Kinetically constrained spin models: a short survey

Fabio Martinelli*

September 28, 2025

Abstract. Kinetically constrained models (KCMs) are interacting particle systems introduced in the '80s by condensed matter physicists to have accessible stochastic models with glassy-type dynamics. The key mechanism behind the complex evolution of these otherwise simple models is the so-called dynamical facilitation, a feature embedded into the models via appropriate kinetic constraints. KCMs are tightly related to bootstrap percolation, a widely studied monotone cellular automaton, and they have recently inspired the construction of quantum spin chains to explore many-body localization in the absence of disorder and ergodicity breaking transition. In the first half of this survey, I will review some of the general features of KCMs and the detailed behaviour of three of the most popular models. In the second half, I will focus on the KCMs counterpart of the famous "universality problem" for two-dimensional bootstrap percolation, a broad research program in extremal probabilistic combinatorics.

1 Introduction Kinetically constrained models (KCMs) are interacting particle systems on the integer lattice $\mathbb{Z}^d$ (other graphs have also been considered), introduced in the physics literature in the 1980s as accessible simple models to test the effects of *dynamical facilitation* on an otherwise rather simple stochastic dynamics. Dynamical facilitation, namely the fact that local evolution is possible only for certain states of the nearby variables, is a mechanism believed by many to play a key role in the glassy dynamics of a supercooled liquid near the glass transition point, where local rearrangements are slow and hindered. Glassy dynamics refers to the presence of several key features, including anomalously long mixing times, aging, dynamical heterogeneities, and ergodicity-breaking transition. We refer the reader to [36, Chapter 1] and references therein for a much more detailed discussion.

Most of the KCMs are spin systems evolving under a standard, non-conservative Glauber dynamics, which is reversible with respect to a trivial product measure, in the presence of kinetic constraints. A notable exception is represented by the Kob-Andersen model [42, 48], a lattice gas with kinetic constraints. In both cases, the role of kinetic constraints is to encode the facilitation mechanism by mimicking the "caging" effect, where particles are trapped by their neighbors, thereby slowing down the dynamics. More precisely, the constraints allow a spin-flip (or an exchange in the conservative case) iff the current spin or particle configuration in a neighborhood of where the dynamical update is attempted satisfies a certain requirement specific to the model. In the sequel, we will restrict ourselves to KCMs with a Glauber-type dynamics.

Mathematically, despite the reversibility w.r.t. a trivial product measure, the presence of hard (i.e., on/off) constraints makes answering even simple questions about the long-time behaviour of KCMs extremely hard, and it has required rather sophisticated tools and constructions, often inspired by other fields like combinatorics and bootstrap percolation. A large part of the difficulties comes from the lack of any natural monotonicity in the dynamics, and the fact that the only coercive inequality one may hope to apply successfully is the Poincaré inequality. Contrary to, e.g., the stochastic Ising model, Logarithmic Sobolev inequalities and the associated hypercontractivity of the dynamical Markov semigroup are most of the time essentially useless. All the above forced most known results to be restricted to the *stationary process*, with the notable exception of the East model, a paradigmatic KCM with a useful orientation of its constraints (see Section 2.1).

Stochastic dynamics with a facilitation mechanism also play a role in other contexts, unrelated to the liquid/glass transition. Examples are the random walk on the space of upper triangular matrices with entries in $\mathbb{F}_2$ (see [54, 24] and references therein), and Glauber dynamics of monotone and Solid-on-Solid surfaces [13].

*Dipartimento di Matematica e Fisica, Università Roma Tre, fabio.martinelli@uniroma3.it .

Finally, we point out that KCMs, particularly the one-dimensional East model (see Section 2.1), have inspired the construction of quantum spin chains featuring (quasi) many-body localization in the absence of disorder and with a transition from a regime of fast relaxation to a regime with very slow relaxation. We refer the reader to [66, 58, 53, 26] and references therein.

1.1 General construction and connection to bootstrap percolation A generic (spin) KCM is constructed as follows. A configuration ω is defined by assigning to each vertex $x \in \mathbb{Z}^d$ an occupation variable or spin $\omega_x \in \{0,1\}$, corresponding to the vertex x being empty or occupied respectively, and to each vertex x is assigned a local (i.e. depending on finitely many variables) constraint $c_x(\omega) \in \{0,1\}$ *independent* of the occupancy of ω at x. The current configuration ω evolves according to the following simple rule. Independently across the lattice, each vertex is equipped with a rate-one Poisson clock. At each arrival, or clock-ring, of the Poisson clock at x, the constraint $c_x(\omega)$ is checked for the current configuration ω. If $c_x(\omega)$ is satisfied, i.e., it is equal to one, then the occupation variable ω_x is refreshed to occupied with probability $p \in (0,1)$ and to empty with probability $q = 1-p$. Otherwise, nothing happens. A clock-ring that occurs with the current configuration satisfying the local constraint is called a *legal ring*. The vacancy density q is related to the inverse temperature of the model via the relation $q = 1/(1+e^{\beta})$, so that the challenging low-temperature regime corresponds to low equilibrium density of the vacancies. Since $c_x(\omega)$ is independent of ω_x, detailed balance w.r.t. the product Bernoulli(p) measure μ, so that the process started from μ at time $t = 0$ is stationary.

Remark 1.1. A natural and very useful generalization goes as follows. Let (S, ρ) be a finite probability space and let $\omega \in S^{\mathbb{Z}^d}$. Given an "effective particle event" $\mathcal{O} \subset S$, let $\hat{\omega} \in \{0,1\}^{\mathbb{Z}^d}$ be defined by $\hat{\omega}_x = \mathbf{1}_{\{\omega_x \in \mathcal{O}\}}$. Then on $S^{\mathbb{Z}^d}$ one defines the generalized KCM as follows. Given the current configuration ω, each vertex x waits an independent Exp(1)-time and then, iff $c_x(\hat{\omega}) = 1$, the state ω_x is updated to $\omega'_x \in S$ sampled from ρ. Clearly the projection process $\{\hat{\omega}(t)\}_{t\geq 0}$ is a standard KCM with constraints $\{c_x\}_{x\in\mathbb{Z}^d}$ and $p = \rho(\mathcal{O})$.

As far as the exact form of the constraints is concerned, most of the models considered in the physics literature belong to the following general class [12]. Fix an *update family* $\mathcal{U} = \{X_1, \ldots, X_m\}$, that is a finite collection of finite subsets of $\mathbb{Z}^d \setminus \{\mathbf{0}\}$, and choose $c_x(\omega) \equiv c_x^{\mathcal{U}}(\omega)$ as the indicator function of the event that there exists $X \in \mathcal{U}$ such that $\omega_y = 0$ for all $y \in X + x$.

Remark 1.2. The empty vertices or "vacancies" become the *facilitating vertices* for the KCM dynamics. However, facilitation does not imply any pathwise monotonicity, because more legal rings can flip more empty vertices into occupied ones.

Bearing on the above definition, KCMs can also be viewed as a natural non-monotone and stochastic counterpart of *$\mathcal{U}$-bootstrap percolation* (BP for shortness), a well-known class of discrete cellular automata (see [10, 51] and references therein). In $\mathcal{U}$-bootstrap percolation, each vertex is either infected or healthy, corresponding to $\omega_x = 0$ or $\omega_x = 1$ respectively in the spin language. Starting from some (possibly random) initial infection set $A \subset \mathbb{Z}^d$ at time $t = 0$, the infection set A_t at a later integer time t is determined by the following deterministic rule. Infected vertices remain infected, and any healthy vertex at time $t-1$ becomes infected at time t iff the translate by x of one of the sets in $\mathcal{U}$ is completely infected, i.e., it belongs to A_{t-1}. One then defines the final infection set and the infection time of the origin as $[A]_{\mathcal{U}} := \bigcup_{t=1}^{\infty} A_t$ and $\tau_0^{BP} = \inf\{t \geq 0 : A_t \ni \{0\}\}$ respectively. It is then natural to define the critical points

$$q_c^{BP} := \inf\Big\{q : \ \mu\big(A : \ [A]_{\mathcal{U}} = \mathbb{Z}^d\big) = 1\Big\}, \tag{1.1}$$

$$\tilde{q}_c^{BP} := \inf\Big\{q : \ \liminf_{t\to\infty} -\frac{\ln(\mu\big(A : \tau_0^{BP} > t\big)}{\ln t} > 0\Big\}. \tag{1.2}$$

Remark 1.3. Clearly $\tilde{q}_c^{BP} \geq q_c^{BP}$ and it is conjectured that the two critical points coincide in any dimension (cf. [36, Chapter 3.1]).

1.2 Ergodicity and stationary measures Due to the kinetic constraints, it is possible for an initially occupied region of the lattice to remain frozen (blocked) under the KCM dynamics. In particular, ω, ω' belong to the same ergodic component iff $[A(\omega)]_{\mathcal{U}} = [A(\omega')]_{\mathcal{U}}$, where $A(\omega)$ is the vacancy set of ω [36, Corollary 3.7]. Hence, questions concerning ergodicity and the structure of stationary measures are quite delicate.

Let $\Omega = \{0,1\}^{\mathbb{Z}^d}$, let $\mu(f) = \int d\mu(\omega) f(\omega)$ and $\mathrm{Var}_\mu(f) = \mu(f^2) - \mu(f)^2$, let

$$(\mathcal{L} f)(\omega) = \sum_x c_x(\omega)\big(\mu_x(f) - f\big)(\omega)$$

be the generator of the process, and let

$$\mathcal{D}_\mu(f,f) = \langle f, -\mathcal{L} f\rangle_{L^2(\Omega,\mu)} = \frac{1}{2}\sum_x \mu\big(c_x \,\mathrm{Var}_{\mu_x}(f)\big)$$

be the Dirichlet form of the process. Here $f : \Omega \mapsto \mathbb{R}$ is e.g. a local function, μ_x is the single Bernoulli(p) measure at $x \in \mathbb{Z}^d$ and $\mathrm{Var}_{\mu_x}(f)$ is the variance of $f(\omega)$ w.r.t. $\omega_x \sim \mu_x$. Because of reversibility w.r.t. μ, $\mathcal{L}$ is a non-positive self-adjoint operator on $\ell^2(\Omega,\mu)$ and we define its *relaxation time* $T_{\mathrm{rel}}(q;\mathcal{U})$ as the best constant C in the Poincaré inequality

$$\mathrm{Var}_\mu(f) \le C\mathcal{D}_\mu(f,f) \quad \forall \text{ local function } f. \tag{1.3}$$

Let also τ_0 be the hitting time of $\{\omega : \omega_0 = 0\}$ and let $\mathbb{P}_\mu(\cdot)$ be the law of the stationary KCM process. It is a general result [36, Section 3.4] that

$$\mathbb{P}_\mu(\tau_0 > t) \le e^{-tq/T_{\mathrm{rel}}} \quad \text{and} \quad \mathbb{E}_\mu(\tau_0) \le T_{\mathrm{rel}}(q;\mathcal{U})/q. \tag{1.4}$$

The fundamental ergodicity result for KCMs is the following (see [12, Proposition 2.5] and [36, Theorems 3.9 and 3.10]).

THEOREM 1.4. *Let $q > q_c^{BP}$. Then 0 is a simple eigenvalue of the spectrum of $-\mathcal{L}$, $\mathbb{P}_\mu(\tau_0 < +\infty) = 1$ and the KCM process is ergodic. If in addition $q > \tilde{q}_c^{BP}$ then $T_{\mathrm{rel}}(q;\mathcal{U}) < +\infty$. If instead $q < q_c^{BP}$ then $T_{\mathrm{rel}}(q;\mathcal{U}) = +\infty$ and $\mathbb{P}_\mu(\tau_0 = +\infty) > 0$.*

As far as the stationary measures are concerned, much less is known (see [49, Section 2] for a detailed discussion). An old result of [38] proves that any *stationary and translation-invariant* measure ν is also reversible. In $d = 1$, entropy methods prove that translation invariance is not necessary [49, Theorem 2.4]. To identify the stationary measures, one can try to combine these ingredients in two different mechanisms: when the bootstrap percolation spreads out from within a box, and when it invades a box from the outside. For example, we quote the following result [49, Theorem 2.10]. Let Λ_N be the box of N vertices centered at the origin and let $\mathbf{0}$ be the all-empty configuration.

THEOREM 1.5. *Consider a KCM such that $\mu\big([A(\omega)\cap\Lambda_N]_{\mathcal{U}} \restriction_{\Lambda_N} = \mathbf{0} \restriction_{\Lambda_N}\big) \xrightarrow{N\to\infty} 1$. Let ν be a measure satisfying detailed balance and assume $[A(\omega)]_{\mathcal{U}} = \mathbf{0}$ ν-a.s. Then $\nu = \mu$.*

We conclude with a general conjecture put forward in [49]. Let $\mathcal{E}$ be the set of stable configurations, i.e. those ω such that $[A(\omega)]_{\mathcal{U}} = \omega$.

CONJECTURE 1.6. *Let ν be a stationary measure. Then there exists a probability measure ρ_* on $\mathcal{E}$ such that $\nu = \int d\rho_*(\beta)\,\nu_\beta$, where ν_β coincides with μ conditioned on the event $\{\omega : [A(\omega)]_{\mathcal{U}} = \beta\}$.*

2 Three specific models Having defined the general context, we now examine three models, which played an important role in the development of the subject. Due to the limited space, we are forced to neglect other important examples, such as the North-East, the Spiral, and the Kob-Andersen models. For a detailed discussion, we refer the interested reader to [36].

2.1 The East model The East model was originally introduced in [41] in one dimension. Here we present the general d-dimensional definition by choosing the update family $\mathcal{U}$ as the set of points that can be reached from the origin by decreasing one of the coordinates by one unit. The fact that the constraint at x queries only the state of the nearest neighbors of x which are "behind x" implies that the projection of the process onto the variables $\{\omega_x\}_{x\prec 0}$, where $x \prec 0$ iff all the coordinates of x are non-negative, coincides with the East process on the orthant $\{x \in \mathbb{Z}^d : x \prec 0\}$. This fact is a key feature that allows for convenient conditioning on the Poisson clocks and coin tosses that govern the dynamics, leading to sharp results for the stationary and non-stationary processes.

Remark 2.1. Based on the East model, an interesting coarse-grained model for atomic glass formers that is a sort of multicolor version of the East model (cf. [17] for a mathematical analysis), featuring a novel, subtle blocking mechanism between vacancies of different types, was proposed in [25]. We also outline that the East model plays a key role in the universality problem for two-dimensional KCMs. It can successfully model the effective dynamics of *mobile droplets* of rooted supercritical models and of critical models with an infinite number of stable directions (cf. Section 3).

For the above choice of $\mathcal{U}$, it is easy to check that $q_c^{BP} = \tilde{q}_c^{BP} = 0$. Hence, Theorem 1.4 implies that the East process is always ergodic with a finite relaxation time $T_{\rm rel} \equiv T_{\rm rel}(q;\mathcal{U})$. For $d = 1$ the set $\mathcal{E}$ of stable configurations consists of the all empty $\mathbf{0}$, the all filled $\mathbf{1}$, and of the configurations $\{\beta^{(i)}\}_{i\in\mathbb{Z}}$, where $\beta_x^{(i)} = 1$ iff $x < i$. In turn, the stationary measures in $d = 1$ turn out to be convex combinations of the Dirac measures on $\mathbf{0}, \mathbf{1}$, and the product measures $\mu^{(i)}, i \in \mathbb{Z}$, with marginals $\mu^{(i)}(\eta_x = 1)$ equal to $1, 0$, and p depending whether $x < i, x = i$, or $x > i$ respectively. An analogous description of the stationary measures holds in higher dimension. We refer the reader to [49, Theorem 2.12, Theorem 2.13]. Quite anomalous, and with far reaching consequences, is the precise low-temperature scaling of $T_{\rm rel}$.

THEOREM 2.2. $\lim_{q\to 0} \frac{\ln(T_{\rm rel})}{\ln^2(1/q)} = \frac{1}{2d\ln(2)}$.

Remark 2.3. If we recall that at low temperature $q \simeq e^{-\beta}$, we get that $T_{\rm rel} \approx e^{\beta^2/(2d\ln(2))}$, i.e. a super-Arrhenius scaling reminiscent of fragile super-cooled liquids (see e.g. [36, Section 1.1 and Fig 1.1] and refs therein).

In $d = 1$, the result was proved in [12], improving upon a result of [2] and correcting a wrong conjecture made by physicists. In higher dimensions, the proof of the upper bound utilizes renormalization group ideas combined with the one-dimensional result [14], whereas the lower bound is proved by a combinatorial construction of a sharp bottleneck in the state space. The key technical tools in one dimension were the so-called *bisection technique* (see e.g. [36, Chapter 4]) to prove an efficient Poincaré inequality, and the following beautiful combinatorial result [15] (see also [36, Section 4.2.1.1]) which allows to characterize in detail the energy/entropy competition of relevant bottlenecks (akin energy barriers) for the East dynamics.

Imagine starting the East process from a single vacancy e.g, at the origin. By construction, such a vacancy will remain in place forever, and no new vacancies will ever appear to its left. On the other hand, new vacancies can be created to the right of the origin, always fulfilling the kinetic constraints. Let $V(n)$ be the set of configurations that the East process can visit without ever observing more than n vacancies in $\{1, 2, \dots\}$. Let also $\ell(n) = \sup_{\omega\in V(n)} \ell(\omega)$, where $\ell(\omega)$ is the position of the rightmost vacancy of ω. Then

$$\ell(n) = 2^n - 1, \text{ and } |V(n)| \le 2^{\binom{n}{2}} n!, \tag{2.1}$$

with the latter upper bound not far from being tight. Equation (2.1) is a cornerstone for understanding the rich and complex behaviour of the East dynamics (stationary and non-stationary) at low temperatures. The fact that the number of extra vacancies that *must be created* in order to observe a new vacancy at distance L grows like $\log_2(L)$, can be reformulated evocatively by saying that the energy barrier that the process has to overcome grows *logarithmically with the distance* L. In particular, since the equilibrium distance between vacancies is $\approx 1/q$, a natural energy-entropy guess for the relaxation time would be

$$T_{\rm rel} \approx |V(1/q)|^{-1} \big(1/q\big)^{\log_2(1/q)} = 2^{\frac{1}{2}\log_2(1/q)^2(1+o(1))} \quad \text{as } q \to 0,$$

i.e., the scaling proved in the theorem.

Remark 2.4. Physicists assumed the entropy contribution to be negligeable w.r.t. to the energy contribution and concluded that $T_{\rm rel} \approx \big(1/q\big)^{\log_2(1/q)}$.

2.1.1 Non-equilibrium behaviour As first discussed in [60, 61] and then rigorously proved in [19, 20], the low-temperature one-dimensional East model features *aging* and *dynamical heterogeneities*.

Imagine to start the process from a Bernoulli(α) product measure and take $q \ll 1 - \alpha$[1]. If one focuses on a time window $[0, T]$, $T \ll T_{\rm rel}$, the (non-equilibrium) evolution will w.h.p. try to remove the excess of vacancies

[1] In physical jargon, we are considering a sudden quench from high to low temperature

present at time $t = 0$. A vacancy can remove the next vacancy, located at 2^n steps to its right, via a complex hierarchical metastable process involving the creation/annihilation of auxiliary vacancies. This process requires on average a time $\approx 1/q^n$ to take place, but its actual duration, once started, is much shorter, a typical metastable signature. Hence, it is not entirely surprising that the evolution in $[0, T]$ is well approximated by a *hierarchical coalescence process* of lattice intervals separating consecutive vacancies. The latter process can be viewed as a hierarchical sequence of suitably linked coalescence processes, each occurring on its proper timescale $\approx 1/q^n$. The main consequence of the above description is that the vacancy density features a peculiar graph with plateaux, and the two-times autocorrelation function features aging, namely a non-trivial dependence on the initial and final times (see [19, Fig 1]).

The exact low-temperature scaling of the relaxation time given in theorem 2.2 has an interesting consequence on the mixing time of the East model in the finite box $\Lambda_L = \{0, \dots, L\}^d$. Consider an initial configuration ω with a vacancy at the origin and no vacancies outside Λ_L. Under the East evolution, the vacancy at the origin remains frozen, and if we consider the marginal of the process in $\hat{\Lambda}_L := \Lambda_L \setminus \{0\}$, we get an ergodic Markov chain reversible w.r.t. $\mu_L := \otimes_{x \in \hat{\Lambda}_L} \text{Bernoulli(p)}$. Let $d_L(t) = \max'_\omega d_L(t, \omega)$, where $d_L(t, \omega)$ is the total variation distance between the law of the chain at time t and μ_L, and $\max'_\omega$ is the maximum over the choices of ω respecting the above conditions.

THEOREM 2.5 (Mixing time cutoff). *There exists $\delta > 0$ such that, for all $q \in (0, \delta) \cup (1 - \delta, 1)$ there exists $v(q) > 0$ such that*

$$\lim_{\alpha \to +\infty} \liminf_{L \to +\infty} d_L\big(L/v - \alpha L^{2/3}\big) = 1 \quad \textit{and} \quad \lim_{\alpha \to +\infty} \limsup_{L \to +\infty} d_L\big(L/v + \alpha L^{2/3}\big) = 0.$$

Moreover, $v(q)$ is the limiting velocity with which the rightmost vacancy advances towards $+\infty$ in the one dimensional East process with a single vacancy at the origin at time $t = 0$ (see [23, 7]), and it satisfies $v(q) \asymp T_{\text{rel}}^{-d}$ as $q \to 0$.

Remark 2.6. For $d = 1$, the theorem was proved in [23] for all $q \in (0, 1)$ and $L^{2/3}$ replaced by $\sqrt{L}$. For $d > 1$ and low temperature (small q) it was proved in [16], and the key step consisted in showing that the speed of propagation of vacancies in the bulk (e.g. along the $(1, 1, \dots, 1)$ direction) is approximately $T_{\text{rel}}(q)^{-1} \asymp v(q)^{1/d} \gg v(q)$. Thus, most of the mixing time $T_{\text{mix}} \sim L/v$ is used to mix the chain along the coordinate axis, and the result follows from the one-dimensional case. The case $q \simeq 1$ was recently addressed in [11] through a comparison with first passage percolation.

We conclude with a natural but challenging conjecture. Start the process with a single vacancy at the origin and define $S(t) = \{x \in \mathbb{Z}^d : \exists \text{ a legal ring at } x \text{ before } t\}$.

CONJECTURE 2.7 (Shape theorem). *There exists a compact set $\hat{S} \subset \mathbb{R}^d$ such that $\forall\ 0 < \varepsilon < 1$ $\lim_{t \to \infty} \mathbb{P}((1 - \varepsilon)t\hat{S} \subseteq S(t) \subseteq (1 + \varepsilon)t\hat{S}) = 1$.*

The main difficulty in proving the conjecture is the lack of subadditivity, at variance with e.g., first passage percolation, due to the absence of any natural monotonicity of the process.

2.2 The one facilitated Fredrickson-Andersen model The Fredrickson-Andersen FA-kf model was introduced in [21, 22] with update family $\mathcal{U}$ all the k-sets of the $2d$ neighbors of the origin. In other words, the constraint at x is satisfied if x has at least k vacancies among its neighbors. For any $1 \le k \le d$ $q_c^{BP} = \tilde{q}_c^{BP} = 0$ (see [36, Theorem 3.1] and references therein) and hence $T_{\text{rel}} < +\infty$, while $q_c^{BP} = \tilde{q}_c^{BP} = 1$ for $d < k \le 2d$. Here, we examine the case $k = 1, d \ge 1$, while in Section 2.3 we tackle the much more complex case $k = 2, d = 2$.

The low-temperature scaling of T_{rel} was determined correctly in $d = 1, 2$ (up to logarithmic correction in $d = 2$) in [12] using again the bisection technique together with suitable test functions, and nailed down in [59] for $d \ge 3$.

THEOREM 2.8. *As $q \to 0$*

$$T_{\text{rel}} = \begin{cases} \Theta(1/q^3) & \textit{for } d = 1, \\ \Theta(1/q^2) & \textit{for } d \ge 3, \end{cases}$$

while for $d = 2$, $\Omega(1/q^2) \le T_{\text{rel}} \le O(\ln(1/q)/q^2)$.

Remark 2.9. In all cases, we get that the leading scaling of $T_{\rm rel}$ as a function of the inverse temperature β is an Arrhenius scaling $T_{\rm rel} \approx e^{c\beta}$, a signature of strong glass formers (cf. Remark 2.3).

For $d = 1, 2$ but not $d \geq 3$, the theorem agrees (modulo logarithmic corrections) with the following naive wrong heuristics. Specifically, for $q \ll 1$ vacancies are typically at mutual distance $\approx 1/q^{1/d}$, and the motion of an isolated vacancy should be well approximated by a symmetric random walk with jump rate q (at rate q an isolated vacancy creates a new vacancy on one of its neighbors and it is then removed, thanks to the newly created vacancy, at rate $p = 1 - q$). However, this picture, which completely neglects second-order branching effects, suggests a scaling $(1/q)^{1+\frac{2}{d}}$ for $T_{\rm rel}$. Notably, the naive picture was already corrected in [40], where FA-1f was mapped into a diffusion-limited aggregation model. By analyzing the latter via renormalization group methods, $T_{\rm rel}$ was then conjectured to satisfy $T_{\rm rel} = \Theta(1/q^2)$ for any $d \geq 2$.

CONJECTURE 2.10 ([36]). *In two dimensions there is a logarithmic correction,* $T_{\rm rel} = \Theta\big(\frac{\ln(1/q)}{q^2}\big)$.

A quite sophisticated and massive analysis [55, 56], taking full care of the vacancies' branching and coalescence, analyzed the low-temperature scaling of the *mixing time* $T_{\rm mix}$ of the FA1f chain in the discrete torus $\mathbb{T}_n^d$ with $n = \lfloor 1/q \rfloor$ vertices, conditioned to have at least one vacancy. The mixing time, i.e., the worst-case over the initial condition of the time necessary for the chain to be close in total variation distance to the stationary reversible measure [43], is, in general, and particularly for KCM, a much harder quantity to deal with w.r.t. $T_{\rm rel}$. The results of [55, 56] were strengthened and extended to other graphs in [32, Corollary 3.2] by successfully comparing (in the Dirichlet forms sense) FA-1f to CBSEP, a coalescing and branching simple symmetric exclusion process whose logarithmic Sobolev constant can be analyzed in great detail. As we will see in Section 2.3, a suitable "block multistate" generalization of CBSEP plays a key role in the analysis of the sharp low-temperature scaling of the mean infection time $\mathbb{E}_\mu(\tau_0)$ for the two-dimensional FA-2f model.

2.2.1 Non-equilibrium behavior The stable configurations are $\mathbf{0}$ and $\mathbf{1}$ because a single vacancy is able to empty the whole lattice. We can then use Theorem 1.5 to conclude that in one dimension any stationary measure is the convex combination of μ and $\delta_{\mathbf{1}}$. The same holds in higher dimensions for all translation-invariant stationary measures. The real challenge for the non-stationary FA-1f is to prove the following conjecture.

CONJECTURE 2.11. *For any* $q \in (0, 1)$ *and any local function* f

$$\lim_{t\to\infty} |\mathbb{E}_\nu(f(\eta(t))) - \pi(f)| = 0, \tag{2.2}$$

provided that $\nu(\exists$ *an infected vertex*$) = 1$. *Moreover, for natural choices of* ν, *e.g., a non-trivial product measure, the above limit is attained exponentially fast.*

Unfortunately, robust tools to attack the above conjecture, like the space-time multiscale analysis described in [62], are not yet available, and the available results are limited to q larger than a certain threshold (see [8, 18, 44] and [37] for a large class of KCM). Remarkably, in [64, 63] the convergence to μ when ν is a (non-trivial) product Bernoulli measure was proved for the *biased annihilating branching process* (BABP) [52], a process which evolves exactly like the FA-1f process, but with the rate of the Poisson clock attach to $x \in \mathbb{Z}^d$ equal to the number of vacancies among the neighbors of x. The main advantage of BABP w.r.t. FA-1f is that it is *self-dual* and *quasi-dual* to a new, reversible interacting particle system dubbed *double flip process* (DFP) (see below (2.3),(2.4)). In DFP, the state of an *edge of* $\mathbb{Z}^d$ is updated with the restriction that the parity of the number of vacancies on the edge is preserved. If $B(t), B'(t), D(t)$ denote the vacancy sets at time t of two independent copies of BABP and one (independent) copy of DFP starting at time $t = 0$ from B, B', and D, then

$$\mathbb{E}\big((-\frac{1}{\lambda})^{|B'(t)\cap B|}\big) = \mathbb{E}\big((-\frac{1}{\lambda})^{|B(t)\cap B'|}\big), \tag{2.3}$$

$$\mathbb{E}\Big((\frac{1}{y+1})^{D\cap B(t)}(\frac{-1}{y-1})^{D^c\cap B(t)}\Big) = \mathbb{E}\Big((\frac{1}{y+1})^{|B\cap D(t)|}(\frac{-1}{y-1})^{|B\cap D^c(t)|}\Big), \tag{2.4}$$

where $\lambda = q/(1-q), y = \sqrt{\lambda+1}$, and $\mathbb{E}(\cdot)$ denotes the average over the processes and over the possible randomness of the initial vacancy sets B, B' and D. Equation (2.3) expresses the self-duality of BABP, while (2.4) expresses the quasi-duality between BABP and DFP. The nice property of DFP on the infinite lattice is that it has a finite logarithmic Sobolev constant and, therefore, it is exponentially ergodic ([49, Theorem 5.2]). Hence, the r.h.s. of (2.4) when e.g. B is either a finite set or it is distributed according to a product measure, can be easily estimated for large t.

2.3 The two facilitated Fredrickson-Andersen model in $\mathbb{Z}^2$ For space reasons, we will necessarily be quite minimalist, and we refer to [33] for a more detailed survey. The FA-2f model in $\mathbb{Z}^2$ is, in some sense, paradigmatic of the complexity of KCM. While in FA-1f a single vacancy is able to create other vacancies and move around - in the glass-formers language it is a *non-cooperative* model - in FA-2f vacancies need to cooperate to facilitate the dynamics, no finite cluster of vacancies can move around without some external help, and there exist infinitely many blocked configurations (e.g. any two neighboring lines of occupied vertices). This complexity has restricted so far the analysis of the *non-stationary* FA-2f process to q close to one [37], whereas, for $q \ll 1$, i.e., at low temperature, it inherits all the complexity of the 2-*neighbour model*, its bootstrap percolation counterpart.

For the 2-neighbour model the first quantitative (and foundational) statement was proved in [1], stating that, as $q \to 0$, $\tau_0^{BP} = \exp(\Theta(1/q))$ w.h.p. Subsequently, the celebrated paper [39] proved that $\tau_0^{BP} = \exp\left(\frac{\pi^2+o(1)}{18q}\right)$, a result that was further refined in [27, 35] with a proof that $o(1) = \Theta(\sqrt{q})$. The high-level idea behind the scaling of τ_0^{BP} is as follows.

Declare a rectangle R *internally filled* (by the initial set of infections A) if $[A \cap R]_{\mathcal{U}} = R$. Then the key estimate is

$$\mu([A \cap R]_{\mathcal{U}} = R) \approx \exp(-\pi^2/9q) \quad \text{as } q \to 0,$$

for R on the critical scale $poly(1/q)$. Once this estimate is available, then in a box of side $\approx \exp(\pi^2/18q)$ centered at the origin w.h.p., there exists an internally filled rectangle with longest side on the *critical scale*. This rectangle, dubbed *critical droplet*, w.h.p. will grow at roughly constant speed, and it will engulf the origin in a time $\approx \exp(\pi^2/18q)$. Regarding FA-2f, the result that mirrors the one for τ_0^{BP} is as follows [34].

THEOREM 2.12. *As $q \to 0$ the stationary FA-2f model on $\mathbb{Z}^2$ satisfies:*

$$\exp\left(\frac{\pi^2}{9q}\big(1 - O(1)\sqrt{q}\big)\right) \le \tau_0 \le \exp\left(\frac{\pi^2}{9q}\big(1 + \sqrt{q}\log(1/q)^{O(1)}\big)\right) \tag{2.5}$$

in mean and w.h.p. In particular, w.h.p. $\tau_0 = \left(\tau_0^{BP}\right)^{2(1+o(1))}$.

Remark 2.13. The theorem is by no means a corollary of the BP results. While the lower bound does indeed follow from a judicious use of the BP results, the proof of the upper bound is much more involved. It requires guessing an efficient mechanism to create/annihilate vacancies, allowing a legal update at the origin, which has no counterpart in the monotone BP.

A high-level description of the proof strategy is well described in [33, Section 4.3]. Here, I will discuss only the background and some of the main mathematical steps.

If one sets the basic scale to be $1/q$, the first paper claiming an exponential scaling of τ_0 (say w.h.p.) was [57]. There, it was argued that the infection process of the origin is dominated by the motion of *macro-defects*, i.e., rare patches of vacancies of polynomial size and exponentially small density, that move at an exponentially small rate. In [65], this picture was considerably refined and modified in two crucial aspects. Firstly, the macro-defects were identified with the *critical droplets* of the 2-neighbour BP model, so that their density is $\approx \exp(-\pi^2/9q)$. Secondly, the rate with which the macro-defects move was argued to be exponentially small in $1/\sqrt{q}$, i.e., much larger than their density. These two ingredients, combined with the assumption that macro defects diffuse, led to conjecture [65, Section 6.3] a scaling $\approx e^{\pi^2/9q}$ for T_{rel}. The proof scheme of Theorem 2.12 in [34] consists of three key steps confirming the heuristic picture of [65] (but with $\mathbb{E}_\mu(\tau_0)$ replacing T_{rel}).

Firstly, similarly to the critical droplets of BP, the macro defects (or *mobile droplets* in the language of [34]) are defined via a multiscale construction, from scale $1/\sqrt{q}$ to scale $1/q^9$, inspired by the one in [39] but finely tuned to the key requirement that mobile droplets should typically move on a much shorter time scale than the target exponential scale of τ_0. In particular, the multiscale construction exhibits a certain degree of "softness," which allows lower-scale droplets to move freely under the FA-2f dynamics, within a larger-scale droplet. Moreover, the density of the mobile droplets scales like the density of the critical droplets of BP.

Secondly, it is observed that w.h.p. the hitting time τ_0 for the FA-2f stationary process coincides with the same hitting time for the stationary FA-2f continuous-time Markov *chain* on the box Λ centered at the origin and of linear size $2e^{(\pi^2+o(1))/9q}$, restricted to a "mobile" environment. Mobile environment refers to the set of configurations satisfying two distinct requirements: (i) they contain at least one mobile droplet, and (ii) their vacancies are such that a mobile droplet can move in the four directions. The latter, very likely, condition

is guaranteed if in each vertical and horizontal discrete interval of suitable polynomial cardinality, there is a vacancy.

Thirdly, the effective dynamics of the mobile droplets is modeled by a *generalized coalescing and branching exclusion process* (g-CBSEP) whose mixing time can be computed very precisely via the logarithmic Sobolev inequality and coupling argument. A general $g-$CBSEP process (a Markov chain with a finite state space in our case) is defined as follows. Consider a finite connected graph $G = (V, E)$, a finite probability space (S, ρ) together with a distinguished event $\mathcal{O} \subset S$ of positive probability. Let $\omega \in S^V$ be the current state (configuration) of the chain. Then, independently across E the state of an edge $e = (u, v) \in E$ such that $\mathcal{O}_e = \{\omega_u \in \mathcal{O}\} \cup \{\omega_v \in \mathcal{O}\}$ holds, with rate one is updated to a new state (ω'_u, ω'_v) sampled from the product measure $\nu \otimes \nu$ conditioned on $\mathcal{O}_e$. It is immediate to check reversibility w.r.t. the natural product measure $\otimes_{v\in V}\nu$ conditioned on the existence of at least one vertex whose state belongs to $\mathcal{O}$. For FA-2f one first fixes $\ell = \text{poly}(1/q)$ and partitions the finite box $\Lambda \subset \mathbb{Z}^d$ above into mesoscopic boxes $\{B_i\}_{i\in V}$ of linear size ℓ. The graph G will be the natural graph, e.g., of the center of the boxes. If B is a reference copy of the boxes, the set S will consist of all particle configurations in B such that every row/column of B contains a vacancy, and ν will be the measure μ_B conditioned on the above property. Finally, the distinguished event $\mathcal{O}$ will be the presence in B of the mobile droplet.

We now present the heuristic motivation for the choice of the g-CBSEP as the effective dynamics of the mobile droplets. Consider two of the above boxes B_i, B_j forming an edge of the graph, and assume that the environment inside $B_i \cup B_j$ is "good", i.e., it satisfies condition (ii) of a mobile environment. In a good (very likely) environment, vacancies are typically well separated and cannot move macroscopically under the FA-2f dynamics. In particular, a mobile droplet cannot be created if it is absent or destroyed if it is present. The case is different when the environment in B_j is only good, while in B_i there is also a mobile droplet. In that case, the mobile droplet in B_i could create a new mobile droplet in B_j thanks to the good environment there. Thus, if the environment of the boxes is labeled M or G according to whether it is a mobile droplet or only good, one may conjecture that the only (macro state) allowed transitions are: $(G, G) \Leftrightarrow (G, G), (M, G) \Leftrightarrow (M, M), (G, M) \Leftrightarrow (M, M)$. These moves are reminiscent of the FA-1f moves, with M/G playing the role of an effective vacancy/particle (cf. Remark 1.1). In this picture, the effective displacement of the original mobile droplet in B_i, $(M, G) \Rightarrow (G, M)$, that in g-CBSEP represents a single possible transition, can only occur as a concatenation of two steps: $(M, G) \Rightarrow (M, M) \Rightarrow (G, M)$. If we recall that the density of mobile droplets scales as $e^{-\Theta(1/q)}$, we see immediately that the above concatenation requires an exponentially large time scale $e^{\Theta(1/q)}$ to occur, much larger than the conjectured one $\approx e^{\Theta(\sqrt{q})}$.

One of the main advances of [34] is the proof that the fundamental "exchange" or SEP[2] move $(M, G) \Leftrightarrow (G, M)$ of g-CBSEP occurs on a much shorter time scale $e^{O(\ln^3(1/q)/\sqrt{q})}$. Such a result suggest that a more reasonable model of the effective dynamics of the mobile droplets is a process with transitions $(G, G) \Leftrightarrow (G, G)$ (i.e. do nothing), a *fast exchange* move $(M, G) \Leftrightarrow (G, M)$, a *slow branching* move $(M, G) \Rightarrow (M, M)$, and a *fast coalescence* move $(M, M) \Rightarrow (G, M)$ or $(M, M) \Rightarrow (M, G)$. By choosing the rates of the transitions appropriately to retain reversibility w.r.t μ, conditioned to the event that all boxes are good and at least one contains a mobile environment, one ultimately ends up with the g-CBSEP discussed above. The technical result supporting the conclusion that the exchange move $(M, G) \Leftrightarrow (G, M)$ is "fast" is the following delicate conditional Poincaré inequality.

Let Λ_1, Λ_2 be two adjacent boxes of linear size $1/q^9$ (the exponent is fairly arbitrary), and let $\Lambda = \Lambda_1 \cup \Lambda_2$. Let also $\mathcal{M}_\Lambda$ be the event that the environment in Λ is mobile, namely each row/column of $\Lambda_i, i = 1, 2$, has a vacancy and there is a mobile droplet in at least one of the two boxes. Then (see [34, Proposition 4.7]),

$$\text{Var}_\Lambda(f|\mathcal{M}_\Lambda) \le \gamma(\Lambda) \sum_{x\in\Lambda} \mu_\Lambda\big(c_x^\Lambda \,\text{Var}_x(f)|\mathcal{M}_\Lambda\big), \tag{2.6}$$

where $c_x^\Lambda(\omega_\Lambda)$ is the FA-2f constraint c_x computed for the infinite configuration obtained from ω_Λ by filling (i.e. setting to one) all the vertices outside Λ, and $\gamma(\Lambda) \le e^{O(\ln^3(1/q)/\sqrt{q})}$. Equation (2.6) can be interpreted as follows. The FA-2f chain in Λ restricted to the mobile environment $\mathcal{M}_\Lambda$ is irreducible and therefore has a finite relaxation time bounded from above by $\gamma(\Lambda)$. In particular, on the timescale $\gamma(\Lambda)$, the mobile droplet can move from one box to the other.

[2]SEP= Simple Exclusion Process

The proof of (2.6) is based on a powerful multi-scale renormalization scheme dubbed "*matryoshka doll technique*" [36, Section 5.3.3.2], which, at a high level, can be explained as follows. Recall that the mobile droplet, and therefore the event $\mathcal{M}_\Lambda$, is constructed recursively on a sequence of length scales ℓ_n [3]. It is then quite natural to try to bound from above the relaxation time $\gamma(\Lambda)$ of the mobile environment $\mathcal{M}_\Lambda$ via the same quantity in progressively larger volumes $\Lambda_1, \dots, \Lambda_N \equiv \Lambda$, with Λ_n a "box" on scale ℓ_n. The "dynamical intuition" is that, after every time interval $\Theta(\gamma(\Lambda_{n-1}))$, the mobile environment on scale ℓ_{n-1} inside the mobile droplet on scale ℓ_n will have reached equilibrium under the FA-2f dynamics. Therefore, $\gamma(\Lambda_n)$, the relaxation time of the mobile environment on scale ℓ_n, should be at most $T(n) \times \gamma(\Lambda_{n-1})$, with $T(n)$ the time that it takes for the mobile droplet on scale ℓ_n to equilibrate its position inside Λ_n, assuming that at each time the infections inside it are at equilibrium. The latter quantity is bounded using, in various fashions, a general powerful technique for KCMs known as the *auxiliary constrained block dynamics* (see [12] and [34, Section 3.3]).

We conclude with two comments. Firstly, although τ_0 for the stationary process and $T_{\text{rel}}(q;\mathcal{U})$ are closely related (see (1.4)), the precise scaling of $T_{\text{rel}}(q;\mathcal{U})$ is still open, with the best available upper bound being $e^{O(\ln^2(1/q)/q)}$ [50]. The main reason is that it is not yet clear if restricting the FA-2f dynamics to a mobile environment, as done in the proof of Theorem 2.12, is possible. Secondly, although most of the intuition behind the proof of the theorem is based on considerations on the possible scenarios for the *effective dynamics* of the mobile droplets immersed in an *approximately static* good environment, there is no way to implement these ideas directly to the FA-2f evolution. As we said at the beginning, the main technical tool at our disposal is various forms of the Poincaré inequality, and a large part of the difficulty of the proof (and of its presentation to a mathematical audience) is to convert the dynamical intuition into useful Poincaré inequalities.

3 Universality for low-temperature two-dimensional KCM In $d = 2$, the East, FA-1f, and FA-2f models discussed previously, all have $q_c^{BP} = 0$, but they exhibit a dramatically different divergence of τ_0 (in mean or w.h.p.) as $q \to 0$. Such sensitivity to the update family $\mathcal{U}$ leads naturally to the following question.

Question 3.1. Is it possible to group all possible update families U into finitely many classes, in such a way that all members of the same class induce (approximately) the same divergence of τ_0 as $q \downarrow q_c^{BP}$?

Such a question has not been addressed in the physics literature, mostly because physicists lacked a general criterion to predict the different scalings. This fact is particularly unfortunate as the anomalous and sharp divergence of the relevant time scales as $q \downarrow q_c^{BP}$ often renders numerical simulations unable to provide clear-cut and reliable answers. The same question for $\mathcal{U}$-bootstrap percolation was addressed and fully answered in a vast research program in extremal probabilistic combinatorics. We refer the reader to [51], [36, Section 6.2] and references therein for $d = 2$ and to [6, 3, 4] for arbitrary $d \geq 3$.

3.1 Universality in $d = 1$ As a warm-up towards the more complex case $d = 2$, we start by examining the one-dimensional case.

Definition 3.2. *Given a one-dimensional update family $\mathcal{U}$, we say that the positive (resp. negative) direction is* unstable *if $\exists\, U \in \mathcal{U}$ contained in $\{-1, -2, \dots\}$ (resp. $\{1, 2, \dots\}$) and* stable *otherwise.*

Remark 3.3. The stability/instability of a direction determines whether vacancies can reproduce in that direction without using the help of vacancies that sit ahead of them along that direction.

With this definition, the one-dimensional universality classification for KCM says that the East, FA-1f, and FA-2f models cover, in a sense, all possible scenarios (see [36, Section 6.1] and references therein). More precisely,

Theorem 3.4.

1. *If $\mathcal{U}$ has exactly one stable direction like the East model then $q_c^{BP} = 0$ and w.h.p. $\tau_0 = e^{\Theta(\ln^2(1/q)/q)}$,*
2. *if $\mathcal{U}$ has no stable direction like the FA-1f model then $q_c^{BP} = 0$ and w.h.p. $\tau_0 = 1/q^{\Theta(1)}$,*
3. *if $\mathcal{U}$ has two stable directions like FA-2f then $q_c = 1$ and $\mathbb{P}_\mu(\tau_0 = +\infty) > 0$.*

An analogous result holds also for the relaxation time $T_{\text{rel}}(q;\mathcal{U})$.

[3] the exact choice was $\ell_n \simeq \frac{e^{n/\sqrt{q}}}{\sqrt{q}}$, with $n = 1, \dots \lceil 8 \log(1/q)/\sqrt{q} \rceil$.

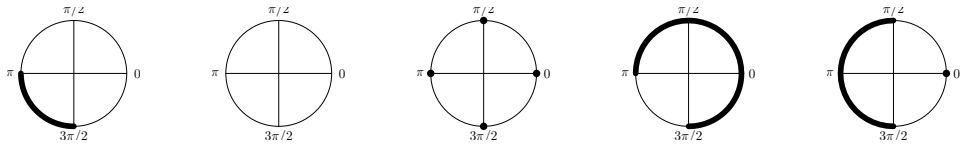

(a) East (b) FA-1f (c) FA-2f (d) North-East (e) Duarte

Figure 3.1: The set of stable directions (in bold) for five popular KCMs. The North-East family is $\{\{(1,0),(0,1)\}\}$, whereas the Duarte family is the collection of the 2-subsets of the North, South, and West neighbors of the origin.

Remark 3.5 (Logarithmic energy barriers). Notice that the dramatic difference between the behaviour of $\mathcal{U}$ with exactly one stable direction and $\mathcal{U}$ with no stable direction is completely absent in the BP setting, where $\tau_0^{BP} \approx \text{poly}(1/q)$ in both cases. The reason is that in BP, there is no penalty for creating new vacancies, whereas for KCMs at low temperatures, new vacancies are energetically penalized. In particular, since the vacancies of the stationary process are typically at mutual distance $\approx 1/q$, the emptying of the origin requires going through unlikely states with an anomalous quantity of vacancies, akin to overcoming an energy barrier. The above result can then be interpreted as follows: if $\mathcal{U}$ has no stable direction, the energy barrier that has to be overcome is $\Theta(1)$, while if $\mathcal{U}$ has exactly one stable direction, the barrier is $\Theta(\ln(1/q))$.

Remark 3.6. As pointed out in [36, problem 6.12], it is a interesting goal to prove that for $\mathcal{U}$ with two unstable directions there exists $\alpha > 0$ such that, w.h.p. as $q \to 0$, $\tau_0 = \Theta(1/q^\alpha)$.

3.2 Universality in $d = 2$ We begin with some quick background from BP. Let $S^1 = \{\mathbf{u} \in \mathbb{R}^2 : \|\mathbf{u}\| = 1\}$ be the unit circle, and call *directions* the elements of S^1. Sometimes we identify a direction with its angle $\theta \in [0, 2\pi)$ with respect to the horizontal axis, measured counterclockwise. Given $\mathbf{u} \in S^1$ let also $\mathbb{H}_\mathbf{u} = \{\mathbf{x} \in \mathbb{R}^2 : \langle \mathbf{x}, \mathbf{u} \rangle < 0\}$. As in $d = 1$, we need the notion of stable directions.

DEFINITION 3.7. *Fix a two-dimensional update family $\mathcal{U}$. A direction $\mathbf{u}$ is unstable if $\exists\ U \in \mathcal{U}$ such that $U \subset \mathbb{H}_\mathbf{u}$ and stable otherwise. A stable direction $\mathbf{u}$ is called* isolated stable *iff there exists no other stable direction in a neighborhood of $\mathbf{u}$.*

With the notion of stable direction at hand, the first-level classification of a two-dimensional update family $\mathcal{U}$ according to BP is as follows.

DEFINITION 3.8. *Let $\mathcal{C}$ be the collection of* open *semicircles of S^1. A two-dimensional update family $\mathcal{U}$ is:*

1. supercritical *if $\exists\ C \in \mathcal{C}$ with no stable direction;*
2. critical *if every $C \in \mathcal{C}$ contains a stable direction and $\exists\ C \in \mathcal{C}$ containing finitely many stable directions;*
3. subcritical *if every $C \in \mathcal{C}$ contains infinitely many stable directions.*

Remark 3.9. The East and the FA-1f models are supercritical, whereas the FA-2f and the Duarte families are critical. The North-East family $\mathcal{U} = \{\{(1,0),(0,1)\}\}$ is subcritical (cf. subsection 3.2).

Supercritical and critical families have $q_c^{BP} = 0$, whereas subcritical families have $q_c^{BP} > 0$ [10, 5]. For example, in the North-East model, $1 - q_c^{BP}$ matches the critical threshold of oriented percolation in $\mathbb{Z}^2$. We need a last fundamental ingredient from BP, namely the notion of the *difficulty* of an update family $\mathcal{U}$.

DEFINITION 3.10. *The difficulty $\alpha(\mathbf{u})$ of a direction $\mathbf{u}$ is defined to be 0 if $\mathbf{u}$ is unstable and to be $+\infty$ if $\mathbf{u}$ is stable but not isolated stable. If $\mathbf{u}$ is isolated stable, consider BP with $\mathbb{H}_\mathbf{u} \cap \mathbb{Z}^2$ as the initial set A of infections. Since $\mathbf{u}$ is stable, no extra sites can be infected. Then $\alpha(\mathbf{u})$ is set equal to the smallest number of extra infections to be added to A in such a way that, with their help, infinitely many new sites become infected. It was proven in [10, Lemma 5.2] that $\alpha(\mathbf{u}) < \infty$ iff $\mathbf{u}$ is isolated stable. The difficulty $\alpha(\mathcal{U})$ of an update family $\mathcal{U}$ is finally defined as $\min_{C \in \mathcal{C}} \max_{\mathbf{u} \in C} \alpha(\mathbf{u})$.*

With this definition, the key result of the BP-universality in $\mathbb{Z}^2$ can be summarized as follows [5, 10, 9].

THEOREM 3.11. *$\mathcal{U}$ is supercritical if $\alpha(\mathcal{U}) = 0$, critical if $1 \le \alpha(\mathcal{U}) < +\infty$, and subcritical otherwise. Moreover, w.h.p. as $q \to 0$, $\tau_0^{BP} = 1/q^{\Theta(1)}$ if $\mathcal{U}$ is supercritical, while $\tau_0^{BP} = e^{\Theta(\ln^\gamma(1/q)/q^\alpha)}$ if $\mathcal{U}$ is critical*

with difficulty α. Here $\gamma = \gamma(\mathcal{U}) \in \{0, 2\}$ depends on a more refined classification of critical families into balanced or unbalanced.

Regarding KCMs, we observe that the East and FA-1f models, both of which are supercritical with $\tau_0^{BP} = \mathrm{poly}(1/q)$ as $q \to 0$, feature quite different scalings of τ_0. Recalling Remark 3.5, it is clear that we need to refine Definition 3.8 to distinguish between these two models.

DEFINITION 3.12. *A two-dimensional supercritical update family $\mathcal{U}$ is called* rooted *if there exist two non-opposite* stable *directions and* unrooted *otherwise. In particular, the East family is rooted, while the FA-1f family is unrooted.*

With this definition, the first-level KCM-universality result in two dimensions reads as follows (see [36, Theorem 6.9] and the references therein).

THEOREM 3.13. *W.h.p. as $q \to 0$,*

1. *if $\mathcal{U}$ is* supercritical *then*
   (a) $\tau_0 = 1/q^{\Theta(1)}$ *if $\mathcal{U}$ is unrooted,*
   (b) $\tau_0 = e^{\Theta(\ln^2(1/q))}$ *if $\mathcal{U}$ is rooted.*
2. *If $\mathcal{U}$ is critical then $\tau_0 = e^{1/q^{\Theta(1)}}$.*

Remark 3.14. Recalling Remark 3.5 we conclude that in the supercritical case the energy barrier that has to be overcome to empty the origin is typically $\Theta(1)$ for unrooted families (as in the FA-1f case), and $\Theta(\log(1/q))$ for rooted ones (as in the East case). In a sense, the rooted case highlights the fundamental differences between two-dimensional supercritical KCMs and BP.

Behind (1.a) vs (1.b) As in the case of the FA-2f model, for all low-temperature supercritical and critical stationary KCM processes, the dominant mechanism to empty the origin should correspond to the arrival close to the origin of a *mobile droplet* which triggers the KCM dynamics there. While there is no unique definition, a mobile droplet is in general characterized by a certain finite set D of vacancies whose geometry depends on the update family $\mathcal{U}$, and whose cardinality may depend on q. For supercritical models a convenient choice of D is any sufficiently large (but not depending on q) rectangle oriented along the mid-point $\mathbf{u}$ of a semicircle C free of stable directions. Because of its orientation, if D is empty, then it is able to infect a suitable translation of itself in the $\mathbf{u}$ direction. For *unrooted* models, the semicircle C can be chosen in such a way that both C and $-C$ are free of stable directions. As a consequence, for unrooted models, the mobile droplet D will be able to infect a suitable translation of itself in *both directions* $\pm\mathbf{u}$. For *rooted* models, the above feature is not true, and the mobile droplet has a preferred propagation direction. The above is evocative of the difference between the one-dimensional FA-1f and East models: a vacancy can duplicate itself forward and backward in the first case, whereas it can only duplicate itself forward in the second case. One then expects that the kind of "orientation" of rooted models induces an East-like *effective dynamics* with parameter $q_{\mathrm{eff}} = q^{|D|}$ of their mobile droplets. Here, East-like means that the presence of a mobile droplet allows the creation (or removal) of another adjacent mobile droplet, but only in a certain direction (or, more precisely, within a limited cone of directions). As seen in Theorem 2.2, the characteristic time scale of the East dynamics is $e^{\Theta(\ln^2(1/q_{\mathrm{eff}}))} = e^{\Theta\left(\ln^2(1/q)\right)}$ because $|D| = \Theta(1)$. That is the scaling (1.b). In the unrooted case, the mobile droplet can move back and forth in a *double cone* of directions and the effective dynamics would be more like an FA-1f dynamics with characteristic time scale (cf. Theorem 2.8) $1/q_{\mathrm{eff}}^{\Theta(1)} = 1/q^{\Theta(1)}$, i.e. the scaling in (1.a). Of course, all the above is only a very high-level intuition which, as shortly discussed in Section 4, can be put to work to suggest a way to establish the key Poincaré inequality (1.3) with the correct scaling as $q \to 0$.

Behind (2) In the critical case the lower bound follows from the sharp asymptotics of τ_0^{BP} and the bound [50, Lemma 4.3], valid for all q small enough, $\mathbb{E}_\mu(\tau_0) \geq \delta \mathbb{E}_\mu(\tau_0^{BP})$ for some $\delta = \delta(\mathcal{U}) \in (0, 1)$. The upper bound follows from (1.4) and an upper bound on $T_{\mathrm{rel}}(q; \mathcal{U})$ obtained using long-range constrained Poincaré inequalities [50] (see Section 4 for a high-level discussion).

The next step in the KCM-universality program is to match the precision of BP for critical families of difficulty α that satisfy $\tau_0^{BP} \approx e^{\ln^{\Theta(1)}(1/q)/q^\alpha}$. In particular, the first main question is whether the scaling of τ_0, on a double logarithmic scale, is still governed by the exponent α.

In the supercritical case, the East family was essentially the ideal "case study" to highlight the differences between KCMs and BP. In the critical case, such a role is played by the Duarte model, a KCM whose update family is the collection of the 2-subsets of the North, South, and West neighbors of the origin. By construction, the stable

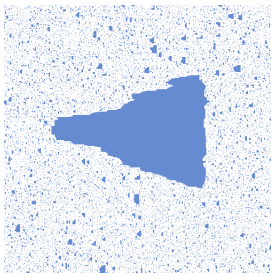

Figure 3.2: A growing droplet under the Duarte bootstrap process (courtesy of P. Smith).

directions are $\theta = 0$ with difficulty $\alpha = 1$ and all the directions in $[\frac{\pi}{2}, \frac{3\pi}{2}]$, each with infinite difficulty. Hence, the Duarte model is a critical (unbalanced) KCM with difficulty $\alpha = 1$, and, w.h.p. as $q \to 0$, $\tau_0^{BP} = e^{\Theta(\ln^2(1/q)/q)}$. For the Duarte KCM it was proved in [46] that $\mathbb{E}_\mu(\tau_0) \le e^{O\left(\ln^4(1/q)/q^2\right)}$ as $q \to 0$. A matching lower bound was subsequently established in [45, Section 5] via an algorithmic proof. Thus, on a double logarithmic scale, the hitting time τ_0 for the Duarte model diverges twice as fast as the corresponding τ_0^{BP}.

The basic intuition behind this remarkable fact is as follows. As in the supercritical case, we need to find a suitable candidate for the mobile droplet and make a reasonable ansatz for its effective dynamics. For critical families $\mathcal{U}$, a convenient choice of the set D characterizing the mobile droplet is a set with a complex geometry depending in a delicate way on the structure of the set of stable directions of $\mathcal{U}$ (see [47, Definition 4.9 and Fig. 5] or [31, Section 3.1.3] according to whether the number of stable directions is infinite or finite), and, contrary to the supercritical case, its cardinality diverges as $q \to 0$ like $\ln^{\Theta(1)}(1/q)/q^\alpha$, where $\alpha = \alpha(\mathcal{U})$ is the difficulty of $\mathcal{U}$.

Back to the Duarte model, its mobile droplets, defined through a complex algorithmic procedure [45, Section 5.2], have cardinality $\Theta\big(\ln(1/q)/q\big)$, and a corresponding effective density $q_{\text{eff}} = q^{\Theta(\ln(1/q)/q)} = e^{-\Theta(\ln^2(1/q)/q)}$. As for supercritical *rooted* families, one expects that the strong "orientation" of the set of stable directions induces an East-like effective dynamics with parameter q_{eff} of the mobile droplet. If that were true, then the characteristic time scale of the Duarte KCM would be $e^{\Theta(\ln^2(1/q_{\text{eff}}))} = e^{\Theta\left(\ln^4(1/q)/q^2\right)}$, i.e., the correct scaling. Notice that, if the mobile droplets could move back and forth in a *double cone* of directions, then the effective dynamics would be more like the FA-1f dynamics with characteristic time scale (cf. Theorem 2.8) $1/q_{\text{eff}}^{\Theta(1)} = e^{\Theta(\ln^2(1/q)/q)}$, i.e. similar to the BP-scaling.

Given the result for the Duarte model, two natural questions arise:

1. do all critical families with an *infinite number of stable directions* feature a divergence of τ_0 twice as fast (on a double logarithmic scale) than that of τ_0^{BP} ?

2. Is having an infinite number of stable directions crucial for the above feature to hold?

The first question was answered positively in [30]. Concerning the second question, consider a critical family with difficulty $\alpha(\mathcal{U}) = 1$ and finitely many stable directions, one at $\theta = 0$ with difficulty $\alpha = 1$, and all the others contained in $[\frac{\pi}{2}, \frac{3\pi}{2}]$, each with difficulty $\alpha \gg 1$. The "orientation" of the set of stable directions is similar to that of the Duarte model, and one could think that, due to the greater difficulty of the stable directions in $[\frac{\pi}{2}, \frac{3\pi}{2}]$, the mobile droplets move again in an East-like fashion in the $\vec{e}_1$ direction. Despite its soundness, this scenario was disproved in [31], correcting a wrong conjecture of [50, Conjecture 3], via the identification of the leading mechanism governing the motion of mobile droplets. It consists of a peculiar hierarchical combination of *mesoscopic East-like motions* resulting in a *macroscopic FA-1f kind of dynamics* (see the illuminating high-level analysis [31, Section 2] for the family $\mathcal{U}$ with four stable directions coinciding with the natural directions of $\mathbb{Z}^2$,

and difficulties $\alpha(\vec{e}_1) = 1$, and $\alpha(-\vec{e}_1) = \alpha(\pm\vec{e}_2) = 2$).
We are now ready for the second-level universality result for KCMs.

THEOREM 3.15. *Let $\mathcal{U}$ be a two-dimensional critical update family of difficulty α. Then, w.h.p. as $q \to 0$, $\tau_0 = e^{\ln^{\Theta(1)}(1/q)/q^\beta}$, where $\beta = \alpha$ if $\mathcal{U}$ has finitely many stable directions and $\beta = 2\alpha$ otherwise.*

Remark 3.16. Quite remarkably, in [28, 29] the logarithmic correction was pinned down very precisely and the collection of critical families classified into seven classes (see [36, Theorem 6.11]).

4 From dynamical intuition to Poincaré inequalities In this final section, we provide a few more details on how to transform suggestive, though unrealistic, dynamical intuition about the stationary KCM evolution into concrete and useful mathematical tools. Specifically, borrowing from [50, 46], we focus on methods to establish upper bounds on $T_{\rm rel}$ for supercritical and critical families. Lower bounds are typically proved either using the estimate $\mathbb{E}_\mu(\tau_0) \geq \delta\, \mathbb{E}_\mu(\tau_0^{BP})$ or via a "look for a bottleneck approach". For supercritical rooted models or critical models with infinitely many stable directions, finding a precise enough bottleneck requires rather sophisticated combinatorial proofs.

As mentioned at the beginning, given a critical update family $\mathcal{U}$, the ultimate goal is to prove a sharp Poincaré inequality for any local function f

$$\mathrm{Var}_\mu(f) \leq C\mathcal{D}_\mu(f,f). \tag{4.1}$$

In [50], it was realized that proving (4.1) requires first developing some tools to prove *constrained Poincaré inequalities* for product measures in a more general setting than the one specific to KCMs. In the sequel, $(S,\hat{\nu})$ will be a finite, positive probability space, and we shall work on the product probability space $(\Omega,\nu) := \otimes_{x\in\mathbb{Z}^d}(S_x,\hat{\nu}_x)$, where $(S_x,\hat{\nu}_x) = (S,\hat{\nu})$ for all $x \in \mathbb{Z}^d$. Often, in applications to KCMs, S is the "particle space" $\{0,1\}^V$, with V a finite subset of $\mathbb{Z}^d$ suitably chosen depending on the structure of the set of stable directions, and $\hat{\nu} = \otimes_{x\in V}\mu_x$. Given $\omega \in \Omega$, we associate to each vertex x a finite collection of constraints $\{c_x^{(i)}(\omega)\}_{i=1}^k$, where each $c_x^{(i)}$ is the indicator function of some event $\mathcal{A}_x^{(i)}$ depending only on the variables $\{\omega_y\}_{y\in\Delta_x^{(i)}}$, with $\Delta_x^{(i)}$ a finite subset of $\mathbb{Z}^d \setminus \{x\}$. Given an exhausting collection $\{V_n\}_{n\in\mathbb{Z}}$ of subsets of $\mathbb{Z}^d$, for any $x \in V_n \setminus V_{n-1}$ the family $\{c_x^{(i)}\}_{i=1}^k$ is assumed to satisfy the key *exterior condition*: for all $x \in V_n \setminus V_{n-1}$ and all $i = 1,\dots,k$, the support set $\Delta_x^{(i)}$ of $c_x^{(i)}$ belongs to $\cup_{j=n+1}^\infty V_j \setminus V_{j-1}$. Then, under a suitable smallness assumption for the probability of failure of the events $\{\mathcal{A}_x^{(i)}\}_{i,x}$ (see [50, Theorem 2]), for any local function f

$$\mathrm{Var}_\nu(f) \leq 4\sum_x \nu\Big(\big(\prod_{i=1}^k c_x^{(i)}\big)\,\mathrm{Var}_{\hat{\nu}_x}(f)\Big). \tag{4.2}$$

Equation (4.2) indicates that the generalized KCM defined by the generator $(\mathcal{L}f)(\omega) = \sum_x c_x(\omega)\big(\hat{\nu}_x(f) - f\big)$ with constraints $c_x = \prod_{i=1}^k c_x^{(i)}$, is ergodic with $T_{\rm rel} \leq 4$.

Remark 4.1. Although the exterior condition is not necessary for (4.2) to hold, it is easy to construct examples of constraints violating the exterior condition for which the r.h.s. of (4.2) is zero for a suitable local function f.

The flexibility of the above setting enables results such as those presented in [50, Theorem 2.6].

THEOREM 4.2. *Consider the two-dimensional 0-1 KCM, where the constraint at site x is defined as the indicator function that the North and East neighbors of x belong to an infinite percolation cluster of vacancies, irrespective of the state at x. For sufficiently small p, this KCM exhibits ergodicity with relaxation time bounded above by 4.*

Remark 4.3. Notice that the kinetic constraints of the above percolation example have infinite support.

Building on this example, we now present (for simplicity, only in $\mathbb{Z}^2$) the most useful consequence of (4.2). Consider two events, G_1 and G_2, with $G_2 \subset G_1$ in $(S,\hat{\nu})$. We will refer to these as the *good* and *super-good* events, respectively. In the applications to KCMs, where $S = \{0,1\}^V$, $\hat{\nu} = \mu_V$, and V is a finite subset of $\mathbb{Z}^2$, $G_1 = S$ and G_2 ensures the presence of a suitable empty rectangle inside V if $\mathcal{U}$ is supercritical. In the critical case, G_1 ensures a good environment in V, whereas G_2 guarantees the presence of a mobile droplet (cf. the discussion in Section 3). We will assume that G_1 has probability p_1 close to one ($p_1 = 1$ in the supercritical case),

whereas G_2 is unlikely with probability p_2. Recall that a (positively) oriented path γ of length n in $\mathbb{Z}^2$ from $x^{(1)}$ to $x^{(n)}$ is a sequence $\gamma = (x^{(1)}, \dots, x^{(n)})$ of vertices with the property that $x^{(k+1)}$ is either the North or the East neighbor of $x^{(k)}$. Given $\omega \in \Omega$, we say that γ is ω-good if ω_x and $\omega_{x\pm\vec{e}_1}$ belong to G_1 for all $x \in \gamma$. The path γ is ω-super-good if it is ω-good and $\omega_x \in G_2$ for some $x \in \gamma$. Finally, for any $x \in \mathbb{Z}^d$, we let

$$\begin{aligned}\mathcal{A}_x = \{\omega :\ &\textit{there exists an oriented, } \omega\textit{-good path } \gamma \textit{ of cardinality } \lfloor 1/p_2^2 \rfloor \textit{ starting at } x+\vec{e}_2, \textit{ the}\\ &\text{smallest (in some natural order) of such paths is } \omega\text{-super-good, and } \omega_{x+\vec{e}_1} \in G_1,\}\end{aligned} \tag{4.3}$$

Due to the orientation of the path γ and the fact that γ starts at $x + \vec{e}_2$, it is easy to verify that the family of constraints $c_x = \mathbf{1}_{\mathcal{A}_x}, x \in \mathbb{Z}^2$, satisfy the exterior condition w.r.t. an exhausting family $\{V_n\}_{n\in\mathbb{Z}}$. Notice also that, under our assumptions on p_1, p_2, the event $\mathcal{A}_x$ is very likely.

Proposition 4.4 ([46, Proposition 3.5]). *There exists $\delta > 0$ such that, if $\max\{p_2, (1-p_1)\ln^2(1/p_2)\} \le \delta$, then , for all local functions f*

$$\mathrm{Var}_\nu(f) \le 4\sum_x \nu\big(\mathbf{1}_{\mathcal{A}_x}\,\mathrm{Var}_{\hat{\nu}_x}(f)\big). \tag{4.4}$$

In the setting suitable for application to KCMs (cf. above), (4.4) is already promising. The variance term $\mathrm{Var}_{\hat{\nu}_x}(f)$ is, in our dynamical intuition, associated with the KCM updates in the box V_x attached to x. This term appears only if the box V_x is connected to a super-good box that contains a mobile droplet. Moreover, the connection is established via an oriented[4] path γ of good boxes, where the environment should be sufficiently favorable for the mobile droplet to travel through and reach V_x. If we assume that the latter feature holds, then the next step is to concretely implement, always at the level of a Poincaré inequality, the picture of the mobile droplet randomly traveling along γ and reaching x. It is here that the East-like or FA-1f-like dynamics appear as possible candidates for the effective dynamics of the droplet.

For any $x \in \mathbb{Z}^2$, let $\mathbb{L}_+(x) = x + \{\vec{e}_1, \vec{e}_2 - \vec{e}_1\}$, $\mathbb{L}_-(x) = x - \{\vec{e}_1, \vec{e}_2 - \vec{e}_1\}$, and let

$$\begin{aligned}\vec{\mathcal{D}}_x(f,f) &= \nu\left(\mathbf{1}_{\{\omega_{x+\vec{e}_2}\in G_2\}}\mathbf{1}_{\{\omega_y\in G_1\ \forall y\in\mathbb{L}_+(x)\}}\,\mathrm{Var}_{\hat{\nu}_x}(f)\right),\\ \vec{\mathcal{E}}_x(f,f) &= \nu\left(\mathbf{1}_{\{\omega_{x+\vec{e}_1}\in G_2\}}\mathbf{1}_{\{\omega_{x-\vec{e}_1}\in G_1\}}\,\mathrm{Var}_{\hat{\nu}_x}\left(f\,|\,G_1\right)\right),\end{aligned} \tag{4.5}$$

$$\begin{aligned}\mathcal{D}_x(f,f) &= \nu\left(\mathbf{1}_{\{\omega_{x+\vec{e}_2}\in G_2\}}\mathbf{1}_{\{\omega_y\in G_1\ \forall y\in\mathbb{L}_+(x)\}}\,\mathrm{Var}_{\hat{\nu}_x}(f)\right) + \nu\left(\mathbf{1}_{\{\omega_{x-\vec{e}_2}\in G_2\}}\mathbf{1}_{\{\omega_y\in G_1\ \forall y\in\mathbb{L}_-(x)\}}\,\mathrm{Var}_{\hat{\nu}_x}(f)\right),\\ \mathcal{E}_x(f,f) &= \nu\left(\mathbf{1}_{\{\omega_{x+\vec{e}_1}\in G_2\}}\mathbf{1}_{\{\omega_{x-\vec{e}_1}\in G_1\}}\,\mathrm{Var}_{\hat{\nu}_x}\left(f\,|\,G_1\right)\right)\right) + \nu\left(\mathbf{1}_{\{\omega_{x-\vec{e}_1}\in G_2\}}\mathbf{1}_{\{\omega_{x+\vec{e}_1}\in G_1\}}\,\mathrm{Var}_{\hat{\nu}_x}\left(f\,|\,G_1\right)\right).\end{aligned} \tag{4.6}$$

Notice that the constraint involving the super-good event G_2 inside the oriented pair $(\vec{\mathcal{D}}_x(f,f), \vec{\mathcal{E}}_x(f,f))$ requires the event G_2 to occur either at $x+\vec{e}_1$ or at $x+\vec{e}_2$, namely next to where the local variance $\mathrm{Var}_{\hat{\nu}_x}(f)$ is computed. Here, "oriented" refers to a direction (East or North) on the lattice. Modulo reversing the direction of the axes, that is exactly the kind of orientation that characterizes the East model. Instead, the same constraint inside the non-oriented pair $(\mathcal{D}_x(f,f), \mathcal{E}_x(f,f))$ requires G_2 to occur at either $x \pm \vec{e}_1$ or $x \pm \vec{e}_2$, a fact that is characteristic of the symmetric orientation of the FA-1f model. This observation justifies dubbing the first pair "*East Dirichlet pair at x*", and the second one "*FA-1f Dirichlet pair at x*".

The next result bounds from above the r.h.s. of (4.4) with a sum over x of either the East or the FA-1f Dirichlet pairs at x. As $p_2 \to 0$, the overall cost of the inequality diverges like $e^{\ln^2(1/p_2)}$ in the first case, and $\mathrm{poly}(1/p_2)$ in the second case. Such a dichotomy is a direct consequence of the very different scalings as $q \to 0$ of the relaxation times of the East and FA-1f models with effective parameter $q_{\mathrm{eff}} = p_2$.

Proposition 4.5 ([46, Part B of Theorem 3.1]). *For any $t \in (0,1)$ there exist $\vec{T}(t)$ and $T(t)$ satisfying $\vec{T}(t) \le e^{O\left(\ln^2(1/t)\right)}$ and $T(t) \le 1/t^{O(1)}$ as $t \to 0$, such that the following holds. There exists $\delta > 0$ such that, for*

[4] the orientation of γ has nothing to do with the next Proposition 4.5 and the discussion before it.

all p_1, p_2 *satisfying* $\max\{p_2, (1-p_1)(\log p_2)^2\} \le \delta$, *and all local functions* f*:*

$$\sum_x \nu\big(\mathbf{1}_{\mathcal{A}_x} \operatorname{Var}_{\hat\nu_x}(f)\big) \le \begin{cases} \vec T(p_2) \sum_{x\in\mathbb{Z}^2} \big(\vec{\mathcal{D}}_x(f,f) + \vec{\mathcal{E}}_x(f,f)\big), \\ \\ T(p_2) \sum_{x\in\mathbb{Z}^2} \big(\mathcal{D}_x(f,f) + \mathcal{E}_x(f,f)\big). \end{cases} \tag{4.7}$$

The proof requires combining conditioning, convexity, and the following high-level idea in a delicate manner. Recall the definition of the event $\mathcal{A}_x$, and consider one term $\nu\big(\mathbf{1}_{\mathcal{A}_x} \operatorname{Var}_{\hat\nu_x}(f)\big)$ in the l.h.s. above. For clarity of presentation, let us pretend to substitute the event $\mathcal{A}_x$ with the (much) simpler event $\mathcal{B}_x = \{\omega :$ at the end point of Γ_x the event G_2 occurs$\}$, where $\Gamma_x = \{x^{(1)}, \dots, x^{(n+1)}\}$ is a deterministic oriented path of length $n+1 = \lfloor 1/p_2^2 \rfloor$ starting at $x^{(1)} = x + \vec e_2$. In doing so, we are: (1) replacing the random path γ entering in $\mathcal{A}_x$ with Γ_x, (2) neglecting the fact that γ is ω-good (a very likely property), and (3) forcing the event G_2 to occur at the end of Γ and not somewhere along it. Approximations (1) and (3) can be dealt with through some careful conditioning, whereas (2) requires working with the modified product measure $\otimes_{x\in\Gamma_x} \hat\nu_x(\cdot \,|\, x$ is good$)$ along the path. If we insist on this simplification, we can first use convexity to enlarge the local variance at x to a variance w.r.t. variables $\{\omega_x, \omega_{x^{(1)}}, \dots, \omega_{x^{(n)}}\}$ and get

$$\nu\big(\mathbf{1}_{\mathcal{B}_x} \operatorname{Var}_{\hat\nu_x}(f)\big) \le \nu\big(\mathbf{1}_{\{\omega_{x^{(n+1)}} \in G_2\}} \operatorname{Var}_{\otimes_{y\in I_x}\hat\nu_y}(f)\big),$$

where $I_x = \{x^{(n)}, x^{(n-1)}, \dots, x^{(1)}, x\}$. At this stage, we can exploit the "boundary condition" $\omega_{x^{(n+1)}} \in G_2$ at the end of Γ_x, to bound from above the variance term $\operatorname{Var}_{\otimes_{y\in I_x}\hat\nu_y}(f)$ using the Poincaré inequality for either the (generalized) East or FA-1f models on I_x, with effective particle event $\mathcal{O} = G_2^c$ (cf. Remark 1.1). It is this step that rigorously implements the intuitive idea that the "effective vacancy event" G_2 (i.e., the presence of a mobile droplet in concrete applications) can move along Γ_x in an East-like or FA-1f-like fashion. If we do so, we obtain

$$\nu\big(\mathbf{1}_{\{\omega_{x^{(n+1)}} \in G_2\}} \operatorname{Var}_{\otimes_{y\in I_x}\hat\nu_y}(f)\big) \le \begin{cases} e^{\Theta(\ln^2(1/p_2))} \sum_{y\in I_x} \mathcal{D}_y^{\text{East}}(f,f), \\ \\ \big(1/p_2\big)^{O(1)} \sum_{y\in I_x} \mathcal{D}_y^{\text{FA1f}}(f,f). \end{cases} \tag{4.8}$$

Here, when $y = x^{(j)}, j = 1, \dots, n$,

$$\begin{aligned} \mathcal{D}_y^{\text{East}}(f,f) &= \nu\big(\mathbf{1}_{\{\omega_{x^{(j+1)}} \in G_2\}} \operatorname{Var}_{\hat\nu_y}(f)\big), \\ \mathcal{D}_y^{\text{FA-1f}}(f,f) &= \nu\big((\mathbf{1}_{\{\omega_{x^{(j+1)}} \in G_2\}} + \mathbf{1}_{\{\omega_{x^{(j-1)}} \in G_2\}}) \operatorname{Var}_{\hat\nu_y}(f)\big). \end{aligned}$$

These two Dirichlet terms are quite similar, at least for the role of the " effective vacancy event" G_2, to the East and FA-1f Dirichlet pairs at y. If we recall that Γ_x is oriented, then $x^{(j+1)}$ can be either $y + \vec e_1$ or $y + \vec e_2$, whereas $x^{(j-1)}$ is either $y - \vec\varepsilon_1$ or $y - \vec\varepsilon_2$, depending on Γ_x. By summing over x either one of the above inequalities and by keeping track of the over-counting cost, we obtain two bounds from above for the expression $\sum_x \nu\big(\mathbf{1}_{\mathcal{B}_x} \operatorname{Var}_{\hat\nu_x}(f)\big)$ which are similar to those of Proposition 4.5.

In applying the Poincaré inequalities (4.7) to KCMs, the final, crucial, step is to convert them into a final Poincaré inequality (4.1). In turn, this requires proving that a full resampling of the block V at x, corresponding to the local variance term $\operatorname{Var}_{\hat\nu_x}(f)$ inside the East/FA-1f Dirichlet pairs, in the presence of good blocks centered at the vertices of $\mathbb{L}_+(x)$ and/or $\mathbb{L}_-(x)$, and a super-good (i.e., with a mobile droplet) block at an oriented neighbor of x in the East case, or at an arbitrary neighbor of x in the FA-1f case, can be simulated (or reproduced) by a sequence of legal single-site updates of the original KCM. Technically, we want to bound from above the East or FA-1f Dirichlet pairs at x by a suitable sum of local Dirichlet terms $\nu(c_y \operatorname{Var}_{\mu_y}(f))$. Of course, the global cost of the latter inequality should be compatible with the target scaling of T_{rel} (see (2.6) for a similar inequality for the FA-2f model). It is here that the full theory of $\mathcal{U}$-bootstrap percolation comes into play. Firstly, it suggests the

complex details of the renormalization construction, namely the choice of the finite set V in $S = \{0,1\}^V$ and of the events G_1, G_2. Secondly, it determines which one of the two inequalities (4.7) one should use. That depends, in fact, whether the mobile droplet (namely the event G_2) is able to move along the ω-good oriented path γ in (4.3) in an East-like or FA-1f-like fashion. While the renormalization construction and the task described above is relatively simple for supercritical models, for critical ones it is significantly more complex and, for finitely many stable directions, required very significant adaptations and new ideas [31] compared to the tools developed in [46].

Acknowledgements KCMs gave me the privilege and the joy to work together with some extraordinary students and young brilliant researchers. Our non-stop discussions about the many facets of the subject have been a great intellectual challenge and a highlight of my personal life for a long time. I am deeply grateful to all my collaborators for that. Special thanks go to C. Toninelli for knocking at my door at IHES in 2004, armed with several simple-looking questions about KCMs and unlimited optimism. I would also like to thank R. Morris for his openness and for sharing his combinatorial insight into the continuous ping-pong between KCMs and bootstrap percolation. Finally, I am grateful to D. Chandler, who very sadly passed away in 2017, and to J.P. Garrahan for their appreciation of our mathematical results.

References

[1] M. Aizenman and J. L. Lebowitz, *Metastability effects in bootstrap percolation*, J. Phys. A, 21 (1988), pp. 3801–3813.

[2] D. Aldous and P. Diaconis, *The asymmetric one-dimensional constrained Ising model: rigorous results*, J. Stat. Phys., 107 (2002), pp. 945–975.

[3] P. Balister, B. Bollobás, R. Morris, and P. Smith, *Subcritical monotone cellular automata*, Random Structures Algorithms, 64 (2024), pp. 38–61.

[4] P. Balister, B. Bollobás, R. Morris, and P. Smith, *The Critical Length for Growing a Droplet*, Mem. Amer. Math. Soc., 310 (2025), pp. v+180.

[5] 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), pp. 7385–7411.

[6] P. Balister, B. Bollobás, R. Morris, and P. Smith, *Universality for monotone cellular automata*, 2022, https://arxiv.org/abs/2203.13806.

[7] O. Blondel, *Front progression in the East model*, Stochastic Process. Appl., 123 (2013), pp. 3430–3465.

[8] 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), pp. 383–406.

[9] B. Bollobás, H. Duminil-Copin, R. Morris, and P. Smith, *Universality for two-dimensional critical cellular automata*, Proc. Lond. Math. Soc. (3), 126 (2023), pp. 620–703.

[10] B. Bollobás, P. Smith, and A. Uzzell, *Monotone cellular automata in a random environment*, Combin. Probab. Comput., 24 (2015), pp. 687–722.

[11] C. Campailla and F. Martinelli, *Cutoff for East models*, 2025, https://arxiv.org/abs/2504.05760.

[12] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Kinetically constrained spin models*, Probab. Theory Related Fields, 140 (2008), pp. 459–504.

[13] P. Caputo, F. Martinelli, and F. Toninelli, *Mixing times of monotone surfaces and SOS interfaces: a mean curvature approach*, Comm. Math. Phys., 311 (2012), pp. 157–189.

[14] 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), pp. 1817–1863.

[15] F. Chung, P. Diaconis, and R. Graham, *Combinatorics for the East model*, Adv. in Appl. Math., 27 (2001), pp. 192–206.

[16] Y. Couzinié and F. Martinelli, *On a front evolution problem for the multidimensional East model*, Electron. J. Probab., 27 (2022), pp. Paper No. 143, 30.

[17] Y. Couzinié, *The multicolour east model*, 2022, https://arxiv.org/abs/2207.09782, https://arxiv.org/abs/2207.09782.

[18] A. Ertul, *Cutoff for the Fredrickson-Andersen one spin facilitated model*, ALEA Lat. Am. J. Probab. Math. Stat., 19 (2022), pp. 511–536.

[19] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *Aging through hierarchical coalescence in the East model*, Comm. Math. Phys., 309 (2012), pp. 459–495.

[20] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *Universality in one-dimensional hierarchical coalescence processes*, Ann. Probab., 40 (2012), pp. 1377–1435.

[21] G. H. Fredrickson and H. C. Andersen, *Kinetic ising model of the glass transition*, Phys. Rev. Lett., 53 (1984), pp. 1244–1247.

[22] G. H. Fredrickson and H. C. Andersen, *Facilitated kinetic ising models and the glass transition*, The Journal of Chemical Physics, 83 (1985), pp. 5822–5831.

[23] S. Ganguly, E. Lubetzky, and F. Martinelli, *Cutoff for the East process*, Comm. Math. Phys., 335 (2015), pp. 1287–1322.

[24] S. Ganguly and F. Martinelli, *Upper triangular matrix walk: cutoff for finitely many columns*, Random Structures Algorithms, 55 (2019), pp. 313–341.

[25] J. Garrahan and D. Chandler, *Coarse-grained microscopic model of glass formers*, Proceedings of the National Academy of Sciences, 100 (2003), pp. 9710–9714.

[26] 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*, Physica A: Statistical Mechanics and its Applications, 504 (2018), pp. 130–154. Lecture Notes of the 14th International Summer School on Fundamental Problems in Statistical Physics.

[27] J. Gravner and A. E. Holroyd, *Slow convergence in bootstrap percolation*, Ann. Appl. Probab., 18 (2008), pp. 909–928.

[28] I. Hartarsky, *Refined universality for critical KCM: upper bounds*, Comm. Math. Phys., 405 (2024), pp. Paper No. 13, 68.

[29] I. Hartarsky and L. Marêché, *Refined universality for critical KCM: lower bounds*, Combin. Probab. Comput., 31 (2022), pp. 879–906.

[30] I. Hartarsky, L. Marêché, and C. Toninelli, *Universality for critical KCM: infinite number of stable directions*, Probab. Theory Related Fields, 178 (2020), pp. 289–326.

[31] I. Hartarsky, F. Martinelli, and C. Toninelli, *Universality for critical KCM: finite number of stable directions*, Ann. Probab., 49 (2021), pp. 2141–2174.

[32] I. Hartarsky, F. Martinelli, and C. Toninelli, *Coalescing and branching simple symmetric exclusion process*, Ann. Appl. Probab., 32 (2022), pp. 2841–2859.

[33] I. Hartarsky, F. Martinelli, and C. Toninelli, *Fredrickson–andersen model in two dimensions**, ESAIM: ProcS, 74 (2023), pp. 108–118.

[34] I. Hartarsky, F. Martinelli, and C. Toninelli, *Sharp threshold for the FA-2f kinetically constrained model*, Probab. Theory Related Fields, 185 (2023), pp. 993–1037.

[35] I. Hartarsky and R. Morris, *The second term for two-neighbour bootstrap percolation in two dimensions*, Trans. Amer. Math. Soc., 372 (2019), pp. 6465–6505.

[36] I. Hartarsky and C. Toninelli, *Kinetically constrained models*, vol. 53 of SpringerBriefs in Mathematical Physics, Springer Nature Switzerland, 2025, https://link.springer.com/book/10.1007/978-3-031-93115-4.

[37] I. Hartarsky and F. L. Toninelli, *Kinetically constrained models out of equilibrium*, Probab. Math. Phys., 5 (2024), pp. 461–489.

[38] R. Holley, *Free energy in a Markovian model of a lattice spin system*, Comm. Math. Phys., 23 (1971), pp. 87–99.

[39] A. E. Holroyd, *Sharp metastability threshold for two-dimensional bootstrap percolation*, Probab. Theory Related Fields, 125 (2003), pp. 195–224.

[40] R. 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 $d_c = 2$*, Journal of Statistical Mechanics: Theory and Experiment, 2006 (2006), p. P03006.

[41] J. Jäckle and S. Eisinger, *A hierarchically constrained kinetic ising model*, Zeitschrift für Physik B Condensed Matter, 84 (1991), pp. 115–124.

[42] W. Kob and H. 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), pp. 4364–4377.

[43] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, Providence, RI, 2009. With a chapter by James G. Propp and David B. Wilson.

[44] L. Marêché, *Exponential convergence to equilibrium in supercritical kinetically constrained models at high temperature*, Markov Process. Related Fields, 26 (2020), pp. 861–884.

[45] L. Marêché, F. Martinelli, and C. Toninelli, *Exact asymptotics for Duarte and supercritical rooted kinetically constrained models*, Ann. Probab., 48 (2020), pp. 317–342.

[46] F. Martinelli, R. Morris, and C. Toninelli, *Universality results for kinetically constrained spin models in two dimensions*, Comm. Math. Phys., 369 (2019), pp. 761–809.

[47] F. Martinelli, R. Morris, and C. Toninelli, *Universality results for kinetically constrained spin models in two dimensions*, Comm. Math. Phys., 369 (2019), pp. 761–809.

[48] 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), pp. 2189–2210.

[49] F. Martinelli, A. Shapira, and C. Toninelli, *Long time behaviour of one facilitated kinetically constrained models: results and open problems*, in preparation, (2025).

[50] F. Martinelli and C. Toninelli, *Towards a universality picture for the relaxation to equilibrium of kinetically constrained models*, Ann. Probab., 47 (2019), pp. 324–361.

[51] R. Morris, *Monotone cellular automata*, in Surveys in combinatorics 2017, vol. 440 of London Math. Soc. Lecture Note Ser., Cambridge Univ. Press, Cambridge, 2017, pp. 312–371.

[52] C. Neuhauser and A. Sudbury, *The biased annihilating branching process*, Adv. in Appl. Probab., 25 (1993), pp. 24–38.

[53] 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), p. 021051.

[54] Y. Peres and A. Sly, *Mixing of the upper triangular matrix walk*, Probab. Theory Related Fields, 156 (2013), pp. 581–591.

[55] N. S. Pillai and A. Smith, *Mixing times for a constrained Ising process on the torus at low density*, Ann. Probab., 45 (2017), pp. 1003–1070.

[56] 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), pp. 1649–1678.

[57] J. Reiter, *Statics and dynamics of the two-spin–facilitated kinetic ising model*, The Journal of Chemical Physics, 95 (1991), pp. 544–554.

[58] K. Royen, S. Mondal, F. Pollmann, and F. Heidrich-Meisner, *Enhanced many-body localization in a kinetically constrained model*, Phys. Rev. E, 109 (2024), p. 024136.

[59] A. Shapira, *A note on the spectral gap of the Fredrickson-Andersen one spin facilitated model*, J. Stat. Phys., 181 (2020), pp. 2346–2352.

[60] P. Sollich and M. Evans, *Glassy time-scale divergence and anomalous coarsening in a kinetically constrained spin chain*, Phys. Rev. Lett., 83 (1999), pp. 3238–3241.

[61] P. Sollich and M. R. Evans, *Glassy dynamics in the asymmetrically constrained kinetic ising chain*, Phys. Rev. E, 68 (2003), p. 031504.

[62] A. Stauffer, *Multiscale analysis of particle systems*, (2020), https://ricerca.matfis.uniroma3.it/users/astauffer/rd.pdf.

[63] A. Sudbury and P. Lloyd, *Quantum operators in classical probability theory. II. The concept of duality in interacting particle systems*, Ann. Probab., 23 (1995), pp. 1816–1830.

[64] A. Sudbury and P. Lloyd, *Quantum operators in classical probability theory. IV. Quasi-duality and thinnings of interacting particle systems*, Ann. Probab., 25 (1997), pp. 96–114.

[65] C. Toninelli, G. Biroli, and D. S. Fisher, *Cooperative behavior of kinetically constrained lattice gas models of glassy dynamics*, J. Stat. Phys., 120 (2005), pp. 167–238.

[66] M. van Horssen, E. Levi, and J. P. Garrahan, *Dynamics of many-body localization in a translation-invariant quantum glass model*, Phys. Rev. B, 92 (2015), p. 100305.