Universality Results for Kinetically Constrained Spin Models in Two Dimensions

Fabio Martinelli, Robert Morris, Cristina Toninelli

UNIVERSALITY RESULTS FOR KINETICALLY CONSTRAINED SPIN MODELS IN TWO DIMENSIONS

FABIO MARTINELLI, ROBERT MORRIS, AND CRISTINA TONINELLI

ABSTRACT. Kinetically constrained models (KCM) are reversible interacting particle systems on $\mathbb{Z}^d$ with continuous time Markov dynamics of Glauber type, which represent a natural stochastic (and non-monotone) counterpart of the family of cellular automata known as $\mathcal{U}$-bootstrap percolation. KCM also display some of the peculiar features of the so-called "glassy dynamics", and as such they are extensively used in the physics literature to model the liquid-glass transition, a major and longstanding open problem in condensed matter physics.

We consider two-dimensional KCM with update rule $\mathcal{U}$, and focus on proving universality results for the mean infection time of the origin, in the same spirit as those recently established in the setting of $\mathcal{U}$-bootstrap percolation. We first identify what we believe are the correct universality classes, which turn out to be different from those of $\mathcal{U}$-bootstrap percolation. We then prove universal upper bounds on the mean infection time within each class, which we conjecture to be sharp up to logarithmic corrections. In certain cases, including all supercritical models, and the well-known Duarte model, our conjecture has recently been confirmed in [27]. In fact, in these cases our upper bound is sharp up to a constant factor in the exponent. For certain classes of update rules, it turns out that the infection time of the KCM diverges much faster than for the corresponding $\mathcal{U}$-bootstrap process when the equilibrium density of infected sites goes to zero. This is due to the occurrence of energy barriers which determine the dominant behaviour for KCM, but which do not matter for the monotone bootstrap dynamics.

CONTENTS

2010 *Mathematics Subject Classification.* Primary 60K35, secondary 60J27.

Key words and phrases. Glauber dynamics, kinetically constrained models, spectral gap, bootstrap percolation, universality.

1. Introduction

Kinetically constrained models (KCM) are interacting particle systems on the integer lattice $\mathbb{Z}^d$, which were introduced in the physics literature in the 1980s in order to model the liquid-glass transition (see e.g. [22,30] for reviews), a major and still largely open problem in condensed matter physics. The main motivation for the ongoing (and extremely active) research on KCM is that, despite their simplicity, they feature some of the main signatures of a super-cooled liquid near the glass transition point.

A generic KCM is a continuous-time Markov process of Glauber type defined as follows. A configuration ω is defined by assigning to each site $x \in \mathbb{Z}^d$ an occupation variable $\omega_x \in \{0,1\}$, corresponding to an empty or occupied site respectively. Each site waits an independent, mean one, exponential time and then, iff a certain local constraint is satisfied by the current configuration ω, its occupation variable is updated to be occupied with rate p and to empty with rate $q = 1-p$. All the constraints that have been considered in the physics literature belong to the following general class [10].

Fix an *update family* $\mathcal{U} = \{X_1, \ldots, X_m\}$, that is, a finite collection of finite subsets of $\mathbb{Z}^d \setminus \{\mathbf{0}\}$. Then ω satisfies the constraint at site x if there exists $X \in \mathcal{U}$ such that $\omega_y = 0$ for all $y \in X + x$. Since each update set belongs to $\mathbb{Z}^d \setminus \{\mathbf{0}\}$, the constraints never depend on the state of the to-be-updated site. As a consequence, the product Bernoulli(p) measure μ is a reversible invariant measure, and the process started at μ is stationary. Despite this trivial equilibrium measure, however, KCM display an extremely rich behaviour which is qualitatively different from that of interacting particle systems with non-degenerate birth/death rates (e.g. the stochastic Ising model). This behaviour includes the key dynamical features of real glassy materials: anomalously long mixing times [1, 10, 26], aging and dynamical heterogeneities [21], and ergodicity breaking transitions corresponding to percolation of blocked structures [22]. Moreover, proving the above results rigorously turned out to be a surprisingly challenging task, in part due to the fact that several of the classical tools typically used to analyse reversible interacting particle systems (e.g., coupling, censoring, logarithmic Sobolev inequalities) fail for KCM.

KCM can be also viewed as a natural non-monotone and stochastic counterpart of $\mathcal{U}$-bootstrap percolation, a well-studied class of discrete cellular automata, see [4,5,8]. For $\mathcal{U}$-bootstrap on $\mathbb{Z}^d$, given a configuration of "infected" sites A_t at time t, infected sites remain infected, and a site v becomes infected at time $t+1$ if the translate by v of one of the sets in $\mathcal{U}$ belongs to A_t. One then defines the final infection set $[A]_{\mathcal{U}} := \bigcup_{t=1}^{\infty} A_t$ and the *critical probability* of the $\mathcal{U}$-bootstrap process on $\mathbb{Z}^d$ to be

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

where $\mathbb{P}_q$ denotes the product probability measure on $\mathbb{Z}^d$ with density q of infected sites. The following key connection between $\mathcal{U}$-bootstrap percolation and KCM has been established by Cancrini, Martinelli, Roberto and Toninelli [10]: KCM processes are ergodic for $q > q_c\big(\mathbb{Z}^d,\mathcal{U}\big)$, and they are not ergodic for $q < q_c\big(\mathbb{Z}^d,\mathcal{U}\big)$. Furthermore, the results of [10] prove that in the ergodic regime time auto-correlations decay exponentially for a large class of KCM including all the models that have been considered in the physics literature (East model, Friedrickson-Andersen models, North-East model). More precisely, for these models the *relaxation time* $T_{\rm rel}(q;\mathcal{U})$ (see Definition 2.9) and the *mean infection time*[1] $\mathbb{E}_\mu(\tau_0)$ (i.e. the mean over the stationary KCM process of the first time at which the origin becomes empty) are finite for $q > q_c\big(\mathbb{Z}^d,\mathcal{U}\big)$ and infinite for $q < q_c\big(\mathbb{Z}^d,\mathcal{U}\big)$. Both from a physical and mathematical point of view, two key questions arise: (i) are the relaxation time and the mean infection time finite if $q > q_c(\mathbb{Z}^d,\mathcal{U})$ for any possible choice of the update families $\mathcal{U}$? (ii) which is the divergence of the time scales $T_{\rm rel}(q;\mathcal{U})$ and $\mathbb{E}_\mu(\tau_0)$ as $q \downarrow q_c(\mathbb{Z}^d,\mathcal{U})$? We will now briefly review some of the known results, which show that KCM exhibit a very large variety of possible scalings depending on the details of the update family $\mathcal{U}$.

We begin by discussing one of the most extensively studied KCM, which was introduced by Jäckle and Eisinger [23]: the so-called *East model*. This model has update family $\mathcal{U} = \big\{\{-\vec{e}_1,\},\dots,\{-\vec{e}_d\}\big\}$, so in the one-dimensional setting $d = 1$ a site can update iff it is the neighbour "to the east" of an empty site. It is not difficult to see that in any dimension $q_c(\mathbb{Z}^d,\mathcal{U}) = 0$. For $d = 1$, it was first proved in [1] that the relaxation time $T_{\rm rel}(q)$ is finite for any $q \in (0,1]$, and it was later shown (see [1, 10, 13]) that it diverges as

$$\exp\left(\big(1+o(1)\big)\frac{\log(1/q)^2}{2\log 2}\right)$$

as $q \downarrow 0$. A similar scaling was later proved in any dimension $d \geqslant 1$, see [14].

Another well-studied KCM, introduced by Friedrickson and Andersen [2], is the k-facilitated model (FA-kf), whose update family consists of the k-sets of nearest neighbours of the origin: a site can be updated iff it has at least k empty nearest neighbours. In this case it was proved in [19, 32] that $q_c(\mathbb{Z}^d,\mathcal{U}) = 0$ for all $1 \leqslant k \leqslant d$, whereas $q_c(\mathbb{Z}^d,\mathcal{U}) = 1$ for all $k > d$. Moreover, the relaxation time $T_{\rm rel}(q)$ diverges as $1/q^{\Theta(1)}$ when $k = 1$ [10, 35], and as a $(k-1)$-times iterated exponential of $q^{-1/(d-k+1)}$ when $2 \leqslant k \leqslant d$ [26]. The above scalings also hold for the mean infection time $\mathbb{E}_\mu(\tau_0)$.

The above model-dependent results (which are, in fact, the only ones that have been proved so far) include a large diversity of possible scalings of the mean infection time, together with a strong sensitivity to the details of the update family $\mathcal{U}$. Therefore, a very natural "universality" question emerges:

Question. *Is it possible to group all possible update families $\mathcal{U}$ into distinct classes, in such a way that all members of the same class induce the same divergence of the mean infection time as q approaches from above the critical value $q_c(\mathbb{Z}^d,\mathcal{U})$?*

[1] The mean infection time is very close to the *persistence time* in the physics literature

Such a general question has not been addressed so far, even in the physics literature: physicists lack a general criterion to predict the different scalings. This fact is particularly unfortunate since, due to the anomalous and sharp divergence of times, numerical simulations often cannot give clear-cut and reliable answers. Indeed, some of the rigorous results recalled above corrected some false conjectures that were based on numerical simulations.

The universality question stated above has, however, being addressed and successfully solved for two-dimensional $\mathcal{U}$-bootstrap percolation (see [4, 5, 8], or [29] for a recent review). The update families $\mathcal{U}$ were classified by Bollobás, Smith and Uzzell [8] into three universality classes: *supercritical*, *critical* and *subcritical* (see Definition 2.2), according to a simple geometric criterion. They also proved in [8] that $q_c(\mathbb{Z}^2, \mathcal{U}) = 0$ if $\mathcal{U}$ is supercritical or critical, and it was proved by Balister, Bollobás, Przykucki and Smith [4] that $q_c(\mathbb{Z}^2, \mathcal{U}) > 0$ if $\mathcal{U}$ is subcritical. For critical update families $\mathcal{U}$, the scaling (as $q \downarrow 0$) of the typical infection time of the origin starting from $\mathbb{P}_q$ was determined very precisely by Bollobás, Duminil-Copin, Morris and Smith [5] (improving bounds obtained in [8]), and various universal properties of the dynamics were obtained.

In this paper we take an important step towards establishing a similar universality picture for two-dimensional KCM with supercritical or critical update family $\mathcal{U}$. Using a geometric criterion, we propose a classification of the two-dimensional update families into universality classes, which is inspired by, but at the same time quite different from, that established for bootstrap percolation. More precisely, we classify a supercritical update family $\mathcal{U}$ as being *supercritical unrooted* or *supercritical rooted* and a critical $\mathcal{U}$ as being α*-rooted* or β*-unrooted*, where $\alpha \in \mathbb{N}$ and $\alpha \leqslant \beta \in \mathbb{N} \cup \{\infty\}$ are called the *difficulty* and the *bilateral difficulty* of $\mathcal{U}$ respectively (see Definitions 2.11 and 2.12). We then prove (see Sections 3-7) the following two main universality results (see Theorems 1 and 2 in Section 2.3) on the mean infection time $\mathbb{E}_\mu(\tau_0)$ and on the relaxation time T_{rel}.

Supercritical KCM. *Let $\mathcal{U}$ be a supercritical two-dimensional update family. Then, for $q > 0$, both $\mathbb{E}_\mu(\tau_0)$ and T_{rel} are finite. And, as $q \to 0$,*

(a) *if $\mathcal{U}$ is unrooted*

$$\mathbb{E}_\mu(\tau_0) \leqslant T_{\mathrm{rel}}/q \leqslant q^{-O(1)};$$

(b) *if $\mathcal{U}$ is rooted*

$$\mathbb{E}_\mu(\tau_0) \leqslant T_{\mathrm{rel}}/q \leqslant \exp\Big(O\big(\log q^{-1}\big)^2\Big).$$

Critical KCM. *Let $\mathcal{U}$ be a critical two-dimensional update family with difficulty α and bilateral difficulty β. Then, for $q>0$, both $\mathbb{E}_\mu(\tau_0)$ and $T_{\rm rel}$ are finite. And, as $q\to 0$,*
(a) if $\mathcal{U}$ is α-rooted

$$\mathbb{E}_\mu(\tau_0)\leqslant T_{\rm rel}/q\leqslant \exp\Big(q^{-2\alpha}\big(\log q^{-1}\big)^{O(1)}\Big);$$

(b) if $\mathcal{U}$ is β-unrooted

$$\mathbb{E}_\mu(\tau_0)\leqslant T_{\rm rel}/q\leqslant \exp\Big(q^{-\beta}\big(\log q^{-1}\big)^{O(1)}\Big).$$

Even though the theorems above only establish universal *upper bounds* on $\mathbb{E}_\mu(\tau_0)$ and $T_{\rm rel}$, we conjecture that our bounds provide the correct scaling up to logarithmic corrections. This has recently been proved for supercritical models in [27]. For critical update families, the bound $\mathbb{E}_\mu(\tau_0)=\Omega(T_\mathcal{U})$ (see [26, Lemma 4.3]), where $T_\mathcal{U}$ denotes the median infection time of the origin for the $\mathcal{U}$-bootstrap process at density q, together with the results of [5] on $T_\mathcal{U}$, provide a matching lower bound for all β-unrooted models with $\alpha=\beta$ (for example, the FA-2f model). In particular, these recent advances combined with the above theorems prove two conjectures that we put forward in [29]. Among the α-rooted models, those which have been considered most extensively in the literature are the Duarte and modified Duarte model (see [6,16,28,33]), for which $\alpha=1$ and $\beta=\infty$. In [27], using very different tools and ideas from those in this paper, a lower bound on $\mathbb{E}_\mu(\tau_0)$ was recently obtained for both models that matches our upper bound, including the logarithmic corrections, yielding $\mathbb{E}_\mu(\tau_0)=\exp\big(\Theta\big(q^{-2}(\log 1/q)^4\big)\big)$.

The above results imply that for all supercritical rooted KCM, and also for the Duarte-KCM, the mean infection time diverges much faster than the median infection time for the corresponding $\mathcal{U}$-bootstrap process, which obeys $T_\mathcal{U}\sim 1/q^{\Theta(1)}$ for supercritical models [8], and $T_\mathcal{U}\sim\exp\big(\Theta\big(q^{-1}(\log 1/q)^2\big)\big)$ for the Duarte model [28]. This is a consequence of the fact that for these KCM the infection time is not well-approximated by the number of updates needed to infect the origin (as it is for bootstrap percolation), but is the result of a much more complex mechanism. In particular, the visits of the process to regions of the configuration space with an anomalous amount of infection (borrowing from physical jargon we may call them "energy barriers") are heavily penalized and require a very long time to actually take place.

Providing an insight into the heuristics and/or the key steps of the proofs at this stage, before providing a clear definition of the geometrical quantities involved, would inevitably be rather vague. We therefore defer these explanations to Section 2.4. We can, however, state two high-level ingredients. The first one consists in identifying, for each class of update families $\mathcal{U}$, an "efficient" (and potentially optimal) dynamical strategy for the difficult (i.e., unlikely) task of infecting the origin. This is necessarily more complex than the growth of the corresponding $\mathcal{U}$-bootstrap process, since an efficient strategy must necessarily feature both infection and healing in order to avoid crossing excessively high energy barriers. The second ingredient consists in using the above strategy as a *guide*,[2] without actually implementing it, for the analytic technique introduced in [26] by two of the authors of the present paper, which allows one to bound the relaxation time $T_{\rm rel}(q; U)$. In [26] this technique was successfully applied to the FA-kf model, with the imagined mechanism for infecting the origin being a large droplet of infected sites moving as a random walk in a suitable (evolving) random environment of sparse infection. Here we have to go well beyond the method of [26], since the random walk picture does not apply to rooted models. Our main novelty is a new and more complex analytic approach to bound $T_{\rm rel}(q, \mathcal{U})$ which is inspired by the East dynamics (see Section 2.4 for more details).

1.1. **Notation.** We gather here (for the reader's convenience) some of the standard notation that we use throughout the paper. First, recall that we write μ for the Bernoulli product measure $\otimes_{x\in\mathbb{Z}^2}\mathrm{Ber}(p)$ on $\mathbb{Z}^2$, where $q = 1 - p$ will always be assumed to be sufficiently small (depending on the update family $\mathcal{U}$).

If f and g are positive real-valued functions of q, then we will write $f = O(g)$ if there exists a constant $C > 0$ (depending on $\mathcal{U}$, but *not* on q) such that $f(q) \leqslant Cg(q)$ for every sufficiently small $q > 0$. We will also write $f(q) = \Omega(g(q))$ if $g(q) = O(f(q))$ and $f(q) = \Theta(g(q))$ if both $f(q) = O(g(q))$ and $g(q) = O(f(q))$.

All constants, including those implied by the notation $O(\cdot)$, $\Omega(\cdot)$ and $\Theta(\cdot)$, are quantities that may depend on the update family $\mathcal{U}$ (and other quantities where explicitly stated) but not on the parameter q. If c_1 and c_2 are constants, then $c_1 \gg c_2 \gg 1$ means that c_2 is sufficiently large, and c_1 is sufficiently large depending on c_2. Similarly, $1 \gg c_1 \gg c_2 > 0$ means that c_1 is sufficiently small, and c_2 is sufficiently small depending on c_1. Finally, we will use the standard notation $[n] = \{1, \dots, n\}$.

2. Universality classes for KCM and main results

In this section we will begin by recalling the main universality results for bootstrap cellular automata. We will then define the KCM process associated to a bootstrap update family, introduce its universality classes, and state our main results about its scaling near criticality. To finish, we will provide an outline of the heuristics behind our main theorems, and a sketch of their proofs.

2.1. **The bootstrap monotone cellular automata and its universality properties.** Let us begin by defining a large class of two-dimensional monotone cellular automata, which were recently introduced by Bollobás, Smith and Uzzell [8].

[2] In this respect our situation shares some similarities with other large deviations problems, where an imagined optimal dynamical strategy has the role of suggesting and motivating several, otherwise mysterious, analytic steps.

Definition 2.1. Let $\mathcal{U} = \{X_1, \ldots, X_m\}$ be an arbitrary finite collection of finite subsets of $\mathbb{Z}^2 \setminus \{\mathbf{0}\}$. The *$\mathcal{U}$-bootstrap process* on $\mathbb{Z}^2$ is defined as follows: given a set $A \subset \mathbb{Z}^2$ of initially *infected* sites, set $A_0 = A$, and define for each $t \geqslant 0$,

$$A_{t+1} = A_t \cup \big\{ x \in \mathbb{Z}^2 \,:\, X + x \subset A_t \text{ for some } X \in \mathcal{U} \big\}. \tag{2.1}$$

We write $[A]_{\mathcal{U}} = \bigcup_{t \geqslant 0} A_t$ for the *closure* of A under the $\mathcal{U}$-bootstrap process.

Thus, a vertex x becomes infected at time $t+1$ if the translate by x of one of the sets in $\mathcal{U}$ (which we refer to as the *update family*) is already entirely infected at time t, and infected vertices remain infected forever. For example, if we take $\mathcal{U}$ to be the family of 2-subsets of the set of nearest neighbours of the origin, we obtain the classical 2-neighbour bootstrap process. One of the key insights of Bollobás, Smith and Uzzell [8] was that, at least in two dimensions, the typical global behaviour of the $\mathcal{U}$-bootstrap process acting on random initial sets should be determined by the action of the process on discrete half-planes.

For each unit vector $u \in S^1$, let $\mathbb{H}_u := \{x \in \mathbb{Z}^2 : \langle x, u \rangle < 0\}$ denote the discrete half-plane whose boundary is perpendicular to u.

Definition 2.2. The set of *stable directions* is

$$\mathcal{S} = \mathcal{S}(\mathcal{U}) = \big\{ u \in S^1 \,:\, [\mathbb{H}_u]_{\mathcal{U}} = \mathbb{H}_u \big\}.$$

The update family $\mathcal{U}$ is:

- *supercritical* if there exists an open semicircle in S^1 that is disjoint from $\mathcal{S}$,
- *critical* if there exists a semicircle in S^1 that has finite intersection with $\mathcal{S}$, and if every open semicircle in S^1 has non-empty intersection with $\mathcal{S}$,
- *subcritical* if every semicircle in S^1 has infinite intersection with $\mathcal{S}$.

The first step towards justifying this trichotomy is given by the following theorem, which was proved in [4, 8]. Recall from (1.1) the definition of $q_c\big(\mathbb{Z}^2,\mathcal{U}\big)$, the critical probability of the $\mathcal{U}$-bootstrap process on $\mathbb{Z}^2$.

Theorem 2.3. *If $\mathcal{U}$ is a supercritical or critical two-dimensional update family, then $q_c\big(\mathbb{Z}^2,\mathcal{U}\big) = 0$, whereas if $\mathcal{U}$ is subcritical then $q_c\big(\mathbb{Z}^2,\mathcal{U}\big) > 0$.*

For supercritical and critical update families, the main question is therefore to determine the scaling as $q \to 0$ of the typical time it takes to infect the origin.

Definition 2.4. The *typical infection time* at density q of an update family $\mathcal{U}$ is defined to be

$$T_{\mathcal{U}} \;=\; T_{q,\mathcal{U}} \;:=\; \inf\left\{ t \geqslant 0 \,:\, \mathbb{P}_q\big(\mathbf{0} \in A_t\big) \geqslant \frac{1}{2} \right\},$$

where (we recall) $\mathbb{P}_q$ indicates that every site is included in A with probability q, independently from all other sites, and A_t was defined in (2.1). We will write $T_{\mathcal{U}}$, omitting the suffix q from the notation, whenever there is no risk of confusion.

In order to state the main result of [5] we need some additional definitions. Let $\mathbb{Q}_1 \subset S^1$ denote the set of rational directions on the circle, and for each $u \in \mathbb{Q}_1$, let ℓ_u^+ be the (infinite) subset of the line $\ell_u := \{x \in \mathbb{Z}^2 : \langle x, u \rangle = 0\}$ consisting of the origin and the sites to the right of the origin as one looks in the direction of u. Similarly, let $\ell_u^- := (\ell_u \setminus \ell_u^+) \cup \{\mathbf{0}\}$ consist of the origin and the sites to the left of the origin. Given a two-dimensional bootstrap percolation update family $\mathcal{U}$, let $\alpha_{\mathcal{U}}^+(u)$ be the minimum (possibly infinite) cardinality of a set $Z \subset \mathbb{Z}^2$ such that $[\mathbb{H}_u \cup Z]_{\mathcal{U}}$ contains infinitely many sites of ℓ_u^+, and define $\alpha_{\mathcal{U}}^-(u)$ similarly (using ℓ_u^- in place of ℓ_u^+).

Definition 2.5. Given $u \in \mathbb{Q}_1$, the *difficulty* of u (with respect to $\mathcal{U}$) is[3]

$$\alpha(u) := \begin{cases} \min\big\{\alpha_{\mathcal{U}}^+(u), \alpha_{\mathcal{U}}^-(u)\big\} & \text{if } \alpha_{\mathcal{U}}^+(u) < \infty \text{ and } \alpha_{\mathcal{U}}^-(u) < \infty \\ \infty & \text{otherwise.} \end{cases}$$

Let $\mathcal{C}$ denote the collection of open semicircles of S^1. The *difficulty* of $\mathcal{U}$ is given by

$$\alpha := \min_{C \in \mathcal{C}} \max_{u \in C} \alpha(u), \tag{2.2}$$

and the *bilateral difficulty* by

$$\beta := \min_{C \in \mathcal{C}} \max_{u \in C} \max\big\{\alpha(u), \alpha(-u)\big\}. \tag{2.3}$$

A critical update family $\mathcal{U}$ is *balanced* if there exists a closed semicircle C such that $\alpha(u) \leqslant \alpha$ for all $u \in C$. It is said to be *unbalanced* otherwise.

Remark 2.6. If $u \in S^1$ is not a stable direction then $[\mathbb{H}_u]_{\mathcal{U}} = \mathbb{Z}^2$ (see [8, Lemma 3.1]), and therefore $\alpha(u) = 0$. Moreover, it was proved in [8, Lemma 5.2] (see also [5, Lemma 2.7]) that if $u \in \mathcal{S}(\mathcal{U})$ then $\alpha(u) < \infty$ if and only u is an isolated point of $\mathcal{S}(\mathcal{U})$. It follows that $\alpha = 0$ for every supercritical update family, and that α is finite for every critical update family. Observe also that $\alpha \leqslant \beta \leqslant \infty$, and that β can be infinite even for a supercritical update family (for example, one can embed the one-dimensional East model in two dimensions). A well-studied critical model with β infinite (and $\alpha = 1$) is the Duarte model (see [6, 16, 28, 33]), which has update family

$$\mathcal{D} = \big\{\{(-1,0),(0,1)\},\{(-1,0),(0,-1)\},\{(0,1),(0,-1)\}\big\}. \tag{2.4}$$

Roughly speaking, Definition 2.5 says that a direction u has finite difficulty if there exists a finite set of sites that, together with the half-plane $\mathbb{H}_u$, infect the entire line ℓ_u. Moreover, the difficulty of u is at least k if it is necessary (in order to infect ℓ_u) to find at least k infected sites that are 'close' to one another. If the open semicircle C with u as midpoint contains no direction of difficulty greater than k, then it is possible for a "critical droplet" of infected sites to grow in the direction of u without ever finding more than k infected sites close together. As a consequence, if the bilateral difficulty is not

[3] In order to slightly simplify the notation, and since the update family $\mathcal{U}$ will always be clear from the context, we will not emphasize the dependence of the difficulty on $\mathcal{U}$.

greater than k, then there exists a direction u (the midpoint of the optimal semicircle in (2.3)) such that a suitable critical droplet is able to grow in *both directions* u *and* $-u$, without ever finding more than k infected sites close together.

We are now in a position to state the main results on the scaling of the typical infection time for supercritical and critical update families. The following bounds were proved in [5] (for critical families) and in [8] (for supercritical families).

Theorem 2.7. *Let $\mathcal{U}$ be a two-dimensional update family. Then, as $q \to 0$,*

(a) if $\mathcal{U}$ is supercritical then

$$T_{\mathcal{U}} \,=\, q^{-\Theta(1)};$$

(b) if $\mathcal{U}$ is critical and balanced with difficulty α, then

$$T_{\mathcal{U}} \,=\, \exp\left(\frac{\Theta(1)}{q^{\alpha}}\right);$$

(c) if $\mathcal{U}$ is critical and unbalanced with difficulty α, then

$$T_{\mathcal{U}} \,=\, \exp\left(\frac{\Theta\big(\log(1/q)\big)^2}{q^{\alpha}}\right).$$

Remark 2.8. Note that in the above result the bilateral difficulty β plays no role. This is because in bootstrap percolation a droplet of empty sites only needs to grow in one direction (as opposed to moving back and forth). For KCM, on the other hand, we will see that the ability to move in two opposite directions will play a crucial role.

2.2. **General finite-range KCM.** In this section we define a class of two-dimensional interacting particle systems known as *kinetically constrained models*. As will be clear from what follows, KCM are intimately connected with bootstrap cellular automata.

We will work on the probability space (Ω, μ), where $\Omega = \{0,1\}^{\mathbb{Z}^2}$ and μ is the product Bernoulli(p) measure, and we will be interested in the asymptotic regime $q \downarrow 0$, where $q = 1-p$. Given $\omega \in \Omega$ and $x \in \mathbb{Z}^2$, we will say that x is "empty" (or "infected") if $\omega_x = 0$. We will say that $f\colon \Omega \mapsto \mathbb{R}$ is a *local function* if it depends on only finitely many of the variables ω_x.

Given a two-dimensional update family $\mathcal{U} = \{X_1, \dots, X_m\}$, the corresponding KCM is the Markov process on Ω associated to the Markov generator

$$(\mathcal{L}f)(\omega) = \sum_{x \in \mathbb{Z}^2} c_x(\omega)\big(\mu_x(f) - f\big)(\omega), \tag{2.5}$$

where $f\colon \Omega \mapsto \mathbb{R}$ is a local function, $\mu_x(f)$ denotes the average of f w.r.t. the variable ω_x, and c_x is the indicator function of the event that there exists an update rule $X \in \mathcal{U}$ such that $\omega_y = 0$ for every $y \in X + x$.

Informally, this process can be described as follows. Each vertex $x \in \mathbb{Z}^2$, with rate one and independently across $\mathbb{Z}^2$, is resampled from $\big(\{0,1\}, \mathrm{Ber}(p)\big)$ iff one of the update rules of the $\mathcal{U}$-bootstrap process at x is satisfied by the current configuration of the empty sites. In what follows, we will sometimes call such an update a *legal update* or *legal spin flip*. It follows (see [10]) that $\mathcal{L}$ is the generator of a reversible Markov process on Ω, with reversible measure μ.

We now define the two main quantities we will use to characterize the dynamics of the KCM process. The first of these is the relaxation time $T_{\rm rel}(q,\mathcal{U})$.

Definition 2.9. We say that $C > 0$ is a Poincaré constant for a given KCM if, for all local functions f, we have

$$\mathrm{Var}(f) \leqslant C\,\mathcal{D}(f), \tag{2.6}$$

where $\mathcal{D}(f) = \sum_x \mu\big(c_x \mathrm{Var}_x(f)\big)$ is the KCM Dirichlet form of f associated to $\mathcal{L}$. If there exists a finite Poincaré constant we then define

$$T_{\rm rel}(q,\mathcal{U}) := \inf\big\{C > 0 \,:\, C \text{ is a Poincaré constant for the KCM}\big\}.$$

Otherwise we say that the relaxation time is infinite.

A finite relaxation time implies that the reversible measure μ is mixing for the semi-group $P_t = e^{t\mathcal{L}}$ with exponentially decaying time auto-correlations [25]. More precisely, in that case $T_{\rm rel}(q,\mathcal{U})^{-1}$ coincides with the best positive constant λ such that,

$$\mathrm{Var}\left(e^{t\mathcal{L}}f\right) \leqslant e^{-2\lambda t}\,\mathrm{Var}(f) \qquad \forall\, f \in L^2(\mu). \tag{2.7}$$

One of the main results of [10] states that $T_{\rm rel}(q,\mathcal{U}) < \infty$ when $q > q_c$ for a large class of KCM including all the models that have been considered in the physics literature (East model, Friedrickson-Andersen models, North-East model).

The second (random) quantity is the hitting time

$$\tau_0 = \inf\big\{t \geqslant 0 \,:\, \omega_0(t) = 0\big\}.$$

In the physics literature the hitting time τ_0 is usually referred to as the *persistence time*, while in the bootstrap percolation framework it would be more conveniently dubbed the *infection time*. For our purposes, the most important connection between the mean infection time $\mathbb{E}_\mu(\tau_0)$ for the stationary KCM process (*i.e.,* with μ as initial distribution) and $T_{\rm rel}(q,\mathcal{U})$ is as follows (see [9, Theorem 4.7]):

$$\mathbb{E}_\mu(\tau_0) \leqslant \frac{T_{\rm rel}(q,\mathcal{U})}{q} \qquad \forall\, q \in (0,1). \tag{2.8}$$

The proof is quite simple. By definition, τ_0 is the hitting time of $A = \big\{\omega :\, \omega_0 = 0\big\}$, and it is a standard result (see, e.g., [3, Theorem 2]) that $\mathbb{P}_\mu(\tau_0 > t) \leqslant e^{-t\lambda_A}$, where

$$\lambda_A = \inf\big\{\mathcal{D}(f) \,:\, \mu(f^2) = 1 \text{ and } f(\omega) = 0 \text{ for every } \omega \in A\big\}.$$

Observe that $\mathrm{Var}(f) \geqslant \mu(A) = q$ for any function f satisfying $\mu(f^2) = 1$ that is identically zero on A. This implies that $\lambda_A \geqslant q/T_{\rm rel}(q,\mathcal{U})$, and so (2.8) follows.

Remark 2.10. If the initial distribution ν of the KCM process is different from the invariant measure μ, then it is only known that $\mathbb{E}_\nu(\tau_0)$ is finite in a couple of specific cases (the d-dimensional East process [11, 12], and the 1-dimensional FA-1f process [7]), even under the assumption that ν is a product Bernoulli(p') measure with $p' \neq p$.

A matching lower bound on $\mathbb{E}_\mu(\tau_0)$ in terms of $T_{\rm rel}(q,\mathcal{U})$ is not known. However, in [26, Lemma 4.3] it was proved that

$$\mathbb{E}_\mu(\tau_0) = \Omega(T_\mathcal{U}). \tag{2.9}$$

2.3. **Universality results.** We are now ready to define precisely the universality classes for KCM with a supercritical or critical update family. We will also restate (in a more precise form) our main results and conjectures on the scaling of $\mathbb{E}_\mu(\tau_0)$ and $T_{\rm rel}$ as $q \to 0$. We begin with the (much easier) supercritical case.

Definition 2.11. A supercritical two-dimensional update family $\mathcal{U}$ is said to be *supercritical rooted* if there exist two non-opposite stable directions in S^1. Otherwise it is called *supercritical unrooted*.

Our first main result, already stated in the Introduction, provides an upper bound on $\mathbb{E}_\mu(\tau_0)$ and $T_{\rm rel}$ for every supercritical two-dimensional update family that is (by the results of [27]) sharp up to the implicit constant factor in the exponent. Recall that if $\mathcal{U}$ is supercritical then $T_\mathcal{U} = q^{-\Theta(1)}$, by Theorem 2.7.

Theorem 1 (Supercritical KCM). *Let $\mathcal{U}$ be a supercritical two-dimensional update family. Then, for $q > 0$, both $\mathbb{E}_\mu(\tau_0)$ and $T_{\rm rel}$ are finite. And, as $q \to 0$,*

(a) *if $\mathcal{U}$ is unrooted*

$$\mathbb{E}_\mu(\tau_0) \leqslant q^{-1}\, T_{\rm rel} \leqslant q^{-O(1)} = \exp\Big(O\big(\log T_\mathcal{U}\big)\Big),$$

(b) *if $\mathcal{U}$ is rooted,*

$$\mathbb{E}_\mu(\tau_0) \leqslant q^{-1}\, T_{\rm rel} \leqslant \exp\Big(O\big(\log q^{-1}\big)^2\Big) = \exp\Big(O\big(\log T_\mathcal{U}\big)^2\Big).$$

We next turn to our bounds for critical update families, the proofs of which will require us to overcome a number of significant technical challenges, in addition to those encountered in the supercritical case. In this setting the distinction between critical unrooted and critical rooted is more subtle, and both the difficulty α and the bilateral difficulty β (see Definition 2.5) play an important role. Recall that for a critical update family the difficulty is finite, but that the bilateral difficulty may be infinite.

Definition 2.12. A critical update family $\mathcal{U}$ with difficulty α and bilateral difficulty β is said to be *α-rooted* if $\beta \geqslant 2\alpha$. Otherwise it is said to be *β-unrooted*.[4]

[4] We warn the attentive reader that when $\alpha < \beta < 2\alpha$ the model is here called β-unrooted, while in [29] it was called α-rooted.

The following theorem is the main contribution of this paper.

Theorem 2 (Critical KCM)**.** *Let $\mathcal{U}$ be a critical two-dimensional update family with difficulty α and bilateral difficulty β. Then, for $q > 0$, both $\mathbb{E}_\mu(\tau_0)$ and $T_{\rm rel}$ are finite. And, as $q \to 0$,*

(a) if $\mathcal{U}$ is α-rooted

$$\mathbb{E}_\mu(\tau_0) \leqslant q^{-1}\ T_{\rm rel} \leqslant \exp\Big(O\Big(q^{-2\alpha}\big(\log q^{-1}\big)^4\Big)\Big) = \exp\Big(\tilde{O}\big(\log T_{\mathcal{U}}\big)^2\Big);$$

(b) if $\mathcal{U}$ is β-unrooted

$$\mathbb{E}_\mu(\tau_0) \leqslant q^{-1}\ T_{\rm rel} \leqslant \exp\Big(O\Big(q^{-\beta}\big(\log q^{-1}\big)^3\Big)\Big) = \exp\Big(\tilde{O}\big(\log T_{\mathcal{U}}\big)^{\beta/\alpha}\Big).$$

It was recently proved in [27] that the upper bounds in Theorem 1 are best possible up to the implicit constant factor in the exponent for all supercritical update families (note that this follows from (2.9) for unrooted models). We conjecture that the bounds for critical models in Theorem 2 are also best possible, though in a slightly weaker sense: up to a polylogarithmic factor in the exponent.

Conjecture 3. *Let $\mathcal{U}$ be a critical two-dimensional update family with difficulty α and bilateral difficulty β. Then, as $q \to 0$,*

(a) if $\mathcal{U}$ is α-rooted

$$\mathbb{E}_\mu(\tau_0) = \exp\Big(q^{-2\alpha}\big(\log q^{-1}\big)^{\Theta(1)}\Big);$$

(b) if $\mathcal{U}$ is β-unrooted

$$\mathbb{E}_\mu(\tau_0) = \exp\Big(q^{-\beta}\big(\log q^{-1}\big)^{\Theta(1)}\Big).$$

The same result hold for $T_{\rm rel}$.

Observe that for α-unrooted update families $\mathcal{U}$ (i.e., families with $\beta = \alpha$), the lower bound in Conjecture 3 follows from Theorem 2.7 and (2.9); in particular Theorem 2 confirms [29, Conjecture 2.4]. If $\mathcal{U}$ is moreover unbalanced, then the upper and lower bounds given by Theorems 2 and 2.7 differ by only a single factor of $\log(1/q)$ (in the exponent), and we suspect that in this case the lower bound is correct, see Remark 6.14.

Conjecture 4. *Let $\mathcal{U}$ be an α-unrooted, unbalanced, critical two-dimensional update family with difficulty α. Then, as $q \to 0$,*

$$\mathbb{E}_\mu(\tau_0) = \exp\Big(\Theta\Big(q^{-\alpha}\big(\log q^{-1}\big)^2\Big)\Big).$$

The same result holds for $T_{\rm rel}$.

We remark that an example of an update family satisfying the conditions of Conjecture 4 is the so-called *anisotropic model* (see, e.g., [17, 18]) whose update family consists of all subsets of size 3 of the set

$$\big\{(-2,0),(-1,0),(1,0),(2,0),(0,1),(0,-1)\big\}.$$

Another model for which Conjecture 3 holds is the Duarte model, defined in (2.4), for which a matching lower bound (this time, up to a *constant* factor in the exponent) was recently proved in [27], confirming (in a strong sense) [29, Conjecture 2.5]. For all other critical models, however, the best known lower bound is that given by Theorem 2.7 and (2.9), and is therefore (we think) very far from the truth.

2.4. **Heuristics and roadmap.** We conclude this section with a high-level description of the intuition behind the proofs of Theorems 1 and 2, together with a roadmap of the actual proof, which is carried out in Sections 3–7.

The first key point to be stressed is that we never actually follow the dynamics of the KCM process itself; instead, we will prove the existence of a Poincaré constant with the correct scaling as $q \to 0$, and use the inequality (2.8) to deduce a bound on the mean infection time. We emphasize that this approach only works for the stationary KCM, that is, the process starting from the stationary measure μ. The second point is that, given that the Dirichlet form of the KCM

$$\mathcal{D}(f) = \sum_{x \in \mathbb{Z}^2} \mu\big(c_x \operatorname{Var}_x(f)\big)$$

is a sum of local variances ($\Leftrightarrow$ spin flips) computed with suitable infection nearby ($\Leftrightarrow$ the constraints c_x), all of our reasoning will be guided by the fact that we need to have some infection ($\Leftrightarrow$ empty sites) next to where we want to compute the variance. Therefore, much of our intuition, and all of the technical tools, have been developed with the aim of finding a way to *effectively* move infection where we need it.

A configuration sampled from μ will always have "mesoscopic" droplets (large patches of infected sites), though these will typically be very far from the origin. The general theory of $\mathcal{U}$-bootstrap percolation developed in [5, 8] allows us to quantify very precisely the critical size of those droplets that (typically) allows infection to grow from them and invade the system. However – and this is a fundamental difference between bootstrap percolation and KCM – it is extremely unlikely for the stationary KCM to create around a given vertex and at a given time a very large cluster of infection. Thus, it is essential to envisage an *infection/healing* mechanism that is able to *move* infection over long distances without creating too large an excess[5] of it.

At the root of our approach lies the notion of a *critical droplet*. A critical droplet is a certain finite set D whose geometry depends on the update family $\mathcal{U}$, and whose characteristic size may depend on q. For supercritical models we can take any sufficiently large (*not* depending on q) rectangle oriented along the mid-point u of a semicircle C free of stable directions. For critical models the droplet D is a more complicated object called a *quasi-stable half-ring* (see Definition 4.9 and Figure 4) oriented along the midpoint u of an open semicircle with largest difficulty either α or β. The long sides of D will have length either $\Theta\big(q^{-\alpha}\log(1/q)\big)$ or $\Theta\big(q^{-\beta}\log(1/q)\big)$ for the α-rooted and

[5]In physical terms an excess of infection is equivalent to an "energy barrier".

β-unrooted cases respectively, while the short sides will always have length $\Theta(1)$. The key feature of a critical droplet for supercritical models (see Section 4.2) is that, if it is empty, then it is able to infect a suitable translate of itself in the u-direction. For unrooted supercritical models the semicircle C can be chosen in such a way that both C *and* $-C$ are free of stable directions. As a consequence, the empty critical droplet will be able to infect a suitable translate of itself in *both* directions $\pm u$.

For critical models the situation changes drastically. An empty critical droplet will not be able to infect freely another critical droplet next to it in the u-direction because of the stable directions which are present in every open semicircle. However, it will be able to do so (in the u-direction if the model is α-rooted, and in the $\pm u$-directions if β-unrooted) provided that it receives some help from a finite number of extra empty sites (in "clusters" of size α or β) nearby. If the size of the critical droplet is chosen as above, then it is straightforward to show that such extra helping empty sites will be present with high probability (see Section 6.1).

Having clarified what a critical droplet is, and under which circumstances it is able to infect nearby sites, we next explain what we mean by "moving a critical droplet". For simplicity we explain the heuristics only for the supercritical case. Imagine that we have a sequence $D_0, D_1, \dots, D_n$ of contiguous, non-overlapping and identical critical droplets such that $D_{i+1} = D_i + d_i u$ for some suitable $d_i > 0$. Suppose first that the model is unrooted and that D_0 is completely infected, and let us write ω_i for the configuration of spins in D_i. Using the infection in D_0 it possible to first infect D_1, then D_2 and then, using reversibility, restore (*i.e.,* heal) the original configuration ω_1 in D_1. Using the infection in D_2 we can next infect D_3 and then, using the infection in D_3, restore ω_2 in D_2 (see the schematic diagram below, where $\emptyset$ stands for an infected droplet)

$$\begin{aligned}\emptyset\,\omega_1\,\omega_2\,\omega_3\dots &\mapsto \emptyset\,\emptyset\,\omega_2\,\omega_3\dots \mapsto \emptyset\,\emptyset\,\emptyset\,\omega_3\dots\\ &\mapsto \emptyset\,\omega_1\,\emptyset\,\omega_3\dots \mapsto \emptyset\,\omega_1\,\emptyset\,\emptyset\dots \mapsto \emptyset\,\omega_1\,\omega_2\,\emptyset\dots\end{aligned}$$

If we continue in this way, we end up moving the original infection in D_0 to the last droplet D_n without having ever created more than two extra infected critical droplets simultaneously. We remark that the sequence described above is reminiscent of how infection moves in the one-dimensional 1-neighbour KCM.

For rooted supercritical models, on the other hand, we cannot simply restore the configuration ω_2 in D_2 using only the infection in D_3 (in the unrooted case this was possible because infection could propagate in both the u and $-u$ directions). As a consequence, we need to follow a more complicated pattern:

$$\begin{aligned}\emptyset\,\omega_1\,\omega_2\,\omega_3\dots &\mapsto \emptyset\,\emptyset\,\omega_2\,\omega_3\dots \mapsto \emptyset\,\emptyset\,\emptyset\,\omega_3\dots\\ &\mapsto \emptyset\,\emptyset\,\emptyset\,\emptyset\dots \mapsto \emptyset\;\;\emptyset\,\omega_2\,\emptyset\dots \mapsto \emptyset\,\omega_1\,\omega_2\,\emptyset\dots,\end{aligned}$$

in which healing is always induced by infection present in the adjacent droplet in the $-u$ direction. This latter case is reminiscent of the one-dimensional East model. In this case, a combinatorial result proved in [15] implies that in order to move the infection to D_n it is necessary to create $\asymp \log n$ *simultaneous* extra infected critical droplets. This logarithmic energy barrier is the reason for the different scaling of $\mathbb{E}_\mu(\tau_0)$ in rooted and unrooted supercritical models (see Theorem 1).

Let us now give a somewhat more detailed outline of our approach. We begin by partitioning $\mathbb{Z}^2$ into 'suitable' rectangular blocks $\{V_i\}_{i\in\mathbb{Z}^2}$ with shortest side orthogonal to the direction u (see Section 4.1). For supercritical models these blocks have sides of constant length, while for critical models they will have length $q^{-\kappa}$ for some constant $\kappa \gg \alpha$, and height equal to that of a critical droplet, so either $\Theta\big(q^{-\alpha}\log(1/q)\big)$ or $\Theta\big(q^{-\beta}\log(1/q)\big)$, depending on the nature of the model. Then, given a configuration $\omega \in \Omega$, we declare a block to be *good* or *super-good* according to the following rules:

- For supercritical models *any* block is good, while for critical models good blocks are those which contain "enough" empty sites to allow an adjacent empty critical droplet to advance in the u (or $\pm u$) direction(s) (see Definition 6.4).
- In both cases, a block is said to be super-good if it is good and also contains an empty (i.e., completely infected) critical droplet.

Good blocks turn out to be very likely w.r.t. μ (a triviality in the supercritical case), and it follows by standard percolation arguments that they form a rather dense infinite cluster. Super-good blocks, on the other hand, are quite rare, with density $\rho = q^{\Theta}(1)$ in the supercritical case, $\rho = \exp\big(-\Theta\big(q^{-\alpha}\log(1/q)^2\big)\big)$ in the critical α-rooted case, and $\rho = \exp\big(-\Theta\big(q^{-\beta}\log(1/q)^2\big)\big)$ for critical β-unrooted models.

We will then prove the existence of a suitable Poincaré constant in three steps, each step being associated to a natural kinetically constrained *block dynamics*[6] on a certain length scale. In each block dynamics the configuration in each block is resampled with rate one (and independently of other resamplings) if a certain constraint is satisfied.

Our first block dynamics forces one of the blocks neighbouring V_i to be at the beginning of an oriented "thick" path γ of good blocks, with length $\approx 1/\rho$, whose last block is super-good. Using the fact that this constraint is very likely, it is possible to prove (see Section 2 in [26]) that the relaxation time of this process is $O(1)$, and moreover (see Proposition 3.5) that the Poincaré inequality

$$\mathrm{Var}(f) \leqslant 4\sum_i \mu\big(\mathbb{1}_{\Gamma_i}\,\mathrm{Var}_i(f)\big) \tag{2.10}$$

holds, where $\mathbb{1}_{\Gamma_i}$ is the indicator of the event that a good path exists for V_i. Though this starting point is similar to the method we develop in [26], for the next two steps of the proof we introduce here a completely different set of tools and ideas in order

[6]See, e.g., Chapter 15.5 of [24] for a introduction to the technique of block dynamics in reversible Markov chains.

to avoid the direct use of canonical paths (which could instead be used in [26] for the special case of the FA-2f model). Indeed for a general model (and especially for rooted models), using canonical paths and evaluating their congestion constants would result in a very heavy and complicated machinery. The next idea is to convert the *long-range* constrained Poincaré inequality (2.10) into a *short-range* one of the form

$$\mathrm{Var}(f) \leqslant C_1(q) \sum_i \mu\big(\mathbb{1}_{SG_i} \mathrm{Var}_i(f)\big), \tag{2.11}$$

in which $\mathbb{1}_{SG_i}$ is the indicator of the event that a suitable collection of blocks *near* V_i are good and one of them is super-good. Which collections of blocks are "suitable", and which one should be super-good, depends on whether the model is rooted or unrooted; we refer the reader to Theorem 3.1 for the details. The main content of Theorem 3.1, which we present in a slightly more general setting for later convenience, is that $C_1(q)$ can be taken equal to the best Poincaré constant (*i.e.,* the relaxation time) of a one-dimensional generalised 1-neighbour or East process at the effective density ρ. Section 3 is entirely dedicated to the task of formalising and proving the above claim.

The final step of the proof is to convert the Poincaré inequality (2.11) into the true Poincaré inequality for our KCM

$$\mathrm{Var}(f) \leqslant C_2(q) \sum_x \mu(c_x \mathrm{Var}_x(f)),$$

with a Poincaré constant $C_2(q)$ which scales with q as required by Theorems 1 and 2. In turn, this requires us to prove that a full resampling of a block in the presence of nearby super-good and good blocks can be simulated (or reproduced) by a sequence of legal single-site updates of the *original* KCM, with a global cost in the Poincaré constant compatible with Theorems 1 and 2. It is here that the results of [5, 8] on the behaviour of the $\mathcal{U}$-bootstrap process come into play. While for supercritical models the task described above is relatively simple (see Section 5), for critical models the problem is significantly more complicated and a suitable generalised East process again plays a key role. A full sketch of the proof can be found in Section 6.1.2, see in particular the proof of Proposition 6.6, and Remark 6.7.

3. Constrained Poincaré inequalities

The aim of this section is to prove a constrained Poincaré inequality for a product measure on $S^{\mathbb{Z}^2}$, where S is a finite set. This general inequality will play an instrumental role in the proof of our main theorems, giving us precise control of the infection time for both supercritical and critical KCM.

In order to state our general constrained Poincaré inequality, we will need some notation. Let $(S, \hat{\mu})$ be a finite positive probability space, and set $\Omega = \big(S^{\mathbb{Z}^2}, \mu\big)$, where $\mu = \otimes_{i \in \mathbb{Z}^2} \hat{\mu}$. A generic element Ω will be denoted by $\omega = \{\omega_i\}_{i \in \mathbb{Z}^2}$. For any local function f we will write $\mathrm{Var}(f)$ for its variance w.r.t. μ and $\mathrm{Var}_i(f)$ for the variance w.r.t. to the variable $\omega_i \in S$ conditioned on all the other variables $\{\omega_j\}_{j\neq i}$. For any $i \in \mathbb{Z}^2$ we set

$$\mathbb{L}^+(i) = i + \{\vec{e}_1, \vec{e}_2 - \vec{e}_1\} \qquad \text{and} \qquad \mathbb{L}^-(i) = i - \{\vec{e}_1, \vec{e}_2 - \vec{e}_1\}.$$

Finally, let $G_2 \subseteq G_1 \subseteq S$ be two events, and set $p_1 := \hat{\mu}(G_1)$ and $p_2 := \hat{\mu}(G_2)$. The main result of this section is the following theorem.

Theorem 3.1. *For any $t \in (0,1)$ there exist $\vec{T}(t), T(t)$ satisfying $\vec{T}(t) \leqslant \exp\big(O\big(\log \frac{1}{t}\big)^2\big)$ and $T(t) \leqslant t^{-O(1)}$ as $t \to 0$, such that the following oriented and unoriented constrained Poincaré inequalities hold.*

(A) *Suppose that $G_1 = S$ and $G_2 \subseteq S$. Then, for all local functions f:*

$$\mathrm{Var}(f) \leqslant \vec{T}(p_2) \sum_{i\in\mathbb{Z}^2} \mu\Big(\mathbb{1}_{\{\omega_{i+\vec{e}_1}\in G_2\}} \mathrm{Var}_i(f)\Big) \tag{3.1}$$

$$\mathrm{Var}(f) \leqslant T(p_2) \sum_{i\in\mathbb{Z}^2} \mu\Big(\mathbb{1}_{\{\{\omega_{i+\vec{e}_1}\in G_2\}\cup\{\omega_{i-\vec{e}_1}\in G_2\}\}} \mathrm{Var}_i(f)\Big). \tag{3.2}$$

(B) *Suppose that $G_2 \subseteq G_1 \subseteq S$. Then there exists $\delta > 0$ such that, for all p_1, p_2 satisfying $\max\big\{p_2, (1-p_1)(\log p_2)^2\big\} \leqslant \delta$, and all local functions f:*

$$\mathrm{Var}(f) \leqslant \vec{T}(p_2)\bigg(\sum_{i\in\mathbb{Z}^2} \mu\Big(\mathbb{1}_{\{\omega_{i+\vec{e}_2}\in G_2\}} \mathbb{1}_{\{\omega_j\in G_1\,\forall j\in\mathbb{L}^+(i)\}} \mathrm{Var}_i(f)\Big) + \sum_{i\in\mathbb{Z}^2} \mu\Big(\mathbb{1}_{\{\omega_{i+\vec{e}_1}\in G_2\}} \mathbb{1}_{\{\omega_{i-\vec{e}_1}\in G_1\}} \mathrm{Var}_i\big(f\,|\,G_1\big)\Big)\bigg), \tag{3.3}$$

$$\mathrm{Var}(f) \leqslant T(p_2)\bigg(\sum_{\varepsilon=\pm1}\sum_{i\in\mathbb{Z}^2} \mu\Big(\mathbb{1}_{\{\omega_{i+\varepsilon\vec{e}_2}\in G_2\}} \mathbb{1}_{\{\omega_j\in G_1\,\forall j\in\mathbb{L}^{\varepsilon}(i)\}} \mathrm{Var}_i(f)\Big) + \sum_{\varepsilon=\pm1}\sum_{i\in\mathbb{Z}^2} \mu\Big(\mathbb{1}_{\{\omega_{i+\varepsilon\vec{e}_1}\in G_2\}} \mathbb{1}_{\{\omega_{i-\varepsilon\vec{e}_1}\in G_1\}} \mathrm{Var}_i\big(f\,|\,G_1\big)\Big)\bigg). \tag{3.4}$$

Remark 3.2. When proving Theorem 1 the starting point will be (3.1) or (3.2), depending on whether the model is rooted or unrooted. Similarly, for critical models we will start the proof of Theorem 2 from (3.3) or (3.4) depending on whether the model is α-rooted or β-unrooted. This choice is dictated by the $\mathcal{U}$-bootstrap process according to the following rule: we will require $V_i \subset [A]_{\mathcal{U}}$ to hold for *any* set A of empty sites such that the indicator function in front of $\mathrm{Var}_i(f)$ is equal to one. We refer the reader to Sections 5 and 6, and in particular to the proof of Lemma 5.2, for more details.

An important role in the proof of the theorem is played by the one-dimensional East and 1-neighbour processes (see, e.g., [10]), and a certain generalization of these processes. For the reader's convenience, we begin by recalling these generalized models.

3.1. **The generalised East and 1-neighbour models.** The standard versions of these two models are ergodic interacting particle systems on $\{0,1\}^n$ with kinetic constraints, which will mean that jumps in the dynamics are facilitated by certain configurations of vertices in state 0. They are both reversible w.r.t. the product measure $\pi = \mathrm{Ber}(\alpha_1) \otimes \cdots \otimes \mathrm{Ber}(\alpha_n)$, where $\mathrm{Ber}(\alpha)$ is the α-Bernoulli measure and $\alpha_1, \ldots, \alpha_n \in (0,1)$.

In the first process, known as the *non-homogeneous East model* (see [20, 23] and references therein), the state ω_x of each point $x \in [n]$ is resampled at rate one (independently across $[n]$) from the distribution $\mathrm{Ber}(\alpha_x)$, provided that $c_x(\omega) = 1$, where

$$c_x(\omega) = \mathbb{1}_{\{\omega_{x+1}=0\}} \qquad \text{and} \qquad \omega_{n+1} := 0.$$

In the second model, known as the *non-homogeneous* 1*-neighbour model* (and also as the FA-1f model [2]), the resampling occurs in the same way, except in this case

$$c_x(\omega) = \max\big\{ \mathbb{1}_{\{\omega_{x-1}=0\}}, \, \mathbb{1}_{\{\omega_{x+1}=0\}} \big\} \qquad \text{where} \qquad \omega_0 := 1 \qquad \text{and} \qquad \omega_{n+1} := 0.$$

It is known [1, 10, 13] that the corresponding relaxation times $T_{\text{East}}(n, \bar\alpha)$ and $T_{\text{FA}}(n, \bar\alpha)$ (where $\bar\alpha = (\alpha_1, \ldots, \alpha_n)$) are finite *uniformly* in n and that they satisfy the following scaling as $q := \min\big\{ 1 - \alpha_x : x \in [n] \big\}$ tends to zero:

$$T_{\text{East}}\big(n, \bar\alpha\big) = q^{-O(\min\{\log n, \, \log(1/q)\})} \qquad \text{and} \qquad T_{\text{FA}}\big(n, \bar\alpha\big) = q^{-O(1)}. \tag{3.5}$$

The proof of (3.5) is deferred to the Appendix. In the proof of Theorem 3.1 we will need to work in the following more general setting.

Consider a finite product probability space of the form $\Omega = \otimes_{x\in[n]}(S_x, \nu_x)$, where S_x is either a finite set or an interval of $\mathbb{R}$, and ν_x is a positive probability measure on S_x. Given $\{\omega_x\}_{x\in[n]} \in \Omega$, we will refer to ω_x as the *the state of the vertex* x. Moreover, for each $x \in [n]$, let us fix a constraining event $S_x^g \subseteq S_x$ with $q_x := \nu_x(S_x^g) > 0$. We consider the following generalisations of the East and FA-1f processes on the space Ω.

Definition 3.3. In the *generalised East chain*, the state ω_x of each vertex $x \in [n]$ is resampled at rate one (independently across $[n]$) from the distribution ν_x, provided that $\vec{c}_x(\omega) = 1$, where

$$\vec{c}_x(\omega) = \mathbb{1}_{\{\omega_{x+1} \in S_{x+1}^g\}}$$

if $x \in \{1, \ldots, n-1\}$, and $\vec{c}_n(\omega) \equiv 1$.

In the *generalised FA-1f chain*, the resampling occurs in the same way, except in this case $c_1(\omega) = \mathbb{1}_{\{\omega_2 \in S_2^g\}}$,

$$c_x(\omega) = \max\big\{ \mathbb{1}_{\{\omega_{x-1} \in S_{x-1}^g\}}, \, \mathbb{1}_{\{\omega_{x+1} \in S_{x+1}^g\}} \big\}$$

if $x \in \{2, \ldots, n-1\}$, and $c_n(\omega) \equiv 1$.

In both cases, set $q := \min_x q_x = \min_x \nu_x(S_x^g)$, and set $\alpha_x := 1 - q_x$ for each $x \in [n]$.

Note that the projection variables $\eta_x = \mathbb{1}_{\{S_x^g\}}$ evolve as a standard East or FA-1f chain, and it is therefore natural to ask whether the relaxation times of these generalised constrained chains can be bounded from above in terms of the relaxation times

$T_{\text{East}}(n,\bar{\alpha})$ and $T_{\text{FA}}(n,\bar{\alpha})$ respectively. The answer is affirmative, and it is the content of the following proposition (cf. [14, Proposition 3.4]), which provides us with Poincaré inequalities for the generalised East and FA-1f chains.

Proposition 3.4. *Let $f:\Omega\mapsto\mathbb{R}$. For the generalised East chain, we have*

$$\operatorname{Var}(f)\leqslant\frac{1}{q}\cdot T_{\mathit{East}}(n,\bar{\alpha})\cdot\sum_{x=1}^{n}\nu\big(\vec{c}_x\operatorname{Var}_x(f)\big),\tag{3.6}$$

and for the generalised FA-1f chain, we have

$$\operatorname{Var}(f)\leqslant\frac{1}{q}\cdot T_{\mathit{FA}}(n,\bar{\alpha})\cdot\sum_{x=1}^{n}\nu\big(c_x\operatorname{Var}_x(f)\big),\tag{3.7}$$

where $\operatorname{Var}_x(\cdot)$ denotes the conditional variance w.r.t. ν_x, given all the other variables.

The proof of this proposition, which is similar to that of [13, Proposition 3.4], is deferred to the Appendix.

3.2. **Proof of Theorem 3.1.** We begin with the proof of part (A), which is a relatively straightforward consequence of Proposition 3.4 and (3.5). The proof of part (B) is significantly more difficult, and we will require a technical result from [26] (see Proposition 3.5, below) and a careful application of Proposition 3.4 (and of convexity) after conditioning on various events.

3.2.1. *Proof of part (A).* Recall that in this setting $G_1=S$ and $G_2\subset S$, where $(S,\hat{\mu})$ is an arbitrary finite positive probability space. Let f be a local function and let $M>0$ be sufficiently large so that f does not depend on the variables at vertices (m,n) with $|m|\geqslant M$. For each $n\in\mathbb{Z}$, let μ_n denote the product measure $\otimes_{m\in\mathbb{Z}}\,\hat{\mu}$ on $S^{\mathbb{Z}\times\{n\}}$, and note that $\mu=\otimes_{n\in\mathbb{Z}}\,\mu_n$. By construction, $\operatorname{Var}_{\mu_n}(f)$ coincides with the same conditional variance computed w.r.t. $\mu_n^M:=\otimes_{m\in\mathbb{Z}\cap[-M,M]}\,\hat{\mu}$.

We apply Proposition 3.4 to the homogeneous product measure μ_n^M with the event G_2 as event S_x^g for all $x\in\{-M,\dots,M\}$. Note that $q_x=\hat{\mu}(G_2)=p_2$ for every x, and that $\operatorname{Var}_{(M,n)}(f)=\operatorname{Var}_{(-M,n)}(f)=0$. It follows, using (3.5), that

$$\operatorname{Var}_{\mu_n}(f)\leqslant\vec{T}(p_2)\sum_{m\in\mathbb{Z}}\mu_n\left(\mathbb{1}_{\{\omega_{(m+1,n)}\in G_2\}}\operatorname{Var}_{(m,n)}(f)\right),$$

where $\vec{T}(p_2)=\exp\Big(O\big(\log\frac{1}{p_2}\big)^2\Big)$, and

$$\operatorname{Var}_{\mu_n}(f)\leqslant T(p_2)\sum_{m\in\mathbb{Z}}\mu_n\left(\mathbb{1}_{\{\omega_{(m+1,n)}\in G_2\}\cup\{\omega_{(m-1,n)}\in G_2\}}\operatorname{Var}_{(m,n)}(f)\right),$$

where $T(p_2)=p_2^{-O(1)}$. Using the standard inequality $\operatorname{Var}_\mu(f)\leqslant\sum_{n\in\mathbb{Z}}\mu\big(\operatorname{Var}_{\mu_n}(f)\big)$, the Poincaré inequalities (3.1) and (3.2) follow.

3.2.2. *Proof of part (B).* We next turn to the significantly more challenging task of proving the constrained Poincaré inequalities (3.3) and (3.4). As noted above, in addition to Proposition 3.4 we will require a technical result from [26], stated below as Proposition 3.5. In order to state this result we need some additional notation.

Recall that an *oriented path of length* n in $\mathbb{Z}^2$ is a sequence $\gamma = (i^{(1)}, \ldots, i^{(n)})$ of n vertices of $\mathbb{Z}^2$ with the property that $i^{(k+1)} - i^{(k)} \in \{\vec{e}_1, \vec{e}_2\}$ for each $k \in [n-1]$. We will say that γ starts at $i^{(1)}$, ends at $i^{(n)}$, and that $i \in \gamma$ if $i = i^{(k)}$ for some $k \in [n]$. Moreover, given $\omega \in \Omega$, we will say that γ is

- *ω-good* if $\omega_i \in G_1$ for all $i \in \bigcup_{j \in \gamma} \{j, j + \vec{e}_1, j - \vec{e}_1\}$, and
- *ω-super-good* if it is good and there exists $i \in \gamma$ such that $\omega_i \in G_2$,

where $G_2 \subseteq G_1 \subseteq S$ are the events in the statement of Theorem 3.1.

In what follows it will be convenient to order the oriented paths of length n starting from a given point according to the alphabetical order of the associated strings of n unit vectors from the finite alphabet $\mathcal{X} = \{\vec{e}_1, \vec{e}_2\}$. Next, for each $i \in \mathbb{Z}^2$ we define the key event $\Gamma_i \subset \Omega$, as follows:

(i) there exists an oriented ω-good path γ, of length $L = \lfloor 1/p_2^2 \rfloor$ starting at $i + \vec{e}_2$;
(ii) the smallest such path (in the above order) is ω-super-good;
(iii) $\omega_{i+\vec{e}_1} \in G_1$.

In what follows, and if no confusion arises, we will abbreviate ω-good and ω-super-good to good and super-good respectively. The following upper bound on $\mathrm{Var}(f)$ is very similar to [26, Proposition 3.4], and we therefore defer the proof to the Appendix.

Proposition 3.5. *There exists $\delta > 0$ such that, if $\max\{p_2, (1-p_1)(\log p_2)^2\} \leqslant \delta$, then*

$$\mathrm{Var}(f) \leqslant 4 \sum_{i \in \mathbb{Z}^2} \mu\big(\mathbb{1}_{\Gamma_i} \mathrm{Var}_i(f)\big) \tag{3.8}$$

for every local function f.

We would like to use Proposition 3.4 to bound the right-hand side of (3.8). However, Proposition 3.4 provides us with an upper bound on the variance of a function, whereas the quantity $\mu\big(\mathbb{1}_{\Gamma_i} \mathrm{Var}_i(f)\big)$ is more like the average of a local variance. We will therefore need to use convexity to bound from above the average of a local variance by a full variance. In order to reduce as much as possible the potential loss of such an operation, we first perform a series of conditionings on the measure μ and use convexity only on the final conditional measure.

Roughly speaking, on the event Γ_i we first reveal, for each $j \neq i$ within distance $2/p_2^2$ of the origin, whether or not the event $\{\omega_j \in G_1\}$ holds. Given this information, we know which paths of length L and starting at $i + \vec{e}_2$ are good and we define γ^* as the smallest one in the order defined above. Next, we reveal the *last* $j^* \in \gamma^*$ such that $\{\omega_{j^*} \in G_2\}$. Note that in doing so we do not need to observe whether or not the event

$\{\omega_j \in G_2\}$ holds for any earlier j (i.e., before j^* in γ^*). Finally, defining $\gamma \subset \gamma^*$ to be the part of γ^* before j^*, we reveal ω_j for all $j \in \mathbb{Z}^2$, except for $j = i$ and $j \in \gamma$.

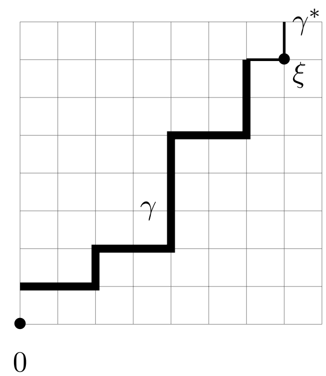

FIGURE 1. The minimal good path γ^*, the position of the first super-good vertex ξ encountered while traveling backward along γ^*, and the subpath $\gamma \subset \gamma^*$ (thick black) connecting $\vec{e}_2$ to a neighbour of ξ.

At the end of this process we are left with a (conditional) probability measure ν on $S^{\gamma\cup\{i\}}$. We will then apply convexity and Proposition 3.4 to this measure. We now detail the above procedure.

Proof of part (B) of Theorem 3.1. Let $\delta > 0$ be given by Proposition 3.5, and assume that the events $G_2 \subseteq G_1 \subseteq S$ satisfy $\max\left\{p_2, (1-p_1)(\log p_2)^2\right\} \leqslant \delta$. By Proposition 3.5, we have

$$\mathrm{Var}(f) \leqslant 4 \sum_{i\in\mathbb{Z}^2} \mu\big(\mathbb{1}_{\Gamma_i} \mathrm{Var}_i(f)\big) \tag{3.9}$$

for every local function f. We will bound each term of the sum in (3.9). Using translation invariance, it will suffice to consider the term $i = (0,0)$.

For each $\omega \in \Gamma_{(0,0)}$, let $\gamma^* = \gamma^*(\omega)$ denote the smallest ω-good oriented path of length L starting from $\vec{e}_2$, and note that γ^* is ω-super-good, since $\omega \in \Gamma_{(0,0)}$. Let $\xi = \xi(\omega) \in \gamma^*$ be the first super-good vertex encountered while travelling along γ^* backwards, *i.e.,* from its last point to its starting point $\vec{e}_2$. Finally, let γ be the portion of γ^* starting at $\vec{e}_2$ and ending at the vertex preceding ξ in γ^*.

We next perform the series of conditionings on the measure μ that were described informally above. Let Λ be the box of side-length $4/p_2^2$ centred at the origin. We first condition on the event $\Gamma_{(0,0)}$ and on the σ-algebra generated by the events

$$\big\{\{\omega_j \in G_1\} : j \in \Lambda \setminus \{(0,0)\}\big\}.$$

Note that, since we are conditioning on the event $\Gamma_{(0,0)}$, these events determine γ^*. Next we condition on the position of ξ on γ^*; this determines the path $\gamma = (i^{(1)}, \ldots, i^{(n)})$.

Finally we condition on all of the variables ω_j with $j \notin \gamma \cup \{(0,0)\}$. Let ν be the resulting conditional measure and observe that $(S^{\gamma\cup\{(0,0)\}}, \nu)$ is a product probability space of the form $\otimes_{j\in\gamma\cup\{(0,0)\}}(S_j, \nu_j)$, with $(S_{(0,0)}, \nu_{(0,0)}) = (S, \hat{\mu})$ and $(S_j, \nu_j) = \big(G_1, \hat{\mu}(\cdot \,|\, G_1)\big)$ for each $j \in \gamma$. Notice that

$$\mu\big(\mathbb{1}_{\Gamma_{(0,0)}} \operatorname{Var}_{(0,0)}(f)\big) = \mu\Big(\mathbb{1}_{\Gamma_{(0,0)}} \nu\big(\operatorname{Var}_{\nu_{(0,0)}}(f)\big)\Big) \leqslant \mu\big(\mathbb{1}_{\Gamma_{(0,0)}} \operatorname{Var}_{\nu}(f)\big), \tag{3.10}$$

because $\nu\big(\operatorname{Var}_{\nu_{(0,0)}}(f)\big) \leqslant \operatorname{Var}_\nu(f)$, by convexity.

We can now bound $\operatorname{Var}_\nu(f)$ from above by applying Proposition 3.4 to the measure $\nu = \otimes_{j\in\gamma\cup\{(0,0)\}}(S_j, \nu_j)$, with the super-good event G_2 as the constraining event S_j^g. Observe that $\nu\big(S^g_{(0,0)}\big) = \hat{\mu}(G_2) = p_2$ and $\nu\big(S^g_j\big) = \hat{\mu}\big(G_2 \,|\, G_1\big) = p_2/p_1$ for each $j \in \gamma$. The first Poincaré inequality (3.6) in Proposition 3.4 therefore gives

$$\mu\big(\mathbb{1}_{\Gamma_{(0,0)}} \operatorname{Var}_\nu(f)\big) \leqslant \vec{T}(p_2) \cdot \mu\Bigg(\mathbb{1}_{\Gamma_{(0,0)}} \sum_{i\in\gamma\cup\{(0,0)\}} \nu\Big(\mathbb{1}_{\{\omega_{m(i)}\in G_2\}} \operatorname{Var}_{\nu_i}(f)\Big)\Bigg), \tag{3.11}$$

where $m(i)$ is the next point on the path γ^* after i (*i.e.*, $m(i)$ is either $m(i) = i + \vec{e}_1$ or $m(i) = i + \vec{e}_2$) and

$$\vec{T}(p_2) \leqslant \frac{1}{p_2} \sup\big\{T_{\text{East}}(n, \bar{\alpha}) : n \leqslant L\big\} \leqslant p_2^{-O(\log(1/p_2))},$$

by (3.5). Recall that in Definition 3.3 the constraint for the last point is identically equal to one (this is in order to guarantee irreducibility of the chain), and observe that this condition holds in the above setting because, by construction, $\omega_\xi \in G_2$.

Finally, we claim that (3.11) implies that

$$\begin{aligned}\mu\big(\mathbb{1}_{\Gamma_{(0,0)}} \operatorname{Var}_\nu(f)\big) \leqslant \vec{T}(p_2) \sum_{i\in\Lambda} \Bigg(&\mu\Big(\mathbb{1}_{\{\omega_{i+\vec{e}_1}\in G_2\}} \mathbb{1}_{\{\omega_{i-\vec{e}_1}\in G_1\}} \operatorname{Var}_i(f \,|\, G_1)\Big) \\ &+ \mu\Big(\mathbb{1}_{\{\omega_{i+\vec{e}_2}\in G_2\}} \mathbb{1}_{\{\omega_j\in G_1\,\forall j\in\mathbb{L}^+(i)\}} \big(\operatorname{Var}_i(f) + \operatorname{Var}_i(f \,|\, G_1)\big)\Big)\Bigg).\end{aligned} \tag{3.12}$$

Indeed, note that $\operatorname{Var}_{\nu_{(0,0)}}(f) = \operatorname{Var}_{(0,0)}(f)$ and that $\operatorname{Var}_{\nu_i}(f) = \operatorname{Var}_i(f \,|\, G_1)$ for each $i \in \gamma$, and recall that, by construction, $\omega_{i+\vec{e}_1}, \omega_{i-\vec{e}_1} \in G_1$ for every $i \in \gamma$. Therefore, for each $i \in \gamma$, if $m(i) = i + \vec{e}_1$ then $\omega_{i-\vec{e}_1} \in G_1$, and if $m(i) = i + \vec{e}_2$ then $\omega_j \in G_1$ for each $j \in \mathbb{L}^+(i) = i + \big\{\vec{e}_1, \vec{e}_2 - \vec{e}_1\big\}$. Moreover, the event $\Gamma_{(0,0)}$ implies that $\omega_j \in G_1$ for each $j \in \mathbb{L}^+((0,0))$. Therefore every term of the right-hand side of (3.11) is included in the right-hand side of (3.12), and hence (3.11) implies (3.12), as claimed.

Now, combining (3.12) with (3.9) and (3.10), and noting that $\operatorname{Var}_i(f) \geqslant p_1 \operatorname{Var}_i(f \,|\, G_1)$ and that $|\Lambda| \leqslant p_2^{-O(1)}$, we obtain

$$\begin{aligned}\operatorname{Var}(f) \leqslant p_1^{-1} p_2^{-O(1)} \vec{T}(p_2) \sum_{i\in\mathbb{Z}^2} \Bigg(&\mu\Big(\mathbb{1}_{\{\omega_{i+\vec{e}_1}\in G_2\}} \mathbb{1}_{\{\omega_{i-\vec{e}_1}\in G_1\}} \operatorname{Var}_i\big(f \,|\, G_1\big)\Big) \\ &+ \mu\Big(\mathbb{1}_{\{\omega_{i+\vec{e}_2}\in G_2\}} \mathbb{1}_{\{\omega_j\in G_1\,\forall j\in\mathbb{L}^+(i)\}} \operatorname{Var}_i(f)\Big)\Bigg),\end{aligned}$$

which implies the oriented Poincaré inequality (3.3), as required.

The proof of the unoriented inequality (3.4) is almost the same, except we will use the second Poincaré inequality (3.7) in Proposition 3.4, instead of (3.6). To spell out the details, we obtain

$$\mu\big(\mathbb{1}_{\Gamma_{(0,0)}} \operatorname{Var}_\nu(f)\big) \leqslant T(p_2) \cdot \mu\bigg(\mathbb{1}_{\Gamma_{(0,0)}} \sum_{i \in \gamma \cup \{(0,0)\}} \nu\Big(c_i \operatorname{Var}_{\nu_i}(f)\Big)\bigg), \tag{3.13}$$

where c_i is the indicator of the event that G_2 holds for at least one of the neighbours of i on the path γ^*, and

$$T(p_2) \leqslant \frac{1}{p_2} \sup_{n \leqslant L} T_{\text{FA}}(n, \bar{\alpha}) = p_2^{-O(1)},$$

by (3.5). Note that the constraint for the last point is again identically equal to one since $\omega_\xi \in G_2$. It follows (cf. (3.12)) that

$$\begin{aligned}\mu\big(\mathbb{1}_{\Gamma_{(0,0)}} \operatorname{Var}_\nu(f)\big) \leqslant T(p_2) \sum_{i \in \Lambda} \sum_{\varepsilon = \pm 1} \bigg(&\mu\Big(\mathbb{1}_{\{\omega_{i+\varepsilon\vec{e}_1} \in G_2\}} \mathbb{1}_{\{\omega_{i-\varepsilon\vec{e}_1} \in G_1\}} \operatorname{Var}_i\big(f \,|\, G_1\big)\Big) \\ &+ \mu\Big(\mathbb{1}_{\{\omega_{i+\varepsilon\vec{e}_2} \in G_2\}} \mathbb{1}_{\{\omega_j \in G_1 \, \forall j \in \mathbb{L}^\varepsilon(i)\}} \big(\operatorname{Var}_i(f) + \operatorname{Var}_i\big(f \,|\, G_1\big)\big)\Big)\bigg),\end{aligned} \tag{3.14}$$

since $\omega_{i+\vec{e}_1}, \omega_{i-\vec{e}_1} \in G_1$ for every $i \in \gamma$, and the event $\Gamma_{(0,0)}$ implies that $\omega_j \in G_1$ for each $j \in \mathbb{L}^+((0,0))$. In particular, note that if $i \in \gamma$ and $i + \vec{e}_2 \in \gamma$, then $\omega_j \in G_1$ for each $j \in \mathbb{L}^+(i) = \mathbb{L}^-(i + \vec{e}_2) = i + \big\{\vec{e}_1, \vec{e}_2 - \vec{e}_1\big\}$. Therefore, as before, every term of the right-hand side of (3.13) is included in the right-hand side of (3.14).

Finally, combining (3.14) with (3.9) and (3.10), and since $\operatorname{Var}_i(f) \geqslant p_1 \operatorname{Var}_i(f \,|\, G_1)$ and $|\Lambda| \leqslant p_2^{-O(1)}$, we obtain

$$\begin{aligned}\operatorname{Var}(f) \leqslant p_1^{-1} p_2^{-O(1)} T(p_2) \sum_{i \in \mathbb{Z}^2} \sum_{\varepsilon = \pm 1} \bigg(&\mu\Big(\mathbb{1}_{\{\omega_{i+\varepsilon\vec{e}_1} \in G_2\}} \mathbb{1}_{\{\omega_{i-\varepsilon\vec{e}_1} \in G_1\}} \operatorname{Var}_i\big(f \,|\, G_1\big)\Big) \\ &+ \mu\Big(\mathbb{1}_{\{\omega_{i+\varepsilon\vec{e}_2} \in G_2\}} \mathbb{1}_{\{\omega_j \in G_1 \, \forall j \in \mathbb{L}^\varepsilon(i)\}} \operatorname{Var}_i(f)\Big)\bigg),\end{aligned}$$

which gives the unoriented Poincaré inequality (3.4), as claimed, and hence completes the proof of Theorem 3.1. □

4. Renormalization and spreading of infection

In this section we shall define the setting to which we will apply Theorem 3.1 in order to bound from above the relaxation time, and hence the mean infection time, of supercritical and critical KCM. We will begin with a very brief informal description, before giving (in Section 4.1) the precise definition. We will then, in Sections 4.2 and 4.3, state two results from the theory of bootstrap percolation that will play an instrumental role in the proofs of Theorems 1 and 2.

Our basic strategy is to partition the lattice $\mathbb{Z}^2$ into disjoint rectangular "blocks" $\{V_i\}_{i\in\mathbb{Z}^2}$, whose size is adapted to the bootstrap update family $\mathcal{U}$. To each block V_i we associate a block random variable ω_i, which is just the collection of i.i.d. $0/1$ Bernoulli(p) variables $\{\omega_x\}_{x\in V_i}$ attached to each vertex of the block. In order to avoid confusion we will always use the letters $i, j, \dots$ for the labels of quantities associated to blocks, and the letters $x, y, \dots$ for the labels of the quantities associated to vertices of $\mathbb{Z}^2$. We will apply Theorem 3.1 to the block variables $\{\omega_i\}_{i\in\mathbb{Z}^2}$.

4.1. **A concrete general setting.** Let v and $v^\perp$ be orthogonal rational directions in the first and second quadrant of $\mathbb{R}^2$ respectively. Let $\vec{v}$ be the vector joining the origin to the first site of $\mathbb{Z}^2$ in direction v, and similarly for $\vec{v}^\perp$. Let $n_1 \geqslant n_2$ be (sufficiently large) even integers, and set

$$R := \big\{x \in \mathbb{R}^2 \,:\, x = \alpha n_1 \vec{v} + \beta n_2 \vec{v}^\perp,\ \alpha, \beta \in [0,1)\big\}. \tag{4.1}$$

The finite probability space $(S, \hat{\mu})$ appearing in Section 3 will always be of the form $S = \{0,1\}^V$, where $V = R \cap \mathbb{Z}^2$, and $\hat{\mu}$ is the Bernoulli(p) product measure. Observe that the probability space $(S^{\mathbb{Z}^2}, \mu)$ is isomorphic to $\Omega = \{0,1\}^{\mathbb{Z}^2}$ equipped with the Bernoulli(p) product measure which, with a slight abuse of notation, we will continue to denote by μ. For our purposes, a convenient isomorphism between the two probability spaces is given by a kind of tilted "brick-wall" partition of $\mathbb{Z}^2$ into disjoint copies of the basic block V (see Figure 2). To be precise, for each $i = (i_1, i_2) \in \mathbb{Z}^2$, set $V_i := R_i \cap \mathbb{Z}^2$, where $R_i := R + (i_1 + i_2/2)n_1\vec{v} + i_2 n_2 \vec{v}^\perp$.

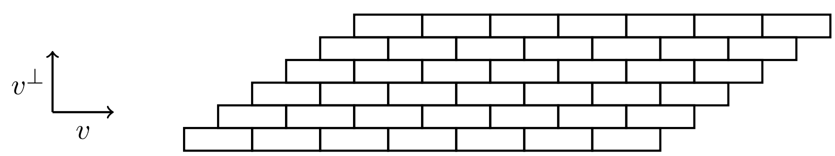

FIGURE 2. The partition into blocks $V_i,\ i \in \mathbb{Z}^2$

In this partition the "northern" and "southern" neighbouring blocks of V_i (*i.e.,* the blocks corresponding to $(i_1, i_2 \pm 1)$) are shifted in the direction $\vec{v}$ by $\pm n_1/2$ w.r.t. V_i. With this notation, and given $\omega \in S^{\mathbb{Z}^2}$, it is then convenient to think of the variable $\omega_i \in S$ as being the collection $\{\omega_x\}_{x\in V_i} \in \{0,1\}^{V_i}$. The local variance term $\mathrm{Var}_i(f)$ (*i.e.,* the variance of f w.r.t. the variable ω_i given all the other variables $\{\omega_j\}_{j\neq i}$), which appears in the various constrained Poincaré inequalities in the statement of Theorem 3.1, is then equal to the variance $\mathrm{Var}_{V_i}(f)$ w.r.t. the i.i.d. Bernoulli(p) variables $\{\omega_x\}_{x\in V_i}$, given all of the other variables $\{\omega_y\}_{y\in\mathbb{Z}^2\setminus V_i}$.

From now on, ω will always denote an element of $\{0,1\}^{\mathbb{Z}^2}$ and, given $\Lambda \subset \mathbb{R}^2$, we will write ω_Λ for the collection of i.i.d. random variables $\{\omega_x\}_{x\in\Lambda\cap\mathbb{Z}^2}$, and μ_Λ for their joint product Bernoulli(p) law. We will say that Λ is *empty* (or *empty in* ω) if ω is identically equal to 0 on $\Lambda \cap \mathbb{Z}^2$, and similarly that Λ is *filled* (or *completely occupied*) if ω is identically equal to 1 on $\Lambda \cap \mathbb{Z}^2$.

We now turn to the definitions of the good and super-good events $G_2 \subset G_1 \subseteq S$. The good event G_1 will depend on the update family $\mathcal{U}$, and will (roughly speaking) approximate the event that the block V_i can be "crossed" in the $\mathcal{U}$-bootstrap process with the help of a constant-width strip connecting the top and bottom of V_i. For supercritical models this event is trivial, and therefore G_1 is the entire space S; for critical models, on the other hand, G_1 will require the presence of empty vertices inside V obeying certain model-dependent geometric constraints (see Definition 6.4, below). The super-good event G_2 for supercritical models will simply require that V is empty. For critical models it will require that G_1 holds, and additionally that there exists an empty subset $\mathcal{R}$ of V, called a *quasi-stable half-ring* (see Definitions 4.9 and 6.4, and Figure 4) of (large) constant width, and height equal to that of V. We emphasize that the parameters n_1, n_2 will be chosen (depending on the model) so that the probabilities p_1 and p_2 of the events G_1 and G_2 (respectively) satisfy the key condition

$$\lim_{q \to 0} \max \Big\{ p_2, \big(1 - p_1\big)\big(\log p_2\big)^2 \Big\} = 0$$

that appears in part (B) of Theorem 3.1.

4.2. **Spreading of infection: the supercritical case.** We are now almost ready to state the property of $\mathcal{U}$-bootstrap percolation (proved by Bollobás, Smith and Uzzell [8]) that we will need when $\mathcal{U}$ is supercritical, i.e., when there exists an open semicircle $C \subset S^1$ that is free of stable directions. If $\mathcal{U}$ is rooted, then we may choose $-v$ (in the construction of the rectangle R and of the partition $\{V_i\}_{i \in \mathbb{Z}^2}$ described in Section 4.1) to be the midpoint of any such semicircle; if $\mathcal{U}$ is unrooted, on the other hand, then C can be chosen in such a way that $-C$ also has no stable directions, and we can choose v to be the midpoint of any such semicircle.

Recall that $[V_i]_{\mathcal{U}}$ denotes the closure of $V_i = R_i \cap \mathbb{Z}^2$ under the $\mathcal{U}$-bootstrap process. The following result, proved in [8], states that a large enough rectangle can infect the rectangle to its "left" (i.e., in direction $-v$) under the $\mathcal{U}$-bootstrap process, and if $\mathcal{U}$ is unrooted then it can also infect the rectangle to its "right" (i.e., in direction v).

Proposition 4.1. *Let $\mathcal{U}$ be a supercritical two-dimensional update family. If n_1 and n_2 are sufficiently large, then the following hold:*

- (i) *If $\mathcal{U}$ is unrooted, then $V_{(-1,0)} \cup V_{(1,0)} \subset [V_{(0,0)}]_{\mathcal{U}}$.*
- (ii) *If $\mathcal{U}$ is rooted, then $V_{(-1,0)} \subset [V_{(0,0)}]_{\mathcal{U}}$.*

Remark 4.2. By definition, in the rooted case the semicircle $-C$ contains some stable directions. Thus, $V_{(1,0)} \not\subset [V_{(0,0)}]_{\mathcal{U}}$.

The proof of Proposition 4.1 in [8] is non-trivial, and required some important innovations, most notably the notion of "quasi-stable directions" (see Definition 4.5, below).

We will therefore give here only a brief sketch, explaining how one can read the claimed inclusions out of the results of [8]

Sketch proof of Proposition 4.1. Both parts of the proposition are essentially immediate consequences of the following claim: if R is a sufficiently large rectangle with two sides parallel to $w \in S^1$, and the semicircle centred at w is entirely unstable, then $[R]_{\mathcal{U}}$ contains *every* element of $\mathbb{Z}^2$ that can be reached from R by travelling in direction w. This claim follows from [8, Lemma 5.5], since in this setting all of the quasi-stable directions in $\mathcal{S}'_U$ (see [8, Section 5.3]) are unstable (since they are contained in the semicircle centred at w), and if u is unstable then the empty set is a u-block (see [8, Definition 5.1]). We refer the reader to [8, Sections 5 and 7] for more details. □

4.3. **Spreading of infection: the critical case.** We next turn to the more complicated task of precisely defining the good and super-good events for critical update families. In this subsection we will lay the groundwork for the precise definitions of these events (which we defer until Section 6, see Definition 6.4) by recalling some definitions from [5, 8], and introducing the key new objects needed for the proof of Theorem 2, which we call "quasi-stable half-rings" (see Definition 4.9 and Figure 4, below). Throughout this subsection, we will assume that $\mathcal{U}$ is a critical update family with difficulty $\alpha \in [1, \infty)$ and bilateral difficulty $\beta \in [\alpha, \infty]$ (see Definition 2.5). Recall that we say that $\mathcal{U}$ is α-rooted if $\beta \geqslant 2\alpha$, and that $\mathcal{U}$ is β-unrooted otherwise.

We begin by noting an important property of the set of stable directions $\mathcal{S}(\mathcal{U})$.

Lemma 4.3. *If $\beta < \infty$ then $\mathcal{S}(\mathcal{U})$ consists of a finite number of isolated, rational directions. Moreover, if $\mathcal{U}$ is β-unrooted and $\alpha(u^*) = \max\big\{\alpha(u) : u \in \mathcal{S}(\mathcal{U})\big\}$, then $\alpha(u) \leqslant \beta$ for every $u \in \mathcal{S}(\mathcal{U}) \setminus \{u^*, -u^*\}$.*

Proof. By [8, Theorem 1.10], $\mathcal{S}(\mathcal{U})$ is a finite union of rational closed intervals of S^1, and by [8, Lemma 5.2] (see also [5, Lemma 2.7]), if $u \in \mathcal{S}(\mathcal{U})$ is a rational direction, then $\alpha(u) < \infty$ if and only if u is an isolated point of $\mathcal{S}(\mathcal{U})$. Thus, if one of the intervals in $\mathcal{S}(\mathcal{U})$ is not an isolated point, then there exist two non-opposite stable directions in S^1, each with infinite difficulty, and so $\beta = \infty$.

Now, suppose that $\mathcal{U}$ is β-unrooted, and that $u \in \mathcal{S}(\mathcal{U})$ satisfies $\alpha(u) > \beta$ and $u \notin \{u^*, -u^*\}$. Then u and u^* are non-opposite stable directions in S^1, each with difficulty strictly greater than β, which contradicts the definition of β. □

In particular, if $\mathcal{U}$ is β-unrooted then Lemma 4.3 guarantees the existence of an open semicircle C such that $(C \cup -C) \cap \mathcal{S}(\mathcal{U})$ consists of finitely many directions, each with difficulty at most β. The next lemma provides a corresponding property for α-rooted models.

Lemma 4.4. *If $\mathcal{U}$ is α-rooted, then there exists an open semicircle C such that $C \cap \mathcal{S}(\mathcal{U})$ consists of finitely many directions, each with difficulty at most α.*

Proof. By Definition 2.5, there exists an open semicircle C such that each $u \in C$ has difficulty at most α. Since $\mathcal{U}$ is critical (and hence α is finite), it follows from [8, Lemma 5.2] (cf. the proof of Lemma 4.3) that each $u \in C$ is either unstable, or an isolated element of $\mathcal{S}(\mathcal{U})$, and hence $C \cap \mathcal{S}(\mathcal{U})$ is finite, as claimed. □

Let us fix (for the rest of the subsection) an open semicircle C, containing finitely many stable directions, and such that the following holds:

- if $\mathcal{U}$ is α-rooted then $\alpha(v) \leqslant \alpha$ for each $v \in C$;
- if $\mathcal{U}$ is β-unrooted then $\alpha(v) \leqslant \beta$ for each $v \in C \cup -C$.

Let us also choose C such that its mid-point u belongs to $\mathbb{Q}_1$, and denote by $\pm u^\perp$ the boundary points of C. When drawing pictures we will always think of C as the semicircle $(-\pi/2, \pi/2)$, though we emphasize that we do not assume that u is parallel to one of the axes of $\mathbb{Z}^2$. We remark that the values of $\alpha(u^\perp)$ and $\alpha(-u^\perp)$ will not be important: we will only need to use the fact that they are both finite.

We are now ready to define one of the key notions from [8], the set of quasi-stable directions. These are directions that are not (necessarily) stable, but which nevertheless it is useful to treat as if they were. For any $v \in S^1$, let us write $\hat{v}$ for the direction in S^1 that is symmetric to v w.r.t. the mid-point u of C.

Definition 4.5 (Quasi-stable directions)**.** We say that a direction $v \in \mathbb{Q}_1$ is quasi-stable if either v or $\hat{v}$ is a member of the set

$$\{u\} \cup \mathcal{S}(\mathcal{U}) \cup \bigg(\bigcup_{X \in \mathcal{U}} \bigcup_{x \in X} \{v \in S^1 : \langle v, x \rangle = 0\} \bigg).$$

Observe that there are only finitely many quasi-stable directions in C (and, if $\beta < \infty$, only finitely many in S^1). The key property of the family of quasi-stable directions is given by the following lemma, which allows us to empty the sites near the corners of "quasi-stable half-rings" (see Definition 4.9, below). Recall that we write ℓ_v for the discrete line $\{x \in \mathbb{Z}^2 : \langle x, v \rangle = 0\}$.

Lemma 4.6 ([8, Lemma 5.3])**.** *For every pair v, v' of consecutive quasi-stable directions there exists an update rule X such that $X \subset (\mathbb{H}_v \cup \ell_v) \cap (\mathbb{H}_{v'} \cup \ell_{v'})$.*

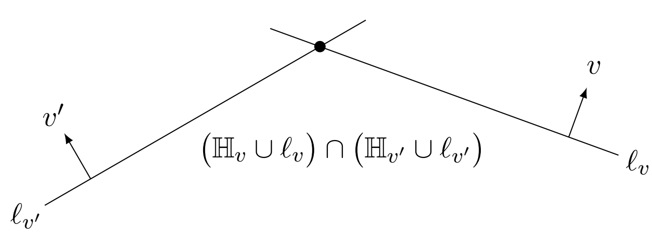

Proof. The statement was proved in [8] (see also [5, Lemma 3.5]) for the family $\mathcal{S}(\mathcal{U}) \cup \big(\bigcup_{X \in \mathcal{U}} \bigcup_{x \in X} \{v \in S^1 : \langle v, x\rangle = 0\}\big)$ of quasi-stable directions, and it therefore holds for any superset of this family. □

In order to define quasi-stable half-rings, we first need to introduce some additional notation:

Definition 4.7. Let $v \in \mathbb{Q}_1$ with $\alpha(v) \leqslant \alpha$. A *$v$-strip* S is any closed parallelogram in $\mathbb{R}^2$ with long sides perpendicular to v and short sides perpendicular to $u^\perp$.

- The $+$-boundary and $-$-boundary of S, denoted $\partial_+(S)$ and $\partial_-(S)$ respectively, are the sides of S with outer normal v and $-v$.
- The external boundary $\partial^{\text{ext}}(S)$ is defined as that translate of $\partial_+(S)$ in the v-direction which captures for the first time a new lattice point not already present in S.
- Given $\lambda > 0$, we define $\partial^{\text{ext}}_\lambda(S)$ as the portion of $\partial^{\text{ext}}(S)$ at distance λ from its endpoints (see Figure 3).

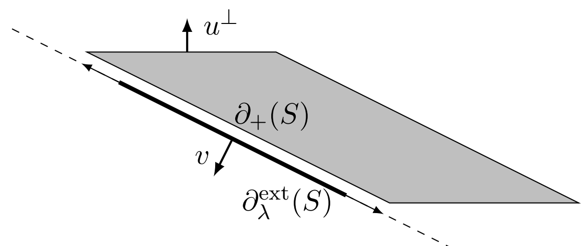

FIGURE 3. A v-strip S, the $+$-boundary of S, the external boundary (solid segment), and its subset $\partial^{\text{ext}}_\lambda(S)$ (thick solid segment)

If v is a stable direction, then a v-strip needs some "help" from other infected sites in order to infect its external boundary (in the $\mathcal{U}$-bootstrap process). Our next ingredient (also first proved in [8]) provides us with a set that suffices for this purpose.

Let v be a quasi-stable direction with difficulty $\alpha(v) \leqslant \alpha$, and let $Z_v \subset \mathbb{Z}^2$ be a set of cardinality α such that $[\mathbb{H}_v \cup Z_v]_{\mathcal{U}} \cap \ell_v$ is infinite. (In the language of [5], Z_v is called a *voracious set*.) The following lemma (see [8, Lemma 5.5] and [5, Lemma 3.4]) states that if S is a sufficiently large v-strip, then a bounded number of translates of Z_v, together with $S \cap \mathbb{Z}^2$, are sufficient to infect $\partial^{\text{ext}}_\lambda(S)$ for some $\lambda = O(1)$.

Lemma 4.8. *There exist $\lambda_v > 0$, $T_v = \{a_1, \dots, a_r\} \subset \ell_v$ and $b \in \ell_v$ such that the following holds. If S is a sufficiently large v-strip such that $\partial^{\text{ext}}(S) \cap \mathbb{Z}^2 \subset \ell_v$, then*

$$\partial^{\text{ext}}_{\lambda_v}(S) \cap \mathbb{Z}^2 \subset \big[(S \cap \mathbb{Z}^2) \cup (Z_v + a_1 + k_1 b) \cup \cdots \cup (Z_v + a_r + k_r b)\big]_{\mathcal{U}} \tag{4.2}$$

for every $k_1, \dots, k_r \in \mathbb{Z}$ such that $a_i + k_i b \in \partial^{\text{ext}}_{\lambda_v}(S)$ for every $i \in [r]$.

Let us fix, for each quasi-stable direction $v \in C$, a constant $\lambda_v > 0$, a set $T_v = \{a_1, \ldots, a_r\} \subset \ell_v$ and a site $b \in \ell_v$ given by Lemma 4.8. If S is a sufficiently large v-strip such that $\partial^{\mathrm{ext}}(S) \cap \mathbb{Z}^2 \subset \ell_v + x$ for some $x \in \mathbb{Z}^2$, then we will refer to any set of the form

$$\big((Z_v + a_1 + k_1 b) \cup \cdots \cup (Z_v + a_r + k_r b)\big) + x, \tag{4.3}$$

with $a_i + k_i b + x \in \partial^{\mathrm{ext}}_{\lambda_v}(S)$ for every $i \in [r]$, as a *helping set* for S.

We are finally ready to define the key objects we will use to control the movement of empty sites in a critical KCM, the *quasi-stable half-rings*. These are non-self-intersecting polygons, obtained by patching together suitable v-strips corresponding to quasi-stable directions (see Figure 4). Recall from Definition 4.5 that, by construction, the set of quasi-stable directions in C is symmetric w.r.t. the midpoint u of C.

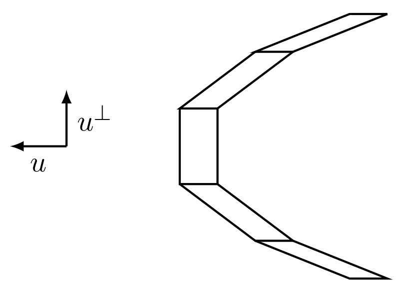

FIGURE 4. A quasi-stable half-ring.

Definition 4.9 (Quasi-stable half-rings). Let $(v_1, \ldots, v_m)$ be the quasi-stable directions in C, ordered in such a way that v_i and v_{i+1} are consecutive directions for any $i \in [m-1]$, and v_{i-1} comes before v_i in clockwise order. Let S_{v_i} be a v_i-strip with length ℓ_i and width w_i. We say that $\mathcal{R} := \bigcup_{i=1}^m S_{v_i}$ is a *quasi-stable half-ring* of width w and length ℓ if the following holds:

(i) $w_i = w$ and $\ell_i = \ell$ for each $i \in [m]$;
(ii) $S_{v_i} \cap S_{v_j} = \emptyset$, unless v_i and v_j are consecutive directions, in which case the two strips share exactly one of their short sides and no other point.

We can finally formulate the "spreading of infection" result that we will need later. Given a quasi-stable half-ring $\mathcal{R}$, we will write $\mathcal{R}^*$ for the quasi-stable half-ring $\mathcal{R} + su$, where $s > 0$ is minimal such that $\big(\mathcal{R}^* \setminus \mathcal{R}\big) \cap \mathbb{Z}^2$ is non-empty. Also, for any set $U \subset \mathbb{Z}^2$, let us write $[A]^U_{\mathcal{U}}$ for the closure of A under the $\mathcal{U}$-bootstrap process restricted to U.

Proposition 4.10. *There exists a constant $\lambda = \lambda(\mathcal{U}) > 0$ such that following holds. Let $\mathcal{R}$ be a quasi-stable half-ring of width w and length ℓ, where $w, \ell \geqslant \lambda$. Let U be the set of points of $\mathbb{Z}^2$ within distance λ of $\mathcal{R} \cup \mathcal{R}^*$, and let Z_i be a helping set for S_{v_i} for each $i \in [m]$. Then*

$$\mathcal{R}^* \cap \mathbb{Z}^2 \subset \big[(\mathcal{R} \cap \mathbb{Z}^2) \cup Z_1 \cup \cdots \cup Z_m\big]^U_{\mathcal{U}}.$$

Proof. This is a straightforward consequence of Lemmas 4.6 and 4.8. To see this, note first that, by Lemma 4.8, the closure of $\big(\mathcal{R} \cap \mathbb{Z}^2\big) \cup Z_1 \cup \cdots \cup Z_m$ under the $\mathcal{U}$-bootstrap process contains all points of $\mathcal{R}^* \cap \mathbb{Z}^2$ except possibly those that lie within distance $O(1)$ of a corner of $\mathcal{R}$. Moreover, the path of infection described in the proof of Lemma 4.8 in [5, 8] only uses sites within distance $O(1)$ of the v-strip S. Thus, if λ is chosen large enough, we have $\partial^{\text{ext}}_{\lambda/4}(S_{v_i}) \cap \mathbb{Z}^2 \subset \big[\big(\mathcal{R} \cap \mathbb{Z}^2\big) \cup Z_i\big]_{\mathcal{U}}^U$ for each $i \in [m]$.

Now, by Lemma 4.6, it follows that the set $\big[\big(\mathcal{R} \cap \mathbb{Z}^2\big) \cup Z_i \cup Z_{i+1}\big]_{\mathcal{U}}^U$ contains the remaining sites of $\partial^{\text{ext}}(S_{v_i}) \cap \mathbb{Z}^2$ and $\partial^{\text{ext}}(S_{v_{i+1}}) \cap \mathbb{Z}^2$ that lie within distance $\lambda/4$ of the intersection of S_{v_i} and $S_{v_{i+1}}$. Indeed, these sites can be infected one by one, working towards the corner, using sites in $\mathcal{R} \cup \partial^{\text{ext}}_{\lambda/4}(S_{v_i}) \cup \partial^{\text{ext}}_{\lambda/4}(S_{v_{i+1}})$. Since this holds for each $i \in [m-1]$, it follows that the whole of $\mathcal{R}^* \cap \mathbb{Z}^2$ is infected, as claimed. □

Given a quasi-stable half-ring $\mathcal{R}$ of width w, we will write $\mathcal{R}'$ for the quasi-stable half-ring $\mathcal{R} + wu$, i.e., the minimal translate of $\mathcal{R}$ in the u-direction such that $\mathcal{R} \cap \mathbb{Z}^2$ and $\mathcal{R}' \cap \mathbb{Z}^2$ are disjoint.

Corollary 4.11. *There exists a constant $\lambda = \lambda(\mathcal{U}) > 0$ such that following holds. Let $\mathcal{R}$ be a quasi-stable half-ring of width w and length ℓ, and suppose that $w \geqslant \lambda$ and $\ell \geqslant \lambda$. Let U be the set of points of $\mathbb{Z}^2$ within distance λ of $\mathcal{R} \cup \mathcal{R}'$, and let $A \subset U$ be such that for any quasi-stable direction v, and any v-strip S_v such that $\partial^{\text{ext}}(S_v) \cap \mathcal{R}'$ has length at least ℓ, there exists a helping set for S_v in A. Then*

$$\mathcal{R}' \cap \mathbb{Z}^2 \subset \big[\big(\mathcal{R} \cap \mathbb{Z}^2\big) \cup A\big]_{\mathcal{U}}^U.$$

Proof. By construction, each v_i-strip of $\mathcal{R}$ has a helping set in $\mathcal{R}'$. Therefore, by Proposition 4.10, the $\mathcal{U}$-bootstrap process restricted to U is able to infect the quasi-stable half-ring $\mathcal{R}^*$. We then repeat with $\mathcal{R}$ replaced by $\mathcal{R}^*$, and so on, until the entire quasi-stable half-ring $\mathcal{R}'$ has been infected. □

Observe that, under the additional assumption that each quasi-stable direction v has a helping set contained in ℓ_v, we may choose A to be a subset of $\mathcal{R}'$, but that in general we may (at some stage) need a helping set not contained in $\mathcal{R}'$ in order to advance in the u-direction.

Remark 4.12. Later on, we will also need the above results in the slightly different setting in which the first v_1-strip entering in the definition of $\mathcal{R}$ is longer than the others, while all of the other v_j-strips, $j \neq 1$ have the same length. In this case we will refer to $\mathcal{R}$ as an *elongated* quasi-stable half-ring. For simplicity we preferred to state Proposition 4.10 in the slightly less general setting above, but exactly the same proof applies if $\mathcal{R}$ is an elongated quasi-stable half-ring.

5. Supercritical KCM: proof of Theorem 1

In this section we shall prove Theorem 1, which gives a sharp (up to a constant factor in the exponent) upper bound on the mean infection time for a supercritical KCM. We will first (in Section 5.1) give a detailed proof in the case that $\mathcal{U}$ is unrooted, and then (in Section 5.2) explain briefly how the proof can be modified to prove the claimed bound for rooted models.

5.1. **The unrooted case.**

Let $\mathcal{U}$ be a supercritical, unrooted, two-dimensional update family; we are required to show that there exists a constant $\lambda = \lambda(\mathcal{U})$ such that

$$\mathbb{E}_\mu(\tau_0) \leqslant q^{-\lambda}$$

for all sufficiently small $q > 0$. To do so, recall first from (2.8) that $\mathbb{E}_\mu(\tau_0) \leqslant T_{\mathrm{rel}}(q,\mathcal{U})/q$, and therefore, by Definition 2.9, it will suffice to prove that

$$\mathrm{Var}(f) \leqslant q^{-\lambda} \sum_x \mu\big(c_x \, \mathrm{Var}_x(f)\big) \tag{5.1}$$

for some $\lambda = \lambda(\mathcal{U}) > 0$ and all local functions f, where c_x denotes the kinetic constraint for the KCM, i.e., c_x is the indicator function of the event that there exists an update rule $X \in \mathcal{U}$ such that $\omega_y = 0$ for each $y \in X + x$. We will deduce a bound of the form (5.1) from Theorem 3.1 and Proposition 4.1.

Recall the construction and notation described in Sections 4.1 and 4.2; in particular, recall the definitions of the blocks V_i, of the parameters n_1 and n_2 (which determine the side lengths of the basic rectangle R), and the choice of v as the midpoint of an open semicircle $C \subset S^1$ such that the set $C \cup -C$ contains no stable directions. As anticipated in Section 4.1, the choice of the good and super-good events $G_2 \subset G_1 \subseteq S$ entering in Theorem 3.1, is, in this case, extremely simple.

Definition 5.1. If $\mathcal{U}$ is a supercritical two-dimensional update family, then:
(a) every block V_i satisfies the *good event* G_1 for $\mathcal{U}$ (*i.e.,* $G_1 = S$);
(b) a block V_i satisfies the *super-good event* G_2 for $\mathcal{U}$ if and only if it is empty.

Let us fix the parameters n_1 and n_2 to be $O(1)$, but sufficiently large so that Proposition 4.1 holds. It follows that if $V_{(0,0)}$ is super-good, then the blocks $V_{(-1,0)}$ and $V_{(1,0)}$ (its nearest neighbours to the left and right respectively) lie in the closure under the $\mathcal{U}$-bootstrap process of the empty sites in $V_{(0,0)}$. In particular,

$$t^{\pm} = \min\big\{ t > 0 \,:\, A_t \supseteq V_{(\pm 1,0)} \big\},$$

are both finite, where A_t is the set of sites infected after t steps of the $\mathcal{U}$-bootstrap process, starting from $A_0 = V_{(0,0)}$ (see Definition 2.1). With foresight, define

$$\Lambda := \big(A_{t^-} \setminus V_{(0,0)}\big) + n_1 \vec{v}, \tag{5.2}$$

and note that $\Lambda \cap V_{\vec{e}_1} = \emptyset$ and $V_{(0,0)} \subset \Lambda$.

Proof of part (a) *of Theorem 1.* The first step is to apply Theorem 3.1 to the probability space $(S^{\mathbb{Z}^2}, \mu)$ described in Section 4.1, in which each 'block' variable $\omega_i \in S$ is given by the collection $\{\omega_x\}_{x \in V_i} \in \{0,1\}^{V_i}$ of i.i.d. Bernoulli(p) variables. Recall that $p_1 = \hat{\mu}(G_1)$ and $p_2 = \hat{\mu}(G_2)$ are the probabilities of the good and super-good events, respectively, and note that, in our setting, $p_1 = 1$ and $p_2 \geqslant q^{O(n_1 n_2)} = q^{O(1)}$. It follows, using (3.2), that

$$\mathrm{Var}(f) \leqslant \frac{1}{q^{O(1)}} \sum_{i \in \mathbb{Z}^2} \mu\left(\mathbb{1}_{\{\text{either } V_{i+\vec{e}_1} \text{ or } V_{i-\vec{e}_1} \text{ is empty}\}} \mathrm{Var}_{V_i}(f) \right) \tag{5.3}$$

for all local functions f, where $\mathrm{Var}_{V_i}(f)$ denotes the variance with respect to the variables $\{\omega_x\}_{x \in V_i}$, given all the other variables $\{\omega_y\}_{y \in \mathbb{Z}^2 \setminus V_i}$.

To deduce (5.1), it will suffice (by symmetry) to prove an upper bound on the right-hand side of (5.3) of the form

$$\mu\left(\mathbb{1}_{\{V_{\vec{e}_1} \text{ is empty}\}} \mathrm{Var}_{V_{(0,0)}}(f) \right) \leqslant \frac{1}{q^{O(1)}} \sum_{x \in \Lambda \cup V_{\vec{e}_1}} \mu\big(c_x \, \mathrm{Var}_x(f) \big) \tag{5.4}$$

for the set Λ defined in (5.2), since the elements of $\Lambda \cup V_{\vec{e}_1}$ are all within distance $O(1)$ from the origin, and so we may then simply sum over all $i \in \mathbb{Z}^2$.

To prove (5.4), the first step is to observe that, by the convexity of the variance, and recalling that $\Lambda \cap V_{\vec{e}_1} = \emptyset$ and $V_{(0,0)} \subset \Lambda$, we have

$$\mu\Big(\mathbb{1}_{\{V_{\vec{e}_1} \text{ is empty}\}} \mathrm{Var}_{V_{(0,0)}}(f) \Big) \leqslant \mu\left(\mathbb{1}_{\{V_{\vec{e}_1} \text{ is empty}\}} \mathrm{Var}_{\Lambda}(f) \right). \tag{5.5}$$

To conclude we appeal to the following result which, for later purposes, we formulate in a slightly more general setting than is needed here. In what follows, for any $\omega \in \Omega$ and $U \subset \mathbb{Z}^2$, we shall write $[\omega]_{\mathcal{U}}^U$ for the closure of the set $\big\{ x \in \mathbb{Z}^2 : \, \omega_x = 0 \big\}$ under the $\mathcal{U}$-bootstrap process restricted to U.

Lemma 5.2. *Let $A, B \subset \mathbb{Z}^2$ be disjoint sets, and let $\mathcal{E}$ be an event depending only on ω_B. Suppose that there exists a set $U \supset A \cup B$ such that $B \subset [\omega]_{\mathcal{U}}^U$ for any $\omega \in \{0,1\}^U$ for which A is empty and $\omega_B \in \mathcal{E}$. Then*

$$\mu\Big(\mathbb{1}_{\{A \text{ is empty}\}} \mathrm{Var}_B\left(f \,|\, \mathcal{E} \right) \Big) \leqslant |U| q^{-|U|} \frac{2}{pq} \sum_{x \in U} \mu\big(c_x \, \mathrm{Var}_x(f) \big) \tag{5.6}$$

for any local function f.

Before proving the lemma we conclude the proof of part (a) of Theorem 1. We apply the lemma with $A = V_{\vec{e}_1}$, $B = \Lambda$, $U = A \cup B$ and $\mathcal{E}$ the trivial event, i.e., $\mathcal{E} = \Omega_B$. Indeed, by construction (see (5.2)), $\Lambda \subset \big[V_{\vec{e}_1} \big]_{\mathcal{U}}^U$. Thus (5.6) becomes

$$\mu\left(\mathbb{1}_{\{V_{\vec{e}_1} \text{ is empty}\}} \mathrm{Var}_{\Lambda}(f) \right) \leqslant |U| q^{-|U|} \frac{2}{pq} \sum_{x \in U} \mu\big(c_x \, \mathrm{Var}_x(f) \big). \tag{5.7}$$

Since $|U| = O(1)$, and using (5.5), we conclude that for all $i \in \mathbb{Z}^2$,

$$\mu\left(\mathbb{1}_{\{V_{i+\vec{e}_1} \text{ is empty}\}} \operatorname{Var}_{V_i}(f)\right) \leqslant \frac{1}{q^{O(1)}} \sum_{x \in U_i} \mu\big(c_x \operatorname{Var}_x(f)\big),$$

where U_i is the analogue of U for the block V_i.

As noted above, summing over $i \in \mathbb{Z}^2$ and using (5.3), we obtain the Poincaré inequality (5.1) with constant $q^{-O(1)}$, and by (2.8) and Definition 2.9 it follows that there exists a constant $\lambda = \lambda(\mathcal{U})$ such that

$$\mathbb{E}_\mu(\tau_0) \leqslant \frac{T_{\text{rel}}(q, \mathcal{U})}{q} \leqslant q^{-\lambda},$$

for all sufficiently small $q > 0$, as required. Since the bootstrap infection time $T_{\mathcal{U}}$ of a supercritical update family satisfies $T_{\mathcal{U}} = q^{-\Theta(1)}$, it also follows that $\mathbb{E}_\mu(\tau_0) \leqslant T_{\mathcal{U}}^{O(1)}$. □

Proof of Lemma 5.2. Observe first that, for any $\omega \in \Omega$,

$$\operatorname{Var}_B\big(f \,|\, \mathcal{E}\big)(\omega_{\mathbb{Z}^2 \setminus B}) \leqslant \frac{1}{\mu_B(\mathcal{E})} \sum_{\eta_B \in \mathcal{E}} \mu_B(\eta_B)\Big(f\big(\eta_B, \omega_{\mathbb{Z}^2 \setminus B}\big) - f\big(0, \omega_{\mathbb{Z}^2 \setminus B}\big)\Big)^2, \qquad (5.8)$$

since $\mathbb{E}\big[(X-a)^2\big]$ is minimized by taking $a = \mathbb{E}[X]$, where $\big(0, \omega_{\mathbb{Z}^2 \setminus B}\big)$ denotes the configuration that is equal to $\omega_{\mathbb{Z}^2 \setminus B}$ outside B, and empty inside B.

We will break each term on the right-hand side of (5.8) into the sum of single spin-flips using the $\mathcal{U}$-bootstrap process as follows. Fix $\omega \in \Omega$ such that A is empty, and fix $\eta_B \in \mathcal{E}$. Using the assumption of the lemma, we claim that there exists a path $\gamma \equiv (\omega^{(0)}, \ldots, \omega^{(m)})$ in Ω such that:

(i) $\omega^{(0)} = (\eta_B, \omega_{\mathbb{Z}^2 \setminus B})$ and $\omega^{(m)} = (0, \omega_{\mathbb{Z}^2 \setminus B})$;
(ii) the length m of γ satisfies $m \leqslant 2|U|$;
(iii) for each $k = 1, \ldots, m$, there exists a vertex $x^{(k)} \in U$ such that
- the configuration $\omega^{(k)}$ is obtained from $\omega^{(k-1)}$ by flipping the value at $x^{(k)}$;
- this flip is legal, i.e., $c_{x^{(k)}}\big(\omega^{(k-1)}\big) = 1$.

We construct γ in two steps: first we empty all of B, and possibly some of $U \setminus B$; then we reconstruct $\omega_{\mathbb{Z}^2 \setminus B}$ without changing ω_B. To spell out the details, observe first that, since $B \subset \big[\big(\eta_B, \omega_{U \setminus B}\big)\big]_{\mathcal{U}}^{U}$, there exists a sequence of legal flips in U connecting $\big(\eta_B, \omega_{\mathbb{Z}^2 \setminus B}\big)$ to a configuration with $A \cup B$ empty. By choosing a minimal such sequence, we may assume that all of the flips are from occupied to empty, and therefore that this first part of the path has length at most $|U|$.

Now, to reconstruct $\omega_{\mathbb{Z}^2 \setminus B}$, we simply run the same sequence backwards, except without performing the steps inside B. Note that all of these flips are legal, since skipping the steps inside B only creates additional empty sites, and that this second part of the path also has length at most $|U|$, as required.

It follows, using Cauchy–Schwarz, that

$$\Big(f\big(\eta_B,\omega_{\mathbb{Z}^2\setminus B}\big)-f\big(0,\omega_{\mathbb{Z}^2\setminus B}\big)\Big)^2 \leqslant m\sum_{k=1}^{m} c_{x^{(k)}}\big(\omega^{(k-1)}\big)\Big(f\big(\omega^{(k)}\big)-f\big(\omega^{(k-1)}\big)\Big)^2$$
$$\leqslant 2|U|\frac{1}{\mu_*}\frac{1}{pq}\sum_{x\in U}\sum_{\eta\in\{0,1\}^U}\mu_U(\eta)c_x(\eta,\omega_{\mathbb{Z}^2\setminus U})\,pq\Big(f\big(\eta^{(x)},\omega_{\mathbb{Z}^2\setminus U}\big)-f\big(\eta,\omega_{\mathbb{Z}^2\setminus U}\big)\Big)^2,$$

for any ω in which A is empty, and any $\eta_B\in\mathcal{E}$, where $\mu_*=\min_{\eta\in\{0,1\}^U}\mu_U(\eta)=q^{|U|}$, and $\eta^{(x)}$ denotes the configuration obtained from η by flipping the spin at x. Notice that the right-hand side does not depend on η_B, and that $pq\big(f\big(\eta^{(x)},\omega_{\mathbb{Z}^2\setminus U}\big)-f\big(\eta,\omega_{\mathbb{Z}^2\setminus U}\big)\big)^2$ is the local variance $\mathrm{Var}_x(f)$ computed for the configuration $\omega\equiv(\eta,\omega_{\mathbb{Z}^2\setminus U})$.

Hence, using (5.8), we obtain

$$\mathbb{1}_{\{A\text{ is empty}\}}\,\mathrm{Var}_B\,\big(f\,|\,\mathcal{E}\big)(\omega_{\mathbb{Z}^2\setminus B})\leqslant\frac{2|U|q^{-|U|}}{pq}\sum_{x\in U}\mu_U\big(c_x\,\mathrm{Var}_x(f)\big)(\omega_{\mathbb{Z}^2\setminus U})$$

for any $\omega\in\Omega$, and inequality (5.6) follows by averaging using the measure μ. □

5.2. **The rooted case.** Let $\mathcal{U}$ be a supercritical, rooted, two-dimensional update family, let $C\subset S^1$ be a semicircle with no stable directions and recall that, thanks to (2.8), it will suffice to prove a Poincaré inequality (cf. (5.1)) with constant $q^{-O(\log(1/q))}$. To prove this we will follow almost exactly the same route of the unrooted case, with the same definition of the blocks V_i and of the good and super-good events. We will therefore only give a very brief sketch of the proof in this new setting.

The main difference w.r.t. the unrooted case is that now the opposite semicircle $-C$ will necessarily contain some stable directions. This forces us to use the oriented Poincaré inequality (3.1) from Theorem 3.1 instead of the unoriented one (3.2), because in this case (see Proposition 4.1 and Remark 4.2) a super-good block is able to infect the block to its left but not the block to its right, *i.e.,* $V_{(-1,0)}\subset[V_{(0,0)}]_{\mathcal{U}}$ but $V_{(1,0)}\not\subset[V_{(0,0)}]_{\mathcal{U}}$.

Proof of part (b) *of Theorem 1.* We again apply Theorem 3.1 to the probability space $(S^{\mathbb{Z}^2},\mu)$ described in Section 4.1, but we use (3.1) instead of (3.2). Recalling that $p_1=1$ and $p_2=q^{O(1)}$, we obtain

$$\mathrm{Var}(f)\leqslant\frac{1}{q^{O(\log(1/q))}}\sum_{i\in\mathbb{Z}^2}\mu\left(\mathbb{1}_{\{V_{i+\vec{e}_1}\text{ is empty}\}}\,\mathrm{Var}_{V_i}(f)\right)\tag{5.9}$$

for all local functions f. As before, using translation invariance, we only examine the $i=(0,0)$ term in the above sum. We claim that

$$\mu\left(\mathbb{1}_{\{V_{\vec{e}_1}\text{ is empty}\}}\,\mathrm{Var}_{V_{(0,0)}}(f)\right)\leqslant\frac{1}{q^{O(1)}}\sum_{x\in U}\mu\big(c_x\,\mathrm{Var}_x(f)\big)\tag{5.10}$$

for $U = V_{\vec{e}_1} \cup \Lambda$, where Λ is the set defined in (5.2). However, the proof of (5.10) is identical to that of (5.4), since Proposition 4.1 implies that $V_{(0,0)}$ can be entirely infected by $V_{\vec{e}_1}$. We therefore obtain the Poincaré inequality

$$\operatorname{Var}(f) \leqslant \frac{1}{q^{O(\log(1/q))}} \sum_x \mu\big(c_x \operatorname{Var}_x(f)\big) \tag{5.11}$$

for all local functions f. Thus $T_{\mathrm{rel}}(q,\mathcal{U}) \leqslant q^{-O(\log(1/q))}$, and hence

$$\mathbb{E}_\mu(\tau_0) \leqslant \frac{T_{\mathrm{rel}}(q,\mathcal{U})}{q} \leqslant q^{-O(\log(1/q))} = T_{\mathcal{U}}^{O(\log T_{\mathcal{U}})},$$

as required, because $T_{\mathcal{U}} = q^{-\Theta(1)}$. □

6. Critical KCM: proof of Theorem 2 under a simplifying assumption

In this section we shall prove Theorem 2 under the following additional assumption (see below): every stable direction v with finite difficulty has a voracious set that is a subset of the line ℓ_v. By doing so, we avoid some technical complications (mostly related to the geometry of the quasi-stable half-ring) which might obscure the main ideas behind the proof. The changes necessary to treat the general case are spelled out in detail in Section 7.

Assumption 6.1. *For any stable direction $u \in \mathcal{S}$ with finite difficulty $\alpha(u)$, there exists a set $Z_u \subset \ell_u$ of cardinality $\alpha(u)$ such that $\big[\mathbb{H}_u \cup Z_u\big]_{\mathcal{U}} \cap \ell_u$ is infinite.*

As in Section 5, our main task will be to establish a suitable upper bound on the relaxation time $T_{\mathrm{rel}}(\mathcal{U};q)$. In Section 6.1 we will first analyse the α-rooted case and the starting point will be the constrained Poincaré inequality (3.3); the proof the β-unrooted case (see Section 6.2) will be essentially the same, the main difference being that (3.3) will be replaced by (3.4).

6.1. **α-rooted update families.** Let $\mathcal{U}$ be a critical, α-rooted, two-dimensional update family, and recall from Definition 2.12 that $\mathcal{U}$ has difficulty α, and bilateral difficulty at least 2α. The properties of $\mathcal{U}$ that we will need below have already been proved in Section 4.3; they all follow from the fact (see Lemma 4.4) that there exists an open semicircle C such that $C \cap \mathcal{S}(\mathcal{U})$ consists of finitely many directions, each with difficulty at most α. In particular, we will make crucial use of Corollary 4.11.

We will prove that, if Assumption 6.1 holds, then there exists a constant $\lambda = \lambda(\mathcal{U})$ such that

$$\mathbb{E}_\mu(\tau_0) \leqslant \frac{T_{\mathrm{rel}}(q,\mathcal{U})}{q} \leqslant \exp\Big(\lambda \cdot q^{-2\alpha}\big(\log(1/q)\big)^4\Big)$$

for all sufficiently small $q > 0$. Note that the first inequality follows from (2.8), and so, by Definition 2.9, it will suffice to prove that

$$\operatorname{Var}(f) \leqslant \exp\Big(\lambda \cdot q^{-2\alpha}\big(\log(1/q)\big)^4\Big) \sum_{x \in \mathbb{Z}^2} \mu\big(c_x \operatorname{Var}_x(f)\big) \tag{6.1}$$

for some $\lambda = \lambda(\mathcal{U})$ and all local functions f. We will deduce a bound of the form (6.1) starting from (3.3) and using Corollary 4.11.

Remark 6.2. We are not able to use the unoriented constrained Poincaré inequality (3.4) in place of the oriented inequality (3.3) in the proof of (6.1) because there exist α-rooted models (the Duarte model [16] is one such example) with $\beta = \infty$ such that, for any choice of the side-lengths n_1 and n_2 of the blocks V_i, and of the good and super-good events $G_2 \subset G_1$ satisfying the condition $(1-p_1)(\log p_2)^2 = o(1)$, the $\mathcal{U}$-bootstrap process is not guaranteed to be able to infect the block V_i using only that facts that the block $V_{i-\vec{e}_1}$ is infected and that some nearby blocks V_j are good. For update families with $2\alpha < \beta < \infty$, it *is* possible to apply (3.4) for certain choices of (n_1, n_2, G_1, G_2), but the best Poincaré constant we are able to obtain in this way is roughly $\exp\left(q^{-\beta}\right)$, which is much larger than the one we prove using (3.3).

6.1.1. *The geometric setting and the good and super-good events.* Recall the construction and notation described in Sections 4.1 and 4.3; in particular, recall that $V = R \cap \mathbb{Z}^2$, where R is a rectangle in the rotated coordinates $(v, v^\perp)$, and $u = -v$ is the midpoint of an open semicircle $C \subset S^1$ in which every stable direction has difficulty at most α. As in Section 4, when drawing figures we will think of u as pointing to the left. We will choose the parameters n_1 and n_2 (which determine the side-lengths of R) depending on q; to be precise, set

$$n_1 = \lfloor q^{-3\kappa} \rfloor \qquad \text{and} \qquad n_2 = \lfloor \kappa^4 q^{-\alpha} \log(1/q) \rfloor,$$

where $\kappa = \kappa(\mathcal{U})$ is a sufficiently large constant.

In order to define the good and super-good events G_1 and G_2, we need to define some structures which we call *κ-stairs*, which will provide us with a way of transporting infection 'vertically'. Let us call the set of points of V lying on the same line parallel to u (resp. $u^\perp$) a *row* (resp. *column*) of V, and let us order the rows from bottom to top and the columns from left to right. Let a and b be (respectively) the number of rows and columns of V, and observe that, since v is a rational direction, we have $a = \Theta(n_2)$ and $b = \Theta(n_1)$, where the implicit constants depend only on the update family $\mathcal{U}$. We will say that a set of vertices is an *interval of V* if it is the intersection with V of a line segment in $\mathbb{R}^2$. Recall that $\kappa = \kappa(\mathcal{U}) > 0$ was fixed above.

Definition 6.3. We say that a collection $\mathcal{M} = \left\{M^{(1)}, \dots, M^{(a)}\right\}$ of disjoint intervals of V of size 2κ forms an *upward κ-stair* with steps $M^{(1)}, M^{(2)}, \dots$ if:

(i) for each $i \in [a]$, the i^{th}-step $M^{(i)}$ belongs to the i^{th}-row of V;
(ii) the i^{th}-step is "to the left" of the j^{th}-step if $i < j$. More precisely, let $\left(M_\ell^{(i)}, M_r^{(i)}\right)$ be the abscissa (in the $(v, v^\perp)$-frame) of the leftmost and rightmost elements (respectively) of the i^{th}-step. Then $M_r^{(i)} < M_\ell^{(j)}$ whenever $i < j$.

We refer the reader to Figure 6 for a picture of an upward κ-stair.

We are now ready to define the good and super-good events. Let us say that a quasi-stable half-ring $\mathcal{R}$ *fits in* the block V_i if the top and bottom sides of $\mathcal{R}$ are contained in the top and bottom sides of R_i, and note that this determines the length ℓ of $\mathcal{R}$, which moreover satisfies $\ell \geqslant n_2/m$ (see Definition 4.9). Let $(v_1, \dots, v_m)$ be the quasi-stable directions in C (see Definition 4.5), and recall the definitions of a v-strip S_v (see Definition 4.7) and of a helping set Z for S_v (see immediately after Lemma 4.8). Note that Assumption 6.1 implies that, for any $j \in [m]$ and v_j-strip S_{v_j}, we may choose the voracious set Z_{v_j} so that the helping sets for S_{v_j} are subsets of $\partial^{\text{ext}}(S_{v_j})$.

Definition 6.4 (Good and super-good events)**.**

(1) The block $V_i = R_i \cap \mathbb{Z}^2$ satisfies the *good* event G_1 iff:
   (a) for each quasi-stable direction $v \in C$ and every v-strip S such that the length of the segment $\partial^{\text{ext}}(S) \cap R_i$ is at least $n_2/(4m)$, there exists an empty helping set $Z \subset \partial^{\text{ext}}(S) \cap V_i$ for S;
   (b) there exists an empty upward κ-stair within the leftmost quarter of V_i.

(2) The block V_i satisfies the *super-good* event G_2 iff it satisfies the good event G_1, and moreover there exists an empty quasi-stable half-ring $\mathcal{R}$ of width κ, that fits in V_i and is entirely contained in the rightmost quarter of R_i.

Next we prove that the hypothesis for the part (B) of Theorem 3.1 holds in the above setting if κ is sufficiently large.

Lemma 6.5. *Let $p_1 := \hat{\mu}(G_1)$ and $p_2 := \hat{\mu}(G_2)$. There exists a constant $\kappa_0(\mathcal{U}) > 0$ such that, for any $\kappa > \kappa_0(\mathcal{U})$,*

$$\lim_{q \to 0} \big(1 - p_1\big)\big(\log(1/p_2)\big)^2 = 0$$

Proof. First, let's bound the probability that there is no empty helping set $Z \subset V_i$ for a given v-strip S (where v is a quasi-stable direction in C) such that $\partial^{\text{ext}}(S) \cap R_i$ has length at least $n_2/(4m)$. Observe that we can choose $\Omega(n_2)$ potential values for each k_j in (4.3) such that the corresponding sets $Z_v + a_j + k_j b$ are pairwise disjoint subsets of $\partial^{\text{ext}}(S) \cap V_i$ (using Assumption 6.1), and that each such translate of Z_v is empty with probability q^α. (Here the implicit constant depends only on $\mathcal{U}$.) The probability that S has no empty helping set is therefore at most

$$O\Big(\big(1 - q^\alpha\big)^{\Omega(n_2)}\Big) \leqslant q^{\kappa^3}$$

if κ is sufficiently large and $q \ll 1$. There are at most $O(n_1^2 n_2^2)$ choices for the quasi-stable direction $v \in C$ and for the intersection of the v-strip S with V_i. Thus, by the union bound, part (a) of the definition of G_1 holds with probability at least $1 - q^{\kappa^2}$.

To bound the probability of part (b), observe that an interval of V of size $\Theta(n_1/n_2)$ contains an empty interval of V of size 2κ with probability at least

$$1 - \big(1 - q^{2\kappa}\big)^{\Theta(n_1/n_2)} \geqslant 1 - \exp\big(- q^{-\kappa + 2\alpha}\big),$$

if $q \ll 1$, and therefore the probability that V contains an empty upward κ-stair is (by the union bound) at least

$$1 - O(n_2) \exp\big(- q^{-\kappa + 2\alpha} \big) \geqslant 1 - q^{\kappa^2} \tag{6.2}$$

if κ is sufficiently large and $q \ll 1$. It follows that

$$1 - p_1 \, = \, 1 - \hat{\mu}(G_1) \, \leqslant \, 2 \cdot q^{\kappa^2}.$$

Moreover, the probability that there exists an empty quasi-stable half-ring $\mathcal{R}$ of width κ that fits in V_i is at least $q^{O(n_2)}$, so (by the FKG inequality) we have

$$\log(1/p_2) \, \leqslant \, O(n_2) \log(1/q) \leqslant O\Big(q^{-\alpha} \big(\log(1/q) \big)^2 \Big),$$

where the implicit constant is independent of q. It follows that, if κ is sufficiently large, then $\big(1 - p_1\big)\big(\log(1/p_2) \big)^2 \to 0$ as $q \to 0$, as required. □

Let us fix, from now on, the constant κ to be sufficiently large so that Lemma 6.5 applies. In particular, by Theorem 3.1, the constrained Poincaré inequality (3.3) holds for any local function f, i.e.,

$$\begin{aligned} \mathrm{Var}(f) \leqslant \vec{T}(p_2) \bigg(\sum_{i \in \mathbb{Z}^2} \mu \Big(\mathbb{1}_{\{\omega_{i+\vec{e}_2} \in G_2\}} \mathbb{1}_{\{\omega_j \in G_1 \, \forall j \in \mathbb{L}^+(i)\}} \mathrm{Var}_{V_i}(f) \Big) \\ + \sum_{i \in \mathbb{Z}^2} \mu \Big(\mathbb{1}_{\{\omega_{i+\vec{e}_1} \in G_2\}} \mathbb{1}_{\{\omega_{i-\vec{e}_1} \in G_1\}} \mathrm{Var}_{V_i} \big(f \,|\, G_1 \big) \Big) \bigg), \end{aligned} \tag{6.3}$$

with

$$\vec{T}(p_2) \, = \, e^{O(\log(p_2)^2)} \, = \, \exp\Big(O\big(q^{-2\alpha} \log(1/q)^4 \big) \Big).$$

As in the supercritical setting (see Section 5), our strategy will be to bound each of the sums in the r.h.s. of (6.3) from above in terms of the Dirichlet form $\mathcal{D}(f)$ of our KCM. To do so, it will suffice to bound from above, for a fixed (and arbitrary) local function f, the following two generic terms:

$$I_1(i) := \mu \Big(\mathbb{1}_{\{\omega_{i+\vec{e}_1} \in G_2\}} \mathbb{1}_{\{\omega_{i-\vec{e}_1} \in G_1\}} \mathrm{Var}_{V_i} \big(f \,|\, G_1 \big) \Big),$$

and

$$I_2(i) := \mu \Big(\mathbb{1}_{\{\omega_{i+\vec{e}_2} \in G_2\}} \mathbb{1}_{\{\omega_j \in G_1 \, \forall j \in \mathbb{L}^+(i)\}} \mathrm{Var}_{V_i}(f) \Big),$$

see Figure 5. Using translation invariance it suffices to consider only the case $i = (0,0)$, so let us set $I_1 \equiv I_1((0,0)$ and $I_2 \equiv I_2((0,0))$.

Define $W_1 = V_{(0,0)} \cup V_{(-1,0)} \cup V_{(1,0)}$ and $W_2 = V_{(0,0)} \cup V_{(-1,0)} \cup V_{(1,0)} \cup V_{(-1,1)} \cup V_{(0,1)}$. We will prove the following upper bounds on I_1 and I_2.

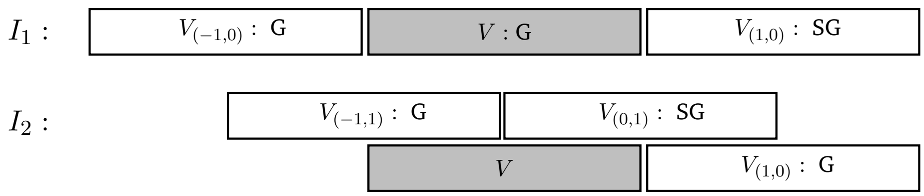

FIGURE 5. In I_1 the block $V \equiv V_{(0,0)}$ is conditioned to be good (G), while the blocks $V_{(-1,0)}$ and $V_{(1,0)}$ are good and super-good (SG) respectively. Recall that $\mathbb{L}^+((0,0)) = \{(1,0),(-1,1)\}$, so in I_2 the blocks $V_{(1,0)}$ and $V_{(-1,1)}$ are good, the block $V_{(0,1)}$ is super-good, and V is unconditioned.

Proposition 6.6. *For each $j \in \{1,2\}$, there exists a $O(1)$-neighbourhood $\hat{W}_j$ of W_j such that*

$$I_j \leqslant \exp\Big(O\big(q^{-\alpha}\log(1/q)^3\big)\Big) \sum_{x \in \hat{W}_j} \mu\big(c_x \operatorname{Var}_x(f)\big).$$

Observe that, combining Proposition 6.6 with (6.3), and noting that $|\hat{W}_j| = q^{-O(1)}$, we immediately obtain a final Poincaré inequality of the form (6.1), i.e.,

$$\operatorname{Var}(f) \leqslant \exp\Big(O\big(q^{-2\alpha}\log(1/q)^4\big)\Big) \sum_{x \in \mathbb{Z}^2} \mu\big(c_x \operatorname{Var}_x(f)\big),$$

as required. It will therefore suffice to prove Proposition 6.6.

6.1.2. *The core of the proof of Proposition 6.6.* Before giving the full technical details of the proof of the proposition, we first explain the high-level idea we wish to exploit. Fix $j \in \{1,2\}$, set $W := W_j$, and fix $\omega \in \Omega$ such that the restriction of ω to W satisfies the requirement of the good and super-good environment of the blocks (see Figure 5). The key idea is to cover the block $V = V_{(0,0)}$ with a collection of pairwise disjoint "fibers" $\hat{F}_1, \dots, \hat{F}_{N+1}$, each of which is a quasi-stable half-ring, for some $N \leqslant |V|$ depending on ω. For each fiber $\hat{F}_i$, the set $F_i := \hat{F}_i \cap \mathbb{Z}^2$ is a subset of W of cardinality $O(n_2)$ with the following key properties (which we will define precisely later):

(a) the fiber F_{N+1} is empty;
(b) in each fiber F_i a certain "helping" event H_i, depending only on the restriction of ω to F_i, and implied by our assumption on the goodness[7] of the blocks in W, occurs;
(c) the helping event H_i has the following property: the $\mathcal{U}$-bootstrap process restricted to a $O(1)$-neighbourhood of the set $F_i \cup F_{i+1}$ is able to infect F_i for any ω such that F_{i+1} is empty and H_i occurs.

[7] It is worth emphasizing here that H_i only requires the blocks to be good, rather than super-good, and therefore holds with high probability.

To be concrete, let us consider the term I_1. In this case we will take F_{N+1} to be $\mathcal{R} \cap \mathbb{Z}^2$, where $\mathcal{R}$ is the rightmost empty quasi-stable half-ring of width κ that fits in $V_{(1,0)}$, which exists by our assumption that $V_{(1,0)}$ is a super-good block. The other fibers $F_1, \dots, F_N$ will be suitable disjoint translates of F_{N+1} in the u-direction, satisfying $V \subset \Lambda := \bigcup_{i=1}^N F_i$. The helping event H_i will require the presence in F_i of suitable helping sets for each quasi-stable direction; we remark that the key requirement that H_i depends only on the restriction of ω to F_i is a consequence of Assumption 6.1. Finally, the third condition (c) above will follow from Corollary 4.11. A similar construction will be used for the term I_2, but the fibers will be slightly more complicated, see Figure 6.

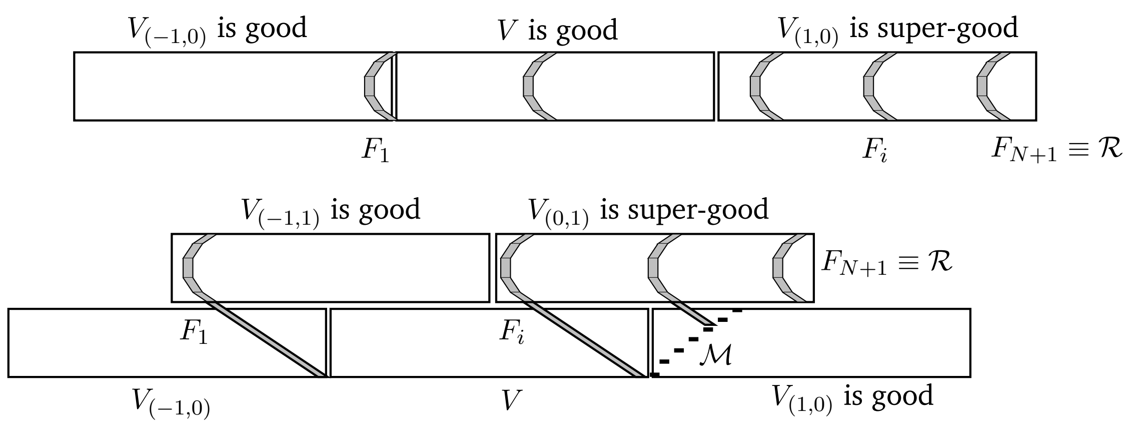

FIGURE 6. The top picture shows the local neighbourhood W_1 of the block $V = V_{(0,0)}$; in this case the fibers are simply the disjoint translates of the rightmost empty quasi-stable half-ring $\mathcal{R}$ in the last quarter of $V_{(1,0)}$. The bottom picture shows the local neighbourhood W_2; in this case the fibers are not all equal: they grow as they 'descend' the steps of the upward κ-stair $\mathcal{M}$ (the little horizontal intervals). Each fiber becomes an elongated version of the rightmost empty half-ring $\mathcal{R}$.

Let us write ν_i for the Bernoulli(p) product measure on $S_i = \{0,1\}^{F_i}$ conditioned on the event H_i. The main step in the proof is the following bound on I_j for $j \in \{1,2\}$:

$$I_j \leqslant \frac{1}{p_1} \cdot \mu\Big(\mathbb{1}_{\{F_{N+1} \text{ is empty}\}} \operatorname{Var}_\nu(f)\Big), \tag{6.4}$$

where $\operatorname{Var}_\nu(\cdot)$ is the variance computed w.r.t. the product measure $\nu = \otimes_{i=1}^N \nu_i$. Before proving (6.4) (see Section 6.1.3, below), let us show how to use Proposition 3.4 and Lemma 5.2 to deduce Proposition 6.6 from it.

Proof of Proposition 6.6, assuming (6.4). Consider the generalized East chain on the space $\otimes_{i=1}^N (S_i, \nu_i)$ with constraining event $S_i^g = \big\{F_i \text{ is empty}\big\}$ (see Definition 3.3).

Note that the East constraint for the last fiber F_N is always satisfied because F_{N+1} is empty, and that the parameters $\{q_i\}_{i=1}^N$ of the generalized East process satisfy

$$q_i = \nu_i\big(S_i^g\big) \geqslant q^{O(n_2)} = \exp\Big(-O\big(q^{-\alpha}\log(1/q)^2\big)\Big).$$

Noting that $N \leqslant |V| = q^{-O(1)}$, it follows from (3.5) that

$$T_{\text{East}}\big(n,\bar{\alpha}\big) \leqslant \exp\Big(O\big(q^{-\alpha}\log(1/q)^3\big)\Big). \tag{6.5}$$

Hence, applying Proposition 3.4 to bound $\mathrm{Var}_\nu(f)$ from above, and recalling (6.4) and that $\Lambda = \bigcup_{i=1}^N F_i$, we obtain

$$\begin{aligned} I_j &\leqslant \frac{1}{p_1}\cdot\mu\Big(\mathbb{1}_{\{F_{N+1}\text{ is empty}\}}\mathrm{Var}_\nu(f)\Big) \\ &\leqslant e^{O(q^{-\alpha}\log(1/q)^3)}\mu\Bigg(\mathbb{1}_{\{F_{N+1}\text{ is empty}\}}\sum_{i=1}^N\nu\Big(\mathbb{1}_{\{F_{i+1}\text{ is empty}\}}\mathrm{Var}_{\nu_i}(f)\Big)\Bigg) \\ &= e^{O(q^{-\alpha}\log(1/q)^3)}\mu\Bigg(\mathbb{1}_{\{F_{N+1}\text{ is empty}\}}\sum_{i=1}^N\mu_\Lambda\Big(\mathbb{1}_{\{F_{i+1}\text{ is empty}\}}\mathrm{Var}_{\nu_i}(f)\Big)\Bigg), \end{aligned} \tag{6.6}$$

where the final inequality follows from the definition of ν_i, and property (b) of the fibers, which implies that the event $H_1\cap\cdots\cap H_N$ has probability at least $p_1^3 = 1-o(1)$ (since it is implied by the goodness of three blocks).

Recall that, by property (c) of the fibers, F_i is contained in the closure, under the $\mathcal{U}$-bootstrap process restricted to a $O(1)$-neighbourhood U_i of the set $F_i\cup F_{i+1}$, of any set of empty sites containing F_{i+1} and for which the event H_i holds. We may therefore apply Lemma 5.2, with $A := F_{i+1}$, $B := F_i$, $\mathcal{E} := H_i$ and $U := U_i$, to obtain

$$\mu_\Lambda\Big(\mathbb{1}_{\{F_{i+1}\text{ is empty}\}}\mathrm{Var}_{\nu_i}(f)\Big) \leqslant O(n_2)q^{-O(n_2)}\sum_{x\in U_i}\mu_\Lambda\big(c_x\,\mathrm{Var}_x(f)\big), \tag{6.7}$$

since $|F_i| = O(n_2)$. Inserting (6.7) into (6.6) we obtain

$$I_j \leqslant e^{O(q^{-\alpha}\log(1/q)^3)}\sum_{x\in\hat{W}_j}\mu\big(c_x\,\mathrm{Var}_x(f)\big)$$

for each $j\in\{1,2\}$, and some $O(1)$-neighbourhood $\hat{W}_j$ of W_j, as required. □

Remark 6.7. We remark that our use of the generalized East chain (rather than the generalised FA-1f chain) in the proof above was necessary (since for α-rooted models Proposition 4.10 can only be used to move infection in the u-direction), and also harmless (since in either case the bound we obtain is of the form $\exp\big(\widetilde{O}(q^{-\alpha})\big)$, which is much smaller than $\exp\big(q^{-2\alpha}\big)$). In the proof for β-unrooted models we will also use the generalized East chain, however, even though in that case we can move infection in both the u- and $-u$-directions, and doing so costs us a factor of $\log(1/q)$ in the exponent for models with $\beta=\alpha$. This is because the method we use in this paper does not appear to easily allow us to use the generalised FA-1f chain in this setting.

In order to conclude the proof of the proposition, it remains to construct in detail the fibers for each case and to prove the basic inequality (6.4).

6.1.3. *Construction of the fibers and the proof of* (6.4). We will first define the helping events and prove (6.4) in the (easier) case $j = 1$. Recall that

$$I_1 = \mu\left(\mathbb{1}_{\{\omega_{\vec{e}_1} \in G_2\}} \mathbb{1}_{\{\omega_{-\vec{e}_1} \in G_1\}} \operatorname{Var}_V\left(f \,|\, G_1\right)\right),$$

where $V = V_{(0,0)}$, and that $\omega_{\vec{e}_1} \in G_2$ implies that there exists an empty quasi-stable half-ring $\mathcal{R}$ of width κ that fits in $V_{(1,0)}$ and is entirely contained in the rightmost quarter of $R_{(1,0)}$, and recall that this determines the length ℓ of $\mathcal{R}$, and that $\ell \geqslant n_2/m$. By translating $\mathcal{R}$ slightly (without changing the set $\mathcal{R} \cap \mathbb{Z}^2$) if necessary, we may also assume that there are no sites of $\mathbb{Z}^2$ on the boundary of $\mathcal{R}$ and in the interior of $R_{(1,0)}$. Let us also choose κ so that the vector κu has integer coordinates. Now, for each such quasi-stable half-ring $\mathcal{R}$, set

$$N = N(\mathcal{R}) := \min\left\{j : \mathcal{R} + j\kappa u \subset V_{(-1,0)}\right\}$$

and define $F_i = F_i(\mathcal{R}) := \hat{F}_i \cap \mathbb{Z}^2$, where

$$\hat{F}_i = \hat{F}_i(\mathcal{R}) := \mathcal{R} + (N + 1 - i)\kappa u,$$

for each $1 \leqslant i \leqslant N+1$. Note that $V_{(0,0)} \subset \bigcup_{i=1}^N F_i$, and that (by our choice of κ) there are no sites of $\mathbb{Z}^2$ on the boundary of $\hat{F}_i$ in the interior of $R_{(-1,0)} \cup R_{(0,0)} \cup R_{(1,0)}$.

Definition 6.8. For each $\mathcal{R}$ and $i \in [N]$, let H_i denote the event that for each quasi-stable direction $v \in C$ and every v-strip S such that the segment $\partial^{\text{ext}}(S) \cap \hat{F}_i$ has length at least $n_2/(2m)$, there exists an empty helping set $Z \subset F_i$ for S.

Notice that in the above definition we do not require the v-strip S to be contained in $\hat{F}_i$. Observe that if the blocks $V_{(-1,0)}$, $V_{(0,0)}$ and $V_{(1,0)}$ are all good, then the event H_i occurs for every $i \in [N]$. Now define $H_{\mathcal{R}}$ to be the event that $\mathcal{R}$ is (up to translates preserving the set $\mathcal{R} \cap \mathbb{Z}^2$) the rightmost empty quasi-stable half-ring in $R_{(1,0)}$, and observe that, conditional on $H_{\mathcal{R}}$, the events $\{H_i\}_{i=1}^N$ are independent. Moreover, by Corollary 4.11, and since κ is sufficiently large, if F_{i+1} is empty and H_i occurs, then the $\mathcal{U}$-bootstrap process restricted to a $O(1)$-neighbourhood of the set $F_i \cup F_{i+1}$ is able to infect F_i. The fibers $\{F_i\}_{i=1}^{N+1}$ therefore satisfy conditions (a), (b) and (c) of Section 6.1.2. Recall that we write $\Lambda = \bigcup_{i=1}^N F_i$. We make the following claim, which implies (6.4) in the case $j = 1$:

Claim 6.9.

$$I_1 \leqslant \frac{1}{p_1} \sum_{\mathcal{R}} \mu\Big(\mathbb{1}_{H_{\mathcal{R}}} \operatorname{Var}_\Lambda\left(f \;\middle|\; H_1 \cap \cdots \cap H_N\right)\Big). \tag{6.8}$$

Note that the sum in the claim is over equivalence classes of quasi-stable half-rings $\mathcal{R}$, where two half-rings are equivalent if they have the same intersection with $\mathbb{Z}^2$.

Proof of Claim 6.9. We first claim that

$$I_1 \leqslant \frac{1}{p_1} \sum_{\mathcal{R}} \mu\Big(\mathbb{1}_{H_{\mathcal{R}}} \mathbb{1}_{\{\omega_{\pm \vec{e}_1} \in G_1\}} \mu_V \Big(\mathbb{1}_{\{\omega_0 \in G_1\}} \big(f - a\big)^2 \Big)\Big), \tag{6.9}$$

where $\omega_0 \equiv \omega_{V_{(0,0)}}$ and, for any $\omega \in H_{\mathcal{R}}$, we set

$$a = a\big(\omega_{\mathbb{Z}^2 \setminus \Lambda}\big) := \mu_\Lambda\big(f \mid H_1 \cap \cdots \cap H_N\big),$$

noting that, on the event $H_{\mathcal{R}}$, the set Λ and the fibers become deterministic. To prove (6.9) we use Definition 6.4, which implies that if $V_{(1,0)}$ is super-good then it is also good, and also that the event $H_{\mathcal{R}}$ holds for some $\mathcal{R}$, and the standard inequality $\mathrm{Var}(X) \leqslant \mathbb{E}\big[(X-a)^2\big]$, which holds for any $a \in \mathbb{R}$ and any random variable X.

Recalling that if the blocks $V_{(-1,0)}$, $V_{(0,0)}$ and $V_{(1,0)}$ are all good, then the event H_i occurs for every $i \in [N]$, it follows from (6.9) that

$$\begin{aligned} I_1 \;&\leqslant\; \frac{1}{p_1} \sum_{\mathcal{R}} \mu\Big(\mathbb{1}_{H_{\mathcal{R}}}\, \mu_\Lambda \Big(\mathbb{1}_{H_1 \cap \cdots \cap H_N} \big(f - a\big)^2 \Big)\Big) \\ &\leqslant\; \frac{1}{p_1} \sum_{\mathcal{R}} \mu\Big(\mathbb{1}_{H_{\mathcal{R}}}\, \mathrm{Var}_\Lambda\left(f \mid H_1 \cap \cdots \cap H_N\right)\Big), \end{aligned}$$

where the last inequality follows from our choice of a and the trivial inequality

$$\mu_\Lambda\Big(\mathbb{1}_{H_1 \cap \cdots \cap H_N} \big(f - a\big)^2 \Big) \leqslant \mu_\Lambda\Big(\big(f-a\big)^2 \mid H_1 \cap \cdots \cap H_N \Big).$$

This proves the claim, and hence (6.4) in the case $j = 1$. ◻

We now turn to the analysis of the term

$$I_2 = \mu\left(\mathbb{1}_{\{\omega_{\vec{e}_2} \in G_2\}} \mathbb{1}_{\{\omega_j \in G_1\, \forall j \in \mathbb{L}^+\}} \mathrm{Var}_V(f) \right).$$

In this case we need to modify the definition of the fibers F_i in order to take into account the different local neighbourhood W_2 of $V_{(0,0)}$ and the different good and super-good environment in W_2 (see Figures 5 and 6).

First, let us define $H_{\mathcal{R}}$ to be the event that $\mathcal{R}$ is (up to translates preserving the set $\mathcal{R} \cap \mathbb{Z}^2$) the rightmost empty quasi-stable half-ring of width κ that fits in $V_{(0,1)}$, and observe that the length ℓ of $\mathcal{R}$ satisfies $\ell \geqslant n_2/m$, and that the event $\omega_{\vec{e}_2} \in G_2$ implies that $H_{\mathcal{R}}$ holds for some $\mathcal{R}$ in the rightmost quarter of $R_{(0,1)}$. As before, we may choose $\mathcal{R}$ so that there are no sites of $\mathbb{Z}^2$ on its boundary in the interior of $R_{(0,1)}$.

The fibers $\{F_i\}_{i=1}^{N+1}$ will be similar to those used above to bound I_1, but some of the v_1-strips (which form the bottom portion of each fiber) will be elongated as the fibers "descend" the upward κ-stair in $V_{(1,0)}$, see Figure 6. (Recall that we call these objects *elongated quasi-stable half-rings*.) To be precise, let us write $L(\mathcal{R})$ for the two-way infinite v_1-strip of width κ that contains the v_1-strip of $\mathcal{R}$, and define

$$N = N(\mathcal{R}) := \min\bigg\{ j : V_{(0,0)} \subset \bigcup_{i=1}^{j} \big(L(\mathcal{R}) + i\kappa u \big) \bigg\}.$$

Now, recall that $a = \Theta(n_2)$ is the number of rows of V, and recall Definition 6.3. Let $\mathcal{M} = \{M^{(1)}, \ldots, M^{(a)}\}$ be an upward κ-stair contained in the leftmost quarter of $V_{(1,0)}$, and define the fibers $\hat{F}_i = \hat{F}_i(\mathcal{R}, \mathcal{M})$ recursively as follows:

(a) $\hat{F}_{N+1} := \mathcal{R}$;

(b) For each $i \in [N]$ set $\hat{F}_i' := \hat{F}_{i+1} + \kappa u$. Now define:

  (i) $\hat{F}_i$ to be an elongated version of $\hat{F}_i'$ such that $(\hat{F}_i \setminus \hat{F}_i') \cap \mathbb{Z}^2$ is a subset of a step of $\mathcal{M}$, if such an elongated quasi-stable half-ring exists;

  (ii) $\hat{F}_i := \hat{F}_i'$ otherwise.

As before, we set $F_i = F_i(\mathcal{R}, \mathcal{M}) := \hat{F}_i \cap \mathbb{Z}^2$ for each $1 \leqslant i \leqslant N+1$. Let us write $H_{\mathcal{M}}$ for the event that $\mathcal{M}$ is the first (in some arbitrary ordering) empty upward κ-stair contained in the leftmost quarter of $V_{(1,0)}$. We can now define the helping events.

Definition 6.10. For each $\mathcal{R}$ and $\mathcal{M}$, and each $i \in [N]$, let H_i denote the event that for each quasi-stable direction $v \in C$ and every v-strip S such that the segment

$$\partial^{\text{ext}}(S) \cap \hat{F}_i \cap \big(R_{(-1,1)} \cup R_{(0,1)}\big)$$

has length at least $n_2/(2m)$, there exists an empty helping set $Z \subset F_i$ for S.

Observe that if the blocks $V_{(-1,1)}$ and $V_{(0,1)}$ are both good, then the event H_i occurs for every $i \in [N]$. Moreover, conditional on the event $H_{\mathcal{R}} \cap H_{\mathcal{M}}$, the events $\{H_i\}_{i=1}^N$ are independent and, by Corollary 4.11 (see Remark 4.12), if F_{i+1} is empty and the events $H_{\mathcal{M}}$ and H_i occur, then the $\mathcal{U}$-bootstrap process restricted to a $O(1)$-neighbourhood of the set $F_i \cup F_{i+1}$ is able to infect F_i. It follows that if the event $H_{\mathcal{R}} \cap H_{\mathcal{M}}$ occurs, then the fibers $\{F_i\}_{i=1}^{N+1}$ satisfy conditions (a), (b) and (c) of Section 6.1.2.

We make the following claim, which implies (6.4) in the case $j = 2$:

Claim 6.11.

$$I_2 \leqslant \frac{1}{p_1} \sum_{\mathcal{R},\mathcal{M}} \mu\Big(\mathbb{1}_{H_{\mathcal{R}}} \mathbb{1}_{H_{\mathcal{M}}} \operatorname{Var}_\Lambda\big(f \mid H_1 \cap \cdots \cap H_N\big)\Big). \tag{6.10}$$

The proof of Claim 6.11 is identical to that of Claim 6.9. As discussed above, this completes the proof of the Proposition 6.6, and hence of Theorem 2 in the case where $\mathcal{U}$ is α-rooted and Assumption 6.1 holds.

6.2. **The β-unrooted case.** In this section we assume that the bilateral difficulty β of the updating rule $\mathcal{U}$ is smaller than 2α. We will prove that, if Assumption 6.1 holds, then there exists a constant $\lambda = \lambda(\mathcal{U})$ such that

$$\mathbb{E}_\mu(\tau_0) \leqslant \frac{T_{\text{rel}}(q,\mathcal{U})}{q} \leqslant \exp\Big(\lambda \cdot q^{-\beta}\big(\log(1/q)\big)^3\Big)$$

for all sufficiently small $q > 0$. Note that the first inequality follows from (2.8), and so, by Definition 2.9, it will suffice to prove that

$$\operatorname{Var}(f) \leqslant \exp\Big(\lambda \cdot q^{-\beta}\big(\log(1/q)\big)^3\Big) \sum_{x \in \mathbb{Z}^2} \mu\big(c_x \operatorname{Var}_x(f)\big) \tag{6.11}$$

for some $\lambda = \lambda(\mathcal{U})$ and all local functions f. We will deduce a bound of the form (6.11) from the *unoriented* constrained Poincaré inequality (3.4) and Corollary 4.11.

Recall from Section 4.3 that $C \subset S^1$ is an open semicircle such that $\alpha(v) \leqslant \beta$ for each $v \in C \cup -C$, and that we let u be the mid-point of C. Similarly to Section 6.1, we set

$$n_1 = \lfloor q^{-3\kappa} \rfloor \qquad \text{and} \qquad n_2 = \lfloor \kappa^4 q^{-\beta} \log(1/q) \rfloor, \tag{6.12}$$

where $\kappa = \kappa(\mathcal{U})$ is a sufficiently large constant.

We need to slightly modify the definition of the good and super-good events G_1 and G_2 as follows. Let $(v_1, \dots, v_m)$ be the quasi-stable directions in C, and let $(v_1, \dots, v_{m'})$ be the quasi-stable directions in $-C$ (see Definition 4.5). As in Section 6.1, it follows by Assumption 6.1 that we may choose the voracious sets so that the helping sets for S_v are subsets of $\partial^{\text{ext}}(S_v)$ for each quasi-stable direction $v \in C \cup -C$.

Definition 6.12 (Good and super-good events)**.**

(1) The block $V_i = R_i \cap \mathbb{Z}^2$ satisfies the *good* event G_1 iff:
  (a) for each quasi-stable direction $v \in C$ and every v-strip S such that the length of the segment $\partial^{\text{ext}}(S) \cap R_i$ is at least $n_2/(4m)$, there exists an empty helping set $Z \subset \partial^{\text{ext}}(S) \cap V_i$ for S;
  (b) for each quasi-stable direction $v \in -C$ and every v-strip S such that the length of the segment $\partial^{\text{ext}}(S) \cap R_i$ is at least $n_2/(4m')$, there exists an empty helping set $Z \subset \partial^{\text{ext}}(S) \cap V_i$ for S;
  (c) there exist two empty upward κ-stairs, one within the leftmost quarter of V_i, and one within the rightmost quarter of V_i.

(2) The block V_i satisfies the *super-good* event G_2 iff it satisfies the good event G_1, and moreover there exist two empty quasi-stable half-rings $\mathcal{R}^+$ and $\mathcal{R}^-$, of width κ, that both fit in V_i, with $\mathcal{R}^+$ relative to C and entirely contained in the rightmost quarter of R_i, and with $\mathcal{R}^-$ relative to $-C$ and entirely contained in the leftmost quarter of R_i.

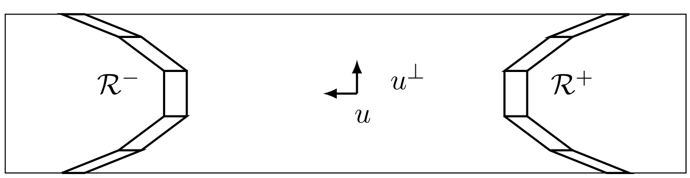

FIGURE 7. The two quasi-stable half-rings $\mathcal{R}^\pm$. For simplicity they have been drawn as mirror images of one another, although in general the quasi-stable directions do not necessarily have this property.

It is easy to check that, with the new definition of the good and super-good events, Lemma 6.5 still holds. It follows, by Theorem 3.1, that the unconstrained Poincaré inequality (3.4) holds for any local function f, i.e.,

$$\operatorname{Var}(f) \leqslant T(p_2)\Bigg(\sum_{\varepsilon=\pm 1} \sum_{i\in\mathbb{Z}^2} \mu\left(\mathbb{1}_{\{\omega_{i+\varepsilon\vec{e}_2}\in G_2\}} \mathbb{1}_{\{\omega_j\in G_1\ \forall j\in \mathbb{L}^{\varepsilon}(i)\}} \operatorname{Var}_i(f)\right) + \sum_{\varepsilon=\pm 1} \sum_{i\in\mathbb{Z}^2} \mu\left(\mathbb{1}_{\{\omega_{i+\varepsilon\vec{e}_1}\in G_2\}} \mathbb{1}_{\{\omega_{i-\varepsilon\vec{e}_1}\in G_1\}} \operatorname{Var}_i\left(f \mid G_1\right)\right) \Bigg). \quad (6.13)$$

with

$$T(p_2) \,=\, p_2^{-O(1)} \,=\, \exp\Big(O\big(q^{-\beta}\log(1/q)^2\big)\Big).$$

As in Section 6.1, using translation invariance it will suffice to bound from above, for a fixed (and arbitrary) local function f, the following four generic terms:

$$I_1^{\pm}(i) := \mu\left(\mathbb{1}_{\{\omega_{i\pm\vec{e}_1}\in G_2\}} \mathbb{1}_{\{\omega_{i\mp\vec{e}_1}\in G_1\}} \operatorname{Var}_{V_i}(f \mid G_1)\right),$$

and

$$I_2^{\pm}(i) := \mu\left(\mathbb{1}_{\{\omega_{i\pm\vec{e}_2}\in G_2\}} \mathbb{1}_{\{\omega_j\in G_1\ \forall j\in \mathbb{L}^{\pm}(i)\}} \operatorname{Var}_{V_i}(f)\right).$$

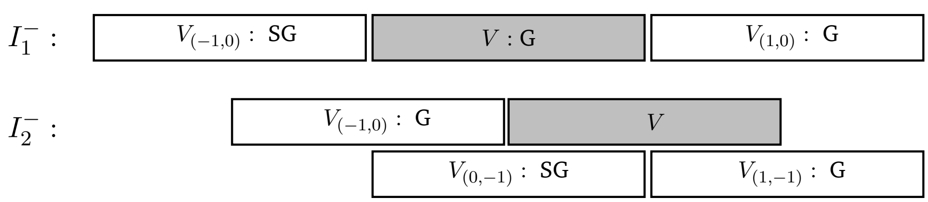

FIGURE 8. In I_1^- the block $V \equiv V_{(0,0)}$ is conditioned to be good (G), while the blocks $V_{(1,0)}$ and $V_{(-1,0)}$ are good and super-good (SG) respectively. In I_2^- the blocks $V_{(-1,0)}$ and $V_{(1,-1)}$ are good, the block $V_{(0,-1)}$ is super-good, and V is unconditioned.

Define $W_1^+ = W_1^- = V_{(0,0)} \cup V_{(-1,0)} \cup V_{(1,0)}$, and $W_2^+ = V_{(0,0)} \cup V_{(-1,0)} \cup V_{(1,0)} \cup V_{(-1,1)} \cup V_{(0,1)}$ and $W_2^- = V_{(0,0)} \cup V_{(1,0)} \cup V_{(-1,0)} \cup V_{(1,-1)} \cup V_{(0,-1)}$. The following upper bounds on $I_1^{\pm}$ and $I_2^{\pm}$ (cf. Proposition 6.6) follow exactly as in Section 6.1.

Proposition 6.13. *For each $j \in \{1,2\}$, there exist $O(1)$-neighbourhoods $\hat{W}_j^{\pm}$ of $W_j^{\pm}$ such that*

$$I_j^{\pm} \,\leqslant\, \exp\Big(O\big(q^{-\beta}\log(1/q)^3\big)\Big) \sum_{x\in \hat{W}_j^{\pm}} \mu\big(c_x \operatorname{Var}_x(f)\big).$$

Sketch proof of Proposition 6.13. The terms I_1^+ and I_2^+ can be treated exactly as the terms I_1 and I_2 analysed in the previous section, because the new good and super-good events imply the good and super-good events for the α-rooted case. We may therefore repeat the proof of Proposition 6.6, with the only difference being that n_2 is now as defined in (6.12), to obtain the claimed bounds on I_1^+ and I_2^+.

For the new terms, I_1^- and I_2^- (which are illustrated in Figure 8), the argument is exactly the same after a rotation of π of the coordinate axes. Indeed, a good block now contains suitable empty helping sets for the quasi-stable directions in $-C$ (as well as C), and an empty upward κ-stair in the rightmost quarter (as well as the leftmost), and a super-good block contains an empty quasi-stable half-ring relative to $-C$ in the leftmost quarter (as well as one relative to C in the rightmost quarter). Such a rotation therefore transforms I_1^- and I_2^- into I_1^+ and I_2^+, and so the proof of the claimed bounds is once again identical to that of Proposition 6.6. □

Remark 6.14. As noted in Remark 6.7, our application of the generalized East chain in the proof above cost us a factor of $\log(1/q)$ in the exponent. More precisely, this log-factor was lost in step (6.5) of the proof of Proposition 6.13, when (roughly speaking) we passed through an energy barrier corresponding to the simultaneous existence of about $\log(1/q)$ empty quasi-stable half-rings in a single block. As stated precisely in Conjecture 4, we expect that (at least for models with $\beta = \alpha$) the true relaxation time does not contain this additional factor of $\log(1/q)$.

Combining Proposition 6.13 with (6.13), and noting that $|\hat{W}_j^{\pm}| = q^{-O(1)}$, we obtain a final Poincaré inequality of the form (6.11), i.e.,

$$\operatorname{Var}(f) \leqslant \exp\Big(O\big(q^{-\beta}\log(1/q)^3\big)\Big) \sum_{x\in\mathbb{Z}^2} \mu\big(c_x \operatorname{Var}_x(f)\big),$$

as required. This completes the proof of Theorem 2 for update families $\mathcal{U}$ such that Assumption 6.1 holds. □

7. Critical KCM: removing the simplifying assumption

In this section we explain how to modify the proof given in Section 6 in order to avoid using Assumption 6.1. Since the argument is essentially identical for α-rooted and β-unrooted families, for simplicity we will restrict ourselves to the α-rooted case.

Our solution requires a slight change in the geometry of the quasi-stable half-ring. In what follows we will always work in the frame $(-u, u^{\perp})$, where u is the midpoint of the semicircle C given by Lemma 4.4 (cf. Sections 4.3 and 6.1).

Recall from Definition 4.7 the definitions of the $+$- and $-$-boundaries of a v-strip S. The following key definition is illustrated in Figure 9a.

Definition 7.1 (Generalised quasi-stable half-rings). Let $(v_1, \dots, v_m)$ be the quasi-stable directions in C, ordered as in Definition 4.9, and let $\mathcal{R}$ be a quasi-stable half-ring of width w and length ℓ relative to C. For each quasi-stable direction $v \in C$, let S_v be the v-strip in $\mathcal{R}$, and let $\hat{S}_v^l$ and $\hat{S}_v^r$ be the (unique) v-strips of width $w/3$ and length $\ell/3$ satisfying the following properties:

(i) $\hat{S}_v^l$ and $\hat{S}_v^r$ each share exactly one corner with S_v; moreover each of these corners lies at the "top" of S_v when working in the frame $(-u, u^\perp)$.

(ii) $\partial_-(\hat{S}_v^l) \subset \partial_+(S_v)$ and $\partial_-(\hat{S}_v^r) \subset \partial_-(S_v)$.

Set

$$S_v^g := \big(S_v \setminus \hat{S}_v^r\big) \cup \hat{S}_v^l,$$

and set

$$\mathcal{R}^g := \bigcup_{i=1}^m S_{v_i}^g.$$

We call $\mathcal{R}^g$ the *generalised version of* $\mathcal{R}$, and define the "core" of $\mathcal{R}^g$ to be the set $\mathcal{R}^g \cap \mathcal{R}$.

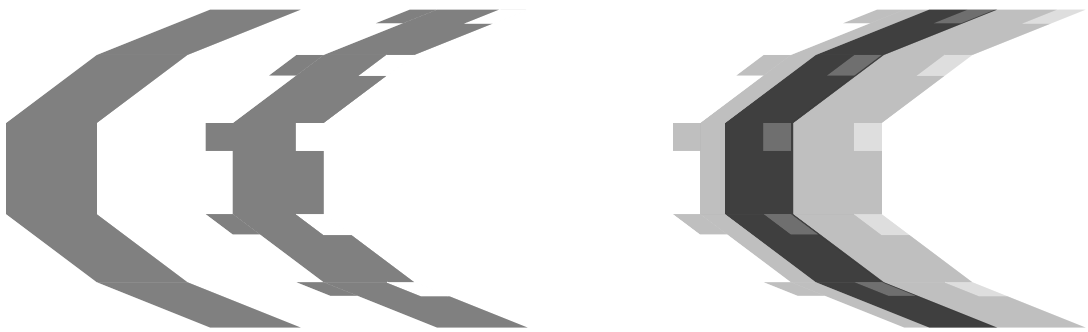

(A) A quasi-stable half-ring $\mathcal{R}$ (left) and its generalised version $\mathcal{R}^g$ (right). The core of $\mathcal{R}^g$ is $\mathcal{R}^g$ stripped off all the left protuberances (*i.e.,* the sets $\{\hat{S}_{v_i}^l\}_{i=1}^n$) .

(B) The union of $\mathcal{R}^g$ and its translate by κ in the u-direction (in light grey) together with the quasi-stable half-ring $\mathcal{R}'$ of smaller width (dark grey).

FIGURE 9. A generalised quasi-stable half-ring.

Recall now the following two key ingredients of the proof given in the previous section (see Section 6.1.2) under the simplifying Assumption 6.1:

(i) a sufficiently large empty quasi-stable half-ring $\mathcal{R}$ is able to completely infect its translate $\mathcal{R} + \kappa u$, provided that a certain "helping" event occurs;

(ii) the helping event depends only on the configuration inside $\mathcal{R}$.

Here we prove a similar result for the generalised quasi-stable half-rings *without* the simplifying assumption. We first define the helping event, cf. Definition 6.8.

Definition 7.2. Given a quasi-stable half-ring $\mathcal{R}$ of length ℓ and width κ, we define $H(\mathcal{R})$ to be the event that for each quasi-stable direction $v \in C$ and every v-strip S of length ℓ with $\partial_+(S) \subset \mathcal{R}$, there exists an empty helping set for S in $\mathcal{R}^g$.

If $H(\mathcal{R})$ holds, then we will say that $\mathcal{R}^g$ is *helping*. We will modify (see below) the good and super-good events G_1 and G_2 (see Definition 6.4) so that they guarantee that this helping event occurs, and choose the constant $\kappa = \kappa(\mathcal{U}) > 0$ (as in Section 6.1) so that the conclusion of Lemma 6.5 holds, and so that κu has integer coordinates. We will also choose our (generalised) quasi-stable half-rings so that there are no sites of $\mathbb{Z}^2$ on their boundary, except on the top and bottom boundaries of the rectangles R_i.

Lemma 7.3. *Let $\mathcal{R}$ be a quasi-stable half-ring of length ℓ and width κ, and let $\mathcal{R}^g$ be the generalised version of $\mathcal{R}$. Assume that the core of $\mathcal{R}^g$ is empty and that $\mathcal{R}^g$ and its translate $\mathcal{R}^g + \kappa u$ are both helping. Then there exists a $O(1)$-neighbourhood U of $\mathcal{R}^g \cup \big(\mathcal{R}^g + \kappa u\big)$ such that the $\mathcal{U}$-bootstrap process restricted to U is able to infect the core of $\mathcal{R}^g + \kappa u$.*

Proof. The lemma is a straightforward consequence of Proposition 4.10, using the geometry of the generalised quasi-stable half-rings. To spell out the details (cf. the proof of Corollary 4.11), fix $\mathcal{R}$ as in the lemma, and let $\mathcal{R}'$ be any quasi-stable half-ring of length ℓ and width $\kappa/3$ such that:

(a) $\mathcal{R}' = \mathcal{R} + \lambda u$ for some $\lambda \geqslant 0$, and
(b) $\mathcal{R}' \subset \mathcal{R}^g \cup \big(\mathcal{R}^g + \kappa u\big)$,

see Figure 9b. We claim that, for every quasi-stable direction $v \in C$, there exists an empty helping set in $\mathcal{R}^g \cup \big(\mathcal{R}^g + \kappa u\big)$ for the v-strip S'_v of $\mathcal{R}'$. Indeed, this follows from the fact that $\mathcal{R}^g$ and $\mathcal{R}^g + \kappa u$ are both helping, since (by construction) either $\partial_+(S'_v) \subset \mathcal{R}$ or $\partial_+(S'_v) \subset \mathcal{R} + \kappa u$.

Now, since the core of $\mathcal{R}^g$ is empty, it follows, by Proposition 4.10, that there exists a $O(1)$-neighbourhood U of $\mathcal{R}^g \cup \big(\mathcal{R}^g + \kappa u\big)$ such that the $\mathcal{U}$-bootstrap process restricted to U can advance in the u-direction, and infect the core of $\mathcal{R}^g + \kappa u$, as claimed. □

Given the above lemma, the proof of Theorem 2 proceeds exactly as the one given in Section 6, with only two main changes:

(a) the fibers $\{F_i\}_{i=1}^N$ are no longer the quasi-stable half-rings (or their elongated version) but rather the generalised quasi-stable half-rings (or their elongated version);
(b) when defining the generalised East process for the fibers, the constraining event S_i^g (see Definition 3.3), which in Section 6 was simply $S_i^g = \{F_i \text{ is empty}\}$, now becomes $S_i^g = \{\text{the core of } F_i \text{ is empty}\}$.

We leave the (straightforward) task of verifying the details to the reader.

APPENDIX A.

1.1. **Proof of Proposition 3.4.** We will follow closely the proof of a very similar result proved in [14, Proposition 3.4]. Let $\{P_t\}_{t\geqslant 0}$ be the Markov semigroup associated to either the generalised East chain or the generalised FA-1f chain. Using reversibility, it follows (see, e.g., [31, Theorem 2.1.7]) that

$$\lim_{t\to\infty} -\frac{1}{t}\log\Big(\max_{\omega}\big\|P_t(\omega,\cdot)-\nu(\cdot)\big\|_{\mathrm{TV}}\Big)=\frac{1}{T_{\mathrm{rel}}}, \tag{A.1}$$

where $\|\cdot\|_{\mathrm{TV}}$ denotes the total variation distance. We now claim that for every function $f\colon \Omega\to\mathbb{R}$ with $\|f\|_\infty\leqslant 1$,

$$\big\|P_t f-\nu(f)\big\|_\infty \leqslant C(n,q)e^{-t/t^*} \tag{A.2}$$

for some $0<C(n,q)<\infty$ and either $t^*\leqslant T_{\text{East}}(n,\bar\alpha)/q$ or $t^*\leqslant T_{\text{FA}}(n,\bar\alpha)/q$, depending on which of the two models we are considering. Clearly (A.1) and (A.2) imply that $T_{\mathrm{rel}}\leqslant t^*$ and (recalling Definition 2.9) the proposition follows.

To prove (A.2), let $\tau_x(\omega)$ be the time of the first legal ring at x (that is, the first time that the state of x is resampled) when the starting configuration is ω. Then, for any function $f\colon \otimes_{x\in[n]} S_x\mapsto\mathbb{R}$ with $\nu(f)=0$, we write

$$\begin{aligned}\|P_t f\|_\infty \leqslant \max_\omega\bigg\{\Big|\mathbb{E}_\omega\Big(f\big(\omega(t)\big)\cdot\mathbb{1}_{\{\tau_x(\omega)<t\ \forall x\}}\Big)\Big| \\ +\|f\|_\infty\cdot n\cdot\max_{x\in[n]}\mathbb{P}_\omega\Big(\tau_x(\omega)>t\Big)\bigg\},\end{aligned} \tag{A.3}$$

where $\mathbb{P}_\omega(\cdot)$ and $\mathbb{E}_\omega(\cdot)$ denote the law and associated expectation of the chain $\{\omega(t)\}_{t\geqslant 0}$ with $\omega(0)=\omega$.

If $\eta(\omega)=\{\eta_x(\omega)\}_{x\in[n]}$ denotes the collection of the 0-1 variables $\eta_x=\mathbb{1}_{\{\omega_x\in S_x^g\}}$ and $\hat\tau_x(\eta)$ is the hitting time of the set $\{\eta' :\ \eta'_x\neq\eta_x\}$, then $\{\tau_x(\omega)>t\}\subset\{\hat\tau_x(\eta(\omega))>t\}$, and hence $\mathbb{P}_\omega\big(\tau_x(\omega)>t\big)\leqslant\mathbb{P}_\omega\big(\hat\tau_x(\eta(\omega))>t\big)$. Notice that $\eta(t)\equiv\eta(\omega(t))$ is itself a Markov chain whose law $\tilde{\mathbb{P}}_\eta(\cdot)$ coincides with that of either the non-homogeneous East chain or the non-homogeneous FA-1f chain, depending on the chain described by P_t. Therefore, $\mathbb{P}_\omega\big(\hat\tau_x(\eta)>t\big)=\tilde{\mathbb{P}}_\eta\big(\hat\tau_x(\eta)>t\big)$, where $\eta\equiv\eta(\omega)$. Letting $\tilde\nu=\mathrm{Ber}(\alpha_1)\otimes\cdots\otimes\mathrm{Ber}(\alpha_n)$, we have that $\tilde\nu$ is the reversible measure for the η-chain and that

$$\begin{aligned}\tilde{\mathbb{P}}_\eta\big(\hat\tau_x(\eta)>t\big)\ &\leqslant\ \frac{1}{\min_\eta\tilde\nu(\eta)}\sum_{\eta'}\tilde\nu(\eta')\tilde{\mathbb{P}}_{\eta'}\big(\hat\tau_x(\eta')>t\big)\\ &\leqslant\ \begin{cases}2q^{-n}\exp\big(-tq/T_{\text{East}}(n,\bar\alpha)\big) & \text{for the East process,}\\ 2q^{-n}\exp\big(-tq/T_{\text{FA}}(n,\bar\alpha)\big) & \text{for the FA-1f process,}\end{cases}\end{aligned}$$

where the factor q^{-n} comes from $\tilde\nu(\eta)\geqslant q^n$ and the exponential bounds follow from [9, Theorem 4.4]. In particular, the inverse of the exponential rate of decay (in t) of the second term in the r.h.s. of (A.3) is smaller than $T_{\text{East}}(n,\bar\alpha)/q$ or $T_{\text{FA}}(n,\bar\alpha)/q$, depending on which of the two models we are considering.

We now analyse the first term in the r.h.s. of (A.3). Conditionally on the event $\bigcap_x \{\tau_x(\omega) < t\}$ and on the vector $\eta(t) \in \{0,1\}^n$, the variables $\{\omega_x(t) : x \in [n]\}$ are independent with law $\otimes_{x\in[n]}\nu_x(\,\cdot \mid \eta_x(t))$. Thus, if $g(\eta') := \nu(f \mid \eta')$, then

$$\begin{aligned}\mathbb{E}_\omega\Big(f\big(\omega(t)\big)\cdot \mathbb{1}_{\{\tau_x(\omega)<t\ \forall x\}}\Big) &= \mathbb{E}_\omega\Big(g\big(\eta(t)\big)\cdot \mathbb{1}_{\{\tau_x(\omega)<t\ \forall x\}}\Big)\\ &= \tilde P_t\, g(\eta) - \mathbb{E}_\omega\Big(g\big(\eta(t)\big)\cdot \mathbb{1}_{\{\max_x \tau_x(\omega)>t\}}\Big)\end{aligned}$$

where $\tilde P_t\, g(\eta) \equiv \tilde{\mathbb{E}}_\eta\big(g\big(\eta(t)\big)\big) = \mathbb{E}_\omega\big(g\big(\eta(t)\big)\big)$. The last term in the r.h.s. above can be analysed exactly as the second term in the r.h.s. of (A.3). Moreover, by the Cauchy–Schwarz inequality and (2.7), the first term satisfies

$$\|\tilde P_t\, g\|_\infty \leqslant \frac{1}{\min_\eta \tilde\nu(\eta)}\operatorname{Var}_{\tilde\nu}\big(\tilde P_t\, g\big)^{1/2} \leqslant \frac{1}{q^n}e^{-\lambda t}\operatorname{Var}_{\tilde\nu}(g)^{1/2},$$

where λ is either $T_{\text{East}}(n,\bar\alpha)^{-1}$ or $T_{\text{FA}}(n,\bar\alpha)^{-1}$ depending on the chosen model. This proves (A.2), and hence the proposition. □

1.2. **Proof of the scaling** (3.5). Recall that $q := \min\{1-\alpha_x : x \in [n]\}$, and let $T_{\text{East}}(n,q)$ and $T_{\text{FA}}(n,q)$ be the relaxation times of the homogenous East and FA-1f chains on $[n]$ with parameters $\alpha_x = 1-q$ for each $x \in [n]$. It was proved in [10, 13] that

$$T_{\text{East}}(n,q) = q^{-O(\min\{\log n,\,\log(1/q)\})} \quad\text{and}\quad T_{\text{FA}}(n,q) = q^{-O(1)}.$$

Thus, it will suffice to prove that

$$T_{\text{East}}(n,\bar\alpha) = \frac{1}{q}\cdot T_{\text{East}}(n,q) \quad\text{and}\quad T_{\text{FA}}(n,\bar\alpha) = \frac{1}{q}\cdot T_{\text{FA}}(n,q).$$

For simplicity we only treat the East case, since the FA-1f case follows by exactly the same arguments.

Consider the generalized East chain on $\Omega = [0,1]^n$ in which each vertex $x \in [n]$, with rate one and independently across $[n]$, is resampled from the uniform measure on $[0,1]$ if either $x \leqslant n-1$ and $\omega_{x+1} \geqslant 1-q$, or $x = n$. The chain is reversible w.r.t. the uniform measure ν on Ω and, by Proposition 3.4, we have

$$\operatorname{Var}_\nu(f) \leqslant \frac{1}{q}\cdot T_{\text{East}}(n,q)\cdot \sum_{x=1}^n \nu\big(\vec c_x \operatorname{Var}_x(f)\big) \tag{A.4}$$

for every function $f\colon \Omega \mapsto \mathbb{R}$, since $\nu\big(\omega_x \geqslant 1-q\big) = q$ for each $x \in [n]$. (Recall that $\vec c_x(\omega) = \mathbb{1}_{\{\omega_{x+1}\geqslant 1-q\}}$ if $x \leqslant n-1$, and that $\vec c_n(\omega) \equiv 1$.)

Now, let $\eta = \{\eta_x\}_{x\in[n]}$ with $\eta_x := \mathbb{1}_{\{\omega_x<\alpha_x\}}$, and, for an arbitrary function $g\colon \{0,1\}^n \mapsto \mathbb{R}$, set $f(\omega) := g\big(\eta(\omega)\big)$. Note that $\eta_x \leqslant \mathbb{1}_{\{\omega_x\leqslant 1-q\}}$ (by the definition of q), and that the law of the variables η w.r.t. ν is the product Bernoulli measure $\pi = \mathrm{Ber}(\alpha_1) \otimes \cdots \otimes \mathrm{Ber}(\alpha_n)$. Therefore, applying (A.4) to f, we obtain

$$\mathrm{Var}_\pi(g) = \mathrm{Var}_\nu(f) \leqslant \frac{1}{q} \cdot T_{\text{East}}(n,q) \cdot \left(\sum_{x=1}^{n-1} \pi\Big(\mathbb{1}_{\{\eta_{x+1}=0\}} \mathrm{Var}_x(g) \Big) + \pi\big(\mathrm{Var}_n(g) \big) \right).$$

The right-hand side of this inequality is exactly $C \cdot \mathcal{D}(g)$, where $C = 1/q \cdot T_{\text{East}}(n,q)$ and $\mathcal{D}(g)$ is the Dirichlet form of g associated to the generator of the non-homogenous East model. Since g was an arbitrary function, it follows by Definition 2.9 that $T_{\text{East}}(n,\bar\alpha) = 1/q \cdot T_{\text{East}}(n,q)$, as required. As noted above, the proof that $T_{\text{FA}}(n,\bar\alpha) = 1/q \cdot T_{\text{FA}}(n,q)$ is identical. □

1.3. **Proof of Proposition 3.5.** We will deduce the proposition from [26, Theorem 1]. The deduction is almost exactly the same as that of [26, Proposition 3.4], but for completeness we give the details. Set $\ell = \lceil \log(1/p_2) \rceil$, $L = \lfloor 1/p_2^2 \rfloor$, and for each $i \in \mathbb{Z}^2$, define

$$C_i(\ell) = \bigcup_{k=0}^{\ell} \big\{ i + \vec{e}_2 + k\vec{e}_1 \big\}.$$

Let also $\mathcal{P}_i(\ell, L)$ be the family of oriented paths starting in $C_i(\ell)$ and of length L. We define two families of events $\big\{A_i^{(1)}, A_i^{(2)}\big\}_{i\in\mathbb{Z}^d}$ as follows:

$$A_i^{(1)} = \big\{ \omega_j \in G_1 \text{ for all } j \in C_i(\ell) \cup \{i+\vec{e}_1\} \cup \{i+\vec{e}_2-\vec{e}_1\} \big\},$$
$$A_i^{(2)} = \big\{ \text{there exists a good path in } \mathcal{P}_i(\ell,L) \text{ and the smallest good one is super-good} \big\},$$

where, if there is more than one smallest good path, then we choose the leftmost one.

Observe that $A_i^{(1)} \cap A_i^{(2)} \subset \Gamma_i$, since $A_i^{(1)}$ implies that the smallest good path in $\mathcal{P}_i(\ell,L)$ starts at $i + \vec{e}_2$,[8] and hence is equal to the smallest path in the definition of Γ_i. We now want to apply [26, Theorem 1] to the two families of constraints $\big\{c_i^{(k)}\big\}_{i\in\mathbb{Z}^2}$, where $c_i^{(k)} := \mathbb{1}_{\{A_i^{(k)}\}}$ for each $k \in \{1,2\}$. To do so, we need to check the following two conditions:

- (a) there exists a two-way infinite sequence of sets $(\ldots, V_{-2}, V_{-1}, V_0, V_1, V_2, \ldots)$, with $V_n \subset V_{n+1}$ for every $n \in \mathbb{Z}$ and $\bigcup_n V_n = \mathbb{Z}^2$, such that if $i \notin V_n$, then the event $A_i^{(k)}$ is independent of the collection of variables $\big\{\omega_i : i \in V_{n+1}\big\}$;
- (b) there exists a family $\big\{\lambda_I : \emptyset \neq I \subset \{1,2\}\big\}$ of positive constants such that the key condition [26, equation (2.1)] holds.

To see (a), let the sets V_n be all translations of the closed half-space

$$\overline{\mathbb{H}}_{(1,2)} = \big\{ x \in \mathbb{Z}^2 : \langle x, (1,2) \rangle \leqslant 0 \big\}$$

[8]This follows from the observation that the word (of length L) obtained from $W \in \{\vec{e}_1, \vec{e}_2\}^L$ by adding $\vec{e}_1$ at the start and removing the final letter is at most W in alphabetical order.

by elements of $\mathbb{Z}^2$ (ordered in the obvious way). Now, observe that if $i \notin V_n$ then $V_{n+1} \subset \overline{\mathbb{H}}_{(1,2)} + i$, and the event $A_i^{(k)}$ is indeed independent of the variables in $\overline{\mathbb{H}}_{(1,2)} + i$.

To prove (b), set $\lambda_I = 1$ for every non-empty set $I \subset \{1,2\}$, and note that the event $A_i^{(1)}$ depends on $\ell + 3$ variables, and that $A_i^{(2)}$ depends on at most $(L+\ell)^2$ variables. It follows that there exists a constant $\hat{\delta} > 0$ such that [26, equation (2.1)] holds if

$$\ell\Big(1 - \mu\big(A_i^{(1)}\big)\Big) + (L+\ell)^2\Big(1 - \mu\big(A_i^{(2)}\big)\Big) \leqslant \hat{\delta}. \tag{A.5}$$

We now claim that if the constant δ of Proposition 3.5 is chosen to be sufficiently small, then (A.5) holds. In order to prove this, it is enough to observe that, by the union bound,

$$1 - \mu\big(A_i^{(1)}\big) \leqslant (\ell+3)(1-p_1),$$

and that

$$\begin{aligned} 1 - \mu\big(A_i^{(2)}\big) \leqslant\; & \mu\big(\text{there is no good path in } \mathcal{P}_i(\ell, L)\big) \\ & + \max_{\gamma \in \mathcal{P}_i(\ell,L)} \mu\big(\gamma \text{ is not super-good} \,\big|\, \gamma \text{ is good}\big) \\ \leqslant\; & e^{-m(p_1)\ell} + \big(1-p_2\big)^L, \end{aligned}$$

with $\lim_{p_1 \to 1} m(p_1) = \infty$, by a standard Peierls bound and by the FKG inequality. In conclusion, if $\delta > 0$ is sufficiently small then we may apply [26, Theorem 1], which gives

$$\mathrm{Var}(f) \leqslant 4 \sum_{i \in \mathbb{Z}^2} \mu\Big(\mathbb{1}_{\{A_i^{(1)} \cap A_i^{(2)}\}} \mathrm{Var}_i(f)\Big) \leqslant 4 \sum_{i \in \mathbb{Z}^2} \mu\Big(\mathbb{1}_{\Gamma_i} \mathrm{Var}_i(f)\Big),$$

where the final inequality holds because $A_i^{(1)} \cap A_i^{(2)} \subset \Gamma_i$.

□

Acknowledgments. This work has been supported by the ERC Starting Grant 680275 "MALIG", ANR-15-CE40-0020-01 and by the PRIN 20155PAWZB "Large Scale Random Structures". RM is also partially supported by CNPq (Proc. 303275/2013-8) and by FAPERJ (Proc. 201.598/2014).

REFERENCES

[1] D. Aldous and P. Diaconis, *The asymmetric one-dimensional constrained Ising model: rigorous results*, J. Stat. Phys. **107** (2002), no. 5-6, 945–975.

[2] H. C. Andersen and G. H. Fredrickson, *Kinetic Ising Model of the Glass Transition*, Phys. Rev. Lett. **53** (1984), no. 13, 1244–1247.

[3] A. Asselah and P. Dai Pra, *Quasi-stationary measures for conservative dynamics in the infinite lattice*, Ann. Prob. **29** (2001), no. 4, 1733–1754.

[4] P. Balister, B. Bollobás, M. J. Przykucki, and P. Smith, *Subcritical $\mathcal{U}$-bootstrap percolation models have non-trivial phase transitions*, Trans. Amer. Math. Soc. **368** (2016), 7385–7411.

[5] B. Bollobás, H. Duminil-Copin, R. Morris, and P. Smith, *Universality of two-dimensional critical cellular automata*, to appear in *Proc. London Math. Soc.* (2016), available at `arXiv.org:1406.6680`.

[6] B. Bollobás, H. Duminil-Copin, R. Morris, and P. Smith, *The sharp threshold for the Duarte model*, Ann. Prob. **45** (2017), 4222–4272.

[7] O. Blondel, N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Fredrickson-Andersen one spin facilitated model out of equilibrium*, Markov Proc. Rel. Fields **19** (2013), 383–406.

[8] B. Bollobás, P. Smith, and A. Uzzell, *Monotone cellular automata in a random environment*, Combin. Probab. Comput. **24** (2015), no. 4, 687–722.

[9] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Facilitated spin models: recent and new results*, Methods of contemporary mathematical statistical physics, Lecture Notes in Math., vol. 1970, Springer, Berlin, 2009, pp. 307–340.

[10] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Kinetically constrained spin models*, Prob. Theory Rel. Fields **140** (2008), no. 3-4, 459–504.

[11] 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.

[12] P. Chleboun, A. Faggionato, and F. Martinelli, *Mixing time and local exponential ergodicity of the East-like process in $\mathbb{Z}^d$*, Annales de la Faculté des Sciences de Toulouse: Mathématiques, Série 6, **24** (2015), no. 4, 717–743.

[13] P. Chleboun, A. Faggionato, and F. Martinelli, *Time scale separation and dynamic heterogeneity in the low temperature East model*, Commun. Math. Phys. **328** (2014), 955-993.

[14] P. Chleboun, A. Faggionato, and F. Martinelli, *Relaxation to equilibrium of generalized East processes on $\mathbb{Z}^d$: Renormalisation group analysis and energy-entropy competition*, Ann. Prob. **44** (2016), no. 3, 1817–1863.

[15] F. Chung, P. Diaconis, and R. Graham, *Combinatorics for the East model*, Adv. Appl. Math. **27** (2001), no. 1, 192–206.

[16] J.A.M.S. Duarte, *Simulation of a cellular automaton with an oriented bootstrap rule*, Phys. A. **157** (1989), no. 3, 1075–1079.

[17] H. Duminil-Copin and A.C.D. van Enter, *Sharp metastability threshold for an anisotropic bootstrap percolation model*, Ann. Prob. **41** (2013), 1218–1242.

[18] H. Duminil-Copin, A.C.D van Enter, and T. Hulshof, *Higher order corrections for anisotropic bootstrap percolation* (2016), available at `arXiv.org:1611.03294`.

[19] A.C.D. van Enter, *Proof of Straley's argument for bootstrap percolation*, J. Stat. Phys. **48** (1987), 943–945.

[20] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *The East model: recent results and new progresses*, Markov Proc. Rel. Fields **19** (2013), 407–458.

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

[22] J. P. Garrahan, P. Sollich, and C. Toninelli, *Kinetically constrained models*, in "Dynamical heterogeneities in glasses, colloids, and granular media" (Eds.: L. Berthier, G. Biroli, J.-P. Bouchaud, L. Cipelletti and W. van Saarloos), Oxford Univ. Press (2011).

[23] J. Jäckle and S. Eisinger, *A hierarchically constrained kinetic Ising model*, Z. Phys. B: Condensed Matter **84** (1991), no. 1, 115–124.

[24] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, 2008.

[25] T. M. Liggett, *Interacting particle systems*, Springer-Verlag, New York, 1985.

[26] F. Martinelli and C. Toninelli, *Towards a universality picture for the relaxation to equilibrium of kinetically constrained models*, to appear in Ann. Prob. (2018), available at `arXiv.org:1701.00107`.

[27] L. Marêché, F. Martinelli, and C. Toninelli, *Energy barriers and the infection time for the kinetically constrained Duarte model* (in preparation).

[28] T. S. Mountford, *Critical length for semi-oriented bootstrap percolation*, Stoch. Proc. Appl. **56** (1995), no. 2, 185–205.

[29] R. Morris, *Bootstrap percolation, and other automata*, Europ. J. Combin. **66** (2017), 250–263.

[30] F. Ritort and P. Sollich, *Glassy dynamics of kinetically constrained models*, Adv. Physics **52** (2003), no. 4, 219–342.

[31] L. Saloff-Coste, *Lectures on finite Markov chains* (P. Bernard, ed.), Lecture Notes in Mathematics, vol. 1665, Springer Berlin Heidelberg, 1997.

[32] R. Schonmann, *On the behaviour of some cellular automata related to bootstrap percolation*, Ann. Prob. **20** (1992), 174–193.

[33] R. Schonmann, *Critical Points of Two-Dimensional Bootstrap Percolation-Like Cellular Automata*, J.Stat.Phys. **58** (1990), 1239–1244.

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

[35] N. S. Pillai and A. Smith, *Mixing times for a constrained Ising process on the torus at low density*, Ann. Prob. **45** (2017), no. 2, 1003–1070.

Dipartimento di Matematica e Fisica, Università Roma Tre, Largo S.L. Murialdo 00146, Roma, Italy
E-mail address: martin@mat.uniroma3.it

IMPA, Estrada Dona Castorina 110, Jardim Botâico, Rio de Janeiro, 22460-320, Brazil
E-mail address: rob@impa.br

Laboratoire de Probabilités, Modélisation et Statistique, CNRS-UMR 7599 Universités Paris VI-VII 4, Place Jussieu F-75252 Paris Cedex 05 France
E-mail address: cristina.toninelli@upmc.fr