{1}_{\Lambda_n}\}$. By (4.3), if $\mathcal{A}$ occurs at time 0, then there exists $t \leqslant \tau_0$ such that $\omega_{\Lambda_n}(t) \notin V(n)$. But, exiting $V(n) \times \Omega_{\mathbb{Z} \backslash \Lambda_n}$ requires visiting $V(n, n) \times \Omega_{\mathbb{Z} \backslash \Lambda_n}$.

Let $T = (n^2 q)^{-n/2}$ and let N denote the number of updates (legal or illegal) at sites $x \in \Lambda_n$ up to time T. Then N has the Poisson law with parameter nT. For $i \in \{1, \dots, N\}$, let θ_i denote the times of these updates. Observe that, by stationarity, $\omega(\theta_i)$ is distributed according to μ (see [116, Claim 3.11] for a formal proof). Putting this together, we get

$$\mathbb{P}_\mu(\tau_0 \leqslant T) \leqslant 1 - \mu(\mathcal{A}) + \mathbb{P}_\mu\left(\bigcup_{i=1}^{N} \{\omega_{\Lambda_n}(\theta_i) \in V(n, n)\}\right)$$

$$\leqslant 1 - \mu(\mathcal{A}) + \mathbb{P}(N \geqslant 2nT) + 2nT\mu_{\Lambda_n}(V(n, n)), \qquad (4.5)$$

that is, if $\tau_0 \leqslant T$, then we start outside $\mathcal{A}$, or there are many updates, or at some update the configuration is in $V(n, n)$.

It remains to bound (4.5). Firstly, for $q \to 0$ we have

$$\mu(\mathcal{A}) = (1 - q)^{2^n} \geqslant (1 - q)^{1/(q \log(1/q))} \to 1.$$

Furthermore, by the Bienaymé-Chebyshev inequality, $\mathbb{P}(N \geqslant 2nT) \to 0$, as $nT \to \infty$, which is the case when $q \to 0$. Finally, by (4.4),

$$2nT\mu_{\Lambda_n}(V(n, n)) = 2nT|V(n, n)|q^n(1 - q)^{2^n - 1 - n} \leqslant 2nTn!2^{\binom{n}{2}}q^n$$

$$\leqslant 2n\left(n^2 q\right)^{-n/2} en(n/e)^n 2^{n^2/2}q^n \leqslant 2en^2 e^{-n} \to 0.$$

Inserting these bound in (4.5), we obtain that $\mathbb{P}_\mu(\tau_0 > T) \to 1$ as $q \to 0$. This concludes the proof for the empytying time, since

$$\frac{\log T}{(\log(1/q))^2} \to \frac{1}{2\log 2}.$$

The analogous lower bound for T_{rel} follows directly from (3.14).

4.2.2 Upper bound: the bisection technique

The upper bound of Theorem 4.4 is our first encounter with the bisection technique. It was introduced in [44], drawing inspiration from [154, Proposition 3.5]. It is not only very useful for the study of KCM, but has also been applied in other settings [24, 51].

4.2.2.1 Two-block dynamics

The reader may have noticed that up to now we have not really proved any Poincaré inequality. We have only been reducing one inequality to another one we already know via renormalisation and canonical paths. The next lemma is, in a sense, the only Poincaré inequality we prove from scratch in this monograph. Morally, it deals with the East model on only two sites with empty boundary condition. This being a Markov process with only 4 states, which even happens to be a birth-death chain, one could compute the spectrum of its generator explicitly by hand. However, we state the result directly for the generalised version of the East model, as this is the version that is useful in renormalisation arguments. While the result is originally from [44, Proposition 4.5], we rather give the proof from [113, Lemma 1.3.8], which is more probabilistic.

Lemma 4.7 (Two-block dynamics) *Let $(\mathbb{X}, \pi)$ be the product of two finite probability spaces $(\mathbb{X}_1, \pi_1)$ and $(\mathbb{X}_2, \pi_2)$. Let $\text{Var}_1(f) = \text{Var}_\pi(f(X_1, X_2)|X_2)$ and similarly for $\text{Var}_2(f)$. Fix a nonempty event $\mathcal{X} \subset \mathbb{X}_1$. Then for any $f : \mathbb{X} \to \mathbb{R}$*

$$\text{Var}_\pi(f) \leqslant \frac{\mathbb{E}_\pi(\text{Var}_1(f) + \mathbb{1}_{\mathcal{X}} \text{Var}_2(f))}{1 - \sqrt{1 - \pi(\mathcal{X})}} \leqslant \frac{2}{\pi(\mathcal{X})} \mathbb{E}_\pi \left(\text{Var}_1(f) + \mathbb{1}_{\mathcal{X}} \text{Var}_2(f) \right).$$

A way to interpret this is as a Poincaré inequality (i.e. bound on the relaxation time) for a continuous time Markov chain which updates X_1 at rate 1 and updates X_2 at rate 1, provided that $\mathcal{X}$ occurs. In fact, the relaxation time bound $1/(1 - \sqrt{1 - \pi(\mathcal{X})})$ is optimal.

Proof Couple two copies of the chain described above, by attempting the same updates in both. For this, use a graphical representation as in Section 2.3 attempting updates at X_1 and X_2 with rate 1, but deeming those in X_2 illegal if $\mathcal{X}$ does not occur. The two chains clearly coalesce as soon as we update X_1 so that $\mathcal{X}$ occurs and then immediately update X_2. Consider (legal or illegal) updates on X_2 preceded by an update at X_1. Their number up to time T is $\lfloor N/2 \rfloor$ with N a Poisson random variable with mean T. Each one succeeds in coupling the chains independently with probability $\pi(\mathcal{X})$. It is elementary to check that $\mathbb{E}(\lambda^N) = e^{-T(1-\lambda)}$ for any

$\lambda \in (0, \infty)$. Thus, the probability that the two chains are not equal at time T is at most

$$\mathbb{E}\left[(1 - \pi(X))^{\lfloor N/2 \rfloor}\right] \leqslant \frac{1}{(1 - \pi(X)) \exp(T(1 - \sqrt{1 - \pi(X)}))}.$$

Classical results on Markov chains [144, Proposition 4.7, Corollary 12.6, Remark 13.13][1] then give that $T_{\text{rel}} \leqslant 1/(1 - \sqrt{1 - \pi(X)})$, as desired. Finally, for the second inequality we use the Bernoulli inequality: $x \leqslant 2(1 - \sqrt{1 - x})$ for any $x \in [0, 1]$. $\qquad\square$

If we apply Lemma 4.7 to $\mathbb{X}_1 = \mathbb{X}_2 = \{0, 1\}$, $\pi_1 = \pi_2 = \mu_q$ and $X = \{0\}$, we obtain that the relaxation time of the East KCM with boundary condition $\mathbf{0}_{\mathbb{Z} \setminus \{0,1\}}$ is at most $1/(1 - \sqrt{1 - q})$. When $q \to 0$ this is approximately $2/q$. However, the great advantage of the first bound in Lemma 4.7 is that its prefactor tends to 1 as $\pi(X) \to 1$. Thus, we can hope to apply this result to progressively larger volumes and more likely events X and obtain a relaxation time bound uniform in the volume. This is the main idea behind the upper bound in Theorem 4.4, which we discuss next.

4.2.2.2 Bisection technique

Recall that by Proposition 3.11, in order to prove the upper bound of Theorem 4.4, it suffices to bound the finite volume relaxation times T_{rel}^{Λ} uniformly as the volume diverges. The bisection technique consists in an iterative application of Lemma 4.7. Rather than presenting the somewhat technical and artificial-looking proof of [44], let us take a more instructive approach to see how the proof is conceived.

Basic idea

Let $\Lambda_k = \{1, \ldots, 2^k\}$ and $\Lambda'_k = \Lambda_{k+1} \setminus \Lambda_k$ for any $k \geqslant 0$. Clearly, $T_{\text{rel}}^{\Lambda_0} = 1$. We further seek to relate $T_{\text{rel}}^{\Lambda_k}$ and $T_{\text{rel}}^{\Lambda_{k+1}}$. Fix $k \geqslant 0$ and apply Lemma 4.7 with $\mathbb{X}_1 = \Omega_{\Lambda'_k}$, $\mathbb{X}_2 = \Omega_{\Lambda_k}$, $\pi_1 = \mu_{\Lambda'_k}$, $\pi_2 = \mu_{\Lambda_k}$ and $X = \{\omega_{2^k+1} = 0\} \subset \mathbb{X}_1$. This gives

$$\text{Var}(f) \leqslant \frac{\mu(\text{Var}_{\Lambda'_k}(f) + \mathbb{1}_X \text{Var}_{\Lambda_k}(f))}{1 - \sqrt{\varepsilon_k}} \tag{4.6}$$

for any $f : \Omega_{\Lambda_{k+1}} \to \mathbb{R}$, where $\varepsilon_k = 1 - \pi_1(X) = 1 - q$. For the first term above, (2.11) and translation invariance directly give

$$\text{Var}_{\Lambda'_k}(f) \leqslant T_{\text{rel}}^{\Lambda_k} \mathcal{D}_{\mathbf{0}_{\mathbb{Z} \setminus \Lambda'_k}}(f). \tag{4.7}$$

---

[1] For continuous time Markov chains the spectral radius in [144, Corollary 12.6] is replaced by $e^{-1/T_{\text{rel}}}$.

Yet, the fact that the East update family $\mathcal{U} = \{\{1\}\}$ only looks to the right gives that for any $x \in \Lambda'_k$ we have $c_x^{\mathbf{0}_{\mathbb{Z}\backslash\Lambda'_k}} = c_x^{\mathbf{0}_{\mathbb{Z}\backslash\Lambda_{k+1}}}$, so

$$\mathcal{D}_{\mathbf{0}_{\mathbb{Z}\backslash\Lambda'_k}}(f) = \sum_{x\in\Lambda'_k} \mu\left(c_x^{\mathbf{0}_{\mathbb{Z}\backslash\Lambda_{k+1}}} \mathrm{Var}_x(f)\right) \tag{4.8}$$

by (2.7). For the second term in (4.6), we similarly have

$$\mathbb{1}_{\mathcal{X}}\mathcal{D}_{\mathbf{0}_{\mathbb{Z}\backslash\Lambda_k}}(f) \leq \sum_{x\in\Lambda_k} \mu\left(c_x^{\mathbf{0}_{\mathbb{Z}\backslash\Lambda_{k+1}}}(f)\right), \tag{4.9}$$

since $\mathcal{X}$ guarantees precisely the presence of an empty site as in the boundary condition $\mathbf{0}_{\mathbb{Z}\backslash\Lambda_k}$. Combining (4.6)-(4.9) and (2.10), we obtain

$$T_{\mathrm{rel}}^{\Lambda_{k+1}} \leq \frac{T_{\mathrm{rel}}^{\Lambda_k}}{1 - \sqrt{\varepsilon_k}}. \tag{4.10}$$

Iterating the above relation, and recalling Proposition 3.11, we get

$$T_{\mathrm{rel}} \leq \frac{2}{q} \prod_{k=0}^{\infty} \frac{1}{1 - \sqrt{\varepsilon_k}}. \tag{4.11}$$

Spread the boundary condition

Unfortunately, this first attempt fails, because $\varepsilon_k = 1 - q$ does not decay with k, so the product in (4.11) is infinite. In order to fix this problem, we should define the event $\mathcal{X}$ differently. Namely, fix some $\delta > 0$ small enough and let $\delta_k = \lfloor 2^{k(1-\delta)} \rfloor$. Then let $\mathcal{X} = \{\omega_{2^k+\{1,\dots,\delta_k\}} \neq \mathbf{1}\}$, that is, there is an empty site among the first δ_k sites to the East of Λ_k. With this choice $\varepsilon_k = (1-q)^{\delta_k}$ does decay sufficiently fast for the right hand side of (4.11) to be finite. In fact, one can compute that the product is at most $q^{-C}(1/q)^{\log_2(1/q)/(2-2\delta)}$ for some constant $C > 0$ (see [44, Section 6.1] for more details). Taking δ small, this is exactly the upper bound we want.

However, with this choice of $\mathcal{X}$, (4.9) is no longer valid. To deal with this issue, on $\mathcal{X}$ we can define the random variable

$$\xi(\omega) = \{\max i \leq \delta_k : \omega_{2^k+i} = 0\} \geq 1 \tag{4.12}$$

indicating the position of the rightmost empty site in Λ'_k at distance at most δ_k from the boundary of Λ_k. Then we can rewrite

$$\mu_{\Lambda_{k+1}}\left(\mathbb{1}_\chi \operatorname{Var}_{\Lambda_k}(f)\right) = \sum_{i=1}^{\delta_k} \mu_{\Lambda_{k+1}}\left(\mathbb{1}_{\xi=i}\operatorname{Var}_{\Lambda_k}(f)\right) \tag{4.13}$$

$$= \sum_{i=1}^{\delta_k} \mu_{\Lambda_{k+1}}\left(\mathbb{1}_{\xi=i}\mu_{\{1,\dots,2^k+i-1\}}\left(\operatorname{Var}_{\Lambda_k}(f)\right)\right),$$

since $\mathbb{1}_{\xi=i}$ is independent of $\omega_{\{1,\dots,2^k+i-1\}}$. But, setting $V_i = \{1,\dots,2^k+i-1\}$, we have

$$\mathbb{1}_{\xi=i}\mu_{V_i}\left(\operatorname{Var}_{\Lambda_k}(f)\right) \leqslant \mathbb{1}_{\xi=i}\operatorname{Var}_{V_i}(f) \leqslant T_{\mathrm{rel}}^{V_i}\mathbb{1}_{\xi=i}\mathcal{D}_{\mathbf{0}_{\mathbb{Z}\backslash V_i}}(f), \tag{4.14}$$

since $\xi = i$ guarantees that $\omega_{2^k+i} = 0$, so the boundary condition is indeed empty. Note that in the first inequality above, we used the convexity of the variance that implies that for any volumes A, B it holds

$$\mu_A(\operatorname{Var}_B(f)) \leqslant \operatorname{Var}_{A\cup B}(f). \tag{4.15}$$

Combining (4.13) and (4.14) together with the monotonicity property (3.16), we get

$$\mu_{\Lambda_{k+1}}\left(\mathbb{1}_\chi \operatorname{Var}_{\Lambda_k}(f)\right) \leqslant T_{\mathrm{rel}}^{V_{\delta_k}}\sum_{i=1}^{\delta_k}\mu_{\Lambda_{k+1}}\left(\mathbb{1}_{\xi=i}\mathcal{D}_{\mathbf{0}_{\mathbb{Z}\backslash V_i}}(f)\right)$$

$$\leqslant T_{\mathrm{rel}}^{V_{\delta_k}}\sum_{x\in V_{\delta_k}}\mu_{\Lambda_{k+1}}\left(c_x \operatorname{Var}_x(f)\right),$$

using $\sum_{i=1}^{\delta_k}\mathbb{1}_{\xi=i} \leqslant 1$ and the fact that $\mathbb{1}_{\xi=i}\cdot c_x^{\mathbf{0}_{\mathbb{Z}\backslash V_i}} \leqslant \mathbb{1}_{\omega_{2^k+i}=0}\cdot c_x^{\mathbf{0}_{\mathbb{Z}\backslash V_i}} \leqslant c_x$ for any $i \in \{1,\dots,\delta_k\}$ and $x \in V_i$. Further combining this with (4.6)-(4.8), we obtain

$$\operatorname{Var}(f) \leqslant \frac{T_{\mathrm{rel}}^{V_{\delta_k}}}{1-\sqrt{\varepsilon_k}}\mu\left(\mathcal{D}_{\Lambda_{k+1}}(f) + \sum_{x\in V_{\delta_k}\cap\Lambda_k'}\left(c_x^{\mathbf{0}_{\mathbb{Z}\backslash\Lambda_{k+1}}}\operatorname{Var}_x(f)\right)\right). \tag{4.16}$$

Recalling (2.10), we see that (4.16) is almost the result we seek to prove, (4.10).

Final adjustments

We are only left with mending two technical problems with the previous argument. Firstly, in the right hand side of (4.16), the terms corresponding to $x \in V_{\delta_k} \cap \Lambda_k'$ appear twice (once in the Dirichlet form). In order to solve this, we consider many possible choices of the partition of Λ_{k+1} into Λ_k and Λ_k', keeping the total volume 2^{k+1} and the overlap δ_k fixed. We then average (4.16) over these choices. This yields an additional factor $1 + 1/s_k$ in (4.10), where $s_k = \lfloor 2^{k\varepsilon/3}\rfloor$ is the number of choices we consider.

Secondly, V_{δ_k} is slightly larger than Λ_k and so is the corresponding relaxation time. This issue is solved by choosing the sizes of all Λ_k growing as $2^k + 2^{k(1-\varepsilon/3)}$ rather than 2^k. Once these problems are solved, we get (4.10) (with the additional factor $1 + 1/s_k$) and conclude the proof of the upper bound in Theorem 4.4. The upper bound on τ_0 follows from (3.14).

Interestingly, prior to [44], the conjecture in the physics literature on the exponent was $T_{\text{rel}} \sim q^{\log_2(1/q)}$, with an exponent off by a factor 2. In order to get the correct scaling one has to take into account a subtle balance between the energetic and entropic contributions that, atypically, lie at the same level for the one-dimensional East model. Remarkably, the bisection technique is able to automatically take into account this subtle balance and provide a tight result correcting the conjectured exponent.

4.3 FA-2f

In this section, we briefly discuss the one-dimensional FA-2f update family $\mathcal{U} = \{\{-1, 1\}\}$. From Theorem 3.1, we have $q_c = 1$, because two neighbouring occupied sites remain occupied at all times. Normally, our study of the model would end here, because the phase $q < q_c$ is rather complicated, but in the one-dimensional setting, we are able to say more.

Specifically for FA-2f, one can check that the BP transformation (recall (3.1)) satisfies $[\omega] = \mathcal{B}_{\mathcal{U}}(\omega)$ for any $\omega \in \Omega$. That is, the BP process becomes stationary after one step. Taking Corollary 3.7 into account, the KCM dynamics can be decomposed into independent dynamics on intervals delimited by two occupied sites. On each such interval, we recover what is known as the hard-core Glauber dynamics with fugacity $\lambda = (1 - q)/q$. That is because occupied sites cannot appear next to other occupied sites, while emptying is always possible (within an interval delimited by two occupied sites). There is a rich literature on Glauber dynamics of the hard-core model, particularly on general graphs with bounded degree, but also lattices of higher dimension (see e.g. [144, 188]). However, the one-dimensional lattice is somewhat degenerate from the standard viewpoint and does not appear to have been the subject of much study.

One natural question one could ask is how the system behaves on its *ergodic component*, that is, the set of configurations such that $[\omega] = \mathbf{0}$. One can then still study the $q \to 0$ regime, which can be viewed as quenching the model from inverse temperature $\beta = -\infty$ to β large. In this setting, it is possible to prove that the relaxation time (this time with respect to the Gibbs measure of the hard-core model, see [100] for background, rather than the plain μ_q product measure) is finite for any $q > 0$. This can be obtained, for example using classical techniques such as block dynamics and strong spatial mixing [156–158]. However, as we will see in Section 4.4, this can also be achieved via the bisection technique, which applies more broadly.

4.4 General KCM

We have so far seen two ways in which it may be desirable to generalise KCM. Namely, allowing a state space larger than $\{0, 1\}$ for each site, and working with the dynamics restricted to an ergodic component. Furthermore, for the purposes of studying higher-dimensional models, it is also useful to consider inhomogeneous KCM with site-dependent update families. We next define general KCM incorporating all these features, following [111].

Fix $R > 0$, $q \in (0, 1)$ and $\Lambda \Subset \mathbb{Z}$. For each $x \in \Lambda$, fix a probability space (Ω_x, π_x) with $|\Omega_x| < \infty$ and an event $\mathcal{I}_x \subset \Omega_x$ with $\pi_x(\mathcal{I}_x) \geqslant q$. Let $(\Omega, \pi) = (\prod_{x \in \Lambda} \Omega_x, \bigotimes_{x \in \Lambda} \pi_x)$ be the corresponding product space. A boundary condition is any configuration $\eta \in \{0, 1\}^{\mathbb{Z} \backslash \Lambda}$. Further fix an update family $\mathcal{U}_x$ for each $x \in \Lambda$ so that for any $U \in \mathcal{U}_x$ and $u \in U$ we have $|u| \leqslant R$. The constraint at $x \in \Lambda$ is defined by

$$c_x^{\eta}(\omega) = \max_{U \in \mathcal{U}_x} \prod_{u \in U, x+u \in \Lambda} \mathbb{1}_{\omega_{x+u} \in \mathcal{I}_{x+u}} \prod_{u \in U, x+u \in \mathbb{Z} \backslash \Lambda} (1 - \eta_{x+u})$$

for a configuration $\omega \in \Omega$ and a boundary condition $\eta \in \{0, 1\}^{\mathbb{Z} \backslash \Lambda}$.

Consider the Markov process such that for each site $x \in \Lambda$ such that $c_x = 1$, the state of site x is updated to an independent random variable with law π_x. That is, the process with generator

$$\mathcal{L}(f) = \sum_{x \in \Lambda} c_x^{\eta} \cdot (\pi_x(f) - f).$$

Let C be an arbitrarily chosen irreducible component of Ω for this dynamics and $\mu = \pi(\cdot | C)$. Then define T_{rel} via (2.7) and (2.10). We refer to this Markov process restricted to C as a *general KCM with range R and facilitating parameter q*. The following result was proved by Hartarsky [111] via the bisection technique adapted for going back and forth several times between the two blocks in the two-block Lemma 4.7.

Theorem 4.8 (General KCM upper bound) *There exists $C > 0$ depending only on R such that for every $q \in (0, 1)$,*

$$T_{\text{rel}} \leqslant (2/q)^{C \log \min(|\Lambda|, 2/q)}.$$

In words, Theorem 4.8 states that, for any one-dimensional general KCM with uniformly bounded update rule range and probability of the facilitating state uniformly bounded away from 0 has a finite relaxation time scaling at most like the one of the East model (recall Theorem 4.4). Note that the minimum reflects the fact that the product in (4.11) approaches its limiting value for scales $k \approx \log(1/q)$. Theorem 4.8 can also be extended to infinite volume along the lines of Proposition 3.11, but one needs to be careful in defining the irreducible components (see [111, Observation 3]).

Let us discuss a few useful applications of Theorem 4.8. Firstly, FA-2f on its ergodic component is covered, just like any 1-dimensional (homogeneous binary) KCM. More importantly, we have the following bounds for the generalised FA-1f and East KCM. They were both derived in [155, Proposition 3.4], using the methods discussed in Sections 4.1 and 4.2.2, the second one also following from Theorem 4.8.

Proposition 4.9 (Generalised FA-1f and East upper bound) *Let Λ be a segment and $\eta \in \{0,1\}^{\mathbb{Z}\backslash\Lambda}$ with $\eta_x = 0$ if $x > \max\Lambda$ and $\eta_x = 1$ if $x < \min\Lambda$. Consider a general KCM on Λ and assume it to be homogeneous with $\mathcal{U}_x = \mathcal{U}$ for all $x \in \Lambda$, with range 1, facilitating parameter $q \in (0,1)$ and boundary condition η. Then for some absolute $C > 0$ we have*

$$T_{\mathrm{rel}} \leqslant \begin{cases} (2/q)^C & \mathcal{U} = \{\{-1\},\{1\}\}, \\ (2/q)^{C\log\min(|\Lambda|,2/q)} & \mathcal{U} = \{\{1\}\}. \end{cases}$$

Let us note that one can also prove a polynomial bound on the relaxation time of (generalised) FA-1f on a segment with **1** boundary condition on its ergodic component (that is, all configurations except **1**), see [29].

4.5 Conclusion

Let us review the state of our toolbox after the developments of this chapter (recall Section 3.6).

Test functions. The lower bound of Theorem 4.1 was proved by guessing a non-trivial test function. This technique will not take us any further in the sequel for more sophisticated models, because guessing a suitable function and being able to compute the variance and Dirichlet form are quite implausible.

Canonical paths. The upper bound of Theorem 4.1 relied on more subtle canonical paths than the ones provided by BP in Section 3.2. They reflect the heuristic view of the dynamics of FA-1f. Once we have a good intuition about the dominant relaxation mechanism of a KCM, we could, in principle try to implement it in a canonical path. Unfortunately, this approach quickly goes out of hand as the models get more complicated, since explicitly defining and analysing the paths involved becomes very laborious and quite tricky. We therefore avoid further recourse to canonical paths.

Renormalisation. In the proof of Theorem 4.1, we saw the details of the 1-dimensional renormalisation we already saw in Section 3.4. This technique for proving upper bounds will be developed much further in the next chapter. Now that we have some simple KCM to build on, it will become our bread and butter tool for proving upper bounds.

Combinatorial bottlenecks. This method discussed in Section 4.2.1 will be our method of choice for proving lower bounds on time scales in what follows. The content of Section 4.2.1.2 will require essentially no adaptation. The main difficulty in implementing this approach lies in identifying what needs to happen before the origin can be updated and proving that it is indeed necessary. Finding the correct

bottleneck is usually guided by heuristics of the dominant relaxation mechanism (from upper bounds), estimating the probability of the bottleneck will usually be done using ideas from BP, while entropy tends not to pose problems. Thus, the main issue is proving rough analogues of (4.3) in more advanced settings.

Bisection. The idea used in Section 4.2.2 was to iterate the simple two-block dynamics of Lemma 4.7. Bisection is our primary technique for proving directly that a KCM has finite relaxation time in infinite volume. In the next chapter we will see how to do the same in higher dimensions.

General KCM. Thanks to bisection, we were able to treat one-dimensional KCM in great generality. They are ready to use in renormalisation schemes. Although we will need to introduce some higher-dimensional models with general state space, one-dimensional general KCM will be sufficient for most of our purposes.

Chapter 5
Fredrickson-Andersen 2-spin facilitated model

This chapter uses the setting of the FA-2f model to develop several new tools for determining the emptying time of the origin with high precision as $q \to 0$. We begin by using bisection in higher dimensions to show that this time scale is finite. We then discuss a robust long range Poincaré inequality approach for proving upper bounds. Finally, we assess the sharp threshold of FA-2f, which relies on a robust relation with bootstrap percolation for the lower bound and on the very flexible method of matryoshka dolls for the upper bound. All of these methods generalise to treat various other models.

Recall from Section 2.2 that the FA-2f model's constraint requires at least two empty neighbours in order to change the state of a site. Throughout this section we work in two dimensions for simplicity of notation, but the arguments apply equally well in any dimension. For FA-2f (in two dimensions), the natural geometry is rectangular. It is therefore, convenient to denote by

$$R(a, b) = \{0, \ldots, a-1\} \times \{0, \ldots, b-1\} \subset \mathbb{Z}^2 \tag{5.1}$$

the rectangle of side lengths $a, b \in \mathbb{N}$.

5.1 Bisection in higher dimensions

Let us begin by paying our debt by proving that $T_{\text{rel}} < \infty$ for any $q > 0$ for FA-2f. Recall that in the proof of Theorem 3.10, the only fact whose proof was postponed until now is that for q close enough to 1, the KCM with update family

$$\mathcal{U}' = \{U'\} = \{\{(1, 0), (1, 1), (0, 1)\}\}$$

has a finite relaxation time. In fact, we rather used this result for the version of this KCM with general state space, as in Section 4.4, but the proof is very similar, so we focus on the binary case. In Section 4.2.2 we discussed how to prove this in one

dimension via the bisection technique. We next explain how the method adapts to higher dimensions, following [44]. In fact, the proof works for any $q > q_c(\mathcal{U}')$.

We start with a simple observation, which could be attributed to Schonmann [181].

Observation 5.1 (Oriented percolation correspondence) *Endow $\mathbb{Z}^2$ with the oriented graph structure defined by the edge set $E = \{(\mathbf{x}, \mathbf{x} + \mathbf{u}) : \mathbf{x} \in \mathbb{Z}^2, \mathbf{u} \in U'\}$. Then for any configuration $\omega \in \Omega$, in $\mathcal{U}'$-BP, the emptying time τ_0^{BP} is given by the number of sites in the longest (oriented) path from 0 whose sites are all occupied.*

The proof of this fact by induction on the number of iterations of the BP map $\mathscr{B}_{\mathcal{U}}$ from (3.1) is left as an exercise to the reader. In view of Observation 5.1, for $\mathbf{x} \in A \subset \mathbb{Z}^2$, $B \subset \mathbb{Z}^2$ and $\omega \in \Omega_A$, we write $\mathbf{x} \xrightarrow{A} B$ if there exists a sequence of sites $(\mathbf{x}_i)_{i=0}^l \in A^{l+1}$ occupied in ω with $\mathbf{x}_0 = \mathbf{x}$, $\mathbf{x}_l \in B$ and $(\mathbf{x}_{i-1}, \mathbf{x}_i) \in E$ for all $i \in \{1, \ldots, l\}$.
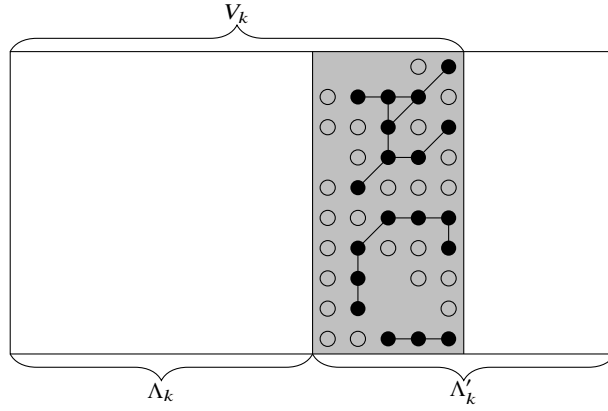

Fig. 5.1: Illustration of the bisection in Section 5.1. The rectangle Λ_{k+1} is partitioned into $\Lambda_k \sqcup \Lambda_k'$. The shaded rectangle V_k' has width δ_k. The full dots in it represent occupied sites in Γ_k, while $\partial \Gamma_k$ consists of the empty sites. In the figure the event $\mathcal{X}_k$ occurs, since the paths do not reach the leftmost column of V_k'.

For any $k \geqslant 0$, consider the rectangles

$$\Lambda_{2k} = R\left(2^k, 2^k\right) \qquad\qquad \Lambda_{2k+1} = R\left(2^{k+1}, 2^k\right)$$

and $\Lambda_k' = \Lambda_{k+1} \setminus \Lambda_k$, which is a translate of Λ_k. Note that the rectangles $(\Lambda_k)_{k \geqslant 0}$ are nested in such a way that each is obtained by stretching the previous rectangle twice either horizontally or vertically (see Figure 5.1). These rectangles will play the role of the intervals Λ_k in Section 4.2.2.2. As in one dimension, for any $k \geqslant 0$, we set $\delta_k = \lfloor 2^{k(1-\delta)} \rfloor$ with some fixed $\delta > 0$ small enough. Further define

$$V'_{2k} = \left(2^k, 0\right) + R\left(\delta_k, 2^k\right) \qquad\qquad V'_{2k+1} = \left(0, 2^k\right) + R\left(2^{k+1}, \delta_k\right)$$

$$\partial_+ V'_{2k} = \left(2^k + \delta_k - 1, 0\right) + R\left(1, 2^k\right) \quad \partial_+ V_{2k+1} = \left(0, 2^k + \delta_k - 1\right) + R\left(2^{k+1}, 1\right)$$

$$\partial_- V'_{2k} = \left(2^k, 0\right) + R\left(1, 2^k\right) \qquad\qquad \partial_- V_{2k+1} = \left(0, 2^k\right) + R\left(2^{k+1}, 1\right)$$

and $V_k = \Lambda_k \cup V'_k$ (see Figure 5.1).

We next need to choose the facilitating event $\mathcal{X}$ in Lemma 4.7, so that its probability gets close to 1 as the scale k increases. Given $k \geqslant 0$, let

$$\mathcal{X}_k = \left\{\omega \in \Omega_{V'_k} : \Gamma_k \cap \partial_- V'_k = \varnothing\right\}, \qquad \Gamma_k = \left\{\mathbf{x} \in V'_k : \mathbf{x} \xrightarrow{V'_k} \partial_+ V'_k\right\}. \tag{5.2}$$

That is, $\mathcal{X}_k$ is the event that no site on the boundary of V'_k is connected to the other boundary via occupied sites in V'_k (see Figure 5.1).

With this definition it is a classical percolation result of Menshikov and Aizenman–Barsky [1, 163] that for some $c > 0$ depending only on $q > q_\mathrm{c}$,

$$\varepsilon_k = 1 - \mu(\mathcal{X}_k) \leqslant \exp\left(-c\delta_k\right) \tag{5.3}$$

(see [73] for a simple proof, [110] for more background on percolation and [121] for more background on oriented percolation). Indeed, taking Observation 5.1 into account, $1 - q_\mathrm{c}$ is the critical parameter of the oriented percolation model on $\mathbb{Z}^2$ with edge set E. Plugging (5.3) into (4.11), we easily obtain that the product there is finite.

Our next task is to show bound the second term in (4.6) as

$$\mu_{V_k}\left(\mathbb{1}_{\mathcal{X}_k} \operatorname{Var}_{\Lambda_k}(f)\right) \leqslant T_\mathrm{rel}^{V_k} \sum_{\mathbf{x} \in V_k} \mu_{V_k}\left(c_{\mathbf{x}}^{\mathbf{0}_{\mathbb{Z}^2 \setminus \Lambda_{k+1}}} \operatorname{Var}_{\mathbf{x}}(f)\right). \tag{5.4}$$

Once (5.4) is established, the proof is concluded by the same final adjustments as in Section 4.2.2.2. In order to prove (5.4), we proceed similarly to the one-dimensional case, but we need to pay attention to the definition of the 'rightmost empty site'. For $\gamma \subset V'_k$, let

$$\bar{\gamma} = \partial_+ V'_k \cup \gamma \cup \left(V'_k \cap (\gamma - U')\right). \tag{5.5}$$

Then $\partial \Gamma_k = \bar{\Gamma}_k \setminus \Gamma_k$ (see Figure 5.1) will play the role of the rightmost empty site ξ in (4.12). By (5.2) and (5.5), Γ_k (and $\bar{\Gamma}_k$) is measurable with respect to $\omega_{\bar{\Gamma}_k}$ and, if $\mathcal{X}_k$ occurs, then $\omega_{\partial \Gamma_k} = \mathbf{0}$. Finally, observe that, by construction, $\partial \Gamma_k$ is a cut-set separating $\partial_+ V'_k$ from Λ_k. Therefore, recalling (4.15), this gives

$$\mu_{V_k}\left(\mathbb{1}_{X_k}\operatorname{Var}_{\Lambda_k}(f)\right) \leqslant \sum_{\gamma \subset V_k' \backslash \partial_- V_k'} \mu_{V_k}\left(\mathbb{1}_{\Gamma_k=\gamma}\operatorname{Var}_{\Lambda_k}(f)\right)$$

$$\leqslant \sum_{\gamma \subset V_k' \backslash \partial_- V_k'} \mu_{V_k}\left(\mathbb{1}_{\Gamma_k=\gamma}\mathbb{1}_{\omega_{\partial\gamma}=\mathbf{0}}\operatorname{Var}_{V_k\backslash\overline{\gamma}}(f)\right)$$

$$\leqslant \sum_{\gamma \subset V_k' \backslash \partial_- V_k'} T_{\mathrm{rel}}^{V_k\backslash\overline{\gamma}} \sum_{\mathbf{x}\in V_k\backslash\overline{\gamma}} \mu_{V_k}\left(\mathbb{1}_{\Gamma_k=\gamma}c_{\mathbf{x}}^{\mathbf{0}_{\mathbb{Z}^2\backslash\Lambda_{k+1}}}\operatorname{Var}_{\mathbf{x}}(f)\right),$$

where we used $\mathbb{1}_{\omega_{\partial\gamma}=\mathbf{0}}c_{\mathbf{x}}(\omega) \leqslant c_{\mathbf{x}}(\omega)$ for all $\mathbf{x} \in V_k \backslash \overline{\gamma}$ and $\omega \in \Omega$. Recalling the the monotonicity property (3.16), we recover (5.4) as desired.

Bisection is good for proving $T_{\mathrm{rel}} < \infty$. However, if we follow the proof of Theorem 3.10 and take into account Theorem 3.1, we obtain the following very poor bound.

Corollary 5.2 (Basic FA-2f upper bound) *For FA-2f in d dimensions there exists* $C > 0$ *such that for all* $q > 0$,

$$T_{\mathrm{rel}} \leqslant \exp^{\circ 2}\left(Cq^{-1/(d-1)}\right).$$

In [44] this was improved to $\exp(Cq^{-5})$ using less brutal canonical paths than the ones suggested in the proof of Theorem 3.10. However, this is still very far from the truth. The next sections examine other techniques for obtaining more accurate bounds.

5.2 Long range renormalisation

In order to improve the bound on the relaxation time of FA-2f provided by Corollary 5.2, we need a new technique by Martinelli and Toninelli [161]. The key idea is to transform our KCM, which has short range but unlikely constraints, into a model with long range but very likely constraints. These constraints require the occurrence of a certain *droplet* far away and of a *good environment*, on which the droplet can move, connecting the droplet location to the origin. We start by proving a *long range Poincaré inequality*, which is the key building block of this technique, in Section 5.2.1. In Section 5.2.2 we discuss how to combine it with renormalisation to obtain results about our model of interest.

5.2.1 A long range constrained Poincaré inequality

Fix a finite probability space $(\mathbb{X}_0, \nu_0)$ and let $(\mathbb{X}, \nu) = (\prod_{\mathbf{x}\in\mathbb{Z}^d}\mathbb{X}_0, \bigotimes_{\mathbf{x}\in\mathbb{Z}^d}\nu_0)$ be the corresponding product space. For each $\mathbf{x} \in \mathbb{Z}^d$ let $\mathbb{N}_{\mathbf{x}}^d := \mathbf{x} + (\{0, 1, \dots\}^d \backslash \{\mathbf{0}\})$. Then fix a finite set $\Delta_0 \subset \mathbb{N}_0^d$. Let $\mathcal{A}_0 \subset \mathbb{X}$ be an event depending only on the

restriction $\eta_{\Delta_0} \in \mathbb{X}_{\Delta_0}$ of the state $\eta \in \mathbb{X}$ to Δ_0. Define $\mathcal{A}_{\mathbf{x}}$ by translating $\mathcal{A}_0$ by $\mathbf{x} \in \mathbb{Z}^d$. For $\mathbf{x} \in \mathbb{Z}^d$, set

$$r_{\mathbf{x}} = \mathbb{1}_{\mathcal{A}_{\mathbf{x}}}, \qquad\qquad \epsilon = 1 - \nu_0(\mathcal{A}_0).$$

Then following result gives a taste of the more general [161, Theorem 2].

Proposition 5.3 (Long range constrained Poincaré inequality) *Assume that*

$$(1 + |\Delta_0|)\epsilon < 1/4, \tag{5.6}$$

then for any local function $f : \mathbb{X} \to \mathbb{R}$ it holds that

$$\mathrm{Var}(f) \leqslant 4 \sum_{\mathbf{x} \in \mathbb{Z}^d} \nu \left(r_{\mathbf{x}} \, \mathrm{Var}_{\mathbf{x}}(f) \right). \tag{5.7}$$

Proof (Sketch) The starting point to prove (5.7) is the inequality [161, Lemma 2.5]

$$\mathrm{Var}(f) \leqslant \sum_{\mathbf{x} \in \mathbb{Z}^d} \nu \left(\mathrm{Var}_{\mathbf{x}} \left(\nu_{\mathbb{N}_{\mathbf{x}}^d}(f) \right) \right), \tag{5.8}$$

which can be obtained using the law of total variance and the product form of ν. Then one can rewrite a generic term in the r.h.s. of (5.8) using the decomposition

$$\mu_{\mathbb{N}_{\mathbf{x}}^d}(f) = \mu_{\mathbb{N}_{\mathbf{x}}^d}(r_{\mathbf{x}} f) + \mu_{\mathbb{N}_{\mathbf{x}}^d}((1 - r_{\mathbf{x}}) f),$$

and repeatedly use Cauchy-Schwarz inequality and convexity of the variance to obtain the final result (5.7). $\square$

Remark 5.4 Proposition 5.3 also plays a key role for the study of KCM on regular trees. It is actually in this context that it was introduced in [160] (also see [47] for a refinement).

In [161], Proposition 5.3 is proved in a more general setting allowing the constraints $r_{\mathbf{x}}$ to be the product of several indicator functions and at the same time transforming (5.6) into a more flexible condition involving the supports and the probabilities of these events. Using this, Proposition 5.6 below is deduced in [161, Proposition 3.4]. Before stating it, we require some more notation.

Definition 5.5 (Good path) We call a sequence of sites $\gamma = (\mathbf{x}_i)_{i \geqslant 0}^k$ with $\mathbf{x}_i - \mathbf{x}_{i-1} \in \{\mathbf{e}_j : j \in \{1, \ldots, d\}\}$ for any $i \in \{1, \ldots, k\}$ an *oriented path of length $k + 1$*. Fix a *good* and a *super good* event $\mathcal{G}, \mathcal{SG} \subset \mathbb{X}_0$. Given a configuration $\eta \in \mathbb{X}$, we say that the oriented path γ is *good*, if $\eta_{\mathbf{x}_i} \in \mathcal{G}$ for all $i \in \{0, \ldots, k\}$ and $\eta_{\mathbf{x}_k} \in \mathcal{SG}$. For any $\mathbf{x} \in \mathbb{Z}^d$ and $K \geqslant 1$, we denote by $\Gamma_{\mathbf{x}}^K$ the event that for every $i \in \{1, \ldots, d\}$ there exists a good oriented path of length at most K starting at $\mathbf{x} + \mathbf{e}_i$.

Proposition 5.6 (Good path constraint dynamics)
There exists $\delta > 0$ such that for all $p_1, p_2 \in [0, 1]$ with $\max(p_2, (1 - p_1)(\log(1/p_2))^2) \leqslant \delta$, the following holds. If $\nu_0(\mathcal{G}) = p_1$ and $\nu_0(\mathcal{SG}) = p_2$, then for any local function $f : \mathbb{X} \to \mathbb{R}$,

$$\mathrm{Var}(f) \leqslant 4 \sum_{\mathbf{x} \in \mathbb{Z}^d} \nu \left(\mathbb{1}_{\Gamma_{\mathbf{x}}^{p_2^{-2}}} \mathrm{Var}_{\mathbf{x}}(f) \right). \tag{5.9}$$

5.2.2 Combining renormalisation and the long range Poincaré inequality

We next explain how to apply Proposition 5.6 to FA-2f. As suggested by the fact that we worked with a general state space in Section 5.2.1, we intend to use renormalisation. Once again, the argument works in any dimension, but we present it in two-dimensions in order to simplify notation. For any $\mathbf{x} \in \mathbb{Z}^2$, the (renormalised) site $\mathbf{x}$ will correspond to the rectangle $R_{\mathbf{x}} = \ell \cdot \mathbf{x} + R(\ell, \ell)$ (recall (5.1)), where $\ell = \lceil 3 \log(1/q)/q \rceil$. We consider the state space $\mathbb{X}_0 = \Omega_{R_0}$ and the corresponding states $\eta_{\mathbf{x}} = \omega_{R_{\mathbf{x}}}$ for $\mathbf{x} \in \mathbb{Z}^2$.

It remains to choose the good and super good events in Definition 5.5. A simple choice is to define $\omega \in \mathcal{SG} \subset \Omega_{R_0}$, if

$$\omega_{R(\ell,\ell) \setminus ((1,1)+R(\ell-2,\ell-2))} = \mathbf{0}. \tag{5.10}$$

That is, a renormalised site is super good, if the perimeter of the corresponding rectangle is empty. The good event is defined by $\omega \in \mathcal{G} \subset \Omega_{R_0}$, if

$$\forall i \in \{0, \ldots, \ell - 1\}, \omega_{(i,0)+R(1,\ell)} \neq \mathbf{1}, \omega_{(0,i)+R(\ell,1)} \neq \mathbf{1}. \tag{5.11}$$

That is a renormalised site is good, if each of its rows and columns contain at least one empty site. The idea behind (5.10) and (5.11) is that the following statements hold (recall the BP closure $[\cdot]$ from (3.2)):

- $\mathcal{SG} \subset \{\omega \in \Omega_{R_0} : [\omega] = \mathbf{0}_{R_0}\}$;
- for any $\mathbf{x} \in \{\mathbf{e}_1, -\mathbf{e}_1, \mathbf{e}_2, -\mathbf{e}_2\}$, if $\omega_{R_0} \in \mathcal{SG}$ and $\omega_{R_{\mathbf{x}}} \in \mathcal{G}$, then $[\omega_{R_0 \cup R_{\mathbf{x}}}] = \mathbf{0}$;
- if $\omega_{R_0} \in \mathcal{SG}$, $\omega_{R_{((-1,0)}} \in \mathcal{G}$ and $\omega_{R_{(0,-1)}} \in \mathcal{G}$, then $[\omega_{[-\ell,\ell)^2}] = \mathbf{0}$.

The reader is invited to verify these deterministic BP claims. They will allow us to transport a super good renormalised site along a good path.

We now have all the ingredients necessary to deal with FA-2f, following [161] to prove the following bound greatly improving on the one from Corollary 5.2.

Theorem 5.7 (FA-2f upper bound up to logarithmic corrections) *For FA-2f in $d \geqslant 2$ dimensions there exists a constant $C > 0$ such that for any q small enough,*

$$T_{\mathrm{rel}} \leqslant \exp \left(\frac{C(\log(1/q))^2}{q^{1/(d-1)}} \right).$$

Proof (Sketch) The reasoning proceeds in three steps. We first apply Proposition 5.6 as described above. We then need to bound the generic term in the right hand side

of (5.9), say for $\mathbf{x} = 0$. The occurrence of the event $\Gamma_0^{p_2^{-2}}$ guarantees good paths of length at most p_2^{-2}.

We can view a good path as a one-dimensional general KCM with FA-1f constraint, as covered by Proposition 4.9. Roughly speaking, this enables us to transform the generic term into a sum of terms of the form

$$\mu \left(\mathbb{1}_{\omega_{R_{\mathbf{y}}} \in \mathcal{SG}} \, \mathrm{Var}_{R_{\mathbf{x}}}(f|\mathcal{G}) \right) \tag{5.12}$$

for $\mathbf{x}$ and $\mathbf{y}$ neighbours along the good path. The prefactor incurred in this transformation is at most p_2^{-C} for some constant $C > 0$. Note that (5.12) only features a short range constraint.

The final step is to bound (5.12) by the Dirichlet form of FA-2f in $R_{\mathbf{x}} \cup R_{\mathbf{y}}$ with $\mathbf{1}$ boundary condition. This can be done using canonical paths. If done brutally, using the legal paths of Section 3.2, this would lead to a prefactor of order $q^{-2\ell^2}$. However, proceeding a little more carefully (only creating one or two empty columns/rows and moving them rather than emptying the entire rectangles), one can easily reduce this cost to $q^{-2\ell}$. Putting everything together and computing $p_2 \geqslant q^{4\ell}$, we obtain the desired result (see [155, 161] for more details). □

The three steps in the last proof can be understood as follows. The first step (applying Proposition 5.6) allows us to reduce the study of the infinite volume KCM to one in a large but finite volume ($1/p_2^2$, which is exponential in $1/q$) containing sufficient empty sites to efficiently empty the origin. The second step goes from the 'global' scale p_2^{-2} to the much smaller 'mesoscopic' scale ℓ. Finally the third step goes from the mesoscopic to the scale of the lattice. In the proof of Theorem 5.7, we used a generalised FA-1f dynamics for the second and third steps, as well as a very simple choice of $\mathcal{SG}$ event. In order to go further, we will need to reconsider all these choices.

5.3 Sharp threshold

Our next goal is a much stronger result that Theorem 5.7 providing a sharp threshold for FA-2f. This is the most precise result for KCM of this type and is due to Hartarsky, Martinelli and Toninelli [120, Theorem 1.3] (also see that reference for a quantitative bound on the second order correction in two dimensions).

Theorem 5.8 (Sharp threshold for FA-2f) *For FA-2f in $d \geqslant 2$ dimensions,*

$$q \log \tau_0 \to d \cdot \lambda(d, 2) \tag{5.13}$$

in $\mathbb{P}_{\mu_q}$-probability, as $q \to 0$, where $\lambda(d, 2)$ is the constant from Theorem 3.1 for j-neighbour BP. Furthermore, $q \log \mathbb{E}_{\mu_q}(\tau_0) \to d \cdot \lambda(d, 2)$.

In other words, for FA-2f we have $\tau_0 = (\tau_0^{\mathrm{BP}})^{d+o(1)}$ as $q \to 0$.

Before moving on to the proof, let us mention that it would be good to improve the statement above to a relaxation time result (the lower bound follows from (3.25) and (5.13)).

Conjecture 5.9 (FA-2f relaxation time) For FA-2f in $d \geqslant 2$, as $q \to 0$, we have

$$q \log T_{\mathrm{rel}} \to d \cdot \lambda(d, 2).$$

5.3.1 Lower bound: combinatorial bottleneck

We start by discussing the lower bound of (5.13). It is a relatively simple consequence of known results in BP and generalises well to other models. Recall that from (3.10) (also recall (3.25)) and Theorem 3.1, we have $\liminf_{q \to 0} q \log \tau_0 \geqslant \lambda(d, 2)$, but we would like to improve this bound. The heuristics behind the improvement is the following combinatorial bottleneck (see [120, Section 2] for more details).

Proof (Sketch of the lower bound of Theorem 5.8) According to Theorem 3.1 and Corollary 3.7, with high probability, the origin cannot be emptied using only empty sites within distance, say, $1/q^3$ from it in the initial configuration. We may therefore consider the first time τ, when the origin can be emptied using only empty sites at distance at most $1/q^3$ from it. That is, setting $\ell = \lceil 1/q^3 \rceil$ and $\Lambda = [-\ell, \ell]^d$, we define

$$\tau = \inf \left\{ t \geqslant 0 : [\omega_\Lambda(t)]^{\mathbf{1}_{\mathbb{Z}^d \setminus \Lambda}} \right\},$$

where $\omega_\Lambda(t)$ is the restriction of the stationary FA-2f process to Λ at time t and we recall (3.2) and (3.4).

The crucial observation is that in $\omega_\Lambda(\tau)$, there is a site $\mathbf{x}$ at the boundary of Λ such that

$$[\omega_\Lambda(\tau)]_0^{\mathbf{1}_{\mathbb{Z}^d \setminus \Lambda}} = 1 \neq \left[\omega_\Lambda^{\mathbf{x}}(\tau)\right]_0^{\mathbf{1}_{\mathbb{Z}^d \setminus \Lambda}} . \tag{5.14}$$

That is, at τ a site at the boundary of Λ becomes empty and this is essential (pivotal in the percolation jargon) to being able to empty the origin inside the box Λ.

Since we are working with the stationary process, we may perform a union bound over the attempted updates in Λ as in Section 4.2.1.2. Thus, it suffices to show that for any $\varepsilon > 0$ and q small enough,

$$\mu_q(\mathcal{A}) \leqslant \exp\left(\frac{-d \cdot \lambda(d, 2) + \varepsilon}{q}\right), \tag{5.15}$$

where $\mathcal{A}$ is the union over $\mathbf{x} \in \partial\Lambda$ of the event in (5.14).

Note that (5.15) only makes reference to 2-neighbour BP, so we have successfully reduced the problem for FA-2f to its BP counterpart. The bound (5.15) is indeed known in BP and is, in fact, the main step in the proof of Theorem 3.1 for this update family. □

In rough terms, we harnessed the fact that, in order for the origin to be 'locally emptiable', a certain 'critical droplet' (similar to the $\mathcal{SG}$ event in Section 5.2.2) needs to be present at the origin at some point in time. We then plugged this combinatorial bottleneck into the standard bound of Section 4.2.1.2.

5.3.2 Coalescing and branching simple symmetric exclusion process

We require one more ingredient as preparation for the upper bound in Theorem 5.8. It comes in the form of a model strongly related to the generalised FA-1f of Section 4.4, but not belonging to the class of KCM we defined in Section 2.3. For the sake of simplicity, we introduce only the binary (non-generalised) model and refer to [119] for the generalised version.

Let $G = (V, E)$ be the box $V = \{1, \ldots, \ell\}^d$ with its usual graph structure, where ℓ is a positive integer. We consider the state space $\Omega = \{0, 1\}^V$ as usual. We define $\Omega_+ = \Omega \setminus \{\mathbf{1}\}$ to be the event that there exists at least one empty site. Similarly, for any edge $e = \{x, y\} \in E$ we refer to $(\omega_x, \omega_y) \in \{0, 1\}^{\{x,y\}}$ as the state of e in ω and write $E_e = \{\omega \in \Omega : \omega_x \omega_y = 0\}$ for the event that e is not occupied (at least one of its vertices is empty). Given $p \in (0, 1)$, let $\pi = \bigotimes_{x \in V} \pi_x$ be the product Bernoulli measure in which each vertex is empty with probability p and let $\mu(\cdot) := \pi(\cdot | \Omega_+)$. Given an edge $e = \{x, y\}$, we write $\pi_e := \pi_x \otimes \pi_y$.

CBSEP is a continuous time Markov chain on Ω_+ for which the state of any edge $e \in E$ such that E_e occurs is resampled with rate one w.r.t. $\pi_e(\cdot | E_e)$. Thus, any edge containing exactly one empty site moves the empty site to the other endpoint of the edge (the *SEP move*) with rate $(1 - p)/(2 - p)$ and creates an extra empty site at the occupied endpoint (the *branching move*) with rate $p/(2 - p)$. Moreover, any edge containing two empty sites occupies one of the two chosen uniformly (the *coalescing move*) with rate $2(1 - p)/(2 - p)$. The chain is readily seen to be reversible with respect to μ and ergodic on Ω_+, because it can reach the configuration with all sites empty. If $c(\omega, \omega')$ denotes the jump rate from ω to ω', the Dirichlet form $\mathcal{D}^{\mathrm{CBSEP}}(f)$ of the chain has the expression

$$\mathcal{D}^{\mathrm{CBSEP}}(f) = \frac{1}{2} \sum_{\omega, \omega' \in \Omega_+} \mu(\omega) c(\omega, \omega') \left(f(\omega') - f(\omega) \right)^2 \qquad (5.16)$$

$$= \sum_{e \in E} \mu(\mathbb{1}_{E_e} \mathrm{Var}_{\pi_e}(f | E_e)).$$

Notice that the branching and coalescing moves of CBSEP are exactly the moves allowed in FA-1f (recall Section 4.1). Moreover, the SEP move can be reconstructed by a branching and a coalescing move. This leads to a comparison between the corresponding Dirichlet forms (see e.g. [144, Section 13.4]).

Although the two models are clearly closely related, we stress that CBSEP has many advantages over FA-1f, making its study simpler. Most notably, CBSEP is *attractive* in the sense of interacting particle systems, that is the natural stochastic order

is preserved by the dynamics (see [146, Section III.2] for background). Furthermore, it is natural to embed in CBSEP a continuous time random walk $(W_t)_{t \geqslant 0}$ on G such that CBSEP has an empty site at W_t for all $t \geqslant 0$. This feature is challenging to reproduce for FA-1f [29]. Finally, it is possible to move an empty site in CBSEP without creating more empty sites, contrary to what is the case in FA-1f (recall Section 4.1).

The main result on CBSEP we will need is the following [120, Proposition 5.2].[1]

Proposition 5.10 (CBSEP relaxation time) *Assume that $d \geqslant 2$ and consider a sequence of box sizes ℓ_n and parameters p_n such that $p_n \ell_n^d \to \infty$ and $p_n \to 0$. Then, for some $C > 0$ and all n large enough,*

$$T_{\mathrm{rel}}^{\mathrm{CBSEP}} \leqslant \frac{C \log^3(1/p_n)}{p_n}.$$

Proof (Sketch) The first step of the proof is to renormalise the model by considering boxes of volume approximately $1/p_n$. This brings us to treating CBSEP with parameter p of order 1 on arbitrary volume and CBSEP on a box of volume approximately $1/p_n$. The former relaxation time is uniformly bounded (see [120, Lemma B.1] and [119, Theorem 1] for details), while the latter is bounded by $C \log^2(1/p_n)/p_n$ thanks to [119, Corollary 3.1].

The proof of the first upper bound follows from the fact that $T_{\mathrm{rel}}^{\mathrm{FA-1f}}$ is bounded for q bounded away from 0 (recall Theorem 3.10) together with a comparison between the Dirichlet forms of CBSEP and FA-1f. The second upper bound may be proved, using canonical paths along the lines of Theorem 4.1 (also recall Conjecture 4.3, whose upper bound is known from [44] for $d = 2$, and its analogue for $d \geqslant 3$ from [184]), see [119, Proposition 4.6].

While this concludes the sketch for CBSEP, its generalised version is more subtle to analyse. Indeed, in [119, Theorem 2] it is shown that one can bound the mixing time (and therefore the relaxation time) of generalised CBSEP, using the mixing time of CBSEP and the cover time of the continuous time simple random walk on the box of interest. The cover time is classically bounded by $\log^2(1/p_n)/p_n$ [144, Chapter 11], but more work is needed to bound the mixing time of CBSEP, as opposed to its relaxation time. The approach of [119] is to prove a logarithmic Sobolev inequality, which classically bounds the mixing time. This is achieved, using deep input from [6, 143]. An alternative approach to bounding the mixing time of FA-1f can be found in [174, 175]. □

---

[1] We record a mistake in [120] leading to the weaker bound stated here. Indeed, in the last but one sentence of the proof of [120, Proposition 5.2] in Appendix B, the bounds on the cover time of the random walk and logarithmic Sobolev constant of CBSEP are not correctly imported from [144, Chapter 11] and [119, Corollary 3.2]. This mistake has no impact on the rest of the paper.

5.3.3 Upper bound

We are now ready to discuss the proof of the upper bound of Theorem 5.8 in dimension $d = 2$, following [120]. On the high level, the proof resembles Section 5.2.2. The first two steps bringing us to a mesoscopic scale are quite general, while reaching the mesoscopic scale from the microscopic one is much more delicate and specific to FA-2f.

5.3.3.1 Reduction to the mesoscopic scale via CBSEP

Recall from Theorem 3.1 that $\lambda(2, 2) = \pi^2/18$. Fixing some $\varepsilon > 0$ and setting

$$t_* = \exp\left(\frac{\pi^2 + \varepsilon}{9q}\right), \tag{5.17}$$

our goal is to prove that

$$\lim_{q \to 0} \mathbb{P}_\mu\left(\tau_0 > t_*\right) = 0. \tag{5.18}$$

We next use finite speed of propagation (recall (3.19)) to show that it is unlikely that, before time t_*, the state of the origin is influenced by the configuration outside the box $B = [-2t_*, 2t_*]^2$. Thus, it suffices to prove (5.18) for FA-2f with boundary condition $\mathbf{1}_{\mathbb{Z}^2 \setminus B}$.

The next step is a renormalisation to the generalised CBSEP model discussed in Section 5.3.2. The renormalisation itself resembles the one of Section 5.2.2, but the choice of super good events is much more delicate. We fix a mesoscopic scale $\ell = \lceil q^{-9} \rceil$ (the exponent is fairly arbitrary) and divide B into boxes $Q_{\mathbf{x}} = \ell \mathbf{x} + R(\ell, \ell)$ (recall (5.1)) for $\mathbf{x} \in \bar{B} = [-\lceil 2t_*/\ell \rceil, \lceil 2t_*/\ell \rceil] \cap \mathbb{Z}^2$. Each renormalised site $\mathbf{x} \in \bar{B}$ will be in one of two states—good and super good. As in Section 5.2.2, $\mathbf{x} \in \bar{B}$ is *good*, if $\omega_{Q_{\mathbf{x}}}$ has at least one empty site on each row and column of $Q_{\mathbf{x}}$. We denote the corresponding event by $\mathcal{G}_{\mathbf{x}}$.

By a union bound, it holds that

$$1 - \mu\left(\bigcap_{\mathbf{x} \in \bar{B}} \mathcal{G}_{\mathbf{x}}\right) \leqslant 5t_*^2(1 - q)^\ell \leqslant 1/t_*^4,$$

owing to our choice of sufficiently large ℓ. Since this probability is so small, as in Section 4.2.1.2 (union bound over attempted updates, using stationarity), it is likely that all renormalised sites in $\bar{B}$ remain good at all times up to t_*.

We will choose the super good event $\mathcal{SG}_{\mathbf{x}}$ in such a way that

$$\mu\left(\mathcal{SG}_{\mathbf{x}}\right) \geqslant \exp\left(-\frac{\pi^2/9 + \varepsilon/2}{q}\right). \tag{5.19}$$

Comparing this with (5.17), we similarly compute that it is likely that, up to time t_*, at all times there is at least one super good renormalised site in $\bar{B}$. Thus, we therefore assume that

$$\mathcal{E} = \bigcap_{\mathbf{x} \in \bar{B}} \mathcal{G}_{\mathbf{x}} \cap \bigcup_{\mathbf{x} \in \bar{B}} \mathcal{SG}_{\mathbf{x}}$$

occurs for all $t \leqslant t_*$. The event $\mathcal{E}$ corresponds to Ω_+ in Section 5.3.2 and super good renormalised sites play the role of empty sites for CBSEP.

By a standard result (see [10, Theorem 2]) similar to (3.14) (also recall (2.10)), but taking into account that we require not exiting $\mathcal{E}$, in order to prove (5.18), it suffices to establish a Poincaré type inequality of the form

$$\frac{\mathcal{D}_{\mathbf{1}_{\mathbb{Z}^2 \setminus B}}(f)}{\mathrm{Var}(f|\mathcal{E})} \geqslant \exp\left(-\frac{\pi^2/9 + 2\varepsilon/3}{q} \right) \tag{5.20}$$

for any function $f : \Omega_B \to \mathbb{R}$ such that $f(\omega) = 0$, if $\omega \notin \mathcal{E}$. We are now ready to apply the Poincaré inequality for generalised CBSEP provided by Proposition 5.10. Recalling (5.19), this yields

$$\mathrm{Var}(f|\mathcal{E}) \leqslant Cq^{-3} \exp\left(-\frac{\pi^2/9 + \varepsilon/2}{q} \right) \sum_{\mathbf{x} \sim \mathbf{y}} \mu\left(\mathbb{1}_{\mathcal{SG}_{\mathbf{x},\mathbf{y}}} \mathrm{Var}_{Q_{\mathbf{x}} \cup Q_{\mathbf{y}}} \left(f | \mathcal{SG}_{\mathbf{x},\mathbf{y}} \right) | \mathcal{E} \right), \tag{5.21}$$

where $\mathcal{SG}_{\mathbf{x},\mathbf{y}} = \mathcal{SG}_{\mathbf{x}} \cup \mathcal{SG}_{\mathbf{y}}$ for neighbours $\mathbf{x} \sim \mathbf{y}$ in $\bar{B}$. Note that, if we had used FA-1f (recall Section 4.1) instead of CBSEP at this point, the exponent above would become roughly $2\pi^2/(9q)$ instead of $\pi^2/(9q)$, which is not enough for proving Theorem 5.8.

Notice that relating $\mathrm{Var}_{Q_{\mathbf{x}} \cup Q_{\mathbf{y}}}(f|\mathcal{SG}_{\mathbf{x},\mathbf{y}})$ to the terms of the Dirichlet form of FA-2f in $Q_{\mathbf{x}} \cup Q_{\mathbf{y}}$ (recall (2.7)) is essentially equivalent to establishing a Poincaré inequality for FA-2f on this mesoscopic volume, conditioned to remain in the event $\mathcal{SG}_{\mathbf{x},\mathbf{y}}$. Indeed, some simple but delicate manipulations (see [120, Claim 5.5]) allow deducing (5.20) from (5.21) and the Poincaré inequality

$$\mathrm{Var}_Q(f|\mathcal{SG}) \leqslant \gamma(Q) \sum_{\mathbf{z} \in Q} \mu\left(c_{\mathbf{z}}^{\mathbf{1}_{\mathbb{Z}^2 \setminus Q}} \mathrm{Var}_{\mathbf{z}}(f)|\mathcal{SG} \right), \tag{5.22}$$

$$\gamma(Q) \leqslant \exp\left(\frac{\varepsilon}{7q} \right)$$

with $Q = Q_{\mathbf{x}} \cup Q_{\mathbf{y}}$ and $\mathcal{SG} = \mathcal{SG}_{\mathbf{x},\mathbf{y}}$, which holds for all $f : \Omega_Q \to \mathbb{R}$. The next section is devoted to the proof of (5.22).

5.3.3.2 Mesoscopic Poincaré inequality: the matryoshka doll technique

The inequality (5.22), which we seek to establish, should be interpreted as stating that, once a critical droplet is present at a given location, it is rather easy to completely change the state there. Note that we have not yet specified our choice of super good

event beyond requiring (5.19) to hold. Our actual choice and the proof of (5.22) go hand in hand and follow a multi-scale renormalisation scheme which we refer to as the *matryoshka doll technique*. For the sake of simplicity of the presentation, we rather present the argument in a way that yields a sharp threshold for slight variants of the model. The additional technical difficulties are dealing with more boundary conditions in (5.22), requiring at least one empty site on every two consecutive lines and implementing a microscopic FA-1f dynamics on the boundary of a droplet. The interested reader can find these features in [120]. The model we will treat is the KCM corresponding to the modified two-neighbour update family obtained by removing the lowest two rules in Figure 2.1c. We refer to it as modified FA-2f and seek to prove (5.22) but with $\mathcal{SG}$ event satisfying (5.19) with $\lambda(2, 2)$ replaced by $\lambda' = \pi^2/6$, which is the correct sharp threshold constant for this model (the matching lower bound is proved exactly as described in Section 5.3.1).

The matryoshka doll technique requires us to keep track of several features of a droplet Λ in parallel:

- geometry, that is, the shape and size of the droplet;
- super good event $\mathcal{SG}(\Lambda) \subset \Omega_\Lambda$, which guarantees that the droplet's state can be resampled efficiently, conditionally on the super good event;
- occurrence probability, that is, $\mu_q(\mathcal{SG}(\Lambda))$;
- relaxation time, that is, the smallest constant $\gamma(\Lambda) \geqslant 1$ such that (5.22) holds.

Given the multi-scale nature of the technique, all of the above are defined or bounded recursively for a sequence of droplets $\Lambda^{(m)}$, starting from a single empty site and reaching the droplet Q from Section 5.3.3.1.

Geometry

In the case of (modified) FA-2f, droplets are simply rectangles. Let $m_0 = \lfloor 1/\sqrt{q} \rfloor$,

$$\ell_m = \begin{cases} m + 1 & m < m_0, \\ \left\lfloor \dfrac{e^{(m-m_0)\sqrt{q}}}{\sqrt{q}} \right\rfloor & m \geqslant m_0, \end{cases} \qquad \Lambda^{(m)} = R\left(\ell_{\lceil m/2 \rceil}, \ell_{\lfloor m/2 \rfloor}\right), \qquad (5.23)$$

for any $m \geqslant 0$. Thus, $\Lambda^{(m)}$ form a nested sequence of rectangular droplets, each second one being a square. At each step only the width or only the length (depending on parity) is increased. The corresponding length scales increase linearly up to $1/\sqrt{q}$ and then exponentially, reaching our final size of interest $\ell = 1/q^9$ after approximately $\frac{17 \log(1/q)}{2\sqrt{q}}$ steps. The exact choice of scales is not of fundamental importance, but some care is needed. The scale $1/\sqrt{q}$ is known to be relevant thanks to BP results (recall (3.7)), but it is also not crucial for our purposes.

Super good event

The definition of the super good event is guided by the idea that lower scale droplets should be allowed to move freely within a larger scale droplet. This is vital for obtaining an efficient relaxation mechanism (that is, a small Poincaré constant $\gamma(Q)$ in (5.22)). For any $\mathbf{x} \in \mathbb{Z}^2$, this intuition leads us to define $\mathcal{SG}(\{\mathbf{x}\}) = \{\omega_{\mathbf{x}} = 0\}$ (recall that $\Lambda^{(0)} = \{0\}$ is just a single site) and, for $m \geqslant 0$,

$$
\mathcal{SG}\left(\mathbf{x} + \Lambda^{(2m+1)}\right) = \bigcup_{s=0}^{\ell_{m+1}-\ell_m} \left(\mathcal{SG}\left((s,0) + \mathbf{x} + \Lambda^{(2m)}\right)\right.
$$
$$
\left. \cap \bigcap_{\substack{t \in \{0,\ldots,\ell_{m+1}-1\} \\ t \notin \{s,\ldots,s+\ell_m-1\}}} \left\{\omega \in \Omega_{\mathbf{x}+\Lambda^{(2m+1)}} : \omega_{(t,0)+\mathbf{x}+R(1,\ell_m)} \neq \mathbf{1}\right\}\right) \quad (5.24)
$$

and similarly for odd scales. In words, at each scale we require that some translate of the lower scale droplet occurs and that each of remaining rows or columns (depending on parity) contains at least one empty site. The definition is illustrated in Figure 5.2.
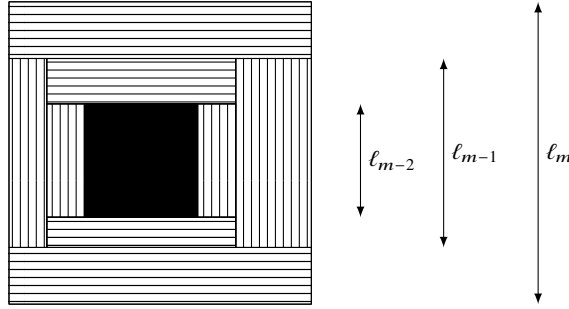

Fig. 5.2: An example structure of the super good event $\mathcal{SG}(\Lambda^{(2m)})$. Only the super good translates of $\Lambda^{(n)}$ for $m - n \in \{0,\ldots,6\}$ are shown. The hatched regions are required to contain at least one empty site per line in the direction of the hatching.

Occurrence probability

We next turn to lower bounding $\mu(\mathcal{SG}(\Lambda^{(m)}))$. The argument goes back to [2] and is only about BP, but we include it, as it is informative. Set $f : (0,\infty) \to (0,\infty) : x \mapsto -\log(1 - e^{-x})$, which is decreasing and convex. Systematically taking $s = 0$ in (5.24), we get that for any $m \in [0, 17\log(1/q)/\sqrt{q}]$,

$$\mu\left(\mathcal{SG}\left(\Lambda^{(m)}\right)\right) \geqslant \prod_{n=0}^{\lceil m/2 \rceil} \left(1 - (1-q)^{\ell_n}\right)^{2(\ell_{n+1} - \ell_n)}$$

$$\geqslant \exp\left(-2 \sum_{n=0}^{\lceil m/2 \rceil} (\ell_{n+1} - \ell_n) f(q\ell_n)\right)$$

$$= \exp\left(-\frac{2}{q} \int_0^\infty f + \frac{o(1)}{q}\right) = \exp\left(-\frac{\pi^2 + o(1)}{3q}\right),$$

where the asymptotic notation is as $q \to 0$. The error in the approximation of the Riemann sum by an integral above is carried out in more detail in [120, Appendix A], while the integral can be computed using the series expansion of $\log(1 - \cdot)$ and the fact that $\zeta(2) = \pi^2/6$. We have thus concluded the proof of the analogue of (5.19) in the context of modified FA-2f.

Relaxation time

Proving (5.22) is the most challenging part of [120]. The proof proceeds by establishing the recursive bound

$$\gamma\left(\Lambda^{(m)}\right) \leqslant e^{C \log^2(1/q)} \gamma\left(\Lambda^{(m-1)}\right). \tag{5.25}$$

for some constant $C > 0$. In turn, proving (5.25) is done via a version of the bisection technique, whose base step is unusually delicate.

For concreteness, we only discuss one parity. Fix some m such that $\ell_m \leqslant 9\log(1/q)/q$. Set $\Lambda' = \Lambda^{(2m+1)} = R(\ell_{m+1}, \ell_m)$ and $\Lambda = \Lambda^{(2m)} = R(\ell_m, \ell_m)$. The bisection will be used to decrease the width difference $\ell_{m+1} - \ell_m$ down to 1 in roughly $\log(\ell_{m+1} - \ell_m)$ steps. For translates of $R(l, \ell_m)$ with $l \in (\ell_m, \ell_{m+1})$, we extend (5.24) by replacing $\ell_{m+1} - \ell_m$ by $l - \ell_m$ and $\Lambda^{(2m+1)}$ by $R(l, \ell_m)$. We consider the rectangles $R^{(k)} = R(\ell_m + d_k, \ell_m)$ with $d_k = \lceil (2/3)^k (\ell_{m+1} - \ell_m) \rceil$ for $k \geqslant 0$. We then seek to prove a recursive bound of the form

$$\gamma\left(R^{(k)}\right) \leqslant a_k \gamma\left(R^{(k+1)}\right) \tag{5.26}$$

with $a_k \leqslant q^{-C}$ for some constant $C > 0$.

In order to prove (5.26), we use an auxiliary dynamics somewhat similar to the two-block one of Lemma 4.7, but using a different mechanism. Contrary to Lemma 4.7, which uses an East mechanism, the dynamics behind Lemma 5.11 below is non-oriented. It has three sites and has two types of constrained updates occurring at rate 1. The first update resamples the first and second sites (blocks) conditionally on some event occurring there before and after the update. The second update is similar for the second and third sites. The first two blocks together correspond to $R^{(k+1)}$, all three blocks form $R^{(k)}$ and the second and third blocks form a translate

of $R^{(k+1)}$ (see Figure 5.3). This leads us to the following lemma adapted from [120, Proposition 3.5].

Lemma 5.11 (Three-block dynamics) *There exists $C > 0$ such that the following holds. Let $(\mathbb{X}, \pi)$ be the product of three finite probability spaces $(\mathbb{X}_i, \pi_i)_{i=1}^3$. Fix some events $\mathcal{A}_1 \subset \mathbb{X}_1$, $\mathcal{A}_3 \subset \mathbb{X}_3$, $\mathcal{B}_{1,2} \subset \mathbb{X}_1 \times \mathbb{X}_2$ and $\mathcal{B}_{2,3} \subset \mathbb{X}_2 \times \mathbb{X}_3$. Set $\mathcal{H} = \mathcal{B}_{1,2} \times \mathcal{A}_3$ and $\mathcal{K} = \mathcal{A}_1 \times \mathcal{B}_{2,3}$. Consider the Dirichlet form*

$$\mathcal{D}(f) = \pi \left(\mathbb{1}_{\mathcal{H}} \operatorname{Var}(f | \mathcal{B}_{1,2} \times \{\omega_3\}) + \mathbb{1}_{\mathcal{K}} \operatorname{Var}(f | \{\omega_1\} \times \mathcal{B}_{2,3}) | \mathcal{H} \cup \mathcal{K} \right)$$

defined for any $f : \mathbb{X} \to \mathbb{R}$. Consider some event $C_2 \subset \mathbb{X}_2$ such that $\mathcal{A}_1 \times C_2 \times \mathcal{A}_3 \subset (\mathcal{H} \cap \mathcal{K})$. Then

$$\operatorname{Var}_{1,2,3}(f | \mathcal{H} \cup \mathcal{K}) \leqslant C \frac{\pi_{1,2}(\mathcal{B}_{1,2}) \pi_{2,3}(\mathcal{B}_{2,3})}{\pi_1(\mathcal{A}_1)(\pi_2(C_2))^2 \pi_3(\mathcal{A}_3)} \mathcal{D}(f) \qquad (5.27)$$

Proof (Sketch) The mechanism behind the proof is the following. Consider two copies of the chain described above with Dirichlet form $\mathcal{D}$, coupled by attempting the same updates. Observe that the following sequence of update attempts guarantees that the two chains reach the same state. First attempt to resample sites 1 and 2 to so that $\mathcal{A}_1 \times C_2$ occurs. Then, before any other update is attempted, update sites 2 and 3 (after the first update, $\mathcal{K}$ necessarily occurs, regardless whether the first update attempt was successful) so that $C_2 \times \mathcal{A}_3$ occurs. Finally, again before any other update is attempted, update sites 1 and 2. Estimating the rate at which such a sequence of updates is attempted yields the desired result (see [120, Proposition 3.5] for more detail). $\qquad \square$
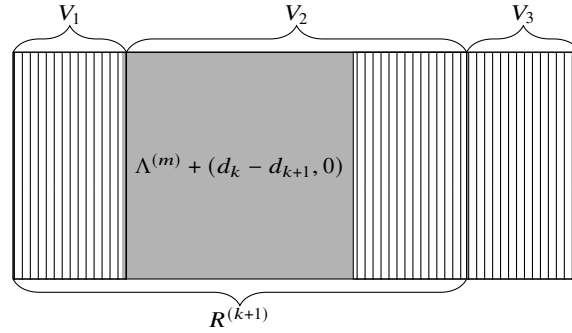

Fig. 5.3: The partition of $R^{(k)}$ into the rectangles V_1, V_2, V_3.

In order to prove (5.26), using Lemma 5.11, we consider the blocks V_1, V_2, V_3 depicted in Figure 5.3. The events $\mathcal{A}_1$ and $\mathcal{A}_3$ require each column to contain at least one empty site. The events $\mathcal{B}_{1,2}$, $\mathcal{B}_{2,3}$ and C_2 are the super good events for $V_1 \cup V_2$, $V_2 \cup V_3$ and V_2, which were already defined. In order to show that the

fraction in (5.27) is at most q^{-C} for some $C > 0$, it remains to observe that

$$\pi_1(\mathcal{A}_1)\pi_2(\mathcal{C}_2) \geqslant \pi_{1,2}(\mathcal{B}_{1,2})/d_k. \tag{5.28}$$

The last inequality itself follows from the observation that the event in (5.24) for each s has equal probability and the position depicted in Figure 5.3 guarantees the occurrence of $\mathcal{H} \cap \mathcal{K}$.

   The above proves (5.26) as desired. However, we are not done proving (5.25). Indeed, iterating (5.26) yields

$$\gamma\left(\Lambda^{(2m+1)}\right) \leqslant e^{C\log^2(1/q)}\gamma\left(R\left(\ell_m + 1, \ell_m\right)\right) \tag{5.29}$$

for some $C > 0$, the rectangle on the right being the thinnest rectangle $R^{(k)}$ we can obtain, which is one column wider than the desired $\Lambda^{(2m)}$. The reason for this is visible in Figure 5.3. Indeed, if $d_k = 1$, we cannot fit a translate of $\Lambda^{(2m)}$ in V_2 in such a way that both V_1 and V_3 remain non-empty. Thus, this base case requires a separate argument contained in [120, Proposition 3.7, Lemma 4.10], which we briefly discuss next.

   The proof that

$$\gamma\left(R(\ell_m + 1, \ell_m)\right) \leqslant q^{-C}\gamma\left(\Lambda^{(2m)}\right) \tag{5.30}$$

for a constant $C > 0$ proceeds along somewhat similar lines to (5.26), but is more subtle. We do use a decomposition as in Figure 5.3 with V_1 and V_3 consisting of a single column. In order to take into account the fact that the remaining V_2 is slightly smaller than $\Lambda^{(2m)}$, we define a "contracted" version of $\mathcal{SG}(\Lambda^{(2m)})$ which requires
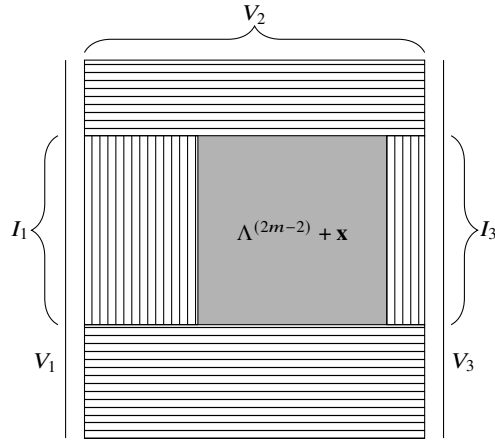

Fig. 5.4: Partition of $R(\ell_m + 1, \ell_m)$ into the rectangle V_2 and the lines V_1 and V_3. The internal structure of the contracted super good event for V_2 is shown. The empty sites in V_1 and V_3 need to be in I_1 and I_3 respectively in order to match it.

as much of its internal structure as it is possible to fit into V_2. That is, we require the super good translate of $\Lambda^{(2m-2)}$ (assuming $m \neq 0$) and lines containing at least one empty site within V_2, mimicking Figure 5.2. This contracted super good event plays the role of C_2 in an appropriate analogue of Lemma 5.11 (the events $\mathcal{A}_1$ and $\mathcal{A}_3$ now need to fit well with the position of the super good translate of $\Lambda^{(2m-2)}$, see Figure 5.4). Finally, using a double iteration of (5.28), we bound the resulting factor as desired, since the position of the super good translate of $\Lambda^{(2m-2)}$ is somewhat uniform inside $\Lambda^{(2m)}$.

Putting (5.29) and (5.30) together, we obtain (5.25). In turn, iterating (5.25) and a trivial bound at scale $m = 0$ gives (5.22) with $\gamma(Q) \leqslant \exp(o(1/q))$, as $q \to 0$, as desired. This concludes the sketch of Theorem 5.8.

5.4 Conclusion

In conclusion, let us summarise the new techniques presented in the present chapter, recalling Section 3.6 and 4.5.

Combinatorial bottlenecks. Thanks to this method, we were able to prove precise lower bounds. The approach of Section 5.3.1 is quite general. It is based on BP results, but more precise than the BP lower bound we saw in (3.10). The intuition behind it is that a critical droplet needs to visit the origin in order to change its state, which does not happen before the time given by the inverse probability of a droplet.

Bisection. In Section 5.1 we saw that the bisection technique can also be applied in higher dimensions by alternating two-block steps in different directions.

Long range renormalisation. This is a new technique in our repertoire. It allows us to focus on a one-dimensional path leading a critical droplet to the origin. This effectively reduces the problem to treating the movement of a critical droplet on its own scale, rather than a much larger one. It has the advantage of yielding upper bounds on T_{rel}, but the one-dimensional character of its path does not allow obtaining sharp thresholds.

Matryoska dolls. This is our state of the art technique for proving sharp upper bounds. It may be viewed as a very adaptable multi-scale renormalisation scheme, whose flexibility will be unleashed in the next chapter. In a sense, bisection in higher dimensions is also an instance of matryoshka dolls. The idea of the method is to recursively prove Poincaré inequalities on progressively larger and larger droplets. At each step, we have full freedom in choosing the auxiliary dynamics. Once a sufficiently large scale is reached via the matryoshka dolls, we conclude by a single step renormalisation, whose auxiliary global dynamics may also be adapted to our needs. A major advance of this technique is that it allows us to prove very precise results incorporating tailored relaxation mechanisms, completely bypassing the need to build any explicit canonical paths. While for modified FA-2f canonical paths reflecting the multi-scale structure of droplets could be envisioned, such approaches very quickly go out of hand in more general settings such as the ones investigated in the next chapter.

Auxiliary dynamics. Let us review the auxiliary dynamics we have plugged into the matryoshka dolls technique so far. In Section 5.1, we always used the simple two-block dynamics of Lemma 4.7. In Section 5.2.2, we employed a one-dimensional FA-1f global dynamics. In Section 5.3.2 we developed a new possible global dynamics given by the CBSEP that was applied to FA-2f in Section 5.3.3.1. Finally, in Section 5.3.3.2, we developed a three-block alternative, Lemma 5.11, to the two-block Lemma 4.7. The three-block dynamics was used repeatedly in one direction and we then switched direction repeatedly, until reaching the desired mesoscopic scale. The key feature of CBSEP and the three-block lemma is that they allow us to move a very unlikely droplet without paying the price of its creation from scratch, but rather just a little bit of internal reshuffling.

Chapter 6
Universality

In this chapter, we consider KCM with completely arbitrary update families. The goal of universality is to identify what kinds of behaviour are possible within this vast collection of models and to efficiently classify all update families with respect to their behaviour. We begin by considering one-dimensional models, which feature three rough universality classes, represented by the three nearest neighbour models studied in Chapter 4. We then provide some background on the two-dimensional BP universality theory, before moving on to two-dimensional KCM. We present the complete rough and refined universality theory for KCM in two dimensions and cover the essential elements of their proofs, building on the techniques exhibited in Chapters 4 and 5.

6.1 KCM universality in one dimension

Let us warm up by considering one-dimensional KCM. We recall that in Theorem 4.1 and 4.4, we saw that the correct scaling of T_{rel} (or τ_0) for FA-1f is $1/q^3$, while it is $\exp((\log(1/q))^2/(2\log 2))$ for the East KCM. For FA-2f, we have $q_c = 1$, as discussed in Section 4.3. The outcome of the universality theory in one dimension is that the only possible behaviours are those of FA-1f, East and FA-2f.

In order to state the universality result, we need to define stable directions, which will determine the universality class.

Definition 6.1 (Stable directions) Fix a one-dimensional update family $\mathcal{U}$. We say that the positive (resp. negative) direction is *unstable*, if there exists $U \in \mathcal{U}$ such that $U \subset \{-1, -2, \dots\}$ (resp. $\{1, 2, \dots\}$) and *stable* otherwise.

Indeed, it is not hard to check that FA-1f, East and FA-2f have zero, one and two stable directions respectively. The stability of a direction governs whether empty sites can reproduce in that direction.

The following result is the rough universality classification in one dimension due to Marêché, Martinelli, Morris and Toninelli [151, 153, 155].

Theorem 6.2 (One-dimensional rough universality) *For a one-dimensional KCM with update family $\mathcal{U}$ we have that*

- *if $\mathcal{U}$ has two unstable directions, then $q_c = 0$ and, for some $C > 0$,*

$$\lim_{q \to 0} \mathbb{P}_\mu \left(1/C \leqslant \frac{\log \tau_0}{\log(1/q)} \leqslant C \right) = 1;$$

- *if $\mathcal{U}$ has one unstable direction, then $q_c = 0$ and, for some $C > 0$,*

$$\lim_{q \to 0} \mathbb{P}_\mu \left(1/C \leqslant \frac{\log \tau_0}{\log^2(1/q)} \leqslant C \right) = 1;$$

- *if $\mathcal{U}$ has no unstable direction, then $q_c = 1$.*

Remark 6.3 The asymptotics in Theorem 6.2 also hold for T_{rel} instead of τ_0.

The proof of this result follows the same lines as what we have already seen, so, rather than repeating it, we only provide the appropriate pointers. The lower bound for update families with two unstable directions follows from the fact that there is typically no empty site at distance much less than $1/q$ from the origin. The corresponding upper bound follows from (3.14), Proposition 4.9 and renormalisation (empty renormalised sites correspond to completely empty intervals of sites with large fixed length). The upper bound for update families with one unstable direction follows from Theorem 4.8. The corresponding lower bound is proved as in Section 4.2.1, again with a renormalisation in order to apply the combinatorial Proposition 4.6, considering a renormalised site empty if it contains at least one empty site (not to be confused with the upper bound renormalisation above). Finally, the result for update families with no unstable directions is immediate, since the state of any sufficiently long interval of occupied sites can never be modified.

It would be interesting to know the sharp asymptotics of $\log \tau_0$ for general $\mathcal{U}$, as in the case of FA-1f and East, but this matter is still open. It is good to note that the corresponding problem for BP is easy (see [113, Proposition 1.3.4]).

6.2 BP universality in two dimensions

Before we move on to two-dimensional universality for KCM, we require some background on the side of BP. Since our focus is on KCM, we take these BP results for granted and refer the interested reader to [165] for a detailed survey of the methods involved.

6.2.1 Rough universality in BP

We start by generalising the definition of stable directions. We use $\| \cdot \|$ and $\langle \cdot, \cdot \rangle$ to denote the Euclidean norm and scalar product. Let $S^1 = \{\mathbf{u} \in \mathbb{R}^2 : \|\mathbf{u}\| = 1\}$ be the unit circle. We call the elements of S^1 *directions*. We consider the open half-plane with outer normal $\mathbf{u} \in S^1$

$$\mathbb{H}_{\mathbf{u}} = \left\{ \mathbf{x} \in \mathbb{R}^2 : \langle \mathbf{x}, \mathbf{u} \rangle < 0 \right\}. \tag{6.1}$$

Definition 6.4 (Stable directions) Fix an update family $\mathcal{U}$. We say that $\mathbf{u} \in S^1$ is *unstable*, if there exists $U \in \mathcal{U}$ such that $U \subset \mathbb{H}_{\mathbf{u}}$ and *stable* otherwise. A direction is called *strongly stable*, if it belongs to the interior of the set of stable directions. A stable direction $\mathbf{u} \in S^1$ is *isolated stable*, if there exists an open interval $I \subset S^1$ such that the only stable direction in I is $\mathbf{u}$.

The stable directions of the models in Figure 2.1 are given in Figure 6.1. Stable directions allow us to define the rough universality classes in two dimensions.

Definition 6.5 (Rough universality partition) Let $C = \{\mathbb{H}_{\mathbf{u}} \cap S^1 : \mathbf{u} \in S^1\}$ denote the set of open semicircles of S^1. An update family $\mathcal{U}$ is:

• *supercritical* if there exists $C \in C$ containing no stable direction. If additionally

  – there exist two non-opposite stable directions, $\mathcal{U}$ is *rooted*;
  – there do not exist two non-opposite stable directions, $\mathcal{U}$ is *unrooted*.

• *critical* if every $C \in C$ contains a stable direction and there exists $C \in C$ containing finitely many stable directions.

• *subcritical* if every $C \in C$ contains infinitely many stable directions. If additionally

  – there exists an unstable direction, $\mathcal{U}$ is *nontrivial*;
  – all directions are stable, $\mathcal{U}$ is *trivial*.

Comparing Definition 6.5 with the one-dimensional case of Theorem 6.2, we see that the new rough universality classes in two dimensions are the critical and subcritical nontrivial ones. The following rough universality theorem for BP is due to Balister, Bollobás, Przykucki, Smith and Uzzell [16, 40].

Theorem 6.6 (Two-dimensional rough universality for BP) *Let $\mathcal{U}$ be a two-dimensional update family. If $\mathcal{U}$ is*

• *supercritical, then $q_c = 0$ and, for some $C > 0$,*

$$\lim_{q \to 0} \mu \left(1/C \leqslant \frac{\log \tau_0^{\mathrm{BP}}}{\log(1/q)} \leqslant C \right) = 1;$$

(a) East (Figure 2.1a)   (b) FA-1f (Figure 2.1b)   (c) FA-2f (Figure 2.1c)

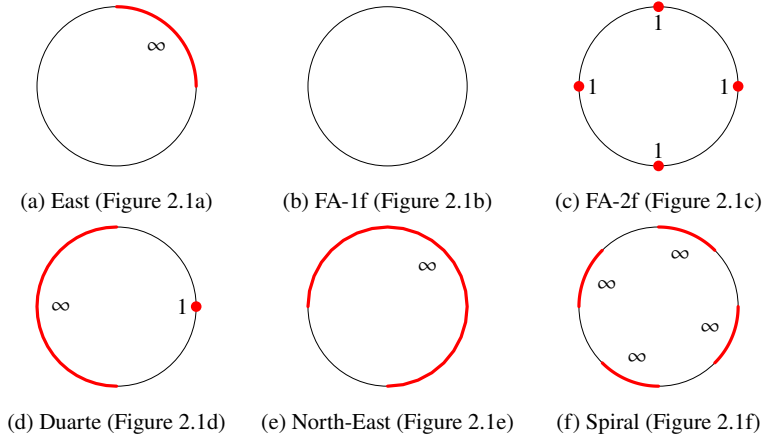(d) Duarte (Figure 2.1d)   (e) North-East (Figure 2.1e)   (f) Spiral (Figure 2.1f)

Fig. 6.1: Stable directions of the two-dimensional update families from Figure 2.1.

- *critical, then $q_c = 0$ and, for some $C > 0$,*

$$\lim_{q \to 0} \mu \left(1/C \leqslant \frac{\log \log \tau_0^{\mathrm{BP}}}{\log(1/q)} \leqslant C \right) = 1;$$

- *subcritical nontrivial, then $0 < q_c < 1$.*
- *subcritical trivial, then $q_c = 1$.*

We note that a complete generalisation of Theorem 6.6 has already been established in arbitrary dimension by Balister, Bollobás, Morris and Smith [13–15], but its KCM counterpart is still missing. The reader interested in subcritical models is also encouraged to consult [108, 112, 114, 122, 200] for different approaches to bounding q_c. While the behaviour of nontrivial subcritical models is very interesting, it is quite challenging and not much is known currently, so we discard them in the sequel. We further refer to [113, Section 1.5.1] for a detailed account of the history of universality in BP and KCM.

6.2.2 Refined universality in BP

The bounds on τ_0^{BP} for critical models in Theorem 6.6 are quite loose due to the iterated logarithm, particularly compared to results for specific models (recall Theorem 3.1). In two dimensions, it is possible to obtain much more precise asymptotics. In order to state them, we require a refinement of the notion of stable direction (see Figure 6.1 for examples). Recall that $|\cdot|$ is the number of empty sites and $[\cdot]$ is the closure from (3.2).

Definition 6.7 (Difficulty) The *difficulty* $\alpha(\mathbf{u})$ of $\mathbf{u} \in S^1$ is

- 0 if **u** is unstable;
- ∞ if **u** is stable, but not isolated stable;
- $\min\{|Z| : Z \subset \mathbb{Z}^2, |[\mathbf{0}_{\mathbb{H}_{\mathbf{u}} \cup Z} \cdot \mathbf{1}_{\mathbb{Z}^2 \setminus (\mathbb{H}_{\mathbf{u}} \cup Z)}]_{\mathbb{Z}^2 \setminus \mathbb{H}_{\mathbf{u}}}| = \infty\}$ otherwise.

The *difficulty* of $\mathcal{U}$ is

$$\alpha = \alpha(\mathcal{U}) = \min_{C \in \mathcal{C}} \max_{\mathbf{u} \in C} \alpha(\mathbf{u}).$$

We say that a direction $\mathbf{u} \in S^1$ is *hard* if $\alpha(\mathbf{u}) > \alpha$. We say that $\mathcal{U}$ is *unbalanced* if there exist two opposite hard directions and *balanced* otherwise.

In words, the difficulty of an isolated stable direction is the smallest number of empty sites needed to empty an infinite number of sites with the help of an empty half-plane with outer normal the direction. The difficulty of the update family is given by the easiest open semi-circle, a semi-circle being as hard as the hardest direction it contains. Comparing Definitions 6.5 and 6.7, it can be shown [39] that an update family is supercritical if $\alpha = 0$, subcritical if $\alpha = \infty$ and critical if $0 < \alpha < \infty$. Among the critical examples of Figure 2.1, FA-2f is balanced and Duarte is unbalanced, both having difficulty 1 (see Figure 6.1).

With these notions, the refined universality result for BP of Bollobás, Duminil-Copin, Morris and Smith [39] is the following.

Theorem 6.8 (Two-dimensional refined universality for BP) *Let $\mathcal{U}$ be a two-dimensional critical update family of difficulty α. Then, for some $C > 0$,*

$$\lim_{q \to 0} \mu \left(1/C \leqslant \frac{q^\alpha \log \tau_0^{\mathrm{BP}}}{(\log(1/q))^\gamma} \leqslant C \right) = 1,$$

where $\gamma = 0$, if $\mathcal{U}$ is balanced, and $\gamma = 2$, if $\mathcal{U}$ is unbalanced.

Naturally, Theorem 6.8 is consistent with Theorem 3.1 for 2-neighbour and Duarte BP.

6.3 KCM universality in two dimensions

6.3.1 Statement

With BP universality at hand, we may turn to KCM. For subcritical models, in view of Theorems 3.9 and 3.10 combined with the BP rough universality Theorem 6.6, we do not have anything new to say. We therefore focus on supercritical and critical models in the next result due to Marêché, Martinelli, Morris and Toninelli [151, 153, 155]. Recall Definition 6.5.

Theorem 6.9 (Two-dimensional rough universality for KCM) *For any two-dimensional KCM with update family $\mathcal{U}$ we have that*

- *if $\mathcal{U}$ is supercritical unrooted, then for some $C > 0$,*

$$\lim_{q \to 0} \mathbb{P}_\mu \left(1/C \leqslant \frac{\log \tau_0}{\log(1/q)} \leqslant C \right) = 1; \qquad (6.2)$$

- *if $\mathcal{U}$ is supercritical rooted, then for some $C > 0$,*

$$\lim_{q \to 0} \mathbb{P}_\mu \left(1/C \leqslant \frac{\log \tau_0}{\log^2(1/q)} \leqslant C \right) = 1; \qquad (6.3)$$

- *if $\mathcal{U}$ is critical, then for some $C > 0$,*

$$\lim_{q \to 0} \mathbb{P}_\mu \left(1/C \leqslant \frac{\log \log \tau_0}{\log(1/q)} \leqslant C \right) = 1. \qquad (6.4)$$

The same asymptotics hold for T_{rel} instead of τ_0.

The proof of Theorem 6.9 will be explained in Section 6.3.2, but before that, let us first discuss refined universality. Recalling Definition 6.7, we only need the following vocabulary in order to define the refined KCM universality classes.

Definition 6.10 (Further refined universality types) A critical two-dimensional update family $\mathcal{U}$ is *rooted*, if there exist two non-opposite hard directions and *unrooted* otherwise. We say that $\mathcal{U}$ is *semi-directed*, if there is exactly one hard direction and *isotropic* if there are no hard directions.

Notice that balanced unrooted update families are either semi-directed or isotropic. The above notions allow us to state the refined universality result obtained over a series of works of Hartarsky, Marêché, Martinelli, Morris and Toninelli [115–118, 155] (see [115] for the final step and a detailed discussion), also relying on the BP refined universality Theorem 6.8 of [39].

Theorem 6.11 (Two-dimensional refined universality for KCM) *Let $\mathcal{U}$ be a two-dimensional critical update family of difficulty α. Then for some $C > 0$,*

$$\lim_{q \to 0} \mathbb{P}_\mu \left(1/C \leqslant \frac{q^{\alpha \cdot \beta} \log \tau_0}{(\log(1/q))^\gamma (\log \log(1/q))^\delta} \leqslant C \right) = 1,$$

where the exponents β, γ and δ are given in Table 6.1, depending on whether $\mathcal{U}$ has finite or infinite number of stable directions; is balanced or unbalanced; is rooted or unrooted; is semi-directed or isotropic, if balanced and unrooted. In particular, $\beta = 2$, if $\mathcal{U}$ has a strongly stable direction, and $\beta = 1$ otherwise.

It is not hard to check that the seven refined universality classes in Theorem 6.11 do exhaust the critical rough universality class of Theorem 6.9. We emphasise that Theorem 6.11 is currently the sharpest result available for any critical KCM with the exception of FA-2f and slight variations thereof, namely modified FA-2f and the Froböse KCM (recall Section 5.3.3.2 and see [91, 107]). As in Theorem 5.8, the upper bounds in Theorem 6.11 are not known to hold for T_{rel}.

β, γ, δ	Infinite stable directions	Finite stable directions	
	Rooted	Rooted	Unrooted
Unbalanced	$2, 4, 0$	$1, 3, 0$	$1, 2, 0$
Balanced	$2, 0, 0$	$1, 1, 0$	S.-dir. $1, 0, 1$ Iso. $1, 0, 0$

Table 6.1: Refined universality exponents for KCM in Theorem 6.11.

Interestingly, an analogue of Theorem 6.11 is not available in one dimension. It would be good to fill this gap by solving the following problem.

Problem 6.12 (Unrooted scaling in one dimension) Let $\mathcal{U}$ be a one-dimensional update family with two unstable directions. Prove existence of and determine $\alpha \in (0, \infty)$ such that

$$\lim_{C \to \infty} \liminf_{q \to 0} \mathbb{P}_\mu \left(1/C \leqslant q^\alpha \tau_0 \leqslant C\right) = 1.$$

It should be noted that, while the analogous problem for BP is an exercise [113, Proposition 1.3.4], in view of Theorem 4.1, Problem 6.12 is not.

6.3.2 Rough universality proofs

We next outline the proof of Theorem 6.9. Firstly, the lower bound of (6.2) is immediate, since there is typically no empty site at distance much smaller than $q^{-1/2}$ from the origin. The lower bound in (6.4) follows from Theorem 6.6 together with (3.10) and (3.25). The upper bounds in (6.2) and (6.3) are proved using Proposition 4.9 and (3.14) together with a simple renormalisation, as for Theorem 6.2. Namely, each empty site corresponds to a suitably oriented rectangle whose sites are all empty. The exact shape of this rectangle is chosen based on the proof of BP rough universality, so that, if it is empty, it is able to reproduce an empty copy of itself (see [40, 155]). The upper bound in (6.4) proceeds like the proof of Theorem 5.7, but using a generalised one-dimensional KCM with East instead of FA-1f constraint, still covered by Proposition 4.9.

We are left with the lower bound in (6.3) which is the only one requiring additional ideas with respect to what we have already seen. The overall scheme remains the same—we seek to establish a combinatorial bottleneck akin to the one of Proposition 4.6 and convert it into the desired bound as in Section 4.2.1.2. However, contrary to the one-dimensional case of Theorem 6.2, the combinatorial bottleneck cannot be deduced from the one-dimensional Proposition 4.6 via renormalisation. We next present a sketch of the proof of the following result of Marêché [151, Theorem 4], which is the crucial ingredient.

Proposition 6.13 (Combinatorial bottleneck for rooted models) *Let $\mathcal{U}$ be a two-dimensional update family which is not supercritical unrooted. There exists an integer*

C > 0 such that the following holds for any integer $n \geqslant 1$. Consider the $\mathcal{U}$-KCM on $\Lambda_n = \{-Cn2^n, \ldots, Cn2^n\}^2$ with boundary condition $\mathbf{0}_{\mathbb{Z}^2 \setminus \Lambda}$. Let $V(n)$ be the set of all configurations that the process can reach from $\mathbf{1}_\Lambda$ via a legal path (recall Definition 3.4) in which all configurations contain at most n empty sites. Then $\omega_0 = 1$ for all $\omega \in V(n)$.
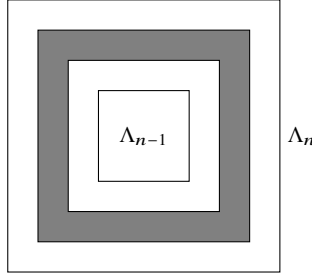

Fig. 6.2: Illustration of the proof of Proposition 6.13. The buffer zone is shaded.

Proof (Sketch) For simplicity, we focus on the two-dimensional East KCM (see Figure 2.1a). The proof proceeds by induction on n claiming that for any $\omega \in V(n)$, $|\omega_{\Lambda_{n-1}}| \leqslant n - 1$. Indeed, iterating this fact, we obtain that for any $\omega \in V(n)$, $|\omega_{\Lambda_1}| \leqslant 1$, so that the single empty site in Λ_1 cannot be at distance more than C from the boundary of Λ_1, so it cannot reach the origin.

In order to show the claim, using the reversibility of legal paths (recall Definition 3.4), we may instead prove that, if we start from $\omega \neq \mathbf{1}_{\Lambda_n}$ such that $\omega_{\Lambda_n \setminus \Lambda_{n-1}} = \mathbf{1}$ and never visit configurations with more than n empty sites, then we cannot reach the $\mathbf{1}_{\Lambda_n}$ configuration. To do this, we prove by a second induction, on the number of steps in the legal path, that the following two conditions remain valid. Firstly, a frame-shaped buffer zone around Λ_{n-1} with no empty site remains intact (see Figure 6.2). Secondly, there always remains an empty site in the internal region encircled by the buffer, so the dynamics cannot reach $\mathbf{1}_{\Lambda_n}$.

We know that so far an empty site remains trapped in the internal region encircled by the buffer, so we only have $n - 1$ empty sites available for disrupting the buffer from the outside, which is impossible by the induction hypothesis on n. Therefore, it suffices to show that we may not disrupt the buffer from the inside either. By projecting the two-dimensional East model on each axis it is clear that no empty site can enter the right and top parts of the buffer from the inside, and the projections of the top and rightmost empty sites in the region inside the buffer need to remain where they were initially. The left part of the buffer (and similarly for the bottom one) cannot be reached from the inside, because at least one empty site needs to remain as far right as the rightmost initial one was, so we only have $n - 1$ empty sites with which to reach the left part of the buffer, which is impossible by induction hypothesis on n. $\square$

With the above sketch in mind, we encourage the reader to consult the full proof in [151], which actually takes only four pages. We also remark that Proposition 6.13 and its proof generalise immediately to any dimension.

6.3.3 Refined universality proofs

6.3.3.1 Lower bounds

We start with the lower bounds in Theorem 6.11, since they are closely related to Proposition 6.13. We follow [116] and refer to that work for the formal proof. Somewhat surprisingly, all seven refined universality classes are governed by the same combinatorial bottleneck, but on different length scales and for different reasons. For the sake of concreteness, we focus on the model from Figure 6.3a, whose difficulty is $\alpha = 1$.

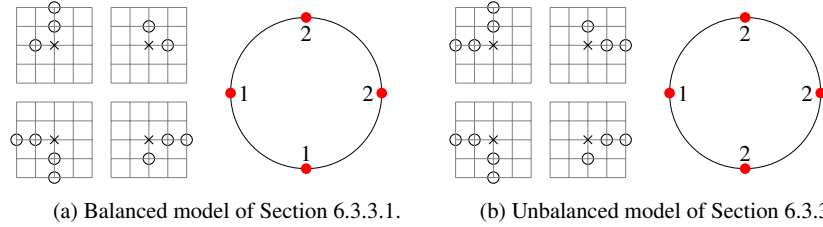(a) Balanced model of Section 6.3.3.1.   (b) Unbalanced model of Section 6.3.3.2.

Fig. 6.3: The update rules, stable directions and difficulties of two example critical rooted update families with difficulty $\alpha = 1$.

Morally speaking, in this model the smallest mobile entity ('droplet') is an empty square of size roughly $1/q$, similarly to FA-2f (recall Chapter 5). Indeed, typically, on its left and bottom sides, one can find an empty site, which allows it to empty the column of sites on its left and the row of sites below it. However, it is essentially impossible for the droplet to grow up or right, as this requires two consecutive empty sites and those are typically only available at distance $1/q^2$ from the droplet. We will only work in the box $\Lambda = \{-1/q^{7/4}, \ldots, 1/q^{7/4}\}^2$, so such couples of empty sites are not available for most columns and rows. Thus, we expect droplets to essentially follow the dynamics of the two-dimensional East KCM (see Figure 2.1a).

On a very high level, we proceed in the same way as in the proof of Proposition 6.13. However, there are several glaring problems in making the above reasoning rigorous. Firstly, much like in Figure 5.2, droplets can be more complex than empty squares of size $1/q$. Thus, one needs to identify an event which says whether a droplet is present and this event should be deterministically necessary for empty sites to spread. Moreover, the event should have probability of the correct order $\exp(-1/q)$, as suggested by the BP Theorem 6.8. It turns out that the notion of

'spanning' introduced in [39] for the proof of the lower bound of Theorem 6.8 for unbalanced models, following [52], is flexible enough for our purposes. Roughly speaking, a droplet (rectangle) is spanned if the empty sites present inside it are sufficient to empty a connected set touching all its sides (see [165, Section 8] for an overview of the BP result and its proof). We call a droplet critical if it has size roughly $1/q$.

If our goal were to only find a single critical droplet, we could proceed as in Section 5.3.1. Indeed, for isotropic and unbalanced unrooted models the proof of Section 5.3.1 or directly (3.10) combined with Theorem 3.1 suffices. However, for other refined universality classes, including the one of the model of Figure 6.3a, we need to work much harder. In order to obtain the exponent $\gamma = 1 > 0$ in Theorem 6.11, we need many droplets. Unfortunately, given a configuration, spanned critical droplets may overlap, so, in order to obtain good bounds on the probability of the configuration, we need to consider disjointly occurring ones (droplets occur disjointly, if they admit disjoint witness sets of empty sites). We may then define the number of spanned critical droplets as the maximal number of disjointly occurring ones. Considering the KCM on Λ with $\mathbf{0}_{\mathbb{Z}^2 \setminus \Lambda}$ boundary condition, starting from a typical initial configuration (which in particular does not contain more than three empty sites close to each other), our aim is to prove that, before we can empty the origin, we need to visit a configuration with at least order $\log(1/q)$ disjointly occurring critical droplets.

If droplets did follow a two-dimensional East dynamics exactly (up to renormalisation), this would follow from Proposition 6.13. But this is not the case. Indeed, by changing their internal structure, droplets may move a bit without creating another droplet, as we saw in Section 5.3.3. Worse yet, they are not really forbidden to move right or up, but simply are not likely to be able to do so wherever they want: it depends on the dynamic environment.

In order to handle these problems, we need the crucial notion of crossing. Consider a vertical strip S of width $1/q^{3/2}$ of our domain Λ. Roughly speaking, we say that S has a crossing if the following two events occur. Firstly, the empty sites in S together with the entire half-plane to the left of S are enough to infect a path from left to right in S. Secondly, S does not contain a spanned critical droplet. Notice that these two events have opposite monotonicity in the configuration. Employing BP tools, it can be shown (see [116, Appendix B]) that the probability of a crossing decays exponentially with the width of S at our scales of interest. In particular, the probability under μ that such a strip S is crossed is of order $\exp(-1/q^{3/2})$. Roughly speaking, the proof proceeds by splitting S into smaller strips which are either crossed by a single spanned droplet of subcritical size, or contain a pair of adjacent empty sites. Having appropriate bounds on the probability of spanned subcritical droplets as a function of their size, one may prove the desired bound on crossing by a union bound over the partition of S into the smaller strips. We note that in the case of infinite number of stable directions, bounding the probability of crossings is quite different [116, Appendix B.2], but remains possible.

Having established such a bound on the probability of crossings, we may incorporate them into the combinatorial bottleneck—we are satisfied if we visit either

a configuration with a crossed strip of width $1/q^{3/2}$, or with $\log(1/q)$ disjointly spanned critical droplets. The lack of crossings allows us to exclude the possibility of a droplet reaching the right side of the vertical strip S without help from the right of S, since the KCM dynamics can never infect more than what bootstrap percolation can (recall Section 3.2).

With these additional inputs, the proof scheme of Proposition 6.13 can be carried out to give the lower bounds in Theorem 6.11. The only difference between refined universality classes comes in the choice of length scales, in the bounds on the probability of spanning droplets and crossing strips and their proofs.

6.3.3.2 Upper bounds

We next turn to the upper bounds in Theorem 6.11. Contrary to the lower bounds, the proofs of upper bounds are highly dependent on the refined universality class. However, there are two classes, for which all the elements of the proof have already been discussed. The weakest upper bound corresponding to $\beta = 2, \gamma = 4, \delta = 0$ in Theorem 6.11, which applies to all models, but is only sharp for unbalanced critical families with infinite number of stable directions, in fact follows from the proof of the rough Theorem 6.9 mentioned in Section 6.3.2 (see [155]). At the other extreme, the upper bound for isotropic models ($\beta = 1, \gamma = 0, \delta = 0$) is proved similarly to Theorem 5.8, using the matryoshka doll technique (recall Section 5.3.3), up to some technical modifications (see [115, Section 5]).

Our next goal is to outline the proof of the upper bound of Theorem 6.11 with $\beta = 1, \gamma = 3, \delta = 0$, which applies to any critical KCM with finite number of stable directions, but is only sharp for unbalanced rooted families with finite number of stable directions. This will clarify the relevance of the absence of any strongly stable directions, which governs the value of the most important exponent β. For the sake of simplicity, we focus on the model depicted in Figure 6.3b.

Let us start with some heuristic considerations before explaining how they can be turned into the proof originating from [118]. Since we are interested in an upper bound, we may choose the notion of droplet in a simple way. Namely, a droplet D is (a translate of) an empty square frame of size $C \log(1/q)/q$ and thickness 2, where C is a suitably large constant (see Figure 6.4). Then typically (under μ), there is an empty site in the column to the left of D allowing us to empty $D - (1, 0)$. However, it is unlikely to find a pair of adjacent empty sites on any of the other sides of D. We conclude that it is easy for D to advance only to the left. An efficient way to perform this leftward motion is given by the legal path for the East KCM from Figure 4.1, where each empty site represents an empty translate of D.

The key idea is that it suffices to perform this East-like motion for a distance of order $1/q^2$, in order to find a pair of adjacent empty sites on the row above the droplet. Once the droplet reaches them, it is able to move one step up. It is then possible to revert the East path to bring the droplet to the original position, but shifted one lattice step up. This procedure effectively yields a step in the hard up direction. We may then iterate this idea, moving upwards in an East-like manner, where each step up is,

in fact, a long East-like path to the left and back. This way, we eventually reach a pair of adjacent empty sites allowing the droplet to move to the right, etc. See Figures 6.4 and 6.5 for an illustration of the mechanism. Based on the above heuristics, we expect droplets to be able to move freely in all directions by creating only about $\log(1/q)$ additional droplets at a time. Hence, we expect the time necessary for droplets to move to be of order $\rho^{-\log(1/q)}$, where $\rho \approx q^{8C\log(1/q)/q}$ is the probability of a single droplet under μ. This gives the desired $\exp(8C\log^3(1/q)/q)$ time scale.

In order to turn the above into a proof, we use the matryoshka doll technique from Section 5.3.3.2. The geometry of the consecutive regions is given in Figure 6.6. The super good event $\mathcal{SG}(\Lambda)$ for $\Lambda = D \cup \mathring{D} \cup B \cup R_0 \cup R_1^+ \cup R_1^-$ requires that:

- the droplet D is empty;
- each column of the *base B* contains an empty site
- each row of the rectangle R_0 contains a pair of adjacent empty sites;
- each column of the rectangles R_1^+ and R_1^- contains a pair of adjacent empty sites.

Thanks to our choice of geometry, it is not hard to check that the latter three events are very likely under μ, so $\mu(\mathcal{SG}(\Lambda)) \approx \rho = q^{|D|} \approx \exp(8C(\log(1/q))^2/q)$. Recalling Section 2.4, we then seek to prove the Poincaré inequality

$$\mu_\Lambda(f|\mathcal{SG}(\Lambda)) \leqslant \gamma(\Lambda)\mathcal{D}_{\mathbf{1}_{\mathbb{Z}^2\backslash\Lambda}}(f), \qquad \gamma(\Lambda) \leqslant e^{(C\log(1/q))^3/q} \qquad (6.5)$$

for any local function $f : \Omega \to \mathbb{R}$. Once (6.5) is proved, concluding the proof of Theorem 6.11 for the model under consideration can be done along the lines of Section 5.3.3.1.

The proof of (6.5) proceeds in a roughly similar way to (5.22), by proving Poincaré inequalities successively for D, $D \cup \mathring{D}$, $D \cup \mathring{D} \cup B,\ldots, \Lambda$. The one for D is trivial, since $\mathcal{SG}(D) = \{\mathbf{0}_D\} \times \Omega_{\mathbb{Z}^2\backslash D}$ is a single configuration if restricted to D. The inequality for $D \cup \mathring{D}$ is proved by dividing $\mathring{D}$ into vertical strips of width 2 and using Proposition 4.9 for generalised FA-1f in one dimension. This yields the relaxation time bound

$$\gamma\left(D \cup \mathring{D}\right) \leqslant q^{C^2\log(1/q)/q},$$

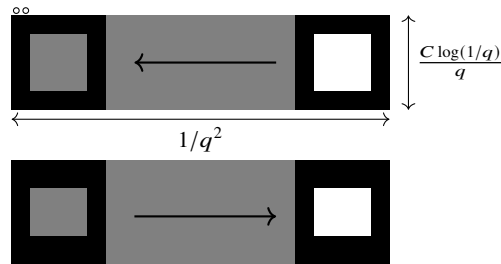where $\gamma(D \cup \mathring{D})$ is defined as in (6.5).

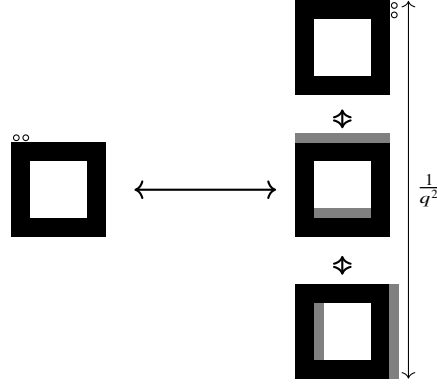Fig. 6.4: The mechanism for the droplet to grow up.

Fig. 6.5: The mechanism for the droplet to grow to the right by making a long excursion up, each of whose steps is a long excursion to the left, as in Figure 6.4.

Proving that $\gamma(D \cup \mathring{D} \cup B) \leqslant \exp(e^{C^2(\log(1/q))^3/q})$ (and similarly for adding the remaining rectangles one by one) is done along the lines of the proof of (5.25). Namely, we use bisection to reduce the length of the base until it reaches 1. However, the factor a_k in the analogue of (5.26) is only bounded by ρ, because we use the original two-block dynamics of Lemma 4.7 rather than the non-oriented three-block variant, Lemma 5.11. This reflects the fact that our heuristics is based on East-like motion rather than CBSEP-like.

Thus, the only remaining ingredient is dealing with adding a single column to the left of $D \cup \mathring{D}$. To do this, we observe that, viewing the empty droplet D as a boundary condition, the problem reduces to dealing with FA-1f on a one-dimensional segment, given that it contains at least one empty site. This was already done in Theorem 4.8. This completes the sketch of the proof of Theorem 6.11 for the model in Figure 6.3b. Let us note that, the fact that Theorem 4.8 covers inhomogeneous one-dimensional
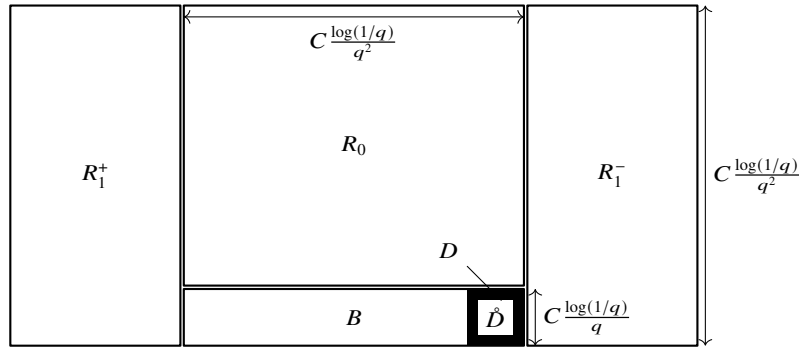

Fig. 6.6: Geometry of the matryoshka dolls in Section 6.3.3.2 for the update family depicted in Figure 6.3b.

KCM on their ergodic components is crucial for dealing with more general update families at this point.

For the remaining unbalanced refined universality class ($\beta = 1, \gamma = 2, \delta = 0$), the proof of [115] is still quite similar, using a CBSEP instead of East-type dynamics in the above argument (Lemma 5.11 versus Lemma 4.7). However, balanced models, particularly with finite number of stable directions, are much more delicate. The reason is that one needs to treat scales below the critical one as well, taking into account the non-trivial internal structure of critical droplets which has a multi-scale form as in Figure 5.2. However, the present situation is more complex than in Section 5.3.3, because some directions are hard, so one needs to use East-like dynamics on some scales and CBSEP-like on others, also carefully choosing directions in which to grow, depending on the scale. When coupled with the not necessarily rectangular geometry required for general models, as well as bounding conditional probabilities of droplets as in Section 5.3.3.2, but in the absence of symmetry, the proof of [115] becomes quite involved. We direct the reader to [115, Section 2] for a detailed description of all mechanisms underlying the proof.

6.4 Conclusion

No new tools were encountered in the present chapter, as compared to previous ones. Instead, we saw how to combine and generalise several of the techniques we were already familiar with, in order to obtain extremely general and precise results. This should clearly showcase the robustness of these methods. For lower bounds we still relied on combinatorial bottlenecks generalising the one for the East model and also incorporating BP ideas. Rough upper bounds used long range renormalisation, while refined ones were proved via the matryoshka dolls technique. One of the takeaways from universality is that a thorough understanding of lower-dimensional models (in our case, one-dimensional) together with respecting the natural geometry and directional preferences of the model can allow one to understand higher dimensional models. The universality viewpoint not only gives a unified framework for understanding the landscape of KCM theory, but also historically supplied the motivation and playground for developing many of the tools presented in the previous chapters.

Chapter 7
Out of equilibrium

In this chapter, we study KCM with initial state not distributed according to the stationary measure. We start by presenting detailed results on the East model, illustrating the kind of results one would like: exponential convergence to equilibrium after a temperature quench, mixing time cutoff in finite volume, etc. We then treat KCM in full generality at the cost of weakening the results. We conclude with open problems, such as bringing the two aspects above together, and some additional out-of-equilibrium settings. This chapter is independent from Chapters 4-6.

In the last three chapters, we discussed properties of the stationary KCM, that is, starting with initial distribution given by the invariant non-trivial product measure $\mu = \mu_q$. While this is the first setting to explore in order to understand these models, it is not the only one of interest. Indeed, both from the mathematics and the physics perspective, it is relevant to study KCM with other initial conditions, that is, out of equilibrium. An initial basic question is to determine under which conditions on $\mathcal{U}$, q and the initial configuration ω, the law of the corresponding infinite volume KCM $(\eta(t))_{t \geqslant 0}$ converges to the equilibrium measure μ, as time goes to infinity. It is natural to expect that, in the ergodic regime, $q > q_c$, provided ω has "enough empty sites" (at the very least, one should have $[\omega] = \mathbf{0}$ in view of Section 3.2) it should hold that for any local function f,

$$\lim_{t \to \infty} |\mathbb{E}_\omega(f(\eta(t))) - \mu(f)| = 0, \tag{7.1}$$

Similarly, if the process is initialised according to a distribution ν, one could expect that if ν has a sufficiently high density of empty sites and $q > q_c$, we should have

$$\lim_{t \to \infty} |\mathbb{E}_\nu(f(\eta(t))) - \mu(f)| = 0. \tag{7.2}$$

A natural choice is $\nu = \mu_{q_0}$ with $q_0 \neq q$. In the physics jargon this is known as a *temperature quench*: abruptly changing the temperature from one value to another (recall (2.1)).[1]

Unfortunately, robust tools to prove (7.1) and (7.2) are not yet available. Indeed, with the notable exception of the East model (see Section 7.2), results are far from being satisfactory and are limited to a restrictive regime of high q (see Section 7.3). The first and foremost reason for this is the fact that, even though the constraints (2.3) are monotone, KCM are not attractive (that is, the product partial order is not preserved by the semi-group of the process, see [146, Sections II.2 and III.2] for background). This is due to the fact that the presence of more empty sites may make certain constraints satisfied and therefore allow certain empty sites to become occupied. Consequently, many of the powerful techniques (e.g. censoring or coupling arguments) which have been developed for the study of other Glauber dynamics (e.g. the contact process and stochastic Ising model, see [145, 146, 154]), fail for KCM. To make matters worse, the usual Holley–Stroock strategy [125] to prove convergence to equilibrium does not apply. Indeed, this approach uses the finiteness of the logarithmic Sobolev constant, which implies hypercontractivity of the semigroup. However, due to the presence of constraints, the logarithmic Sobolev constant is infinite for KCM (see Corollary 3.13) and the technique fails.

Another natural question is how the mixing time (recall (2.19)) scales with the volume $\Lambda^{(n)} = \{0, \dots, n\}^d$, as $n \to \infty$ with $\mathcal{U}$ and q fixed. For $q > q_c$, we expect linear scaling: there exists $C > 0$ such that, for any $\delta \in (0, 1)$ and n large enough depending on δ,

$$t_{\text{mix}}^{(n)}(\delta) := t_{\text{mix}}^{\Lambda^{(n)}}(\delta) \leqslant Cn. \tag{7.3}$$

While lower bounds linear in n are easy to obtain (see Proposition 3.12), proving (7.3) is more challenging and will be addressed in the next sections. Once linear upper and lower bounds are established, it is natural to seek a finer *cutoff* result. Namely, we expect that there exists ν independent of δ such that

$$t_{\text{mix}}^{(n)}(\delta) = \nu n + \epsilon_{n,\delta}, \qquad\qquad \lim_{n \to \infty} \epsilon_{n,\delta}/n = 0. \tag{7.4}$$

Proving this cutoff is essentially equivalent to establishing a limit shape result like those known e.g. for the contact process (see [74]). This question was raised by Kordzakhia and Lalley [139] for the North-East KCM and is readily supported by simulations (see [139, Figure 1]).

---

[1] The extreme case, $q = 1$, is well understood. Indeed, the process becomes a continuous-time version of BP, and essentially behaves like BP.

7.1 Oriented KCM

Before turning to the East model, let us gather a few useful facts which hold more generally for oriented models. For $\mathbf{u} \in \mathbb{R}^d \setminus \{0\}$, we consider the open half-space $\mathbb{H}_{\mathbf{u}} = \{\mathbf{x} \in \mathbb{R}^d : \langle \mathbf{x}, \mathbf{u} \rangle < 0\}$, as in (6.1).

Definition 7.1 (Oriented KCM) Fix a dimension $d \geqslant 1$ and a d-dimensional update family $\mathcal{U}$. We say that the $\mathcal{U}$ is oriented (and that the corresponding KCM is *oriented*) if there exists $\mathbf{v} \in \mathbb{R}^d \setminus \{0\}$ such that $\mathbb{H}_{\mathbf{v}} \supset \bigcup_{U \in \mathcal{U}} U$.

Remark 7.2 (Oriented examples) Among the models introduced in Section 2.2 only the East model and the North-East model are oriented (in every dimension). In both cases a possible choice is $\mathbf{v} = -\sum_{i=1}^{d} \mathbf{e}_i$.

The propositions below state two very handy properties shared by all oriented KCMs:

- dependence propagates only in one direction among well-chosen hyperplanes;
- conditionally on a given site having been legally updated, its occupation variable has its equilibrium distribution.

Proposition 7.3 (Oriented dependence) *Fix an oriented update family $\mathcal{U}$, a site $\mathbf{x} \in \mathbb{Z}^d$ and let $\mathbb{H}_{\mathbf{v}} \supset \bigcup_{U \in \mathcal{U}} U$. Recalling the graphical construction (see Section 2.3), the restriction $(\eta_{\mathbf{x} + \mathbb{H}_v}(t))_{t \geqslant 0}$ of the KCM process is independent of the initial condition, clock rings and coin tosses in $\mathbb{Z}^d \setminus (\mathbf{x} + \mathbb{H}_{\mathbf{v}})$.*

In the case of the one-dimensional East model, Proposition 7.3 states that a site $x \in \mathbb{Z}$ is only influenced by the restriction of the process to its right.

Corollary 7.4 (Exact equilibrium) *Fix an oriented update family $\mathcal{U}$, an initial configuration ω, and a site $\mathbf{x} \in \mathbb{Z}^d$. Let $\mathcal{E}_{\mathbf{x}}(t)$ for $t \geqslant 0$ be the event that there has already been a legal update (see Section 2.3) at $\mathbf{x}$ by time t. Then*

$$\mathbb{P}_\omega(\eta_t(x) = 1 | \mathcal{E}_x(t)) = q.$$

The proofs of Proposition 7.3 and Corollary 7.4 are left as an exercise to the reader.

The following result was proved by Chleboun and Martinelli [59], using an idea similar to Proposition 7.3.

Theorem 7.5 (Quasi-linear mixing for oriented KCM) *Fix a dimension $d \geqslant 1$ and a d-dimensional oriented update family $\mathcal{U}$. For any $q > q_c$ there exists $C > 0$ such that*

$$t_{\mathrm{mix}}^{(n)}(\delta) \leqslant Cn \log n$$

for any $\delta \in (0, 1)$ and n large enough depending on δ.

Though Theorem 7.5 is expected to be sub-optimal (see (7.3) and Conjecture 7.30 below), we recall the main ideas of its simple and instructive proof.

Proof (Sketch) Let $\Lambda = \{1, \ldots, n\}^d$. We proceed iteratively by decomposing Λ into its sections

$$\{\mathbf{x} \in \Lambda : \langle \mathbf{x}, \mathbf{v} \rangle = \lambda\}$$

for $\lambda \in \mathbb{R}$ and $\mathbf{v} \in \mathbb{R}^d \setminus \{0\}$ as in the definition of an oriented update family. Choosing the direction $\mathbf{v} \in \mathbb{Z}^d$, we obtain order n such sections of cardinality at most order n^{d-1}.

Observe that sites on the first hyperplane (corresponding to the smallest value of λ above) are unconstrained thanks to the boundary condition, so a classical coupon collector argument (see e.g. [144, Section 5.3.3]) shows that its mixing time is of order at most $\log n$. The idea is to show that after a time of order $i \log n$, the distribution of the state of the first i hyperplanes is very close to the stationary one. Then the inductive step is performed using the fact that $T_{\mathrm{rel}} < \infty$ whenever $q > \tilde{q}_{\mathrm{c}}$, by Theorem 3.10 and $q_{\mathrm{c}} = \tilde{q}_{\mathrm{c}}$ for oriented update families by [114, Corollary 1.8] (recall Conjecture 3.2, which is established in this case). Indeed, using $T_{\mathrm{rel}} < \infty$ along the lines of the proof of Theorem 3.10, we show that it is likely that, at time of order $\log n$, each site of the last hyperplane has received a legal update, assuming that the initial marginal on the first $i - 1$ sections is exactly (or close to) the equilibrium measure μ_q. Of course, the fact that the equilibrium marginal is preserved over time is true thanks to the oriented nature of the constraint. $\qquad\square$

7.2 East model

7.2.1 Results

In addition to orientation, the East model has other helpful features enabling the proof of rather detailed results outlined next. Let $\mathbb{N} = \{0, 1, \ldots\}$ and $\Delta_{\mathbf{x}}^d = \mathbf{x} - (\mathbb{N}^d \setminus \{0\})$.

Theorem 7.6 (Exponential convergence in all dimensions) *Consider the East model on $\mathbb{Z}^d$ with $q \in (0, 1]$. Fix $\mathbf{x} \in \mathbb{Z}^d$ and an initial configuration $\omega \in \Omega$ such that $\omega_{\mathbf{x}+\mathbb{N}^d} \neq \mathbf{1}_{\mathbf{x}+\mathbb{N}^d}$. Then there exists $m = m(\omega, q) > 0$ and for each local function f with support contained in $\Delta_{\mathbf{x}}^d$ there exists $C = C(f, \omega, q)$ such that for $t > 0$ sufficiently large it holds that*

$$|\mathbb{E}_\omega(f(\eta(t))) - \mu(f)| \leqslant Ce^{-mt}.$$

Remark 7.7 Theorem 7.6, which was proved in one dimension by Cancrini, Schonmann, Martinelli and Toninelli [48] and extended by Marêché [150] to any dimension, implies relaxation to equilibrium in the sense of (7.1) with exponential decay for the minimal possible initial condition, namely as soon as ω has at least one empty site in $\mathbf{x} + \mathbb{N}^d$ for any $\mathbf{x} \in \mathbb{Z}^d$. Using Theorem 7.6, one can also easily prove convergence to equilibrium in the sense of (7.2) with $\nu = \mu_{q_0}$ for any $q_0 \in (0, 1]$ (see [48, Theorem 4.3] and [150, Theorem 2.2 and Remark 2.1]). Furthermore, this convergence occurs exponentially fast, namely there exists $m = m(q_0, q) > 0$ and, for f local, there

exists $C = C(f, q_0, q)$ such that

$$\left| \mathbb{E}_{\mu_{q_0}}(f(\eta(t))) - \mu(f) \right| \leqslant Ce^{-mt}.$$

In dimension one, the dominant term of the time scale of convergence to equilibrium when $q \downarrow 0$ has also been established. It holds that $T_{\text{rel}} \leqslant m^{-1} \leqslant T_{\text{rel}} \log(1/q) = e^{(\log(1/q))^2/(2\log 2)} \log(1/q)$. The lower bound follows by using an argument similar to the one in Section 4.2.1.2, the upper bound may be found in [87, Theorem 3.5]. We also refer to [28, Proposition 4.3] for non-local functions.

Theorem 7.8 (Linear mixing time in all dimensions) *For the East model in any dimension* (7.3) *holds. Furthermore, the result also holds when instead of the completely empty boundary condition we fix any ergodic boundary condition (see Definition 3.8).*

We will provide a proof of Theorem 7.6 in Section 7.2.2 for $d = 1$ and refer the reader to [150] for $d > 1$. Concerning the proof of Theorem 7.8, which is based on similar ingredients, we refer the reader to [56].

Let us now turn to finer results which have been proved only in the one dimensional case. Consider the East model on $\mathbb{Z}$ with initial condition $\omega \in \Omega$ such that $\omega_{-\mathbb{N}} = \mathbf{1}_{-\mathbb{N}\setminus\{0\}} \cdot \mathbf{0}_{\{0\}}$. Let $X_t = \min\{x \in \mathbb{Z} : \eta_x(t) = 0\}$ be the position of the leftmost empty site of $\eta(t)$, which we call the *front*. Notice that X_t makes only nearest neighbour jumps. Indeed, an empty site can appear or disappear only when its left neighbour is empty. Since, once a site is legally updated its occupation variable is forever set to equilibrium (see Corollary 7.4), one could imagine that to the right of the front the distribution is μ. If this were true, the front would move as a biased random walk: negative increments would occur at rate q (the constraint is always satisfied on the site $X_t - 1$), and positive increments at rate $q(1 - q)$ (the occupation variable at the position of the front can be updated to occupied only if $X_t + 1$ is empty). This would yield a speed q^2 to the left. While it is not true that the configuration seen from the front has distribution μ, the following result due to Ganguly, Lubetzky and Martinelli [92], confirms that (as for a biased random walk) the front moves at a negative speed with normal fluctuations and that its concentrated passage times imply cutoff (see (7.4)) with a window $\sqrt{n}$.

Theorem 7.9 (East in one dimension: CLT for the front) *Let $d = 1$ and $\omega \in \Omega$ be such that $\omega_{-\mathbb{N}} = \mathbf{1}_{-\mathbb{N}\setminus\{0\}} \cdot \mathbf{0}_{\{0\}}$. For any $q \in (0, 1)$, there exist constants $\sigma = \sigma(q) > 0$, $v = v(q) < 0$ and $C = C(q) > 0$ such that, the front X_t of the one-dimensional East process with initial condition ω satisfies*

$$\mathbb{P}\left(\lim_{t \to \infty} \frac{X_t}{t} = v \right) = 1, \qquad\qquad |\mathbb{E}(X_t) - vt| \leqslant C,$$

$$\lim_{t \to \infty} \frac{\text{Var}(X_t)}{t} = \sigma^2, \qquad\qquad \lim_{t \to \infty} \frac{X_t - vt}{\sigma\sqrt{t}} = Z$$

in distribution for a standard normal random variable Z.

Corollary 7.10 (Cutoff in one dimension) *For any $q \in (0, 1)$, there exists $C = C(q) > 0$ such that the one-dimensional East model satisfies* (7.4) *with*

$$|\epsilon_{n,\delta}| \leqslant C\phi^{-1}(1 - \delta)\sqrt{n},$$

where ϕ is the cumulative distribution function of the standard normal law.

We refer the reader to [92] for the proof of Theorem 7.9 based on:

- ergodicity of the process seen from the front (see Theorem 7.11 below);
- showing that, after an initial burn-in time, the front increments behave like a stationary sequence of weakly dependent random variables, and applying an ingenious Stein's method argument of Bolthausen [41] to derive the central limit theorem.

Corollary 7.10 follows almost immediately. We further refer the reader to [65] for a first step in the direction of proving (7.4) for East in higher dimensions.

Theorem 7.11 (Ergodicity of the process seen from the front) *Let $d = 1$ and $\omega \in \Omega$ be such that $\omega_{-\mathbb{N}} = \mathbf{1}_{-\mathbb{N}\backslash\{0\}} \cdot \mathbf{0}_{\{0\}}$. Let $(\eta(t))_{t\geqslant 0}$ be the one-dimensional East KCM with initial condition ω. For all $t \geqslant 0$ and $x \in \mathbb{Z}$, let $\tilde{\eta}_x(t) = \eta_{x+X_t}(t)$, defining the process seen from the front. There exists a unique measure $\tilde{\mu} = \tilde{\mu}(q)$ such that $\tilde{\eta}(t) \to \tilde{\mu}$ in distribution as $t \to \infty$. Moreover, there exist $C = C(q)$ and $m = m(q) > 0$, such that for any $x \in \mathbb{Z}$*

$$\|\tilde{\mu} - \mu\|_{[x,\infty)} \leqslant Ce^{-mx},$$

where, for $\Lambda \subset \mathbb{Z}$, $\|\tilde{\mu} - \mu\|_\Lambda$ denotes the total variation distance between the marginals of $\tilde{\mu}$ and μ on Λ.

Remark 7.12 (Front velocity) The velocity of the front v in Theorem 7.9 can be expressed in terms of the invariant measure as $v = -q + (1 - q)\tilde{\mu}(\omega_1 = 0)$.

The main ingredients for the proof of Theorem 7.11 due to Blondel [28], in addition to the techniques needed for Theorem 7.6, are:

- coupling the processes seen from the front, starting with two different initial conditions, in order to prove that their laws converge to the same limit (see [28, Theorem 4.7]);
- using the fact that the front moves at most linearly (by the finite speed argument of Proposition 3.12) and leaves empty sites behind to prove that far from the front the process is almost at equilibrium (see [28, Theorem 4.7]). Here the distinguished zero of Definition 7.14 below plays a key role.

Problem 7.13 (Front measure) Theorems 7.9-7.11 leave various questions unanswered. For instance, can one quantify the correlations between adjacent occupation variables in the invariant measure $\tilde{\mu}$? Can one determine the asymptotics of the velocity v of Theorem 7.9 and Remark 7.12 in the $q \to 0$ limit?

7.2.2 Exponential decay to equilibrium for East in one dimension

Let us start by introducing a key notion, due to Aldous and Diaconis [4].

Definition 7.14 (Distinguished zero) Fix $x \in \mathbb{Z}$ and an initial configuration $\omega \in \Omega$ with $\omega_x = 0$. Call the site x *distinguished*. We set $\xi_0 = x$. The position $\xi_s \in \mathbb{Z}$ of the *distinguished zero* at time $s > 0$ is defined according to the following iterative rule. For all times s strictly smaller than the first legal update t_1 at x, we set $\xi_s = x$, while $\xi_{t_1} = x + 1$. Then we wait for the first legal update t_2 at $x + 1$ after t_1, at which point we set $\xi_{t_2} = x + 2$ and so on.

Note that, almost surely, the trajectory $(\xi_s)_{s \geqslant 0}$ is right-continuous, piece-wise constant, increasing by 1 at each jump and not exploding in finite time. Also note that, by definition of the legal updates, necessarily the state $\eta_{\xi_s}(s)$ of the East process at the position of the distinguished zero is 0 for any $s \geqslant 0$, hence the name.

Proposition 7.15 (Properties of the distinguished zero) *Fix $t > 0$. In the setting of Definition 7.14, conditioning on the knowledge of the trajectory $(\xi_s)_{s \leqslant t}$ and denoting by $0 < t_1 < \cdots < t_{n-1} < t$ the times of the distinguished zero's jumps (with $t_0 = 0$, $t_n = t$), the following holds for all $i \in \{0, \ldots, n - 1\}$*

- *for $s \in [t_i, t_{i+1})$, $\xi_s = x + i$ and the restriction of the process to $\{x, \ldots, x + i - 1\}$ follows an East dynamics with zero boundary condition;*
- *at time t_{i+1}, the configuration at site $x + i$ is updated according to a Bernoulli$(1 - q)$ variable.*

The proof, which is derived using Definition 7.14, Proposition 7.3 and Corollary 7.4, is left to the reader.

Remark 7.16 (Conditional graphical construction) Recall the graphical construction of Section 2.3. Proposition 7.15 implies that, given $(\xi_{t'})_{t' \leqslant t}$, for any (s, y) satisfying $s \leqslant t$ and $y \leqslant \xi_s$, the variable $\eta_y(s)$ is determined by the following "conditional graphical construction". In the time interval $[0, t_1)$, the occupation variables in the interval $\{y, \ldots, x - 1\}$ evolve as an East process with fixed empty site at x, using the clock rings $(t_{z,k})_{z \in \{y, \ldots, x-1\}}$ and coin tosses $(s_{z,k})_{z \in \{y, \ldots, x-1\}}$. In the time interval $[t_1, t_2)$ the same happens with empty boundary condition at $x + 1$ instead of x and so on. Note that, at t_1, the clock at x rings and $\eta_x(t_1)$ takes the value of the corresponding coin toss.

Furthermore, Proposition 7.15 yields the following.

Corollary 7.17 (Equilibrium zone) *Fix $a \leqslant x \in \mathbb{Z}$ and two measures ψ^- and ψ^+ on $\Omega_{a-1-\mathbb{N}}$ and $\Omega_{x+1+\mathbb{N}}$ respectively. Let*

$$\psi = \begin{cases} \psi^- \otimes \mu_{\{a, \ldots, x-1\}} \otimes \delta_0 \otimes \psi^+ & \text{if } a < x \\ \psi^- \otimes \delta_0 \otimes \psi^+ & \text{if } a = x \end{cases}$$

Let $(\eta(t))_{t \geqslant 0}$ be the East process with initial distribution ψ. Fix $t \geqslant 0$. Then the conditional distribution of $\eta_{\{a, \ldots, \xi_t - 1\}}(t)$, given the distinguished zero trajectory $(\xi_s)_{s \leqslant t}$, is the equilibrium one, $\mu_{\{a, \ldots, \xi_t - 1\}}$.

The following proof, due to Cancrini, Martinelli, Schonmann and Toninelli [48], uses as key ingredients Proposition 7.15, Corollary 7.17 and the L^2 convergence to equilibrium guaranteed by Theorem 3.10(iv) and (2.12).

Proof (Theorem 7.6 for d = 1) Assume for simplicity that $\mu(f) = 0$, $x = 0$, $\omega_0 = 0$ and let the support of f be contained in $\{a, \ldots, a'\} \Subset -1 - \mathbb{N}$. Let $b \leqslant 0$ be the position of the first empty site in ω to the right of a'. Make b distinguished and denote by ξ_s its position at time s. Given the trajectory $(\xi_s)_{s \leqslant t}$, let $0 < t_1 < t_2 < \cdots < t_{n-1} < t$ be the times when the distinguished zero jumps, and set $t_0 = 0$, $t_n = t$. For $i \in \{0, \ldots, n-1\}$, let

$$\Xi_i = (\xi_s)_{s \in [t_i, t]} \qquad\qquad V_i = \{a, \ldots, b + i - 1\}.$$

We claim that

$$|\mathbb{E}_\omega (f(\eta(t)))| \leqslant \mathbb{E}_\omega (|\mathbb{E}_\omega(f(\eta(t))|\Xi_0)|)$$

$$\leqslant (\min(q, 1-q))^{-(b-a)} \mathbb{E}_\omega \left(\int \mathrm{d}\mu_{V_0}(\omega') \, |\mathbb{E}_{\omega'}(f(\eta'(t))|\Xi_0)| \right)$$

$$\leqslant (\min(q, 1-q))^{-(b-a)} \mathbb{E}_\omega \left(\sqrt{\int \mathrm{d}\mu_{V_0}(\omega')(\mathbb{E}_{\omega'}(f(\eta'(t))|\Xi_0))^2} \right)$$

$$\leqslant (\min(q, 1-q))^{-(b-a)} \mathbb{E}_\omega \left(\sqrt{\mathrm{Var}_{V_0}\left(g_t^{(0)} \right)} \right), \qquad (7.5)$$

where we let $\eta'(t)$ be the configuration obtained following the conditional graphical construction of Remark 7.16 with $\eta'_{V_0}(0) = \omega' \in \Omega_{V_0}$ and

$$g_t^{(0)}(\omega') = \mathbb{E}_{\omega'}(f(\eta'(t))|\Xi_0).$$

In order to obtain (7.5), we used Remark 7.16 together with the fact that for any $\omega \in \{0, 1\}^{\mathbb{Z}}$ it holds $\mu_{V_0}(\omega|_{V_0}) \geqslant (\min(q, 1-q))^{-(b-a)}$ to obtain the second inequality, Cauchy-Schwarz for the third inequality and, for the last inequality, we used Corollary 7.17 and the hypothesis that the support of f is contained in $[a, b)$, which yield

$$\int \mathrm{d}\mu_{V_0}(\omega')\mathbb{E}_{\omega'}(f(\eta'(t))|(\xi_s)_{s \leqslant t}) = \mu_{V_0}(f) = 0. \qquad (7.6)$$

Let $P_s^{(i)}$ for $s \in [t_i, t_{i+1})$ be the Markov semigroup associated to the East process in the interval V_i with a fixed empty site at the right boundary $b + i$. Then, using Propositions 7.3 and 7.15, we get

$$g_t^{(0)}(\omega') = \sum_{\sigma \in \{0,1\}^{V_0}} \sum_{\sigma' \in \{0,1\}} P_{t_1}^{(0)}(\omega', \sigma)\mu_b(\sigma')g_{t-t_1}^{(1)}(\sigma \cdot \sigma'), \qquad (7.7)$$

where, for $s \geqslant 0$, we define $g_s^{(1)} : \Omega_{V_1} \to \mathbb{R}$ by

$$g_s^{(1)}(\omega'') = \mathbb{E}_{\omega''}(f(\eta''(s))|\Xi_1) \tag{7.8}$$

where $\eta''(s)$ for $s \leqslant t$ denotes the configuration in the interval $\{a, \ldots, \xi_t - 1\}$ obtained starting at time t_1 from the configuration $\omega'' \in \Omega_{V_1}$ and evolving according to the conditional graphical construction described in Remark 7.16 applied to the time interval $(t_1, t_1 + s]$. Therefore, using (2.12) we get

$$
\begin{aligned}
\mathrm{Var}_{V_0}\left(g_t^{(0)}\right) &\leqslant e^{-2t_1/T_{\mathrm{rel}}^{V_0}} \, \mathrm{Var}_{V_0}\left(\sum_{\sigma' \in \{0,1\}} \mu_b(\sigma')\mu_{V_1}\left(g_{t-t_1}^{(1)}\right)\right) \\
&\leqslant e^{-2t_1/T_{\mathrm{rel}}} \, \mathrm{Var}_{V_1}\left(g_{t-t_1}^{(1)}\right)
\end{aligned}
\tag{7.9}
$$

where, as usual, we denote by $T_{\mathrm{rel}}^{V_0}$ (resp. T_{rel}) the relaxation time of the East model on V_0 with empty boundary condition (resp. $\mathbb{Z}$). In order to obtain (7.9) we use convexity of the variance and (3.15). We can now proceed analogously to get

$$\mathrm{Var}_{V_1}\left(g_{t-t_1}^{(1)}\right) \leqslant e^{-2(t_2-t_1)/T_{\mathrm{rel}}} \, \mathrm{Var}_{V_2}\left(g_{t-t_1-t_2}^{(2)}\right) \tag{7.10}$$

and, by induction, we get

$$\mathrm{Var}_{V_0}\left(g_t^{(0)}\right) \leqslant e^{-2t/T_{\mathrm{rel}}} \, \mathrm{Var}_{\{a,\ldots,\xi_t-1\}}(f).$$

Plugging this bound into (7.5) and recalling that $T_{\mathrm{rel}} > 0$ for any $q \in (0, 1]$ for the East model, we finally get

$$|\mathbb{E}_\omega(f(\eta(t)))| \leqslant ce^{-t/T_{\mathrm{rel}}}\mathbb{E}_\omega\left(\sqrt{\mathrm{Var}_{\{a,\ldots,\xi_t-1\}}(f)}\right) \leqslant ce^{-t/T_{\mathrm{rel}}}\|f\|_\infty$$

for some $c = c(q, a, b) > 0$. $\qquad\square$

While the distinguished zero does not generalise to oriented KCM, a version of it was used in [48, Section 4] for an oriented KCM on a tree.

7.3 High vacancy density regime

In this section we focus on results on KCM with arbitrary update family with q close enough to 1. While the same results should hold whenever $q > q_c$, the current techniques do not allow proving this.

7.3.1 Results

The next two theorems, proved by Hartarsky and F. Toninelli [124], establish (7.2) with an exponential convergence when ν is a product measure with a vacancy density $q_0 > q_c$ and a linear upper bound for the mixing time (7.3). Both results apply to all models but are restricted to a high vacancy density regime for q and their proof is delegated to Section 7.3.2. Before stating them we need to introduce the notion of trivial subcritical models.

Definition 7.18 (Trivial subcritical) We say that an update family $\mathcal{U}$ is not *trivial subcritical*, if there exists $U \in \mathcal{U}$ and a direction $\mathbf{v} \in \mathbb{R}^d$ such that $\langle \mathbf{u}, \mathbf{v} \rangle < 0$ for all $\mathbf{u} \in U$.

In [14, Theorem 7.1] it is proved that $q_c = 1$ iff $\mathcal{U}$ is trivial subcritical. Therefore, excluding trivial subcritical models from the following two results is necessary.

Theorem 7.19 (Exponential convergence to equilibrium after a temperature quench) *Fix a dimension $d \geqslant 1$ and an update family $\mathcal{U}$ which is not trivial subcritical. For any $\alpha > 0$ there exist $\varepsilon = \varepsilon(\alpha) > 0$ and $c = c(\alpha) > 0$ such that the following holds for the $\mathcal{U}$-KCM η. For any $q_0 \in [\tilde{q}_c + \alpha, 1]$, $q \in [1 - \varepsilon, 1]$ and local function $f : \Omega \to \mathbb{R}$,*

$$\left| \mathbb{E}_{\mu_{q_0}} (f(\eta(t))) - \mu_q(f) \right| \leqslant \frac{\|f\|_\infty \cdot |\mathrm{supp} f|}{ce^{ct}},$$

where $\mathrm{supp} f$ is the set of sites on whose state the value of f depends.

Theorem 7.20 (Linear mixing at high vacancy density) *Fix a dimension $d \geqslant 1$ and a d-dimensional update family $\mathcal{U}$ which is not trivial subcritical. Then there exist $\varepsilon > 0$ and $C > 0$ such that for any $q \in [1 - \varepsilon, 1]$, inequality (7.3) holds for any $\delta \in (0, 1)$ and n large enough depending on δ.*

It should be noted that Theorem 7.20 also applies to domains of non-hypercubic shape, but is stated as is for the sake of simplicity. We refer the reader to [124, Section 4] for an account of previous works in this direction, in particular [29, 152, 166]. Moving on to more precise results, it only remains to report the following one due to Ertul [83] (also recall Corollary 7.10 for the one dimensional East model).

Theorem 7.21 (Cutoff for one-dimensional FA-1f) *Consider the FA-1f model in one dimension. There exists an explicit $\varepsilon > 0$ such that for any $q \in (1 - \varepsilon, 1]$, there exists $v, \alpha, \beta > 0$ such that (7.4) holds with*

$$-\alpha\sqrt{n} \leqslant \epsilon_{n,\delta} \leqslant \beta\sqrt{n}.$$

Furthermore, the number v in Theorem 7.21 corresponds to twice the speed at which the rightmost empty site, known as *front*, moves in FA-1f on $\{1, 2, \dots\}$ with boundary condition $\mathbf{0}_{\{0,-1,\dots\}}$ (compare with Remark 7.12 for East). Indeed, as for Corollary 7.10, the proof of Theorem 7.21 is based on the identification of the

front speed thanks to the convergence of the process seen from the front [30] (as in Theorem 7.11). Then, considering each interval of occupied sites, one may show that it shrinks at the front speed at both ends, leading to a double speed (see [83] for more details). Nonetheless, deducing Theorem 7.21 from the analogue of Theorem 7.11 is not immediate, as in the case of Corollary 7.10.

7.3.2 Proofs via cooperative contact processes

We now turn to outlining the proofs of Theorems 7.19 and 7.20. In order to simplify the presentation (see [124] for the full details[2]), we focus on Theorem 7.20 in the case of FA-2f in two dimensions (see Figure 2.1c). The only additional ingredient needed in the general case and in Theorem 7.19 are one-scale renormalisations along the lines of Section 3.4.

The proof proceeds in several steps involving a number of interacting particle systems other than KCM. In Section 7.3.2.1, we reduce the mixing time problem for FA-2f to a more complex problem for a simpler model. Namely, studying the space-time connected components of occupied sites in a cooperative contact process known as the sexual contact process. In Section 7.3.2.2, we use a comparison with last passage percolation to replace the arbitrary initial condition by the fully empty one. In Section 7.3.2.3 we discretise time to transform the sexual contact process into a North-East BP with death. In Section 7.3.2.4, we introduce Toom cycles to show that, at q close enough to 1, with empty initial condition, BP with death mostly has empty sites. Finally, in Section 7.3.2.5, we show that long chains of such Toom cycles are also unlikely.

7.3.2.1 Sexual contact process

Let us start by introducing the *sexual contact process* (SCP) of [76]. It is a continuous time Markov process on $\{0, 1\}^{\mathbb{Z}^2}$ (or in finite volume Λ with $\mathbf{0}_{\mathbb{Z}^2 \setminus \Lambda}$ boundary condition as in Section 2.4) which, using a graphical representation similar to the one for KCM from Section 2.3, is defined as follows. Each site $\mathbf{x} \in \mathbb{Z}^2$ waits an independent exponentially distributed time with mean one before attempting to update. At that time, if both $\mathbf{x} + \mathbf{e}_1$ and $\mathbf{x} + \mathbf{e}_2$ are in state 0, the state of $\mathbf{x}$ becomes 0 with probability q and becomes 1 with probability $1 - q$. Otherwise the state of $\mathbf{x}$ remains unchanged with probability q and becomes 1 with the remaining probability $1 - q$. SCP has two notable advantages over FA-2f and one disadvantage. Namely, SCP is attractive and oriented, but its upper invariant measure is not explicit when it is not the trivial $\delta_{\mathbf{1}}$.

The above formulation of the process suggests a canonical coupling with FA-2f, using the same clock rings and Bernoulli variables with parameter q, the same

---

[2] Beware that in [124] the roles of 0 and 1 states are exchanged.

domain $\Lambda = \{1, \ldots, n\}^2$ and boundary condition $\mathbf{0}_{\mathbb{Z}^2 \setminus \Lambda}$. It is an exercise to check the following.

Lemma 7.22 (Sexual contact process comparison) *Under the canonical coupling of FA-2f $(\eta(t))_{t \geqslant 0}$ and SCP $(\zeta(t))_{t \geqslant 0}$, if $\zeta(0) = \mathbf{1}_\Lambda$, then $\eta_\mathbf{x}(t) \leqslant \zeta_\mathbf{x}(t)$ for all $\mathbf{x} \in \Lambda$ and $t \geqslant 0$.*

However, we require a finer relation from [124, Section 7] (also see [145, Section III.1], [43, Section 4.2] and [105, Section 1.3] for similar ideas), since the two processes do not share the invariant measure μ_q. Consider a set of *orange* sites $O_t \subset \Lambda$ defined as follows for $t \geqslant 0$. At time 0, we set $O_0 = \Lambda$. At each clock ring $t \geqslant 0$ at site $\mathbf{x} \in \Lambda$, we obtain O_t from O_{t-} by:

- removing $\mathbf{x}$, if $\zeta_\mathbf{x}(t) = 0$;
- adding $\mathbf{x}$, if $\zeta_\mathbf{x}(t) = 1$ and there is an orange site around $\mathbf{x}$, that is, $O_{t-} \cap \{\mathbf{x}, \mathbf{x} + \mathbf{e}_1, \mathbf{x} + e_2, \mathbf{x} - \mathbf{e}_1, \mathbf{x} - \mathbf{e}_2\} \neq \varnothing$;
- changing nothing otherwise.

The purpose of orange sites is to ensure that FA-2f processes with different initial conditions are coupled outside orange sites.

Lemma 7.23 (Orange set coupling) *Under the canonical coupling of SCP $(\zeta(t))_{t \geqslant 0}$ with initial condition $\mathbf{1}_\Lambda$ and two FA-2f processes $(\eta(t))_{t \geqslant 0}$ and $(\eta'(t))_{t \geqslant 0}$ with different initial conditions, for all $t \geqslant 0$, we have*

$$\{\mathbf{x} \in \Lambda : \eta_\mathbf{x}(t) \neq \eta'_\mathbf{x}(t)\} \subset O_t. \tag{7.11}$$

Proof We proceed by induction on the number of clock rings in Λ. Removing $\mathbf{x}$, if $\zeta_\mathbf{x}(t) = 0$ is justified by Lemma 7.22. Setting $O_t = O_{t-} \cup \{\mathbf{x}\}$ cannot violate (7.11), at time t, if the induction hypothesis is verified at $t-$. Finally, assume that there is no orange site around $\mathbf{x}$. Then by induction hypothesis $\eta(t-)$ and $\eta'(t-)$ are equal around $\mathbf{x}$ and, therefore, $\eta_\mathbf{x}(t) = \eta'_\mathbf{x}(t)$.   □

Owing to Lemma 7.23 and the standard results on coupling and mixing times [144, Corollary 5.5], proving Theorem 7.20 is reduced to showing that with high probability the set of orange sites is empty at time Cn for C large enough. We say that two space-time points $(t, \mathbf{x})$ and $(t', \mathbf{x}')$ are connected in $X \subset [0, \infty) \times \Lambda$, if there exists a sequence of segments $[s_i, t_i] \times \{\mathbf{x}_i\} \subset X$ indexed by $i \in \{0, \ldots, N\}$ such that $[s_i, t_i] \cap [s_{i-1}, t_{i-1}] \neq \varnothing$ and $\mathbf{x}_{i-1}$ and $\mathbf{x}_i$ are neighbours in Λ, such that $(t, \mathbf{x}) \in [s_0, t_0] \times \{\mathbf{x}_0\}$ and $(t', \mathbf{x}') \in [s_N, t_N] \times \{\mathbf{x}_N\}$. By the definition of orange sites, it is clear that $\bigcup_{t \geqslant 0} O_t$ is contained in the connected component C of $\{(t, \mathbf{x}) \in [0, \infty) \times \Lambda : \zeta_\mathbf{x}(t) = 1\}$ containing $\{0\} \times \Lambda$. Thus, it suffices to show that, for $q \geqslant 1 - \varepsilon$,

$$\mathbb{P}(C \subset [0, Cn] \times \Lambda) \geqslant 1 - \delta \tag{7.12}$$

for some constant C independent of $\delta > 0$, provided n is large enough. In order to simplify the exposition, we fix $q = 1 - \varepsilon < 1$ in the rest of the section with ε to be chosen small enough in Section 7.3.2.3.

7.3.2.2 Last passage percolation

Following [124, Section 11], our next goal is to replace the worst initial condition, $\mathbf{1}_\Lambda$, of SCP ζ of Section 7.3.2.1 by the best one, $\mathbf{0}_\Lambda$. For this we use a comparison with a version of last passage percolation, which will play a similar role to the orange set above. Define the *last passage* set $L_t \subset \Lambda$ for $t \geqslant 0$ as follows. Set $L_0 = \Lambda$ and at each clock ring $t \geqslant 0$ at site $\mathbf{x} \in \Lambda$, we obtain L_t from L_{t-} by removing $\mathbf{x}$, if the following conditions are both satisfied:

- at time t the Bernoulli variable with parameter $q = 1 - \varepsilon$ takes the value 0, that is, SCP changes the state of $\mathbf{x}$ to 1 regardless of the current configuration;
- $\{\mathbf{x} + \mathbf{e}_1, \mathbf{x} + \mathbf{e}_2\} \cap L_{t-} = \varnothing$.

Otherwise, we set $L_t = L_{t-}$. The proof of the following observation, similar to Lemma 7.23, is left as an exercise for the reader.

Lemma 7.24 (Last passage percolation coupling) *Under the above coupling, we have*

$$\{\mathbf{x} \in \Lambda : \zeta_\mathbf{x}(t) \neq \zeta'_\mathbf{x}(t)\} \subset L_t$$

for any $t \geqslant 0$, where ζ and ζ' are SCP with initial condition $\mathbf{1}_\Lambda$ and $\mathbf{0}_\Lambda$ respectively.

We next invoke a classical fact about last passage percolation. We note that much more precise results are available, but not more useful for our purposes.

Lemma 7.25 (Linear last passage time) *For some absolute constant $C' > 0$ we have*

$$\mathbb{P}\left(L_{C'n/\varepsilon} \neq \varnothing\right) \leqslant \delta$$

for any n large enough depending on $\delta > 0$.

Proof (Sketch) A proof of Lemma 7.25 generalising to arbitrary dimension was given by Greenberg, Pascoe and Randall [109]. The idea is to introduce an exponential metric on the set of possible values of the last passage set L_t:

$$d(A, B) = \sum_{\mathbf{x} \in A \Delta B} e^{\gamma\langle \mathbf{x}, \mathbf{1}\rangle},$$

where $\gamma > 0$ is a suitably large constant depending only on dimension and Δ denotes the symmetric difference of sets. One can verify that the expected distance of two canonically coupled last passage percolations starting from neighbouring configurations contracts. Applying the path coupling method (see [144, Chapter 14]), this yields that the hitting time of $\varnothing$ is logarithmic in the diameter of the space, which is exponential in n. $\qquad\square$

Combining Lemmas 7.24 and 7.25 and recalling (7.12), it is now sufficient to prove that, with probability $1 - \delta$, the connected components of $\{(t, \mathbf{x}) \in [0, \infty) \times \Lambda : \zeta'_\mathbf{x}(t) = 1\}$ intersecting $\{C'n/\varepsilon\} \times \Lambda$ are contained in $[0, Cn] \times \Lambda$. Taking $C > C'/\varepsilon$ and performing a union bound, it suffices to show that for any space-time point $(t, \mathbf{x})$, its connected component $C_{t,\mathbf{x}}$ satisfies

$$\mathbb{P}(\operatorname{diam}(C_{t,\mathbf{x}}) \geqslant k) \leqslant e^{-ck} \tag{7.13}$$

for some $c > 0$ independent of n and all k large enough. Since SCP is attractive, ζ' has initial condition $\mathbf{0}_\Lambda$ and boundary condition $\mathbf{0}_{\mathbb{Z}^2 \setminus \Lambda}$, it suffices to prove (7.13) for the infinite volume SCP ζ'' with initial condition $\mathbf{0}$. Indeed, by induction on the number of clock rings in Λ in the graphical construction, one can show that for every $t \geqslant 0$ and $\mathbf{x} \in \mathbb{Z}^2$, we have $(\mathbf{0}_{\mathbb{Z}^2 \setminus \Lambda} \cdot \zeta'(t))_{\mathbf{x}} \leqslant \zeta''_{\mathbf{x}}(t)$, which implies that the connected component $C_{t,\mathbf{x}}$ for ζ'' contains the one for ζ'.

7.3.2.3 BP with death

The next step of the proof is a discretisation in time [124, Section 8]. Namely, we fix T large enough but such that $T\varepsilon$ is small. This way, in each time interval of length T, it is likely that SCP ζ'' with initial condition $\mathbf{0}$ attempts to change the state of a given site $\mathbf{x}$ to 0 (and succeeds, if $\mathbf{x} + \mathbf{e}_1$ and $\mathbf{x} + \mathbf{e}_2$ are in state 0), but never attempts to change the state of $\mathbf{x}$, $\mathbf{x} + \mathbf{e}_1$ or $\mathbf{x} + \mathbf{e}_2$ to 1. We declare each space-time point $(m, \mathbf{x}) \in \{0, 1, \dots\} \times \mathbb{Z}^2$ *good*, if the above event occurs for site $\mathbf{x}$ in the time interval $[mT, (m+1)T)$. Note that almost surely, no clocks ring at times $(mT)_{m=0}^\infty$, so these events are 1-dependent and have probability $1 - \varepsilon'$ for $\varepsilon' > 0$ that can be chosen arbitrarily small, if ε is small enough. Using the Liggett-Schonmann-Stacey theorem [147], we can replace this 1-dependent field of indicators of good sites by an independent one with high marginals (at the cost of changing ε').

With the good space-time points at hand, we define the discrete time version ξ of SCP, which we refer to as North-East BP with death parameter ε'. Set $\xi(0) = \mathbf{0}$. For $m \geqslant 1$ and $\mathbf{x} \in \mathbb{Z}^2$, define $\xi_{\mathbf{x}}(m) = 0$, if $(m-1, \mathbf{x})$ is good and $\xi_{\mathbf{x}}(m-1) = 0$, or if $(m-1, \mathbf{x})$ is good and $\xi_{\mathbf{x}+\mathbf{e}_1}(m-1) = \xi_{\mathbf{x}+\mathbf{e}_2}(m-1) = 0$. Otherwise, set $\xi_{\mathbf{x}}(m) = 1$. The name is justified by the fact that, in the absence of non-good space-time points, this process is exactly BP with the North-East update family (see Figure 2.1e). It is not hard to check that, if $\xi_{\mathbf{x}}(m) = \xi_{\mathbf{x}}(m+1) = 0$, then $\zeta''_{\mathbf{x}}(t) = 0$ for all $t \in [mT, (m+1)T)$. In view of this, we consider the set

$$X = \left\{ (m, \mathbf{x}) \in \{0, 1, \dots\} \times \mathbb{Z}^2 : \xi_{\mathbf{x}}(m) = 1 \right\}. \tag{7.14}$$

equipped with all edges of the form $((m, \mathbf{x}), (m', \mathbf{x}'))$ with

$$(m' - m, \mathbf{x}' - \mathbf{x}) \in \{-1, 0, 1\} \times \{\mathbf{0}, \mathbf{e}_1, \mathbf{e}_2, -\mathbf{e}_1, -\mathbf{e}_2\}. \tag{7.15}$$

Hence, it suffices to prove (7.13) for the connected component $C'_{m,\mathbf{x}}$ of any space-time point $(m, \mathbf{x})$ in X.

7.3.2.4 Toom cycles

Before treating connected components $C'_{m,\mathbf{x}}$ in X of (7.14), it is useful to first show that $\mathbb{P}(\xi_{\mathbf{x}}(m) = 1)$ is small for any space-time point $(m, \mathbf{x})$. This is an instance of a

classical result of Toom [200] for stability of cellular automata subjected to random noise. However, since our cellular automaton is BP with update family consisting of a single rule, it is possible to use a simpler argument of Swart, Szabó and Toninelli [193, Section 3.5] presented below. We present the construction of a *Toom cycle* in the form of an algorithm illustrated in Figure 7.1.
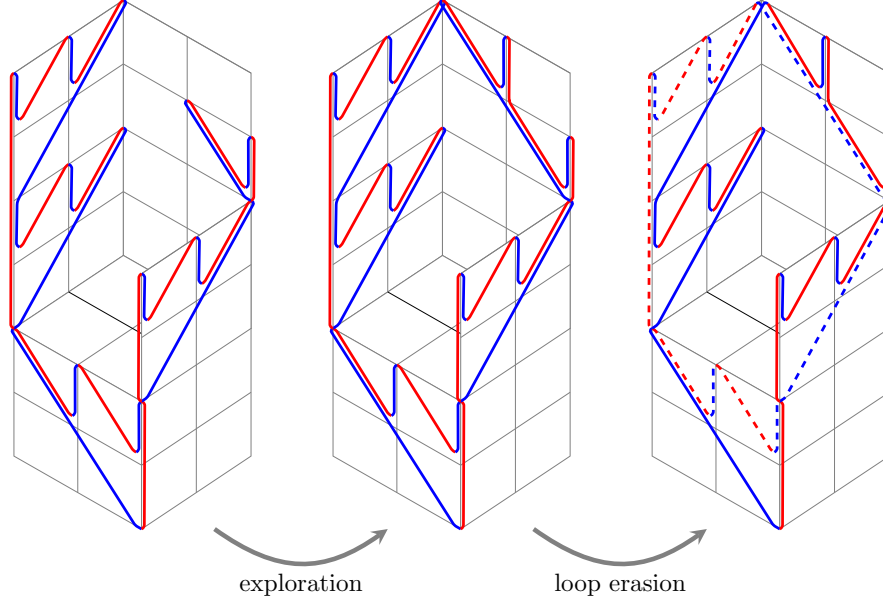

exploration   loop erasion

Fig. 7.1: Illustration of the algorithm used to construct a Toom cycle rooted at the bottom space-time point $(m, \mathbf{x})$. Time increases downwards.

Fix a realisation of the good space-time points such that $\xi_{\mathbf{x}}(m) = 1$. For each $(m', \mathbf{x}')$ such that $\xi_{\mathbf{x}'}(m') = 1$ and $(m' - 1, \mathbf{x}')$ is good, fix some $\mathbf{e}(m', \mathbf{x}') \in \{\mathbf{e}_1, \mathbf{e}_2\}$ such that $\xi_{\mathbf{x}'+\mathbf{e}(m',\mathbf{x}')}(m' - 1) = 1$, which necessarily exists by construction. We construct a sequence of space-time points starting at $(m, \mathbf{x})$. Elements of the sequence such that both the previous and next element have larger time coordinate are called *sinks*. Initially the sequence is the single point $\mathcal{T}_0 = (m, \mathbf{x})$. If $\mathbf{e}(m, \mathbf{x})$ is not defined, we stop and output $\mathcal{T}_0$. Otherwise, we *explore*: replace the point $(m, \mathbf{x})$ by the sequence $\mathcal{T}_1 = (m, \mathbf{x}), (m - 1, \mathbf{x}), (m, \mathbf{x}), (m - 1, \mathbf{x} + \mathbf{e}(m, \mathbf{x})), (m, \mathbf{x})$. For $l \geqslant 1$, given $\mathcal{T}_l$, define $\mathcal{T}_{l+1}$ as follows. Among the sinks $(m', \mathbf{x}')$ of $\mathcal{T}_l$ such that $\mathbf{e}(m', \mathbf{x}')$ is defined, we find the first one that maximises m'. If no such element exists, we stop and output $\mathcal{T}_l$. Otherwise, explore: replace the selected space-time point $(m', \mathbf{x}')$ in $\mathcal{T}_l$ by $(m', \mathbf{x}'), (m' - 1, \mathbf{x}'), (m', \mathbf{x}'), (m' - 1, \mathbf{x}' + \mathbf{e}(m', \mathbf{x}')), (m', \mathbf{x}')$. Denote the result of this exploration operation by $\mathcal{T}_l'$. If the points $(m' - 1, \mathbf{x}')$ and $(m' - 1, \mathbf{x}' + \mathbf{e}(m', \mathbf{x}'))$ are not already present in $\mathcal{T}_l$, we set $\mathcal{T}_{l+1} = \mathcal{T}_l'$. If $(m' - 1, \mathbf{x}')$ appears in $\mathcal{T}_l$, we remove from $\mathcal{T}_l'$ all vertices of the corresponding sub-cycle except one of the two

occurrences of $(m' - 1, \mathbf{x}')$ to obtain $\mathcal{T}_l''$. Otherwise, set $\mathcal{T}_l'' = \mathcal{T}_l'$. Finally, we do the same *loop erasion* operation on $\mathcal{T}_l''$, if $(m' - 1, \mathbf{x}' + \mathbf{e}(m', \mathbf{x}'))$ appears twice in this sequence. The final result defines $\mathcal{T}_{l+1}$. The output of this algorithm is called the *Toom cycle rooted at* $(m, \mathbf{x})$. A number of combinatorial properties of this object are needed. All of their proofs are fairly simple, but fiddly, so we refer to [193] for the details. Firstly, the Toom cycle is well defined, contains its root and its increments belong to

$$\{(-1, 0, 0), (1, 0, 0), (-1, 1, 0), (1, -1, 0), (-1, 0, 1), (1, 0, -1)\}.$$

Consequently, the number of Toom cycles with given root of length l is at most 6^l. Secondly, for each sink $(m', \mathbf{x}')$, the space-time point $(m' - 1, \mathbf{x}')$ is not good and each sink appears only once in the cycle.

   Observe that there are three types of vertices in a Toom cycle: the ones with two neighbours with smaller time coordinate, those with two neighbours with larger time coordinate and the others. We call them *sources*, *sinks* and *internal vertices*. The number n_* of sources and the number of sinks are the same by double counting. One can prove that internal vertices $(m_i, \mathbf{x}_i)$ such that their neighbours are not sinks or sources satisfy

$$((m_{i+1} - m_i), \langle \mathbf{1}, \mathbf{x}_{i+1} - \mathbf{x}_i \rangle) = ((m_i - m_{i-1}), \langle \mathbf{1}, \mathbf{x}_i - \mathbf{x}_{i-1} \rangle) \in \{(-1, 0), (1, -1)\}.$$

When the above quantity is $(1, -1)$ (resp. $(-1, 0)$), we call the internal vertex *blue* (resp. *red*) and denote the number of such vertices n_b (resp. n_r). Examining the time increments, it is clear that $|n_b - n_r| \leqslant 4n_*$. On the other hand, by examining the projected space increments $\langle \mathbf{1}, \mathbf{x}_{i+1} - \mathbf{x}_i \rangle$, we see that $n_b \leqslant 6n_*$. Combining these facts, we obtain that the total length of the Toom cycle is at most $6n_* + n_b + n_r \leqslant 22n_*$.

   As noted above, sinks are not good and distinct, so the probability that all sinks of a Toom cycle are bad is at most $(\varepsilon')^{n_*}$. We are then able to use a union bound

$$\mathbb{P}(\xi_{\mathbf{x}}(m) = 1) \leqslant \varepsilon' + \sum_{l \geqslant 1} 6^{2l}(\varepsilon')^{2l/22} < (\varepsilon')^{1/12},$$

for ε' small enough, taking into account that cycles have even length.

7.3.2.5 Chains of Toom cycles

In Section 7.3.2.4, we showed that for any space-time point $(m, \mathbf{x})$, the probability $\mathbb{P}(\xi_{\mathbf{x}}(m) = 1)$ that it is in state 1 in North-East BP with death parameter ε' and initial condition $\mathbf{0}$ is small. However, following Section 7.3.2.3, we need to prove that the probability of a connected component of such points is exponentially low in the diameter of the component. We already have one exponential bound, namely for a given space-time point, the probability that its Toom cycle has length more than $l \geqslant 1$ is at most $6^l(\varepsilon')^{l/22}$. It therefore remains to treat connected components of Toom cycles. Unfortunately, if Toom cycles for different points intersect at a sink,

we lose independence. In order to deal with this issue, we introduce the following notion of chain (of Toom cycles), following [124, Section 9.3].

Definition 7.26 (Chain) A *chain* is a finite sequence of vertex-disjoint Toom cycles T_i rooted at space-time points $(m_i, \mathbf{x}_i)$ of lengths l_i such that

$$d((m_i, \mathbf{x}_i), (m_{i+1}, \mathbf{x}_{i+1})) \leqslant 7(l_i + l_{i+1}). \tag{7.16}$$

The length of a chain is $\sum_i l_i$.

The next result [124, Lemma 9.15] shows that we can extract a chain from a large connected component of occupied sites.

Lemma 7.27 (Existence of chains) *There exists a constant $C > 0$ (depending only on the dimension, which we fixed equal to 2) such that the following holds. If the connected component $C'_{m,\mathbf{x}}$ of $(m, \mathbf{x})$ in X from (7.14) has diameter at least k, then there exists a chain of length $l \geqslant k/C$ contained in the ball $B_{(m,\mathbf{x})}(Cl)$ of radius Cl centered at $(m, \mathbf{x})$.*

Proof (Sketch) The construction of the chain is obtained algorithmically by pruning the sequence $(T_i)_{i=1}^N$ of Toom cycles rooted at the space-time points of a path from $(m, \mathbf{x})$ to the farthest point in its connected component $C'_{m,\mathbf{x}}$. Initially the Toom cycles in this sequence may intersect. We start by discarding the Toom cycles rooted at space-time points at distance at most $6l_1$ from $(m_1, \mathbf{x}_1) = (m, \mathbf{x})$ except T_1. We inspect the first remaining Toom cycle T_i with $i > 1$ (if any). If $T_1 \cap T_i \neq \varnothing$, we also remove T_1 from our sequence, in which case we observe that necessarily $l_i > 3l_1$, so

$$\{(m_j, \mathbf{x}_j) : j < i\} \subset B_{(m_1, \mathbf{x}_1)}(6l_1) \subset B_{(m_i, \mathbf{x}_i)}(6l_i).$$

If, on the contrary, $T_1 \cap T_i = \varnothing$, we keep both T_1 and T_i for the moment and observe that $d((m_1, \mathbf{x}_1), (m_i, \mathbf{x}_i)) \leqslant 6l_1 + \sqrt{2}$ by construction, since the roots initially form a connected path in the sense of (7.15).

In subsequent steps, we proceed similarly. Namely, we first remove the Toom cycles rooted at space-time points $(m_j \mathbf{x}_j)$, with $j > i$, at distance at most $6l_i$ from $(m_i, \mathbf{x}_i)$. Then, for the first remaining Toom cycle T_j with $j > i$, we remove those among the remaining ones with smaller index which intersect T_j. After the algorithm terminates, denoting the set of remaining indices of Toom cycles by I, one can show that $T_i \cap T_j = \varnothing$ for distinct $i, j \in I$ and

$$\bigcup_{i \in I} B_{(m_i, \mathbf{x}_i)}(6l_i)$$

is a connected set containing the initial path. This allows us to extract the desired chain. See [124] for more details. $\qquad\square$

With Lemma 7.27 at hand, the proof of Theorem 7.20 is nearly complete. We use a union bound over the possible chains. Recall that Toom cycles in a chain are disjoint (so independent) and their probability is at most $(\varepsilon')^{l/22}$ for a chain length l. It therefore remains to show that the number of possible chains of length l is bounded

by C^l for some constant C independent of ε'. Indeed, the number of chains modulo the choice of the roots of the Toom cycles is at most 6^l as discussed above. Finally, the number of choices for the roots can be bounded using (7.16) and the fact that $\mathbb{Z}^2$ does not have super-exponential growth.

Putting the above together, we obtain an exponential bound on $\mathbb{P}(\mathrm{diam}(C'_{m,\mathbf{x}}) \geqslant k)$, concluding the proof of Theorem 7.20 for FA-2f in two dimensions.

7.4 Other out-of-equilibrium results

We next briefly mention a few more works tackling KCM out of equilibrium from angles different from those discussed above.

7.4.1 The biased annihilating branching process

The biased annihilating branching process (BABP) is an interacting particle system closely related to FA-1f, to which several works have been devoted in the 1990s [168, 191, 192]. Here sites are updated to 0 (resp. 1) with rate proportional to the number of neighbouring empty sites. More precisely BABP has generator (2.4) with $c_{\mathbf{x}}$ replaced by $\tilde{c}_{\mathbf{x}}$, where

$$\tilde{c}_{\mathbf{x}}(\omega) := \sum_{i=1}^{d} \sum_{\varepsilon \in \{-1,1\}} (1 - \omega_{\mathbf{x}+\varepsilon\mathbf{e}_i}).$$

Like FA-1f, BABP is reversible w.r.t. μ and not attractive. Furthermore, for all $\omega \in \Omega$,

$$\tilde{c}_{\mathbf{x}}(\omega) = 0 \Leftrightarrow c_{\mathbf{x}}(\omega) = 0$$

with $c_{\mathbf{x}}$ the FA-1 constraint function (recall (2.3) and Section 2.2). Despite these similarities, BABP has some special features not shared by FA-1f, which make it more tractable. In particular, it enjoys a self-duality property (see [191, Eq. 34]) and quasi-duality with another model known as double flip process (see [192, Section 6]). Thanks to these features, which correspond to some special algebraic properties satisfied by the generator, Sudbury [192] established that for BABP in any dimension convergence to equilibrium (i.e. (7.2)) holds when $\nu = \mu_{q_0}$, provided $q_0 \neq 0$. Unfortunately, this result does not seem to provide any insight or tool for proving (7.2) at all vacancy densities for FA-1f.

7.4.2 FA-1f at low density

The FA-1f model has been particularly investigated, since it is the most accessible model beyond the East one. It has the notable advantage of being non-cooperative in the sense that a single empty site can move around and allow the configuration to be resampled. Moreover, as we saw in Section 4.1, empty sites move more or less like random walks, albeit possibly coalescing when they meet and occasionally branching. If one works in a suitably chosen finite volume with q small and only follows the process over a relatively short amount of time, it can be possible to track these random walks and deal with delicate collision events. More precisely, consider a 'critical' volume, that is, a torus (or any other bounded degree finite connected graph) $\Lambda = (\mathbb{Z}/n\mathbb{Z})^d$ of cardinal n^d of order $1/q$. Since the graph is connected, by Corollary 3.7, the FA-1f model is ergodic on $\Omega \setminus \{\mathbf{1}\}$, so we exclude the $\mathbf{1}$ configuration in the sequel. Typically, under μ_q, Λ only contains a bounded number of empty sites and they are well separated. If we follow the process up to a time horizon $1/q^2$, say, we do not expect to see more than three empty sites near each other. Since moving an empty site requires creating another one, this means that only binary collisions or branchings occur and these can be analysed.

Furthermore, the above reasoning can also work out of equilibrium, provided that one can control the length of the initial period of time it takes for the number of empty sites to drop down to order 1. This delicate strategy was employed by Pillai and Smith [174, 175] yielding that the mixing time in the above setting is of order q^{-2}, up to logarithmic corrections, for $d \geqslant 2$. As already discussed in the proof of Proposition 5.10, one can prove the same result, but also for other graphs by establishing a finite-volume logarithmic Sobolev inequality [119] (recall (2.15), Corollary 3.13). In fact, such inequalities already proved useful in [29] for studying FA-1f at high vacancy density on infinite graphs (e.g. $\mathbb{Z}^d$) with initial condition such that there are empty sites at bounded distance from all vertices of the graph. However, in the absence of attractiveness, it is usually difficult to relate finite and infinite volume results.

7.4.3 Large deviations in trajectory space

A radically different viewpoint consists in studying trajectories of KCM rather than their state at a given time. Given, e.g., a finite box $\Lambda = \{1, \ldots, n\}^d$ with suitable boundary condition $\sigma \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ and a time t, the *activity* $\mathcal{A}(t)$ is the total number of times any site changed its state up to time t. One is interested in the large deviation properties of $\mathcal{A}(t)$ as $t \to \infty$. That is, for $a \in [0, \infty)$, one expects that

$$f(a) = \lim_{\varepsilon \to 0} \lim_{n \to \infty} \lim_{t \to \infty} \frac{-\log \mathbb{P}_\mu(|\mathcal{A}(t) - a| < \varepsilon)}{n^d t}$$

exists and would like to study the properties of the large deviation rate f. The constraint of the KCM impacts the function f in that $f(a) = 0$ for all $a < \lim_{n \to \infty} \lim_{\to \infty} \mathbb{E}_\mu(\mathcal{A}(t))/(n^d t)$: low activity is easy to achieve. This can be seen as a constraint-induced dynamical phase transition. We refer the reader to [36, 37, 96, 97, 130, 164], where these and finer matters are investigated for FA-1f and East in one dimension.

7.4.4 Aging for the one-dimensional East

Remark 7.7 guarantees that, for $q \ll 1$, relaxation to equilibrium after a density quench (namely starting from $\nu = \mu_{q_0}$ with a fixed $q_0 \in (0, 1]$) occurs at an exponential rate on a time scale of order $T_{\text{rel}} \sim e^{c|\log q|^2}$ with $c = (2 \log 2)^{-1}$ (see Theorem 4.4). In this section we will discuss a peculiar behaviour that occurs at intermediate times $1 \ll t \ll T_{\text{rel}}$.

Fix $\epsilon, q \in (0, 1)$, and, for $n \geq 1$, set

$$
t_0 = 1, \qquad\qquad t_0^- = 0, \qquad\qquad t_0^+ = \left(\frac{1}{q}\right)^\epsilon,
$$

$$
t_n = \left(\frac{1}{q}\right)^n, \qquad\qquad t_n^- = t_n^{1-\epsilon}, \qquad\qquad t_n^+ = t_n^{1+\epsilon}. \qquad (7.17)
$$

The time intervals $[t_n^-, t_n^+]$ and $[t_n^+, t_{n+1}^-]$ are called n^{th} *active* and n^{th} *stalling* periods respectively. Note that in the limit $q \to 0$ there is a sharp separation of time scales, $t_n/t_{n'} \to 0$ for $n < n'$ and, if $\epsilon \ll 1$, the n^{th} stalling period is much longer than the n^{th} active period.

Suppose that $q \ll 1$ and initialise the process from a Bernoulli distribution at density $q_0 > q$. Most of the non-equilibrium evolution will try to remove the excess of empty sites present initially and will thus be dominated by the coalescence of domains corresponding to the intervals separating two consecutive empty sites. This process must occur in a cooperative way because, in order to remove an empty site, another empty site must be created to its right. Furthermore, recalling that creating an empty site at distance ℓ requires creating $\sim \log_2 \ell$ additional empty sites (see Proposition 4.6) the following heuristic picture emerges

  (i) During the n^{th} active period, only the empty sites with another empty site to their right at distance less than 2^n can be removed;
 (ii) at the end of the n^{th} active period, no empty sites at distance less than $2^n + 1$ are present any more, and there is no empty site that was not present at the beginning of the period;
(iii) during the n^{th} stalling period, nothing happens: none of the empty sites present at the beginning is destroyed and no new empty site is present at the end.

The above heuristics, set out by Evans and Sollich in two physics papers [189, 190], was turned into rigorous results by Faggionato, Martinelli, Roberto and Toninelli

in [85]. This implies that the vacancy density displays a peculiar staircase behavior (see [85, Theorem 2.5(i)]) and the two-time autocorrelation function depends in a non trivial way on the two times and not just on their difference (see [85, Theorem 2.5(i)]). This phenomenon is known as *aging*. In the same paper, sharp results for the statistics of the interval between two consecutive zeros in the different periods were established (see[85, Theorem 2.6]). The key ingredients for these results are:

- proving that the non equilibrium dynamics of the East model starting from a renewal process is well approximated when $q \downarrow 0$ by a hierarchical coalescence process whose rates depend on large deviation probabilities of the East model;
- universality results for the scaling limit of this coalescence process (see [86]).

Problem 7.28 It could be interesting to investigate whether a staircase behaviour for local functions and aging for two-time functions hold also for East in higher dimensions or for other KCM featuring a sharp separation of time scales, in particular for supercritical rooted models where logarithmic energy barrier also occur (see Proposition 6.13).

7.5 Basic open problems

The results discussed in the previous sections leave several basic questions open.

7.5.1 Ergodic regime

Fix a dimension $d \geqslant 1$, a d-dimensional update family $\mathcal{U}$ and $q > q_c(\mathcal{U})$. The following conjectures can be viewed either as an extension to the whole set of $\mathcal{U}$ of the results of Section 7.2 for the East model or as extension to the whole ergodic regime of the high temperature results of Section 7.3.

Conjecture 7.29 (Convergence to equilibrium after a temperature quench) Let ν be the Bernoulli measure μ_{q_0}, $q_0 > q_c$. Then,

   (i) for any local function f, (7.2) holds;
  (ii) furthermore, the convergence occurs exponentially fast.

Conjecture 7.30 (Linear pre-cutoff) There exists C such that(7.3) holds for any $\delta \in (0, 1)$ and n large enough.

Conjecture 7.31 (Cutoff) There exists ν such that (7.4) holds.

7.5.2 Beyond the ergodic regime

Fix a dimension $d \geqslant 1$, a d-dimensional update family $\mathcal{U}$ and $q < q_c(\mathcal{U})$.

Problem 7.32 (Quench from high to low temperature) Let the initial distribution ν be μ_{q_0} with $q_0 > q_c$. Since a.s. the initial state has closure **0**, this will be true at any later time (see Lemma 3.5). However, it is natural to expect a behaviour dominated by the growth of occupied "clusters" that can be unblocked only from their boundary. Is this intuition correct and how fast do these clusters grow? E.g. how does τ_0^{BP} of the current configuration scale with time? Another natural question is whether, locally, the measure converges to the equilibrium density, namely whether $\lim_{t\to\infty} \mathbb{E}_{\mu_{q_0}}(\eta_0(t)) = 1 - q$. In principle, this is still compatible with the fact that the closure is **0** at any time. See however [48, Section 4] for an example of KCM defined on trees (see Section 8.1) for which such a result is ruled out.

The inverse regime, $q_0 < q_c < q$, is often uninteresting, because the set of sites which can be updated may partition into an infinite collection of independent finite Markov chains. However, in some cases, an infinite chain remains. This leads to a setting similar to inhomogeneous KCM, discussed in Section 4.4 and 8.2 below.

7.6 Techniques

Finally, let us review the new techniques encountered in this chapter. In fact, the only one that was not new (recall Section 3.6) is renormalisation used in the proof of the general case of Theorems 7.19 and 7.20.

Distinguished empty site. Definition 7.14 provided this essential tool for the analysis of the one dimensional East model. It can be seen as a boundary condition moving to the right and leaving equilibrium to its left.

Couplings. In Sections 7.3.2.1-7.3.2.3, we saw useful couplings between KCM and various attractive processes: cooperative contact processes, BP with death and last passage percolation, all based on the graphical construction of the models. The drawback of such comparisons is that, in order for them to be useful, the attractive process may need to be supercritical, leading to an artificial limitation to the high vacancy density regime.

Toom cycles. This combinatorial tool for proving stability of BP with death, discussed in Section 7.3.2.4, has broader applications. Namely, suitable extensions can be used for perturbations of arbitrary cellular automata and, in some cases, interacting particle systems [193].

Chains. In order to control connected components rather than individual sites, in Section 7.3.2.5, we relied on chains. This is a rather general percolation method not specific to Toom cycles.

Cutoff strategy. In both Corollary 7.10 and Theorem 7.21 the same strategy was used to prove a linear time cutoff with square root window. While the entire programme has only been implemented for specific models in one dimension, it is, in principle, feasible for any KCM. The first step is proving a (possibly stretched) exponential convergence to equilibrium as in Theorem 7.19. Then one shows a positive speed result as in Theorem 7.20. The third step is showing ergodicity for the process seen

from the front and a law of large numbers for the front position. Finally, a central limit theorem for the front can be sought in order to control the cutoff window.

Chapter 8
Related settings and models

In this final chapter, we conclude by surveying a few additional settings not covered by the models defined in Chapter 2, but strongly related. Indeed, so far we have allowed general update families and, in Chapter 7, non-equilibrium initial conditions. However, we have restricted our attention to the equilibrium measure μ being product, constraints being identical at all sites, dynamics changing the state of a single site at a time, the underlying graph being a d-dimensional lattice. Each of these hypotheses may be revoked and leads to interesting models and questions, many of which have not yet been explored.

8.1 KCM on other graphs

KCM can also be defined on graphs different from $\mathbb{Z}^d$, including arbitrary graphs [45, 119], trees [47, 48, 160, 173], hyperbolic lattices [180] and many more such as Bienaymé–Galton–Watson trees or various models of random graphs waiting to be explored. See e.g. [3] for a possible application of FA-1f to information storage in a sensor network. The most studied case is FA-jf on oriented or non-oriented regular trees of degree $k + 1$. In the non-oriented version, the constraint at site $\mathbf{x}$ is satisfied if at least j of the neighbours of $\mathbf{x}$ in the tree are empty. In the oriented version, at least j empty sites should be among the k children. As for KCM on $\mathbb{Z}^d$, the ergodicity thresholds for the KCM and for the corresponding BP dynamics coincide (see Theorem 3.9). Thanks to the tree structure, it is not difficult to write recursive equations for the critical thresholds [19], yielding that, both for oriented and non-oriented models, we have $q_c = 1$ for $j > k$, $q_c = 0$ for $j = 1$ and $0 < q_c < 1$ for $j \in \{2, \ldots, k\}$. In [160], Martinelli and Toninelli prove that $T_{rel} < \infty$ for all models in the whole ergodic regime, $q > q_c$. In [47], the same authors together with Cancrini and Roberto, analyse the scaling in the critical regime in the case $j = k$ for the oriented model and prove a power law divergence of T_{rel} as $q \downarrow q_c$. An analogous scaling is also conjectured in all cases $j \in \{2, \ldots, k - 1\}$, but has yet to be proven. A fundamental difference between the case $j = k$ and $j \in [2, k - 1]$,

at the root of the difficulties to handle the missing cases, is that the BP transition is continuous for $j = k$ (it essentially correspond to a standard percolation transition), while for all $j \in \{2, \ldots, k-1\}$ it is discontinuous, namely $\mu_{q_c}(\tau_0^{BP} = \infty) > 0$. Using a strategy similar to the proof of Theorem 7.6, for the oriented model with $k = j = 2$, exponential convergence to equilibrium when $q > q_c$ starting from an initial distribution $\mu_{q'}$ with $q' > q_c$ is proven in [48] (see Theorem 4.3 therein together with [160, Theorem 2]). The proof can be readily extended to all oriented models for $j = 1$ or $j = k$. The result is conjectured to hold also for the remaining cases $j \in [2, k]$.

8.2 Inhomogeneous KCM

We briefly considered an inhomogeneous setting in Section 4.4, where the update family defining the constraint is allowed to depend on the site. While our treatment in one dimension is very general, one can also consider such inhomogeneous models in higher dimensions, possibly choosing the update families at random. For instance, one may consider sites $\mathbf{x} \in \mathbb{Z}^d$ having FA-$j_{\mathbf{x}}$f constraint with $j_{\mathbf{x}}$ chosen i.i.d. at random according to some distribution on $\{0, \ldots, 2d\}$. A few such models are studied in [183, 186], but most remain unexplored.

8.3 KCM with interactions

Another natural modification of KCM is to introduce static interactions between occupied sites. This may be achieved by updating each site w.r.t. a measure depending on the current state of other sites. For instance, one could consider the $\mathcal{U}$-constrained Glauber dynamics for the Ising model with inverse temperature β with generator (2.4), where, instead of $\mu_{\mathbf{x}}(\omega_{\mathbf{x}} = 0) = q$, we set

$$\mu_{\mathbf{x}}(\omega_{\mathbf{x}} = 0) = \frac{1}{1 + \exp(\beta \sum_{\mathbf{y} \sim \mathbf{x}} (2\omega_{\mathbf{y}} - 1))},$$

the sum running over nearest neighbours of $\mathbf{x}$ in $\mathbb{Z}^d$. In fact, this was already considered in [89] together with an external magnetic field. More generally Gibbs measures were considered in [45, Section 5]. While the initial motivation behind KCM is to investigate the extent to which glassy phenomenology can be explained by purely dynamical means, in reality interactions are certainly present. It is therefore interesting, but probably challenging, to study such models.

8.4 Plaquette models

Instead of kinetic constraints, *plaquette models* have static interactions as in Section 8.3, but of multi-body type. They were introduced in [148, 169] to show that kinetic constraints can emerge from static interactions at low temperatures (also see [94, 202]). An example is the *square plaquette model* on $\mathbb{Z}^2$. For any $\mathbf{x} \in \mathbb{Z}^2$, the plaquette of $\mathbf{x}$ in this model is the square $P_\mathbf{x} = \{\mathbf{x}, \mathbf{x} + \mathbf{e}_1, \mathbf{x} + \mathbf{e}_2, \mathbf{x} + \mathbf{e}_1 + \mathbf{e}_2\}$. The Gibbs weights are defined by the Hamiltonian

$$- \sum_{\mathbf{x} \in \mathbb{Z}^2} \prod_{\mathbf{y} \in P_\mathbf{x}} (2\omega_\mathbf{y} - 1)$$

and one then considers an unconstrained single site Glauber dynamics. The behaviour of the square plaquette model turns out to be similar to that of FA-1f, while a similar triangular plaquette model with plaquettes of the form $P_\mathbf{x} = \{\mathbf{x}, \mathbf{x} + \mathbf{e}_1, \mathbf{x} + \mathbf{e}_1 + \mathbf{e}_2\}$ is conjectured to behave like the East KCM. Work on these models can be found in [58, 60, 61], following physics studies (see e.g. [98, 128, 129]).

8.5 Conservative models

The physical motivation behind KCM (recall Chapter 1) views sites of $\mathbb{Z}^d$ as mesoscopic volumes whose particle density may be lower or higher, as reflected by the state of the site. However, if we take a microscopic perspective, it is more natural to consider constrained models in which the number of particles is conserved. The first and most classical such models are the Kob–Andersen ones [137]. In KA-jf, one may exchange the states of any two neighbouring sites, $\mathbf{x}$ and $\mathbf{y}$, provided they both have at least $j - 1$ empty neighbours in $\mathbb{Z}^d \setminus \{\mathbf{x}, \mathbf{y}\}$. The case $j = 1$ coincides with the simple symmetric exclusion process (SSEP) and we will disregard it in the following. One can similarly define conservative versions of $\mathcal{U}$-KCM, obtaining the class of *kinetically constrained lattice gases* (KCLG) (see [46, Section 2] for a formal definition of this class). It is immediate to verity that for any $q \in [0, 1]$, μ_q is a reversible measure for these dynamics. As for the non conservative models, we define q_c by (2.13). In [46, Proposition 2.16], a conservative analogue of Theorem 3.9 is established stating that for any KCLG the ergodicity threshold coincides with the *excheangeability threshold* defined as the minimal value of q above which, for μ_q-almost every configuration and for any couple of sites, there exists a legal path exchanging their occupation variables. The first mathematical result on KCLG is due to Biroli, Fisher and Toninelli [194, 197] who proved that $q_c = 0$ for any KA-jf with $j \in \{2, \ldots, d\}$.[1] The key ingredients of this proof are:

- the construction of a set of configurations, the so called *frameable configurations*, which can be connected by a legal path to a configuration with well-chosen bound-

---

[1] The cases $j > d$ trivially lead to $q_c = 1$.

ary state (see [159, Definition 3.3] and [197]) for the dynamics on finite volume
with occupied boundary condition;

- the construction of legal paths connecting any two frameable configurations with
  equal number of empty sites;
- the fact that the μ_q-probability that a configuration is frameable goes sufficiently
  fast to one as the volume increases (see [159, Proposition 3.26]). More precisely ,
  there exist $c_+, c_- > 0$ dependent on d, j such that, setting

$$\Xi_\pm(q, j, d) = \exp^{\circ(j-1)}\left(\frac{c_\pm}{q^{\frac{1}{d-j+1}}}\right),$$

when $q \to 0$ and $L \to \infty$ faster (resp. slower) than $\Xi_+(q, j, d)$ (resp. $\Xi_-(q, j, d)$),
the probability of being frameable goes to one (resp. to zero).

Combining these results with canonical paths, renormalization arguments and tools
borrowed from oriented percolation, Martinelli, Shapira and Toninelli [159, Theorem
1] prove that the spectral gap of the KA-jf models on $\{1, \ldots, L\}^d$ with unconstrained
sources at the boundary in any dimension d and for any $j \in [2, d]$ scales diffusively
(as for SSEP), namely as L^{-2}, and with a density pre-factor of the form $\Xi_\pm(q, j, d)^{-1}$
(while for SSEP there is no density pre-factor).

In [34], Blondel and Toninelli consider the behavior of a tagged particle, and prove
(following the ideas sketched in [194]), that for all $d \geqslant 2$ and for any $j \in [2, d]$,
diffusive behaviour holds at any density $q \in (0, 1)$. More precisely, if we distribute
the configuration according to μ_q, condition on the presence of a particle initially at
the origin, tag it and denote by $\mathbf{X}_t$ its position at time t, for some matrix $D(q)$ such
that $\mathbf{e}_i \cdot D(q)\mathbf{e}_i > 0$ for all $i \in \{1, \ldots, d\}$, it holds that

$$\lim_{\epsilon \to 0} \epsilon \mathbf{X}_{\epsilon^{-2}t} = \sqrt{2D(q)}B_t, \tag{8.1}$$

where B_t is a standard d-dimensional Brownian motion and the convergence holds in
the sense of weak convergence of path measures. This contradicts conjectures based
on numerical simulations [137, 140] claiming the occurrence of a diffusive/non-
diffusive transition. Furthermore Ertul and Shapira prove (see [84, Theorem 2.3])
upper and lower bound for $D(q)$ of the form $\Xi_\pm(q, j, d)^{-1}$ (modulo a logarithmic
correction in the case $d = 2$). The fast shrinking to zero explains why it was
incorrectly inferred from numerical simulation ([137]) that for $d = 3$ and $j = 3$ a
diffusive/non-diffusive transition would occur.

Problem 8.1 Two other natural issues in the conservative setting are

 (i) determining the evolution of macroscopic density profiles, namely establishing
     the hydrodynamic limit, and the fluctuations around these profiles;
(ii) establishing relaxation at equilibrium in infinite volume.

Concerning (i), a natural candidate for the hydrodynamic limit is a parabolic equa-
tion of porous media that degenerates when the density approaches one. As for
fluctuations, it is reasonable to expect they would be Edward–Wilkinson Gaussian

fluctuations as for SSEP [67]. Establishing these results in the presence of constraints is particularly challenging (see [104] where this is achieved for a different KCLG). Concerning (ii), a first result [48] shows that there exists $C(q) > 0$ such that for any local function f it holds that, for all $t > 0$,

$$\mathrm{Var}_{\mu_q}(P_t f) \leqslant \frac{C(q)\|f\|_\infty}{t}.$$

We expect the correct behavior to be of the form $t^{-d/2}$, as for SSEP.

Other kinetically constrained lattice gases can be found in [22, 42, 68, 82, 103, 104, 132, 167, 185]. A model that is currently very actively considered is known as the *facilitated exclusion process*. In this one-dimensional model, a particle is allowed to jump to a neighbouring empty site, provided its other neighbour is occupied. We direct the reader to [12, 31, 32, 66, 79–81, 102, 162] for work on this topic.

8.6 Tracer diffusion

In the previous section we described the tagged particle behavior for KCLG. Though the non-conservative dynamics of KCM is not diffusive, one can define the following similar problem. Consider a stationary KCM evolving from a configuration distributed according to μ_q and inject at time zero a particle (the tracer) at the origin. The tracer moves like a modified random walk attempting to jump at rate one to a site chosen uniformly at random among its nearest neighbours, with the jump being allowed if and only if both the sites occupied by the walker before and after the move are empty (see [33] for a precise definition). Note that the KCM constitutes a dynamical random environment in which the tracer evolves, and is not influenced by the motion of the tracer. Blondel proved (see [33, Proposition 3.1 and 3.2]) that if the underlying KCM has a positive spectral gap, the tracer has a diffusive behavior with a non-degenerate diffusion matrix, namely (8.1) holds. Furthermore for the FA-1f model in any dimension $d \geqslant 1$, it holds that $\mathbf{e}_i \cdot D(q)\mathbf{e}_i \sim q^2$ for $q \downarrow 0$ and $i \in \{1, \ldots d\}$. Instead, for the East model [33, Theorem 3.5] proves that D scales as the spectral gap (modulo power law corrections in q). This corrects the conjecture that had been put forward by physicists ([134, 135]) affirming that for the East model D would scale as $T_{\mathrm{rel}}^{-\xi}$ with $\xi < 1$.

Asymmetric tracers on stationary KCM have also been the object of investigation. See for example [11] for results on a tracer on the one-dimensional East model and with positive (resp. negative) drift when on its current position the occupation variable of the East model is occupied (resp. empty).

8.7 Upper triangular matrix walk

We conclude with a further context in which KCM arise naturally beyond the study of glassy dynamics and interacting particle systems. Let G_n be the group of $n \times n$ upper triangular matrices with entries in the two-element field $\mathbb{F}_2$ and ones on the diagonal. The following Markov chain was considered in [63]. At each step, with probability $1/2$ nothing happens and, for each $i \in \{1, \ldots, n-1\}$, with probability $1/(2n-2)$, we add row $i+1$ to row i. This corresponds to performing a lazy random walk on G_n with generator set $(I + E_{i,i+1})_{i=1}^{n-1}$. If we restrict our attention to column j of the matrix, this Markov chain becomes exactly the East process with parameter $q = 1/2$ on the segment $\Lambda_j = \{1, \ldots, j-1\}$ with boundary condition $\mathbf{0}_{\mathbb{Z} \backslash \Lambda_j}$. We direct the reader to [93, 172] for works on this random walk.

References

[1] M. Aizenman and D. J. Barsky, *Sharpness of the phase transition in percolation models*, Comm. Math. Phys. **108** (1987), no. 3, 489–526 pp. MR874906

[2] M. Aizenman and J. L. Lebowitz, *Metastability effects in bootstrap percolation*, J. Phys. A **21** (1988), no. 19, 3801–3813 pp. MR968311

[3] D. Aldous, *Constrained Ising models and speculative application to sensor networks*, 2005. Presentation available at `https://www.stat.berkeley.edu/~aldous/Talks/persi-ising.pdf`.

[4] D. Aldous and P. Diaconis, *The asymmetric one-dimensional constrained Ising model: rigorous results*, J. Stat. Phys. **107** (2002), no. 5-6, 945–975 pp. MR1901508

[5] D. Aldous and J. A. Fill, *Reversible Markov chains and random walks on graphs*, 2002. Unfinished monograph, recompiled 2014, available at `http://www.stat.berkeley.edu/~aldous/RWG/book.html`.

[6] G. Alon and G. Kozma, *Comparing with octopi*, Ann. Inst. Henri Poincaré Probab. Stat. **56** (2020), no. 4, 2672–2685 pp. MR4164852

[7] P. W. Anderson, *Through the glass lightly*, Science **267** (1995), no. 5204, 1615–1616 pp.

[8] F. Arceri, F. P. Landes, L. Berthier, and G. Biroli, *A statistical mechanics perspective on glasses and aging*, Encyclopedia of Complexity and Systems Science, 2021, 1–68 pp.

[9] D. J. Ashton, L. O. Hedges, and J. P. Garrahan, *Fast simulation of facilitated spin models*, J. Phys. Conf. Ser. **40** (2006), 99–104 pp. MR2299904

[10] A. Asselah and P. Dai Pra, *Quasi-stationary measures for conservative dynamics in the infinite lattice*, Ann. Probab. **29** (2001), no. 4, 1733–1754 pp. MR1880240

[11] L. Avena, O. Blondel, and A. Faggionato, *A class of random walks in reversible dynamic environments: antisymmetry and applications to the East model*, J. Stat. Phys. **165** (2016), no. 1, 1–23 pp. MR3547832

[12] J. Ayre and P. Chleboun, *Mixing times for the facilitated exclusion process*, arXiv e-prints (2024), available at arXiv:2402.18999.

[13] P. Balister, B. Bollobás, R. Morris, and P. Smith, *Universality for monotone cellular automata*, arXiv e-prints (2022), available at arXiv:2203.13806.

[14] P. Balister, B. Bollobás, R. Morris, and P. Smith, *Subcritical monotone cellular automata*, Random Structures Algorithms **64** (2024), no. 1, 38–61 pp. MR4672996

[15] P. Balister, B. Bollobás, R. Morris, and P. Smith, *The critical length for growing a droplet*, Mem. Amer. Math. Soc. (To appear).

[16] P. Balister, B. Bollobás, M. Przykucki, and P. Smith, *Subcritical $\mathcal{U}$-bootstrap percolation models have non-trivial phase transitions*, Trans. Amer. Math. Soc. **368** (2016), no. 10, 7385–7411 pp. MR3471095

[17] J. Balogh, B. Bollobás, H. Duminil-Copin, and R. Morris, *The sharp threshold for bootstrap percolation in all dimensions*, Trans. Amer. Math. Soc. **364** (2012), no. 5, 2667–2701 pp. MR2888224

[18] J. Balogh, B. Bollobás, and R. Morris, *Bootstrap percolation in three dimensions*, Ann. Probab. **37** (2009), no. 4, 1329–1380 pp. MR2546747

[19] J. Balogh, Y. Peres, and G. Pete, *Bootstrap percolation on infinite trees and non-amenable groups*, Combin. Probab. Comput. **15** (2006), no. 5, 715–730 pp. MR2248323

[20] L. Berthier and G. Biroli, *Theoretical perspective on the glass transition and amorphous materials*, Rev. Mod. Phys. **83** (2011), no. 2, 587–645 pp.

[21] L. Berthier, G. Biroli, J.-P. Bouchaud, L. Cipelletti, and W. van Saarloos, *Dynamical heterogeneities in glasses, colloids, and granular media*, Oxford University Press, 2011.

[22] L. Bertini and C. Toninelli, *Exclusion processes with degenerate rates: convergence to equilibrium and tagged particle*, J. Stat. Phys. **117** (2004), no. 3-4, 549–580 pp. MR2099727

[23] C. Bezuidenhout and G. Grimmett, *The critical contact process dies out*, Ann. Probab. **18** (1990), no. 4, 1462–1482 pp. MR1071804

[24] N. Bhatnagar, P. Caputo, P. Tetali, and E. Vigoda, *Analysis of top-swap shuffling for genome rearrangements*, Ann. Appl. Probab. **17** (2007), no. 4, 1424–1445 pp. MR2344312

[25] P. Billingsley, *Ergodic theory and information*, John Wiley & Sons, Inc., New York-London-Sydney, 1965. MR192027

[26] P. Billingsley, *Convergence of probability measures*, 2nd ed., Wiley series in probability and statistics: Probability and statistics, John Wiley & Sons, Inc., New York, 1999. MR1700749

[27] G. Biroli and J. P. Garrahan, *Perspective: the glass transition*, J. Chem. Phys. **138** (2013), no. 12, 12A301, 13 pp.

[28] O. Blondel, *Front progression in the East model*, Stochastic Process. Appl. **123** (2013), no. 9, 3430–3465 pp. MR3071385

[29] O. Blondel, N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Fredrickson-Andersen one spin facilitated model out of equilibrium*, Markov Process. Related Fields **19** (2013), no. 3, 383–406 pp. MR3156958

[30] O. Blondel, A. Deshayes, and C. Toninelli, *Front evolution of the Fredrickson-Andersen one spin facilitated model*, Electron. J. Probab. **24** (2019), Paper No. 1, 32 pp. MR3903501

[31] O. Blondel, C. Erignoux, M. Sasada, and M. Simon, *Hydrodynamic limit for a facilitated exclusion process*, Ann. Inst. Henri Poincaré Probab. Stat. **56** (2020), no. 1, 667–714 pp. MR4059004

[32] O. Blondel, C. Erignoux, and M. Simon, *Stefan problem for a nonergodic facilitated exclusion process*, Prob. Math. Phys. **2** (2021), no. 1, 127–178 pp. MR4404818

[33] O. Blondel and C. Toninelli, *Is there a fractional breakdown of the Stokes-Einstein relation in kinetically constrained models at low temperature?*, EPL **107** (2014), no. 2, 26005, 6 pp.

[34] O. Blondel and C. Toninelli, *Kinetically constrained lattice gases: tagged particle diffusion*, Ann. Inst. Henri Poincaré Probab. Stat. **54** (2018), no. 4, 2335–2348 pp. MR3865675

[35] S. G. Bobkov and P. Tetali, *Modified logarithmic Sobolev inequalities in discrete settings*, J. Theoret. Probab. **19** (2006), no. 2, 289–336 pp. MR2283379

[36] T. Bodineau, V. Lecomte, and C. Toninelli, *Finite size scaling of the dynamical free-energy in a kinetically constrained model*, J. Stat. Phys. **147** (2012), no. 1, 1–17 pp. MR2922756

[37] T. Bodineau and C. Toninelli, *Activity phase transition for constrained dynamics*, Comm. Math. Phys. **311** (2012), no. 2, 357–396 pp. MR2902193

[38] B. Bollobás, H. Duminil-Copin, R. Morris, and P. Smith, *The sharp threshold for the Duarte model*, Ann. Probab. **45** (2017), no. 6B, 4222–4272 pp. MR3737910

[39] B. Bollobás, H. Duminil-Copin, R. Morris, and P. Smith, *Universality for two-dimensional critical cellular automata*, Proc. Lond. Math. Soc. (3) **126** (2023), no. 2, 620–703 pp. MR4550150

[40] B. Bollobás, P. Smith, and A. Uzzell, *Monotone cellular automata in a random environment*, Combin. Probab. Comput. **24** (2015), no. 4, 687–722 pp. MR3350030

[41] E. Bolthausen, *On the central limit theorem for stationary mixing random fields*, Ann. Probab. **10** (1982), no. 4, 1047–1050 pp. MR672305

[42] L. Bonorino, R. de Paula, P. Gonçalves, and A. Neumann, *Hydrodynamics of porous medium model with slow reservoirs*, J. Stat. Phys. **179** (2020), no. 3, 748–788 pp. MR4099999

[43] A. Bušić, J. Mairesse, and I. Marcovici, *Probabilistic cellular automata, invariant measures, and perfect sampling*, Adv. in Appl. Probab. **45** (2013), no. 4, 960–980 pp. MR3161292

[44] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Kinetically constrained spin models*, Probab. Theory Related Fields **140** (2008), no. 3-4, 459–504 pp. MR2365481

[45] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Facilitated spin models: recent and new results*, Methods of contemporary mathematical statistical physics, 2009, 307–340 pp. MR2581609

[46] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Kinetically constrained lattice gases*, Comm. Math. Phys. **297** (2010), no. 2, 299–344 pp. MR2651901

[47] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Mixing time of a kinetically constrained spin model on trees: power law scaling at criticality*, Probab. Theory Related Fields **161** (2015), no. 1-2, 247–266 pp. MR3304752

[48] N. Cancrini, F. Martinelli, R. Schonmann, and C. Toninelli, *Facilitated oriented spin models: some non equilibrium results*, J. Stat. Phys. **138** (2010), no. 6, 1109–1123 pp. MR2601425

[49] R. Candelier, O. Dauchot, and G. Biroli, *Building blocks of dynamical heterogeneities in dense granular media*, Phys. Rev. Lett. **102** (2009), no. 8, 088001, 4 pp.

[50] R. Candelier, A. Widmer-Cooper, J. K. Kummerfeld, O. Dauchot, G. Biroli, P. Harrowell, and D. R. Reichman, *Spatiotemporal hierarchy of relaxation events, dynamical heterogeneities, and structural reorganization in a supercooled liquid*, Phys. Rev. Lett. **105** (2010), no. 13, 135702, 4 pp.

[51] P. Caputo, F. Martinelli, and F. L. Toninelli, *Mixing times of monotone surfaces and SOS interfaces: a mean curvature approach*, Comm. Math. Phys. **311** (2012), no. 1, 157–189 pp. MR2892467

[52] R. Cerf and E. N. M. Cirillo, *Finite size scaling in three-dimensional bootstrap percolation*, Ann. Probab. **27** (1999), no. 4, 1837–1850 pp. MR1742890

[53] R. N. Chacko, F. P. Landes, G. Biroli, O. Dauchot, A. J. Liu, and D. R. Reichman, *Dynamical facilitation governs the equilibration dynamics of glasses*, Phys. Rev. X **14** (2024), no. 3, 031012 pp.

[54] J. Chalupa, P. L. Leath, and G. R. Reich, *Bootstrap percolation on a Bethe lattice*, J. Phys. C **12** (1979), no. 1, L31–L35 pp.

[55] P. Chleboun, A. Faggionato, and F. Martinelli, *Time scale separation and dynamic heterogeneity in the low temperature East model*, Comm. Math. Phys. **328** (2014), no. 3, 955–993 pp. MR3201217

[56] P. Chleboun, A. Faggionato, and F. Martinelli, *Mixing time and local exponential ergodicity of the East-like process in $\mathbb{Z}^d$*, Ann. Fac. Sci. Toulouse Math **24** (2015), no. 4, 717–743 pp. MR3434253

[57] P. Chleboun, A. Faggionato, and F. Martinelli, *Relaxation to equilibrium of generalized East processes on $\mathbb{Z}^d$: renormalization group analysis and energy-entropy competition*, Ann. Probab. **44** (2016), no. 3, 1817–1863 pp. MR3502595

[58] P. Chleboun, A. Faggionato, F. Martinelli, and C. Toninelli, *Mixing length scales of low temperature spin plaquettes models*, J. Stat. Phys. **169** (2017), no. 3, 441–471 pp. MR3711606

[59] P. Chleboun and F. Martinelli, *Mixing time bounds for oriented kinetically constrained spin models*, Electron. Commun. Probab. **18** (2013), Paper No. 60, 9 pp. MR3084571

[60] P. Chleboun and A. Smith, *Mixing of the square plaquette model on a critical length scale*, Electron. J. Probab. **25** (2020), Paper No. 89, 53 pp. MR4136469

[61] P. Chleboun and A. Smith, *Cutoff for the square plaquette model on a critical length scale*, Ann. Appl. Probab. **31** (2021), no. 2, 668–702 pp. MR4254492

[62] F. Chung, P. Diaconis, and R. Graham, *Combinatorics for the East model*, Adv. Appl. Math. **27** (2001), no. 1, 192–206 pp. MR1835679

[63] D. Coppersmith and I. Pak, *Random walk on upper triangular matrices mixes rapidly*, Probab. Theory Related Fields **117** (2000), no. 3, 407–417 pp. MR1774070

[64] Y. Couzinié, *The multicolour East model*, arXiv e-prints (2022), available at arXiv:2207.09782.

[65] Y. Couzinié and F. Martinelli, *On a front evolution problem for the multidimensional East model*, Electron. J. Probab. **27** (2022). MR4505380

[66] H. Da Cunha, C. Erignoux, and M. Simon, *Hydrodynamic limit for a boundary-driven facilitated exclusion process*, arXiv e-prints (2024), available at arXiv:2401.16535.

[67] A. De Masi, N. Ianiro, A. Pellegrinotti, and E. Presutti, *A survey of the hydrodynamical behavior of many-particle systems*, Nonequilibrium phenomena, II, 1984, 123–294 pp. MR757003

[68] R. de Paula, P. Gonçalves, and A. Neumann, *Energy estimates and convergence of weak solutions of the porous medium equation*, Nonlinearity **34** (2021), no. 11, 7872–7915 pp. MR4331252

[69] P. G. Debenedetti and F. H. Stillinger, *Supercooled liquids and the glass transition*, Nature **410** (2001), no. 6825, 259–267 pp.

[70] P. Diaconis and L. Saloff-Coste, *Logarithmic Sobolev inequalities for finite Markov chains*, Ann. Appl. Probab. **6** (1996), no. 3, 695–750 pp. MR1410112

[71] J. A. M. S. Duarte, *Simulation of a cellular automaton with an oriented bootstrap rule*, Phys. A **157** (1989), no. 3, 1075–1079 pp.

[72] R. M. Dudley, *Real analysis and probability*, Cambridge studies in advanced mathematics, vol. 74, Cambridge University Press, Cambridge, 2002. MR1932358

[73] H. Duminil-Copin and V. Tassion, *A new proof of the sharpness of the phase transition for Bernoulli percolation on $\mathbb{Z}^d$*, Enseign. Math. **62** (2016), no. 1, 199–206 pp. MR3605816

[74] R. Durrett, *On the growth of one-dimensional contact processes*, Ann. Probab. **8** (1980), no. 5, 890–907 pp. MR586774

[75] R. Durrett, *Oriented percolation in two dimensions*, Ann. Probab. **12** (1984), no. 4, 999–1040 pp. MR757768

[76] R. Durrett, *Some peculiar properties of a particle system with sexual reproduction*, Stochastic spatial processes (Heidelberg, 1984), 1986, 106–111 pp. MR877768

[77] R. Durrett, *Probability: Theory and examples*, 5th ed., Cambridge series in statistical and probabilistic mathematics, vol. 49, Cambridge University Press, Cambridge, 2019. Originally published by Wadsworth & Brooks/Cole Advanced Books & Software, Pacific Grove, CA (1991). MR3930614

[78] M. D. Ediger, *Spatially heterogeneous dynamics in supercooled liquids*, Annu. Rev. Phys. Chem. **51** (2000), no. 1, 99–128 pp.

[79] C. Erignoux and B. Massoulié, *Cutoff for the transience and mixing time of a SSEP with traps and consequences on the FEP*, arXiv e-prints (2024), available at arXiv:2403.20010.

[80] C. Erignoux, A. Roget, A. Shapira, and M. Simon, *Hydrodynamic behavior near dynamical criticality of a facilitated conservative lattice gas*, arXiv e-prints (2024), available at arXiv:2403.09324.

[81] C. Erignoux and L. Zhao, *Stationary fluctuations for the facilitated exclusion process*, arXiv e-prints (2023), available at arXiv:2305.13853.

[82] W. Ertel, K. Froböse, and J. Jäckle, *Constrained diffusion dynamics in the hard-square lattice gas at high density*, J. Chem. Phys. **88** (1988), no. 8, 5027–5034 pp.

[83] A. Ertul, *Cutoff for the Fredrickson-Andersen one spin facilitated model*, ALEA Lat. Am. J. Probab. Math. Stat. **19** (2022), no. 1, 511–536 pp. MR4394307

[84] A. Ertul and A. Shapira, *Self-diffusion coefficient in the Kob-Andersen model*, Electron. Commun. Probab. **26** (2021), Paper No. 3, 12 pp. MR4218031

[85] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *Aging through hierarchical coalescence in the East model*, Comm. Math. Phys. **309** (2012), no. 2, 459–495 pp. MR2864800

[86] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *Universality in one-dimensional hierarchical coalescence processes*, Ann. Probab. **40** (2012), no. 4, 1377–1435 pp. MR2978129

[87] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *The East model: recent results and new progresses*, Markov Process. Related Fields **19** (2013), no. 3, 407–452 pp. MR3156959

[88] W. Feller, *An introduction to probability theory and its applications. Vol. I*, 3rd ed., John Wiley & Sons, Inc., New York-London-Sydney, 1968. MR228020

[89] G. H. Fredrickson and H. C. Andersen, *Kinetic Ising model of the glass transition*, Phys. Rev. Lett. **53** (1984), no. 13, 1244–1247 pp.

[90] G. H. Fredrickson and H. C. Andersen, *Facilitated kinetic Ising models and the glass transition*, J. Chem. Phys. **83** (1985), no. 11, 5822–5831 pp.

[91] K. Froböse, *Finite-size effects in a cellular automaton for diffusion*, J. Stat. Phys. **55** (1989), no. 5-6, 1285–1292 pp. MR1002492

[92] S. Ganguly, E. Lubetzky, and F. Martinelli, *Cutoff for the East process*, Comm. Math. Phys. **335** (2015), no. 3, 1287–1322 pp. MR3320314

[93] S. Ganguly and F. Martinelli, *Upper triangular matrix walk: cutoff for finitely many columns*, Random Structures Algorithms **55** (2019), no. 2, 313–341 pp. MR3983784

[94] J. P. Garrahan, *Glassiness through the emergence of effective dynamical constraints in interacting systems*, J. Phys.: Condens. Matter **14** (2002), no. 7, 1571–1579 pp.

[95] J. P. Garrahan, *Aspects of non-equilibrium in classical and quantum systems: slow relaxation and glasses, dynamical large deviations, quantum non-ergodicity, and open quantum dynamics*, Phys. A **504** (2018), 130–154 pp. MR3805503

[96] 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** (2007), no. 19, 195702, 4 pp.

[97] 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** (2009), no. 7, 075007, 34 pp. MR2525457

[98] J. P. Garrahan and M. E. J. Newman, *Glassiness and constrained dynamics of a short-range nondisordered spin model*, Phys. Rev. E **62** (2000), no. 6, 7670–7678 pp.

[99] P. Garrahan, P. Sollich, and C. Toninelli, *Kinetically constrained models*, Dynamical heterogeneities in glasses, colloids and granular media and jamming transitions, 2011, 341–369 pp.

[100] H.-O. Georgii, *Gibbs measures and phase transitions*, 2nd ed., De gruyter studies in mathematics, vol. 9, Walter de Gruyter & Co., Berlin, 2011. MR2807681

[101] S. H. Glarum, *Dielectric relaxation of isoamyl bromide*, J. Chem. Phys. **33** (1960), no. 3, 639–643 pp.

[102] S. Goldstein, J. L. Lebowitz, and E. R. Speer, *The discrete-time facilitated totally asymmetric simple exclusion process*, Pure Appl. Funct. Anal. **6** (2021), no. 1, 177–203 pp. MR4213301

[103] P. Gonçalves, *Microscopic dynamics for the porous medium equation*, Dynamics, games and science. II, 2011, 387–392 pp. MR2883293

[104] P. Gonçalves, C. Landim, and C. Toninelli, *Hydrodynamic limit for a particle system with degenerate rates*, Ann. Inst. Henri Poincaré Probab. Stat. **45** (2009), no. 4, 887–909 pp. MR2572156

[105] M. Gottschau, M. Heydenreich, K. Matzke, and C. Toninelli, *Phase transition for a non-attractive infection process in heterogeneous environment*, Markov Process. Related Fields **24** (2018), no. 1, 39–56 pp. MR3793942

[106] J. Gravner and D. Griffeath, *Scaling laws for a class of critical cellular automaton growth rules*, Random walks (Budapest, 1998), 1999, 167–186 pp. MR1752894

[107] J. Gravner and A. E. Holroyd, *Local bootstrap percolation*, Electron. J. Probab. **14** (2009), Paper No. 14, 385–399 pp. MR2480546

[108] L. F. Gray, *Toom's stability theorem in continuous time*, Perplexing problems in probability, 1999, 331–353 pp. MR1703139

[109] S. Greenberg, A. Pascoe, and D. Randall, *Sampling biased lattice configurations using exponential metrics*, Proc. Twent. Annu. ACM-SIAM Symp. Discrete Algorithms, 2009, 76–85 pp. MR2809307

[110] G. Grimmett, *Percolation*, Second edition, Grundlehren der mathematischen Wissenschaften, Springer, Berlin, Heidelberg, 1999. Originally published by Springer, New York (1989). MR1707339

[111] I. Hartarsky, *Bisection for kinetically constrained models revisited*, Electron. Commun. Probab. **26** (2021), Paper No. 60, 10 pp. MR4346864

[112] I. Hartarsky, $\mathcal{U}$-*bootstrap percolation: critical probability, exponential decay and applications*, Ann. Inst. Henri Poincaré Probab. Stat. **57** (2021), no. 3, 1255–1280 pp. MR4291442

[113] I. Hartarsky, *Bootstrap percolation and kinetically constrained models: two-dimensional universality and beyond*, Ph.D. Thesis, Université Paris Dauphine, PSL University, 2022. Available at `https://tel.archives-ouvertes.fr/tel-03527333`.

[114] I. Hartarsky, *Bootstrap percolation, probabilistic cellular automata and sharpness*, J. Stat. Phys. **187** (2022), no. 3, Article No. 21, 17 pp. MR4408459

[115] I. Hartarsky, *Refined universality for critical KCM: upper bounds*, Comm. Math. Phys. **405** (2024), no. 1, Paper No. 13, 68 pp. MR4694420

[116] I. Hartarsky and L. Marêché, *Refined universality for critical KCM: lower bounds*, Combin. Probab. Comput. **31** (2022), no. 5, 879–906 pp. MR4472293

[117] I. Hartarsky, L. Marêché, and C. Toninelli, *Universality for critical KCM: infinite number of stable directions*, Probab. Theory Related Fields **178** (2020), no. 1, 289–326 pp. MR4146539

[118] I. Hartarsky, F. Martinelli, and C. Toninelli, *Universality for critical KCM: finite number of stable directions*, Ann. Probab. **49** (2021), no. 5, 2141–2174 pp. MR4317702

[119] I. Hartarsky, F. Martinelli, and C. Toninelli, *Coalescing and branching simple symmetric exclusion process*, Ann. Appl. Probab. **32** (2022), no. 4, 2841–2859 pp. MR4474521

[120] I. Hartarsky, F. Martinelli, and C. Toninelli, *Sharp threshold for the FA-2f kinetically constrained model*, Probab. Theory Related Fields **185** (2023), no. 3, 993–1037 pp. MR4556287

[121] I. Hartarsky and R. Szabó, *Generalised oriented site percolation*, Markov Process. Related Fields **28** (2022), no. 2, 275–302 pp. MR4560695

[122] I. Hartarsky and R. Szabó, *Subcritical bootstrap percolation via Toom contours*, Electron. Commun. Probab. **27** (2022), Paper No. 55, 13 pp. MR4510850

[123] I. Hartarsky and A. Teixeira, *Bootstrap percolation is local*, arXiv e-prints (2024), available at arXiv:2404.07903.

[124] I. Hartarsky and F. Toninelli, *Kinetically constrained models out of equilibrium*, Probab. Math. Phys. **5** (2024), no. 2, 461–489 pp. MR4749812

[125] R. A. Holley and D. W. Stroock, *Uniform and L^2 convergence in one-dimensional stochastic Ising models*, Comm. Math. Phys. **123** (1989), no. 1, 85–93 pp. MR1002033

[126] A. E. Holroyd, *Sharp metastability threshold for two-dimensional bootstrap percolation*, Probab. Theory Related Fields **125** (2003), no. 2, 195–224 pp. MR1961342

[127] A. E. Holroyd, *The metastability threshold for modified bootstrap percolation in d dimensions*, Electron. J. Probab. **11** (2006), Paper No. 17, 418–433 pp. MR2223042

[128] R. L. Jack, L. Berthier, and J. P. Garrahan, *Static and dynamic length scales in a simple glassy plaquette model*, Phys. Rev. E **72** (2005), no. 1, 016103 pp.

[129] R. L. Jack and J. P. Garrahan, *Caging and mosaic length scales in plaquette spin models of glasses*, J. Chem. Phys. **123** (2005), no. 16, 164508 pp.

[130] 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. Theory Exp. **03** (2006), P03006, 33 pp.

[131] J. Jäckle and S. Eisinger, *A hierarchically constrained kinetic Ising model*, Z. Phys. B Con. Mat. **84** (1991), no. 1, 115–124 pp.

[132] J. Jäckle, K. Froböse, and D. Knödler, *Size dependence of self-diffusion in the hard-square lattice gas*, J. Stat. Phys. **63** (1991), no. 1-2, 249–260 pp.

[133] M. Jeng and J. M. Schwarz, *On the study of jamming percolation*, J. Stat. Phys. **131** (2008), no. 4, 575–595 pp. MR2398943

[134] Y. Jung, J. P. Garrahan, and D. Chandler, *Excitation lines and the breakdown of Stokes-Einstein relations in supercooled liquids*, Phys. Rev. E **69** (2004), no. 6, 061205, 7 pp.

[135] Y. Jung, J. P. Garrahan, and D. Chandler, *Dynamical exchanges in facilitated models of supercooled liquids*, J. Chem. Phys. **123** (2005), no. 8, 084509 pp.

[136] G. Keller, *Equilibrium states in ergodic theory*, London mathematical society student texts, vol. 42, Cambridge University Press, Cambridge, 1998. MR1618769

[137] 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** (1993), no. 6, 4364–4377 pp.

[138] P. M. Kogut and P. L. Leath, *Bootstrap percolation transitions on real lattices*, J. Phys. C **14** (1981), no. 22, 3187–3194 pp.

[139] G. Kordzakhia and S. P. Lalley, *Ergodicity and mixing properties of the northeast model*, J. Appl. Probab. **43** (2006), no. 3, 782–792 pp. MR2274800

[140] J. Kurchan, L. Peliti, and M. Sellitto, *Aging in lattice-gas models with constrained dynamics*, EPL **39** (1997), no. 4, 365–370 pp.

[141] G. F. Lawler and A. D. Sokal, *Bounds on the L^2 spectrum for Markov chains and Markov processes: a generalization of Cheeger's inequality*, Trans. Amer. Math. Soc. **309** (1988), no. 2, 557–580 pp. MR930082

[142] J.-F. Le Gall, *Measure theory, probability, and stochastic processes*, Graduate texts in mathematics, vol. 295, Springer, Cham, 2022. MR4559717

[143] T.-Y. Lee and H.-T. Yau, *Logarithmic Sobolev inequality for some models of random walks*, Ann. Probab. **26** (1998), no. 4, 1855–1873 pp. MR1675008

[144] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, Providence, RI, 2009. With a chapter by J. G. Propp and D. B. Wilson. MR2466937

[145] T. M. Liggett, *Stochastic interacting systems: contact, voter and exclusion processes*, Grundlehren der mathematischen Wissenschaften, vol. 324, Springer, Berlin, Heidelberg, 1999. MR1717346

[146] T. M. Liggett, *Interacting particle systems*, Classics in mathematics, Springer, Berlin, Heidelberg, 2005. Originally published by Springer, New York (1985). MR2108619

[147] T. M. Liggett, R. H. Schonmann, and A. M. Stacey, *Domination by product measures*, Ann. Probab. **25** (1997), no. 1, 71–95 pp. MR1428500

[148] A. Lipowski, *Glassy behaviour and semi-local invariance in Ising model with four-spin interaction*, J. Phys. A: Math. Gen. **30** (1997), no. 21, 7365–7373 pp.

[149] M. K. Mapes, S. F. Swallen, and M. D. Ediger, *Self-diffusion of supercooled o-Terphenyl near the glass transition temperature*, J. Phys. Chem. B **110** (2006), no. 1, 507–511 pp.

[150] L. Marêché, *Exponential convergence to equilibrium for the d-dimensional East model*, Electron. Commun. Probab. **24** (2019), Paper No. 55, 10 pp. MR4003129

[151] L. Marêché, *Combinatorics for general kinetically constrained spin models*, SIAM J. Discrete Math. **34** (2020), no. 1, 370–384 pp. MR4062795

[152] L. Marêché, *Exponential convergence to equilibrium in supercritical kinetically constrained models at high temperature*, Markov Process. Related Fields **26** (2020), no. 5, 861–884 pp. MR4246011

[153] L. Marêché, F. Martinelli, and C. Toninelli, *Exact asymptotics for Duarte and supercritical rooted kinetically constrained models*, Ann. Probab. **48** (2020), no. 1, 317–342 pp. MR4079438

[154] F. Martinelli, *Lectures on Glauber dynamics for discrete spin models*, Lectures on probability theory and statistics (Saint-Flour, 1997), 1999, 93–191 pp. MR1746301

[155] F. Martinelli, R. Morris, and C. Toninelli, *Universality results for kinetically constrained spin models in two dimensions*, Comm. Math. Phys. **369** (2019), no. 2, 761–809 pp. MR3962008

[156] F. Martinelli and E. Olivieri, *Approach to equilibrium of Glauber dynamics in the one phase region. I. The attractive case*, Comm. Math. Phys. **161** (1994), no. 3, 447–486 pp. MR1269387

[157] F. Martinelli and E. Olivieri, *Approach to equilibrium of Glauber dynamics in the one phase region. II. The general case*, Comm. Math. Phys. **161** (1994), no. 3, 487–514 pp. MR1269388

[158] F. Martinelli, E. Olivieri, and R. H. Schonmann, *For 2-D lattice spin systems weak mixing implies strong mixing*, Comm. Math. Phys. **165** (1994), no. 1, 33–47 pp. MR1298940

[159] F. Martinelli, A. Shapira, and C. Toninelli, *Diffusive scaling of the Kob–Andersen model in $\mathbb{Z}^d$*, Ann. Inst. Henri Poincaré Probab. Stat. **56** (2020), no. 3, 2189–2210 pp. MR4116722

[160] F. Martinelli and C. Toninelli, *Kinetically constrained spin models on trees*, Ann. Appl. Probab. **23** (2013), no. 5, 1967–1987 pp. MR3134727

[161] F. Martinelli and C. Toninelli, *Towards a universality picture for the relaxation to equilibrium of kinetically constrained models*, Ann. Probab. **47** (2019), no. 1, 324–361 pp. MR3909971

[162] B. Massoulié, *Cutoff for the mixing time of the facilitated exclusion process*, arXiv e-prints (2024), available at arXiv:2412.04032.

[163] M. V. Menshikov, *Coincidence of critical points in percolation problems*, Dokl. Akad. Nauk SSSR **288** (1986), no. 6, 1308–1311 pp. MR852458

[164] M. Merolle, J. P. Garrahan, and D. Chandler, *Space–time thermodynamics of the glass transition*, Proc. Natl. Acad. Sci. U.S.A. **102** (2005), no. 31, 10837–10840 pp.

[165] R. Morris, *Monotone cellular automata*, Surveys in combinatorics 2017, 2017, 312–371 pp. MR3728111

[166] T. Mountford and G. Valle, *Exponential convergence for the Fredrickson-Andersen one-spin facilitated model*, J. Theoret. Probab. **32** (2019), no. 1, 282–302 pp. MR3908915

[167] Y. Nagahata, *Lower bound estimate of the spectral gap for simple exclusion process with degenerate rates*, Electron. J. Probab. **17** (2012), Paper No. 92, 19 pp. MR2994840

[168] C. Neuhauser and A. Sudbury, *The biased annihilating branching process*, Adv. in Appl. Probab. **25** (1993), no. 1, 24–38 pp. MR1206531

[169] M. E. J. Newman and C. Moore, *Glassy dynamics and aging in an exactly solvable spin model*, Phys. Rev. E **60** (1999), no. 5, 5068–5072 pp.

[170] M. Ozawa and G. Biroli, *Elasticity, facilitation, and dynamic heterogeneity in glass-forming liquids*, Phys. Rev. Lett. **130** (2023), no. 13, 138201, 8 pp.

[171] N. Pancotti, G. Giudice, J. I. Cirac, J. P. Garrahan, and M. C. Bañuls, *Quantum East model: localization, nonthermal eigenstates, and slow dynamics*, Phys. Rev. X **10** (2020), no. 2, 021051, 21 pp.

[172] Y. Peres and A. Sly, *Mixing of the upper triangular matrix walk*, Probab. Theory Related Fields **156** (2013), no. 3-4, 581–591 pp. MR3078280

[173] G. Perrupato and T. Rizzo, *Thermodynamics of the Fredrickson-Andersen model on the Bethe lattice*, Phys. Rev. E **110** (2024), no. 4, 044312, 22 pp. MR4834814

[174] N. S. Pillai and A. Smith, *Mixing times for a constrained Ising process on the torus at low density*, Ann. Probab. **45** (2017), no. 2, 1003–1070 pp. MR3630292

[175] N. S. Pillai and A. Smith, *Mixing times for a constrained Ising process on the two-dimensional torus at low density*, Ann. Inst. Henri Poincaré Probab. Stat. **55** (2019), no. 3, 1649–1678 pp. MR4010947

[176] M. Pollak and I. Riess, *Application of percolation theory to 2D-3D Heisenberg ferromagnets*, Phys. stat. sol. (b) **69** (1975), no. 1, K15–K18 pp.

[177] J. Reiter, F. Mauch, and J. Jäckle, *Blocking transitions in lattice spin models with directed kinetic constraints*, Phys. A **184** (1992), no. 3, 458–476 pp.

[178] F. Ritort and P. Sollich, *Glassy dynamics of kinetically constrained models*, Adv. Phys. **52** (2003), no. 4, 219–342 pp.

[179] L. Saloff-Coste, *Lectures on finite Markov chains*, Lectures on probability theory and statistics (Saint-Flour, 1996), 1997, 301–413 pp. MR1490046

[180] F. Sausset, C. Toninelli, G. Biroli, and G. Tarjus, *Bootstrap percolation and kinetically constrained models on hyperbolic lattices*, J. Stat. Phys. **138** (2010), no. 1-3, 411–430 pp. MR2594903

[181] R. H. Schonmann, *Critical points of two-dimensional bootstrap percolation-like cellular automata*, J. Stat. Phys. **58** (1990), no. 5-6, 1239–1244 pp. MR1049067

[182] R. H. Schonmann, *On the behavior of some cellular automata related to bootstrap percolation*, Ann. Probab. **20** (1992), no. 1, 174–193 pp. MR1143417

[183] A. Shapira, *Kinetically constrained models with random constraints*, Ann. Appl. Probab. **30** (2020), no. 2, 987–1006 pp. MR4108128

[184] A. Shapira, *A note on the spectral gap of the Fredrickson–Andersen one spin facilitated model*, J. Stat. Phys. **181** (2020), no. 6, 2346–2352 pp. MR4179809

[185] A. Shapira, *Noncooperative models of kinetically constrained lattice gases*, arXiv e-prints (2023), available at arXiv:2301.13559.

[186] A. Shapira and E. Slivken, *Time scales of the Fredrickson-Andersen model on polluted $\mathbb{Z}^2$ and $\mathbb{Z}^3$*, arXiv e-prints (2019), available at arXiv:1906.09949.

[187] A. Sinclair and M. Jerrum, *Approximate counting, uniform generation and rapidly mixing Markov chains*, Inform. and Comput. **82** (1989), no. 1, 93–133 pp. MR1003059

[188] A. Sly, *Computational transition at the uniqueness threshold*, 2010 IEEE 51st Annual Symposium on Foundations of Computer Science—FOCS 2010, 2010, 287–296 pp. MR3025202

[189] P. Sollich and M. R. Evans, *Glassy time-scale divergence and anomalous coarsening in a kinetically constrained spin chain*, Phys. Rev. Lett. **83** (1999), no. 16, 3238–3241 pp.

[190] P. Sollich and M. R. Evans, *Glassy dynamics in the asymmetrically constrained kinetic Ising chain*, Phys. Rev. E **68** (2003), no. 3, 031504, 16 pp.

[191] A. Sudbury and P. Lloyd, *Quantum operators in classical probability theory. II. The concept of duality in interacting particle systems*, Ann. Probab. **23** (1995), no. 4, 1816–1830 pp. MR1379169

[192] A. Sudbury and P. Lloyd, *Quantum operators in classical probability theory. IV. Quasi-duality and thinnings of interacting particle systems*, Ann. Probab. **25** (1997), no. 1, 96–114 pp. MR1428501

[193] J. M. Swart, R. Szabó, and C. Toninelli, *Peierls bounds from Toom contours*, arXiv e-prints (2022), available at arXiv:2202.10999.

[194] C. Toninelli and G. Biroli, *Dynamical arrest, tracer diffusion and kinetically constrained lattice gases*, J. Stat. Phys. **117** (2004), no. 1-2, 27–54 pp. MR2098557

[195] C. Toninelli and G. Biroli, *Jamming percolation and glassy dynamics*, J. Stat. Phys. **126** (2007), no. 4-5, 731–763 pp. MR2311884

[196] C. Toninelli and G. Biroli, *A new class of cellular automata with a discontinuous glass transition*, J. Stat. Phys. **130** (2008), no. 1, 83–112 pp. MR2375957

[197] C. Toninelli, G. Biroli, and D. S. Fisher, *Cooperative behavior of kinetically constrained lattice gas models of glassy dynamics*, J. Stat. Phys. **120** (2005), no. 1-2, 167–238 pp. MR2165529

[198] C. Toninelli, G. Biroli, and D. S. Fisher, *Jamming percolation and glass transitions in lattice models*, Phys. Rev. Lett. **96** (2006), no. 3, 035702, 4 pp.

[199] C. Toninelli, G. Biroli, and D. S. Fisher, *Toninelli, Biroli, and Fisher reply:*, Phys. Rev. Lett. **98** (2007), no. 12, 129602, 1 pp.

[200] A. L. Toom, *Stable and attractive trajectories in multicomponent systems*, Multicomponent random systems, 1980, 549–575 pp. MR599548

[201] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić, *Weak ergodicity breaking from quantum many-body scars*, Nature Phys **14** (2018), no. 7, 745–749 pp.

[202] R. M. Turner, R. L. Jack, and J. P. Garrahan, *Overlap and activity glass transitions in plaquette spin models with hierarchical dynamics*, Phys. Rev. E **92** (2015), no. 2, 022115, 13 pp. MR3531148