



Universality for critical KCM: infinite number of stable directions

Ivailo Hartarsky^{1,2} · Laure Maréché³ · Cristina Toninelli²

Received: 29 July 2019 / Revised: 9 April 2020 / Published online: 8 June 2020
© Springer-Verlag GmbH Germany, part of Springer Nature 2020

Abstract

Kinetically constrained models (KCM) are reversible interacting particle systems on \mathbb{Z}^d with continuous-time constrained Glauber dynamics. They are a natural non-monotone stochastic version of the family of cellular automata with random initial state known as \mathcal{U} -bootstrap percolation. KCM have an interest in their own right, owing to their use for modelling the liquid-glass transition in condensed matter physics. In two dimensions there are three classes of models with qualitatively different scaling of the infection time of the origin as the density of infected sites vanishes. Here we study in full generality the class termed ‘critical’. Together with the companion paper by Hartarsky et al. (Universality for critical KCM: finite number of stable directions. arXiv e-prints [arXiv:1910.06782](https://arxiv.org/abs/1910.06782), 2019) we establish the universality classes of critical KCM and determine within each class the critical exponent of the infection time as well as of the spectral gap. In this work we prove that for critical models with an infinite number of stable directions this exponent is twice the one of their bootstrap percolation counterpart. This is due to the occurrence of ‘energy barriers’, which determine the dominant behaviour for these KCM but which do not matter for the monotone bootstrap dynamics. Our result confirms the conjecture of Martinelli et al. (Commun Math Phys 369(2):761–809. <https://doi.org/10.1007/s00220-018-3280-z>, 2019), who proved a matching upper bound.

Keywords Kinetically constrained models · Bootstrap percolation · Universality · Glauber dynamics · Spectral gap

Mathematics Subject Classification Primary 60K35; Secondary 82C22 · 60J27 · 60C05

This work was supported by European Research Council Starting Grant 680275 MALIG and by Agence Nationale de la Recherche-15-CE40-0020-01.

Extended author information available on the last page of the article

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 by Fredrickson and Andersen [16] in order to model the liquid-glass transition (see e.g. [17,31] for reviews), a major and still largely open problem in condensed matter physics [5]. A generic KCM is a continuous-time Markov process of Glauber type characterised by a finite collection \mathcal{U} of finite nonempty subsets of $\mathbb{Z}^d \setminus \{0\}$, its *update family*. 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 $x \in \mathbb{Z}^d$ waits an independent, mean one, exponential time and then, iff there exists $U \in \mathcal{U}$ such that $\omega_y = 0$ for all $y \in U + x$, site x is updated to empty with probability q and to occupied with probability $1 - q$. Since each $U \in \mathcal{U}$ is contained in $\mathbb{Z}^d \setminus \{0\}$, the constraint to allow the update does not depend on the state of the to-be-updated site. As a consequence, the dynamics satisfies detailed balance w.r.t. the product Bernoulli($1 - q$) measure, μ , which is therefore a reversible invariant measure. Hence the process started at μ is stationary.

Both from a physical and from a mathematical point of view, a central issue for KCM is to determine the speed of divergence of the characteristic time scales when $q \rightarrow 0$. Two key quantities are: (i) the *relaxation time* T_{rel} , i.e. the inverse of the spectral gap of the Markov generator (see Definition 2.5) and (ii) the *mean infection time* $\mathbb{E}(\tau_0)$, i.e. the mean over the stationary process started at μ of the first time at which the origin becomes empty. Several works have been devoted to the study of these time scales for some specific choices of the constraints [2,9,12,13,25,27] (see also [17] section 1.4.1 for a non exhaustive list of references in the physics literature). These results show that KCM exhibit a very large variety of possible scalings depending on the update family \mathcal{U} . A question that naturally emerges, and that has been first addressed in [26], is whether it is possible to group all possible update families into distinct *universality classes* so that all models of the same class display the same divergence of the time scales.

Before presenting the results and the conjectures of [26], we should describe the key connection of KCM with a class of discrete monotone cellular automata known as \mathcal{U} -bootstrap percolation (or simply bootstrap percolation) [8]. For \mathcal{U} -bootstrap percolation on \mathbb{Z}^d , given an update family \mathcal{U} and a set A_t of sites *infected* at time t , the infected sites in A_t remain infected at time $t + 1$, and every site x becomes infected at time $t + 1$ if the translate by x of one of the sets in \mathcal{U} is contained in A_t . The set of initial infections A is chosen at random with respect to the product Bernoulli measure with parameter $q \in [0, 1]$, which identifies with μ : for every $x \in \mathbb{Z}^d$ we have $\mu(x \in A) = q$. One then defines the *critical probability* $q_c(\mathbb{Z}^d, \mathcal{U})$ to be the infimum of the q such that with probability one the whole lattice is eventually infected, namely $\bigcup_{t \geq 0} A_t = \mathbb{Z}^d$. A key time scale for this dynamics is the first time at which the origin is infected, τ_{BP} . In order to study this infection time for models on \mathbb{Z}^2 , the update families were classified by Bollobás et al. [8] into three universality classes: *supercritical*, *critical* and *subcritical*, according to a simple geometric criterion (see Definition 2.1). In [8] they proved that $q_c(\mathbb{Z}^2, \mathcal{U}) = 0$ if \mathcal{U} is supercritical or critical,

and it was proved by Balister et al. [4] that $q_c(\mathbb{Z}^2, \mathcal{U}) > 0$ if \mathcal{U} is subcritical. For supercritical update families, [8] proved that $\tau_{BP} = q^{-\Theta(1)}$ w.h.p. as $q \rightarrow 0$, while in the critical case $\tau_{BP} = \exp(q^{-\Theta(1)})$. The result for critical families was later improved by Bollobás et al. [7], who identified the critical exponent $\alpha = \alpha(\mathcal{U})$ such that $\tau_{BP} = \exp(q^{-\alpha+\rho(1)})$.

Back to KCM, if we fix an update family \mathcal{U} and an initial configuration ω and we identify the empty sites with infected sites, a first basic observation is that the clusters of sites that will never be infected in the \mathcal{U} -bootstrap percolation correspond to clusters of sites which are occupied and will never be emptied under the KCM dynamics. A natural issue is whether there is a direct connection between the infection mechanism of bootstrap percolation and the relaxation mechanism for KCM, and, more precisely, whether the scaling of T_{rel} and $\mathbb{E}(\tau_0)$ is connected to the typical value of τ_{BP} when the law of the initial infections is μ . It is not difficult to establish that $\mu(\tau_{BP})$ provides a lower bound for $\mathbb{E}(\tau_0)$ and T_{rel} (see [27, Lemma 4.3] and (10)), but in general, as we will explain, this lower bound does not provide the correct behaviour.

Martinelli et al. [26] proposed that the supercritical class should be refined into *unrooted* supercritical and *rooted* supercritical models in order to capture the richer behavior of KCM. For unrooted models the scaling is of the same type as for bootstrap percolation, $T_{rel} \sim \mathbb{E}(\tau_0) = q^{-\Theta(1)}$ as $q \rightarrow 0$ [26, Theorem 1(a)],¹ while for rooted models the divergence is much faster, $\mathbb{E}(\tau_0) \sim T_{rel} = e^{\Theta((\log q)^2)}$ (see [26, Theorem 1(b)] for the upper bound and [25, Theorem 4.2] for the lower bound).

Concerning the critical class, the lower bound with $\mu(\tau_{BP})$ mentioned above and the results of [8] on bootstrap percolation imply that T_{rel} and $\mathbb{E}(\tau_0)$ diverge at least as $\exp(q^{-\Theta(1)})$. In [26, Theorem 2] an upper bound of the same form was established and a conjecture [26, Conjecture 3] was put forward on the value of the critical exponent ν such that both $\mathbb{E}(\tau_0)$ and T_{rel} scale as $\exp(|\log q|^{O(1)}/q^\nu)$, with ν in general different from the exponent of the corresponding bootstrap percolation process. Furthermore, a toolbox was developed for the study of the upper bounds, leading to upper bounds matching this conjecture for all models. The main issue left open in [26] was to develop tools to establish sharp lower bounds. A first step in this direction was done by Marêché et al. [25] by analyzing a specific critical model known as the Duarte model for which the update family contains all the 2-elements subsets of the North, South and West neighbours of the origin. Theorem 5.1 of [25] establishes a sharp lower bound on the infection and relaxation times for the Duarte KCM that, together with the upper bound in [26, Theorem 2(a)], proves $\mathbb{E}^{Duarte}(\tau_0) = \exp(\Theta((\log q)^4/q^2))$ as $q \rightarrow 0$, and the same result holds for T_{rel} . The divergence is again much faster than for the corresponding bootstrap percolation model, for which it holds $\tau_{BP} = e^{\Theta((\log q)^2/q)}$ w.h.p as $q \rightarrow 0$ [30] (see also [6], from which the sharp value of the constant follows), namely the critical exponent for the Duarte KCM is twice the critical exponent for the Duarte bootstrap percolation.

Both for Duarte and for supercritical rooted models, the sharper divergence of time scales for KCM is due to the fact that the infection time of KCM is not well approximated by the infection mechanism of the monotone bootstrap percolation process, but

¹ For the lower bound of T_{rel} one does not need to use the bootstrap percolation results, as $T_{rel} \geq q^{-\min_{U \in \mathcal{U}} |U|/|\mathcal{U}|}$ by plugging the test function $\mathbb{1}_{\{\omega_0=0\}}$ in Definition 2.5.

is instead the result of a much more complex infection/healing mechanism. Indeed, visiting regions of the configuration space with an anomalous amount of empty sites is heavily penalised and requires a very long time to actually take place. The basic underlying idea is that the dominant relaxation mechanism is an *East-like dynamics* for large *droplets* of empty sites. Here East-like means that the presence of an empty droplet allows to empty (or fill) another adjacent droplet but only in a certain direction (or more precisely in a limited cone of directions). This is reminiscent of the relaxation mechanism for the East model, a prototype one-dimensional KCM for which x can be updated iff $x - 1$ is empty, thus a single empty site allows to create/destroy an empty site only on its right (see [15] for a review on the East model). For supercritical rooted models, the empty droplets that play the role of the single empty sites for East have a finite (model dependent) size, hence an equilibrium density $q_{\text{eff}} = q^{\Theta(1)}$. For the Duarte model, droplets have a size that diverges as $\ell = |\log q|/q$ and thus an equilibrium density $q_{\text{eff}} = q^\ell = e^{-(\log q)^2/q}$. Then a (very) rough understanding of the results of [25,26] is obtained by replacing q with q_{eff} in the time scale for the East model $T_{\text{rel}}^{\text{East}} = e^{\Theta((\log q)^2)}$ [2]. The main technical difficulty to translate this intuition into a lower bound is that the droplets cannot be identified with a rigid structure. In [25] this difficulty for the Duarte model was overcome by an algorithmic construction that allows to sequentially scan the system in search of sets of empty sites that *could* (without violating the constraint) empty a certain rigid structure. These are the droplets that play the role of the empty sites for the East dynamics.

In [26] all critical models which have an *infinite number of stable directions* (see Sect. 2.1), of which the Duarte model is but one example, were conjectured to have a critical exponent $\nu = 2\alpha$, with $\alpha = \alpha(\mathcal{U})$ the critical exponent of the corresponding bootstrap percolation dynamics (defined in Definition 2.2). The heuristic is the same as for the Duarte model, the only difference being that droplets would have in general size $\ell = |\log q|^{O(1)}/q^\alpha$. However, the technique developed in [25] for the Duarte model relies heavily on the specific form of the Duarte constraint and in particular on its oriented nature,² and it cannot be extended readily to this larger class.

In this work, together with the companion paper by Hartarsky et al. [20], we establish in full generality the universality classes for critical KCM, determining the critical exponent for each class.

Here we treat all choices of \mathcal{U} for which there is an *infinite number of stable directions* and prove (Theorem 2.8) a lower bound for T_{rel} and $\mathbb{E}(\tau_0)$ that, together with the matching upper bound of [26, Theorem 2], yields

$$\mathbb{E}(\tau_0) = e^{|\log q|^{O(1)}/q^{2\alpha}}$$

for $q \rightarrow 0$ and the same result for T_{rel} . Our technique is somewhat inspired by the algorithmic construction of [25], however, the nature of the droplets which move in an East-like way is here much more subtle, and in order to identify them we construct an algorithm which can be seen as a significant improvement on the α -covering and u -iceberg algorithms developed in the context of bootstrap percolation [7].

² Note that, since the Duarte update rules contain only the North, South and West neighbours of the origin, the constraint at a site x does not depend on the sites with abscissa larger than the abscissa of x .

In the companion paper [20] we prove for the complementary class of models, namely all critical models with a *finite number of stable directions*, an upper bound that (together with the lower bound from bootstrap percolation) yields instead

$$\mathbb{E}(\tau_0) = e^{|\log q|^{O(1)}/q^\alpha}$$

for $q \rightarrow 0$ and the same result for T_{rel} .

A comparison of our results with Conjecture 3 of [26] is due. The class that we consider here is, in the notation of [26], the class of models with *bilateral difficulty* $\beta = \infty$, hence belong to the α -rooted class defined therein. Therefore, our Theorem 2.8 proves Conjecture 3(a) in this case. We underline that it is not a limitation of our lower bound strategy that prevents us from proving Conjecture 3(a) for the other α -rooted models, namely those with $2\alpha \leq \beta < \infty$. Indeed, as it is proven in the companion paper [20], in this case the conjecture of [26] is not correct, since it did not take into account a subtle relaxation mechanism which allows to recover the same critical exponent as for the bootstrap percolation dynamics.

The plan of the paper is as follows. In Sect. 2 we develop the background for both KCM and bootstrap percolation needed to state our result, Theorem 2.8. In Sect. 3 we give a sketch of our reasoning and highlight the important points. In Sect. 4 we gather some preliminaries and notation. Section 5 is the core of the paper—there we define the central notions and establish their key properties, culminating in the Closure Proposition 5.20. In Sect. 6 we establish a connection between the KCM dynamics and an East dynamics and use this to wrap up the proof of Theorem 2.8. Finally, in Sect. 7 we discuss some open problems.

2 Models and background

2.1 Bootstrap percolation

Before turning to our models of interest, KCM, let us recall recent universality results for the intimately connected bootstrap percolation models in two dimensions. \mathcal{U} -bootstrap percolation (or simply bootstrap percolation) is a very general class of monotone transitive local cellular automata on \mathbb{Z}^2 first studied in full generality by Bollobás et al. [8]. Let \mathcal{U} , called *update family*, be a finite family of finite nonempty subsets, called *update rules*, of $\mathbb{Z}^2 \setminus \{0\}$. Let A , called the set of *initial infections*, be an arbitrary subset of \mathbb{Z}^2 . Then the \mathcal{U} -bootstrap percolation dynamics is the discrete time deterministic growth of infection defined by $A_0 = A$ and, for each $t \in \mathbb{N}$,

$$A_{t+1} = A_t \cup \{x \in \mathbb{Z}^2 : \exists U \in \mathcal{U}, U + x \subset A_t\}.$$

In other words, at any step each site becomes infected if a rule translated at it is already fully infected, and infections never heal. We define the *closure* of the set A by $[A] = \bigcup_{t \geq 0} A_t$ and we say that A is *stable* when $[A] = A$. The set of initial infections A is chosen at random with respect to the product Bernoulli measure μ with parameter $q \in [0, 1]$: for every $x \in \mathbb{Z}^2$ we have $\mu(x \in A) = q$.

Arguably, the most natural quantity to consider for these models is the typical (e.g. mean) value of τ_{BP} , the infection time of the origin.

The combined results of Bollobás et al. [8] and Balister et al. [4] yield a pre-universality partition of all update families into three classes with qualitatively different scalings of the median of the infection time as $q \rightarrow 0$. In order to define this partition we will need a few definitions.

For any unitary vector $u \in S^1 = \{z \in \mathbb{R}^2 : \|z\| = 1\}$ ($\|\cdot\|$ denotes the Euclidean norm in \mathbb{R}^2) and any vector $x \in \mathbb{R}^2$ we denote $\mathbb{H}_u(x) = \{y \in \mathbb{R}^2 : \langle u, y - x \rangle < 0\}$ — the open half-plane directed by u passing through x . We also set $\mathbb{H}_u = \mathbb{H}_u(0)$. We say that a direction $u \in S^1$ is *unstable* (for an update family \mathcal{U}) if there exists $U \in \mathcal{U}$ such that $U \subset \mathbb{H}_u$ and *stable* otherwise. The partition is then as follows.

Definition 2.1 (Definition 1.3 of [8]) An update family \mathcal{U} is

- *supercritical* if there exists an open semi-circle of unstable directions,
- *critical* if it is not supercritical, but there exists an open semi-circle with a finite number of stable directions,
- *subcritical* otherwise.

The main result of [8] then states that in the supercritical case $\tau_{BP} = q^{-\Theta(1)}$ with high probability as $q \rightarrow 0$, while in the critical one $\tau_{BP} = \exp(q^{-\Theta(1)})$. The final justification of the partition in Definition 2.1 was given by Balister et al. [4] who proved that the origin is never infected with positive probability for subcritical models for $q > 0$ sufficiently small, i.e. $q_c(\mathbb{Z}^2, \mathcal{U}) > 0$ if \mathcal{U} is subcritical. From the bootstrap percolation perspective supercritical models are rather simple, while subcritical ones remain very poorly understood (see [19]). Nevertheless, most of the non-trivial models considered before the introduction of \mathcal{U} -bootstrap percolation, including the 2-neighbour model (see [1,22] for further results), fall into the critical class, which is also the focus of our work.

Significantly improving the result of [8], Bollobás et al. [7] found the correct exponent determining the scaling of τ_{BP} for critical families. Moreover, they were able to find $\log \tau_{BP}$ up to a constant factor. To state their results we need the following crucial notion.

Definition 2.2 (Definition 1.2 of [7]) Let \mathcal{U} be an update family and $u \in S^1$ be a direction. Then the *difficulty* of u , $\alpha(u)$, is defined as follows.

- If u is unstable, then $\alpha(u) = 0$.
- If u is an isolated stable direction (isolated in the topological sense), then

$$\alpha(u) = \min\{n \in \mathbb{N} : \exists K \subset \mathbb{Z}^2, |K| = n, |[\mathbb{Z}^2 \cap (\mathbb{H}_u \cup K)] \setminus \mathbb{H}_u| = \infty\}, \quad (1)$$

i.e. the minimal number of infections allowing \mathbb{H}_u to grow infinitely.

- Otherwise, $\alpha(u) = \infty$.

We define the *difficulty* of \mathcal{U} by

$$\alpha(\mathcal{U}) = \inf_{C \in \mathcal{C}} \sup_{u \in C} \alpha(u), \quad (2)$$

where $\mathcal{C} = \{\mathbb{H}_u \cap S^1 : u \in S^1\}$ is the set of open semi-circles of S^1 .

It is not hard to see (Theorem 1.10 of [8], Lemma 2.6 of [7]) that the set of stable directions is a finite union of closed intervals of S^1 and that (Lemmas 2.7 and 2.10 of [7]) (1) also holds for unstable and *strongly stable* directions, that is directions in the interior of the set of stable directions (but not for *semi-isolated stable* directions i.e. endpoints of non-trivial stable intervals). Furthermore (see [7, Lemma 2.7], [8, Lemma 5.2]), $1 \leq \alpha(u) < \infty$ if and only if u is an isolated stable direction, so that \mathcal{U} is critical if and only if $1 \leq \alpha(\mathcal{U}) < \infty$. As a final remark we recall that, contrary to determining whether an update family is critical, finding $\alpha(\mathcal{U})$ is a NP-hard question [21].

We are now ready to describe the universality results. A weaker form of the result of [7] is that $\tau_{BP} = \exp(q^{-\alpha(\mathcal{U})+o(1)})$ with high probability as $q \rightarrow 0$. For the full result however, we need one last definition.

Definition 2.3 A critical update family \mathcal{U} is *balanced* if there exists a closed semi-circle C such that $\max_{u \in C} \alpha(u) = \alpha(\mathcal{U})$ and *unbalanced* otherwise.

Then [7] provides that for balanced models $\tau_{BP} = \exp(\Theta(1)/q^{\alpha(\mathcal{U})})$ with high probability as $q \rightarrow 0$, while for unbalanced ones $\tau_{BP} = \exp(\Theta((\log q)^2)/q^{\alpha(\mathcal{U})})$. These are the best general estimates currently known. We refer to [28,29] for recent surveys on these results as well as on sharper results for some specific models.

2.2 Kinetically constrained models

Returning to KCM, let us first define the general class of KCM introduced by Cancrini et al. [9] directly on \mathbb{Z}^2 . Fix a parameter $q \in [0, 1]$ and an update family \mathcal{U} as in the previous section. The corresponding KCM is a continuous-time Markov process on $\Omega = \{0, 1\}^{\mathbb{Z}^2}$ which can be informally defined as follows. A configuration ω is defined by assigning to each site $x \in \mathbb{Z}^2$ an *occupation variable* $\omega_x \in \{0, 1\}$ corresponding to an *empty* (or *infected*) and *occupied* (or *healthy*) site respectively. Each site waits an independent exponentially distributed time with mean 1 before attempting to update its occupation variable. At that time, if the configuration is completely empty on at least one update rule translated at x , i.e. if $\exists U \in \mathcal{U}$ such that $\omega_y = 0$ for all $y \in U + x$, then we perform a *legal update* or *legal spin flip* by setting ω_x to 0 with probability q and to 1 with probability $1 - q$. Otherwise the update is discarded. Since the constraint to allow the update never depends on the state of the to-be-updated site, the product measure μ is a reversible invariant measure and the process started at μ is stationary. More formally, the KCM is the Markov process on Ω with generator \mathcal{L} acting on local functions $f: \Omega \mapsto \mathbb{R}$ as

$$(\mathcal{L}f)(\omega) = \sum_{x \in \mathbb{Z}^2} c_x(\omega) (\mu_x(f) - f)(\omega), \tag{3}$$

for any $\omega \in \Omega$, where $\mu_x(f)$ denotes the average of f when the occupation variable at x has law $\text{Ber}(1 - q)$ and the other occupation variables are set to $\{\omega_y\}_{y \neq x}$, and

c_x is the indicator function of the event that there exists $U \in \mathcal{U}$ such that $U + x$ is completely empty, i.e. $\omega_{U+x} \equiv 0$. We refer the reader to chapter I of [24], where the general theory of interacting particle systems is detailed, for a precise construction of the Markov process and the proof that \mathcal{L} is the generator of a reversible Markov process $\{\omega(t)\}_{t \geq 0}$ on Ω with reversible measure μ .

The corresponding Dirichlet form is defined as

$$\mathcal{D}(f) = \sum_{x \in \mathbb{Z}^2} \mu(c_x \text{Var}_x(f)), \tag{4}$$

where $\text{Var}_x(f)$ denotes the variance of the local function f with respect to the variable ω_x conditionally on $\{\omega_y\}_{y \neq x}$. The expectation with respect to the stationary process with initial distribution μ will be denoted by $\mathbb{E} = \mathbb{E}_\mu^{q, \mathcal{U}}$. Finally, given a configuration $\omega \in \Omega$ and a site $x \in \mathbb{Z}^2$, we will denote by ω^x the configuration obtained from ω by *flipping* site x , namely by setting $(\omega^x)_x = 1 - \omega_x$ and $(\omega^x)_y = \omega_y$ for all $y \neq x$. For future use we also need the following definition of legal paths, that are essentially sequences of configurations obtained by successive legal updates.

Definition 2.4 (*Legal path*) Fix an update family \mathcal{U} , then a *legal path* γ in Ω is a finite sequence $\gamma = (\omega_{(0)}, \dots, \omega_{(k)})$ such that, for each $i \in \{1, \dots, k\}$, the configurations $\omega_{(i-1)}$ and $\omega_{(i)}$ differ by a legal (with respect to the choice of \mathcal{U}) spin flip at some vertex $v = v(\omega_{(i-1)}, \omega_{(i)})$.

As mentioned in Sect. 1, our goal is to prove sharp bounds on the characteristic time scales of critical KCM. Let us start by defining precisely these time scales, namely the relaxation time T_{rel} (or inverse of the spectral gap) and the mean infection time $\mathbb{E}(\tau_0)$ (with respect to the stationary process).

Definition 2.5 (*Relaxation time T_{rel}*) Given an update family \mathcal{U} and $q \in [0, 1]$, we say that $C > 0$ is a *Poincaré constant* for the corresponding KCM if, for all local functions f , we have

$$\text{Var}_\mu(f) = \mu(f^2) - \mu(f)^2 \leq C \mathcal{D}(f). \tag{5}$$

If there exists a finite Poincaré constant, we define

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

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 (see e.g. [3, Section 2.1]).

Definition 2.6 (*Infection time τ_0*) The random time τ_0 at which the origin is first infected is given by

$$\tau_0 = \inf \{t \geq 0 : \omega_0(t) = 0\},$$

where we adopt the usual notation letting $\omega_0(t)$ be the value of the configuration $\omega(t)$ at the origin, namely $\omega_0(t) = (\omega(t))_0$.

The East model We close this section by defining a specific example of KCM on \mathbb{Z} , the East model of Jäckle and Eisinger [23], which will be crucial to understand our results (KCM on \mathbb{Z} are defined in the same way as KCM on \mathbb{Z}^2). It is defined by an update family composed by a single rule containing only the site to the left of the origin (-1) . In other words, site x can be updated iff $x - 1$ is empty. For this model both T_{rel} and $\mathbb{E}(\tau_0)$ scale as $\exp\left(\frac{(\log q)^2}{2 \log 2}\right)$ as $q \rightarrow 0$, see [2,9,12].³ One of the key ingredients behind this scaling is the following combinatorial result [32] (see [14, Fact 1] for a more mathematical formulation).

Proposition 2.7 *Consider the East model on $\{1, \dots, M\}$ defined by fixing $\omega_0 = 0$ at all time. Then any legal path γ connecting the fully occupied configuration (namely ω s.t. $\omega_x = 1$ for all $x \in \{1, \dots, M\}$) to a configuration ω' such that $\omega'_M = 0$ goes through a configuration with at least $\lceil \log_2(M + 1) \rceil$ empty sites.*

This logarithmic ‘energy barrier’, to employ the physics jargon, and the fact that at equilibrium the typical distance to the first empty site is $M = \Theta(1/q)$ are responsible for the divergence of the time scales as roughly $1/q^{\lceil \log_2(M+1) \rceil} = e^{\Theta((\log q)^2)}$.

2.3 Result

In this paper we study critical KCM with an infinite number of stable directions or, equivalently, with a non-trivial interval of stable directions. Recall that \mathbb{E} denotes the expectation with respect to the stationary KCM process.

Theorem 2.8 *Let \mathcal{U} be a critical update family with an infinite number of stable directions. Then there exists a sufficiently large constant $C > 0$ such that*

$$\mathbb{E}(\tau_0) \geq \exp\left(1 / \left(Cq^{2\alpha(\mathcal{U})}\right)\right),$$

as $q \rightarrow 0$ and the same asymptotics holds for T_{rel} .

This theorem combined with the upper bound of Martinelli et al. [26, Theorem 2(a)], determines the critical exponent of these models to be 2α in the sense of Corollary 2.9 below. We thus complete the proof of universality and Conjecture 3(a) of [26] for these models.⁴

Corollary 2.9 *Let \mathcal{U} be a critical update family with an infinite number of stable directions. Then*

$$q^{2\alpha(\mathcal{U})} \log \mathbb{E}(\tau_0) = (-\log q)^{O(1)}$$

³ Actually these references focus on the study of T_{rel} . A matching upper bound for $\mathbb{E}(\tau_0)$ follows from (10). The lower bound for $\mathbb{E}(\tau_0)$ follows easily from the lower bound for $\mathbb{P}(\tau_0 > t)$ with $t = \exp(\log(q)^2/2 \log 2)$ obtained in the proof of Theorem 5.1 of [11].

⁴ The conjecture involuntarily asks for a positive power of $\log q$, which we do not expect to be systematically present (see Conjecture 7.1).

as $q \rightarrow 0$ and the same holds for T_{rel} .

Universality for the remaining critical models is proved in a companion paper by Hartarsky et al. [20] and, in particular, Conjecture 3(a) of [26] is disproved for models other than those covered by Theorem 2.8. It is important to note that Theorem 2.8 significantly improves the best known results for all models with the exception of the recent result of Marêché et al. [25] for the Duarte model. Indeed, the previous bound had exponent α , and was proved via the general (but in this case far from optimal) lower bound with the mean infection time for the corresponding bootstrap percolation model [27, Lemma 4.3].

3 Sketch of the proof

In this section we outline roughly the strategy to derive our main result, Theorem 2.8. The hypothesis of infinite number of stable directions provides us with an interval of stable directions. We can then construct stable ‘droplets’ of shape as in Fig. 3 (see Definitions 5.5 and 5.6), where we recall from Sect. 2.1 that a set is stable if it coincides with its closure. Thus, if all infections are initially inside a droplet, this will be true at any time under the KCM dynamics. The relevance and advantage of such shapes come from the fact that only infections situated to the left of a droplet can induce growth left. This is manifestly not feasible without the hypothesis of having an interval of stable directions. It is worth noting that these shapes, which may seem strange at first sight, are actually very natural and intrinsically present in the dynamics. Indeed, such is the shape of the stable sets for a representative model of this class – the modified 2-neighbour model with one (any) rule removed, that is the three-rule update family with rules $\{(-1, 0), (0, 1)\}, \{(-1, 0), (0, -1)\}, \{(0, -1), (1, 0)\}$ (it can also be seen as the modified Duarte model with an additional rule). The stable sets in this case are actually Young diagrams.

We construct a collection of such droplets covering the initial configuration of infections, so that it gives an upper bound on the closure. To do this, we devise an improvement of the α -covering algorithm of Bollobás et al. [7]. It is important for us not to overestimate the closure as brutally. Indeed, a key step and the main difficulty of our work is the Closure Proposition 5.20, which roughly states that the collections of droplets associated to the closure of the initial infections is equal to the collection for the initial infections. This is highly non-trivial, as in order not to overshoot in defining the droplets, one is forced to ignore small patches of infections (larger than the ones in [7]), which can possibly grow significantly when we take the closure for the bootstrap percolation process and especially so if they are close to a large infected droplet. In order to remedy this problem, we introduce a relatively intrinsic notion of ‘crumb’ (see Definition 5.1) such that its closure remains one and does not differ too much from it. A further advantage of our algorithm for creating the droplets over the one of [7] is that it is somewhat canonical, with a well-defined unique output, which has particularly nice ‘algebraic’ description and properties (see Remark 5.10). Another notable difficulty we face is systematically working in roughly a half-plane

(see Remark 5.21 for generalisations) with a fully infected boundary condition, but we manage to extend our reasoning to this setting very coherently.

Finally, having established the Closure Proposition 5.20 alongside standard and straightforward results like an Aizenmann-Lebowitz Lemma 5.13 and an exponential decay of the probability of occurrence of large droplets (Lemma 5.15), we finish the proof via the following approach, inspired by the one developed by Marêché et al. [25] for the Duarte model. The key step here (see Sect. 6) is mapping the KCM legal paths to those of an East dynamics via a suitable renormalisation. Roughly speaking, we say that a renormalised site is empty if it contains a large droplet of infections. However, for the renormalised configuration to be mostly invariant under the original KCM dynamics, we rather look for the droplets in the closure of the original set of infections instead. This is where the Closure Proposition 5.20 is used to compensate the fact that the closure of equilibrium is not equilibrium. In turn, this mapping together with the combinatorial result for the East model recalled in Sect. 2.2 (Proposition 2.7), yield a bottleneck for our dynamics corresponding to the creation of $\log(1/q_{\text{eff}})$ droplets, where $1/q_{\text{eff}}$ is the equilibrium distance between two empty sites in the renormalized lattice, and $q_{\text{eff}} \sim e^{-1/q^\alpha}$. This provides for the time scales the desired lower bound $q_{\text{eff}}^{\log(q_{\text{eff}})} \sim e^{1/q^{2\alpha}}$ of Theorem 2.8. The last part of the proof follows very closely the ideas put forward in [25] for the Duarte model. However, in [25], there was no need to develop a subtle droplet algorithm since, owing to the oriented character of the Duarte constraint, droplets could simply be identified with some large infected vertical segments. It is also worth noting that, thanks to the less rigid notion of droplets that we develop in the general setting, some of the difficulties faced in [25] for Duarte are no longer present here.

4 Preliminaries and notation

Let us fix a critical update family \mathcal{U} with an infinite number of stable directions for the rest of the paper. We will omit \mathcal{U} from all notation, such as $\alpha(\mathcal{U})$.

Directions The next lemma establishes that one can make a suitable choice of 4 stable directions, which we will use for all our droplets. At this point the statement should look very odd and technical, but it simply reflects the fact that we have a lot of freedom for the choice and we make one which will simplify a few of the more technical points in later stages. Nevertheless, this is to a large extent not needed besides for concision and clarity.

A direction $u \in S^1$ is called *rational* if $\tan u \in \mathbb{Q} \cup \{\infty\}$.

Lemma 4.1 *There exists a set of rational stable directions $\mathcal{S} = \{u_1, u_2, v_1, v_2\}$ with difficulty at least α such that (see Fig. 1)*

- The directions appear in counterclockwise order u_1, u_2, v_1, v_2 .
- No $u \in \mathcal{S}$ is a semi-isolated stable direction.
- u_{3-i} belongs to the cone spanned by v_i and u_i for $i \in \{1, 2\}$ i.e. the strictly smaller interval among $[v_i, u_i]$ and $[u_i, v_i]$ contains u_{3-i} .

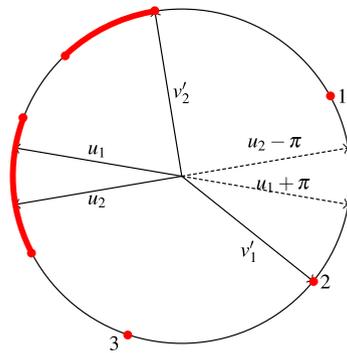


Fig. 1 Illustration of Lemma 4.1 and its proof. Thickened arcs represent intervals of strongly stable directions. Solid dots represent isolated and semi-isolated stable directions. The difficulties of the isolated stable directions are indicated next to them and yield that the difficulty of the model is $\alpha = 2$. The directions chosen in Lemma 4.1 are the solid vectors $u_1, u_2, v_1 = v_1'$ and a direction v_2 in the strongly stable interval ending at v_2' sufficiently close to v_2' . Note that the definition of v_2' (and v_1') disregards stable directions with difficulty smaller than α as present on the figure

- 0 is contained in the interior of the convex envelope of \mathcal{S} .
- Either $u_2 < v_1 - \pi/2$ or $u_1 > v_2 + \pi/2$.
- $(\mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}) \cap \mathbb{Z}^2$ is stable or, equivalently, $\nexists U \in \mathcal{U}, U \subset \mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}$.
- the directions

$$\begin{aligned} u' &= (u_1 + u_2)/2, \\ u'_1 &= (3u_1 + u_2)/4, \\ u'_2 &= (u_1 + 3u_2)/4 \end{aligned}$$

are rational.

Proof Since \mathcal{U} has an infinite number of stable directions and they form a finite union of closed intervals with rational endpoints [8, Theorem 1.10], there exists a non-empty open interval I''' of stable directions. Further note that the set J of directions u such that there exists a rule $U \in \mathcal{U}$ and $x \in U$ with $\langle x, u \rangle = 0$ is finite, so one can find a non-trivial closed subinterval $I'' \subset I'''$ which does not intersect J . The directions u_1 and u_2 will be chosen in I'' , which clearly implies that they are strongly stable and thus with infinite difficulty. Moreover, if there exists $U \in \mathcal{U}$ with $U \subset \mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}$, by stability of u_2 , we have $U \cap (\mathbb{H}_{u_1} \setminus \mathbb{H}_{u_2}) \neq \emptyset$, which contradicts $I'' \cap J = \emptyset$.

Since \mathcal{U} is critical it does not have two opposite strongly stable directions, so there is no strongly stable direction in $I'' + \pi$. If there are any (isolated or semi-isolated) stable directions in $I'' + \pi$, we can further choose a non-trivial open subinterval $I' \subset I''$, for which this is not the case (there is a finite number of isolated and semi-isolated stable directions). Let $\pi > \delta > 0$ be such that the angle between any two consecutive directions of difficulty at least α is at most $\pi - \delta$ (it is well defined by (2)). We then choose a non-trivial closed subinterval $I' \supset I = [u_1, u_2]$ with u_1 rational and $u'_1 = (3u_1 + u_2)/4$ rational and with $0 < u_2 - u_1 < \delta < \pi$. It easily follows

from the sum and difference formulas for the tangent function that u' , u'_2 and u_2 are also rational.

Let

$$v'_1 = \max\{v \in (u_2, u_1 + \pi) : \alpha(v) \geq \alpha\},$$

$$v'_2 = \min\{v \in (u_2 - \pi, u_1) : \alpha(v) \geq \alpha\}.$$

These both exist, since $I + \pi$ does not contain stable directions, both $(u_2, u_2 + \pi)$ and $(u_1 - \pi, u_1)$ contain directions with difficulty at least α by (2) and the set of such directions is closed. If v'_1 is not semi-isolated, we set $v_1 = v'_1$ and similarly for v_2 . Otherwise, we choose a rational strongly stable direction sufficiently close to v'_1 as v_1 and similarly for v_2 . We claim that this choice satisfies all the desired conditions. Indeed, all directions in \mathcal{S} are stable non-semi-isolated rational with difficulty at least α and the last but one condition was already verified.

One does have that u_1 is in the cone spanned by v_2 and u_2 , which is implied by $v_2 \in (u_2 - \pi, u_1)$ and similarly for u_2 , so the third condition is also verified. If $v'_2 - v'_1 \geq \pi$, then there is an open half circle contained in (v'_1, v'_2) with no direction of difficulty at least α , which contradicts (2), so $v_2 - v_1 < \pi$ and the same holds for $u_1 - v_2, u_2 - u_1$ and $v_1 - u_2$ by the definition of v'_1 and v'_2 , the fact that v_1 and v_2 are sufficiently close to them and the fact that I was chosen smaller than π . Thus 0 is in the convex envelope of \mathcal{S} .

Finally, if one has both $v_1 - u_2 \leq \pi/2$ and $u_1 - v_2 \leq \pi/2$, then one obtains $v'_2 - v'_1 > \pi - \delta$, since I is smaller than δ . However, v'_1 and v'_2 are consecutive directions of difficulty at least α , which contradicts the definition of δ . □

Notation For the rest of the paper we fix directions $\mathcal{S} = \{u_1, u_2, v_1, v_2\}$ as in Lemma 4.1 and assume without loss of generality that $u_2 < v_1 - \pi/2$.

Let us fix large constants

$$1 \ll C_1 \ll C'_2 \ll C_2 \ll C_3 \ll C'_4 \ll C_4 \ll C_5,$$

each of which can depend on previous ones as well as on \mathcal{U} and \mathcal{S} . We will also use asymptotic notation whose constants can depend on \mathcal{U} and \mathcal{S} , but not on C_1 or the other constants above. All asymptotic notation is with respect to $q \rightarrow 0$, so we assume throughout that $q > 0$ is sufficiently small.

For any two sets $K, \partial \subset \mathbb{R}^2$ we define $[K]_\partial = [(K \cup \partial) \cap \mathbb{Z}^2] \setminus \partial$.

Finally, we make the convention that throughout the article all distances, balls and diameters are Euclidean unless otherwise stated. We say that a set $X \subset \mathbb{R}^2$ is *within distance* δ of a set $Y \subset \mathbb{R}^2$ if $d(x, Y) \leq \delta$ for all $x \in X$ where d is the Euclidean distance.

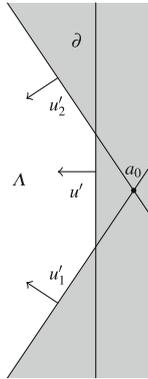


Fig. 2 The open domain ∂ defined in (6) is shaded, while its complement Λ is not. The lines are the boundaries of the three half-planes defining ∂ . Note that if $a_0 \notin \mathbb{H}_{u'}$, then Λ becomes simply a cone

5 Droplet algorithm

In this section we define our main tool—the droplet algorithm. It can be seen as a significant improvement on the α -covering and u -iceberg algorithms [7, Definitions 6.6 and 6.22], many of whose techniques we adapt to our setting.

We will work in an infinite domain Λ defined as follows (see Fig. 2). Fix some vector $a_0 \in \mathbb{R}^2$ and let

$$\begin{aligned} \partial &= \mathbb{H}_{u'} \cup \mathbb{H}_{u'_1}(a_0) \cup \mathbb{H}_{u'_2}(a_0), \\ \Lambda &= \mathbb{R}^2 \setminus \partial, \end{aligned} \tag{6}$$

where the directions u' , u'_1 and u'_2 are those defined in Lemma 4.1. In other words, Λ is a cone with sides perpendicular to u'_1 and u'_2 cut along a line perpendicular to u' . The reader is invited to simply think that ∂ is a half-plane directed by u' , which will not change the reasoning.

5.1 Clusters and crumbs

Let Γ be the graph with vertex set \mathbb{Z}^2 but with $x \sim y$ if and only if $\|x - y\| \leq C_2$. Let Γ' be defined similarly with C_2 replaced by C'_2 . Given a finite $K \subset \Lambda \cap \mathbb{Z}^2$, we say that $\kappa \subset K$ is a *connected component* of K in Γ if the subgraph of Γ induced by the vertex set κ is connected and there do not exist vertices $x \in K \setminus \kappa$ and $y \in \kappa$ such that $x \sim y$ in Γ .

Crumbs For a given finite set $K \subset \Lambda \cap \mathbb{Z}^2$ of infections we would like to have a notion of a connected component being ‘big’ or ‘small’. ‘Small’ components will be dubbed ‘crumbs’ and will play a negligible perturbative role in the bootstrap percolation process, by inducing only ‘very localised’ growth and being ‘well isolated’ from the rest of the infections. A sufficient condition for this, as identified in [7], is that $|\kappa| < \alpha$.

However, contrary to what was the case in [7], we need the notion of ‘crumb’ to be stable under the closure (with respect to the bootstrap percolation process), i.e. the closure of a ‘crumb’ to still be a ‘crumb’. We thus identify as ‘crumb’ any component, which is the closure of a set of size less than α . Also taking into account the boundary, this leads us to the following notion.

Definition 5.1 (*Crumb*) Fix a finite set $K \subset \Lambda \cap \mathbb{Z}^2$ and let κ be a connected component of K in Γ . We say that κ is a *crumb* for K if the following conditions hold.

- For all $x \in \kappa$ we have $d(x, \partial) > C_2$.
- There exists a set $P_\kappa \subset \mathbb{Z}^2$ such that $[P_\kappa] \supset \kappa$ and $|P_\kappa| = \alpha - 1$.

First properties of crumbs It follows from the definition that a crumb κ for K is at distance more than C_2 from $\partial \cup (K \setminus \kappa)$. Moreover, the closure of a crumb is within bounded distance from the crumb, as we shall see in Corollary 5.17 (see Fig. 5a). Also, crumbs have diameters much smaller than C_3 , as we shall see in Corollary 5.17. The proofs of this corollary and Observation 5.16, which it follows from, are both independent of the rest of the argument and are only postponed for convenience. Nevertheless, we allow ourselves to use these (easy) results ahead of their proofs.

These properties justify and quantify the intuition that crumbs are ‘small’, that they only grow ‘locally’, and it is clear that (if we disregard the boundary) the closure of a crumb is a crumb.

Modified crumbs Unfortunately, if K is the union of two crumbs at distance slightly larger than C_2 , it is not necessarily true that $[K]$ is still composed of crumbs (recall that, albeit locally, crumbs can grow under the bootstrap percolation process), which can be disastrous. This is the reason for introducing ‘modified crumbs’ with $C'_2 \ll C_2$, so that in the scenario above all connected components of $[K]$ in Γ' are ‘modified crumbs’ (there may now be more than two of them).

Definition 5.2 (*Modified crumb*) We define a *modified crumb* by replacing in Definition 5.1 Γ by Γ' and C_2 by C'_2 .

In the sequel we will encounter more ‘modified’ notions and constants (like C'_2). These will be applied to K equal to the closure $[K']_\partial$ of some K' , which is our initial set of infections. Our ultimate goal is to ensure that simply using these modified notions based on (much smaller) modified constants will compensate the closure operation.

Clusters We next consider connected components which are not crumbs. Since they can be very large (particularly so if we are working with the closure of a set), we cut them up into (possibly overlapping) pieces termed ‘clusters’, which have bounded size. Roughly speaking, a ‘cluster’ is any ‘big, but not too big’ connected set of infections.

Definition 5.3 (*Cluster*) Fix a finite set $K \subset \Lambda \cap \mathbb{Z}^2$. Let κ be a connected component of K in Γ which is not a crumb. We say that a subset C of κ is a *cluster* for K if the following conditions hold.

- $\text{diam}(C) \leq C_3$.
- C is connected in Γ (i.e. C is a connected component of C in Γ).
- Either $C = \kappa$ or for all $x \in \kappa \setminus C$ and $y \in C$ such that $x \sim y$ in Γ we have $\text{diam}(C \cup \{x\}) > C_3$.

A cluster is called *boundary cluster* if it is at distance at most C_2 from ∂ . For a cluster C we denote by $Q(C)$ the smallest open quadrilateral with sides perpendicular to \mathcal{S} containing the set $\{x \in \mathbb{R}^2 : d(x, C) < C_4\}$.

We similarly define *modified cluster* and *modified boundary cluster* by replacing Γ by Γ' and C_2 by C'_2 . For a cluster or modified cluster C we denote by $Q'(C)$ the smallest open quadrilateral with sides perpendicular to \mathcal{S} containing the set $\{x \in \mathbb{R}^2 : d(x, C) < C'_4\}$.

Identifying clusters and crumbs In order to identify the clusters and crumbs of K , one may proceed as follows. Determine the connected components of K in Γ and consider each of them separately. For a given component κ first check if it is at distance at most C_2 from ∂ . If so, then it is not a crumb and will give rise to clusters. If not, then check if κ is the closure of at most $\alpha - 1$ sites. If this second verification succeeds, then κ is determined to be a crumb and, as mentioned above, it must have diameter much smaller than C_3 .

If κ is thus determined not to be a crumb, we proceed to identify its clusters. If $\text{diam}(\kappa) \leq C_3$, then there is a single cluster— κ —and we are done. If not, we construct the clusters of κ by the following algorithm. Initialise the set $C = \emptyset$. If there exists $y \in \kappa \setminus C$ such that $C \cup \{y\}$ is connected in Γ and has diameter at most C_3 , then replace C by $C \cup \{y\}$ and repeat. If several such y exist, then we do this for each possible y in parallel. The clusters containing x are all possible sets C obtained via this algorithm to which no y can be added.

In particular, this provides us with a partition of K into well separated crumbs, single clusters equal to their corresponding connected component and sets of overlapping clusters whose union is a connected component of diameter larger than C_3 .

First properties of clusters Following the algorithm above, we obtain some basic properties of clusters.

Observation 5.4 *Let C be a non-boundary cluster or non-boundary modified cluster for a finite $K \subset \Lambda \cap \mathbb{Z}^2$. Then $|C| \geq \alpha$.*

Proof Let κ be the connected component of K in Γ containing C . If $\text{diam}(\kappa) \leq C_3$, then $C = \kappa$ and κ would be a crumb if we had $|\kappa| \leq \alpha - 1$, by taking $P_\kappa \supset \kappa$. If, on the contrary, $\text{diam}(\kappa) > C_3$, then $\text{diam}(C) \geq C_3 - C_2$ (by the third condition of Definition 5.3) and we can choose C_3 large enough to have $\frac{C_3 - C_2}{C_2} \geq \alpha$. □

Finally, for every cluster C we have $\text{diam}(C) \leq C_3$, so C intersects at most $2^{5C_3^2}$ other clusters. Also, $Q(C) \supset [C]$, since $Q(C) \cap \mathbb{Z}^2 \supset C$ is stable. Furthermore, we have $\text{diam}(Q(C)) = \Theta(C_4)$, as $\text{diam}(C) \leq C_3$. Analogous statements hold for modified clusters.

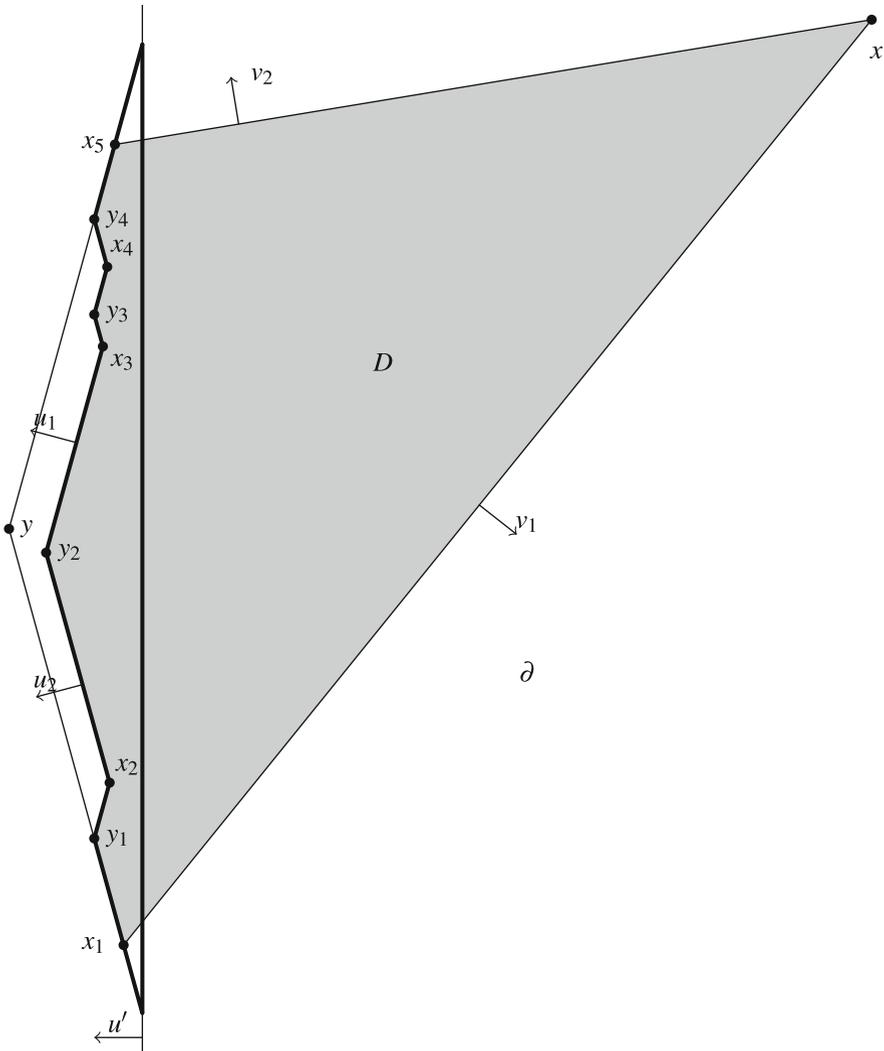


Fig. 3 The shaded region D is a distorted Young diagram (DYD) as in Definition 5.5. The larger quadrilateral with vertices x, x_1, y and x_5 is $Q(D)$. Note that $Q(D)$ can degenerate into a triangle, but we call it a quadrilateral nevertheless. On the figure $|D|$ is the length of the v_1 side, but this is not always the case. The thickened region is the cut distorted Young diagram (CDYD) $C(D)$ of D . The vertical line is the boundary between Λ on its left and ∂ on its right

5.2 Distorted Young diagrams

We now define the shape that our ‘droplets’ will have, which resembles Young diagrams.⁵ The following definitions are illustrated in Fig. 3.

⁵ For the 3-rule model alluded to in Sect. 3 stable sets consist precisely of Young diagrams and the directions \mathcal{S} provided by Lemma 4.1 can be arbitrarily close to the four axis directions, yielding Young diagrams.

Definition 5.5 (DYD) A distorted Young diagram (DYD) is a subset of \mathbb{R}^2 of the form

$$(\mathbb{H}_{v_1}(x) \cap \mathbb{H}_{v_2}(x)) \cap \bigcap_{i \in I} (\mathbb{H}_{u_1}(x_i) \cup \mathbb{H}_{u_2}(x_i)) \tag{7}$$

for a finite set I , some set $X = \{x_i : i \in I\}$ of vectors $x_i \in \mathbb{R}^2$ and $x \in \mathbb{R}^2$. The vectors x_i and x are uniquely defined up to redundancy (and up to the convention that all x_i are on the topological boundary of the DYD). Alternatively, a DYD can also be defined by

$$(\mathbb{H}_{v_1}(x) \cap \mathbb{H}_{v_2}(x)) \cap \bigcup_{i \in I} (\mathbb{H}_{u_1}(y_i) \cap \mathbb{H}_{u_2}(y_i)), \tag{8}$$

where y_i are the convex corners of the diagram rather than the concave ones.

For any DYD D we denote by y the vector such that

$$\langle y, u_j \rangle = \sup_{a \in D} \langle a, u_j \rangle = \max_{i \in I} \langle y_i, u_j \rangle$$

for $j \in \{1, 2\}$. We further denote

$$Q(D) = \mathbb{H}_{u_1}(y) \cap \mathbb{H}_{u_2}(y) \cap \mathbb{H}_{v_1}(x) \cap \mathbb{H}_{v_2}(x),$$

i.e. the minimal open quadrilateral containing D with sides directed by \mathcal{S} . In these terms, for any cluster or modified cluster C we have that $Q(C)$ and $Q'(C)$ are DYD, $Q(Q(C)) = Q(C)$ and $Q(Q'(C)) = Q'(C)$.

Definition 5.6 (CDYD) A cut distorted Young diagram (CDYD) is a subset of \mathbb{R}^2 of the form

$$\Lambda \cap (\mathbb{H}_{u_1}(y) \cap \mathbb{H}_{u_2}(y)) \cap \bigcap_{i \in I} (\mathbb{H}_{u_1}(x_i) \cup \mathbb{H}_{u_2}(x_i))$$

for a finite set I and some vectors $x_i \in \mathbb{R}^2$ and $y \in \Lambda$. Alternatively, one can write

$$\Lambda \cap \bigcup_{i \in I} (\mathbb{H}_{u_1}(y_i) \cap \mathbb{H}_{u_2}(y_i)),$$

where $y_i \in \Lambda$ are the convex corners.

For a DYD, D , we denote by $C(D)$ the CDYD defined by the same x_i and y or the same y_i . We extend the notation $C(D)$ to CDYD by setting $C(D) = D$ if D is a CDYD. Note that by Lemma 4.1 all DYD and CDYD are stable for the bootstrap percolation dynamics (restricted to Λ). Also pay attention to the fact that CDYD are not necessarily connected, contrary to DYD.

Definition 5.7 (*Size*) For a DYD D we set $\pi(D) = \{x \in \mathbb{R}^2 : \exists y \in D, \langle y, v_1 + \pi/2 \rangle = x\}$ to be its *projection* (parallel to v_1) and $|D| = \sup \pi(D) - \inf \pi(D)$ to be its *size* – the length of the projection. For a CDYD D we denote its *size* $\text{diam}(D)/C_1$ by $|D|$.

Note that if D is a DYD, then $|D| = |Q(D)|$ by Lemma 4.1 and the assumption we made that $u_2 < v_1 - \pi/2$. Furthermore, for all DYD $\text{diam}(D) = \Theta(|D|)$ again by Lemma 4.1 with constants depending only on \mathcal{S} . One should be careful with the meaning of size for disconnected CDYD, but it will not cause problems, as all CDYD arising in our forthcoming algorithm are connected.

Observation 5.8 *Note that for any $d \geq 1$ the number of discretised DYD and CDYD (i.e. intersections of a DYD or CDYD with \mathbb{Z}^2) containing a fixed point $a \in \mathbb{R}^2$ of diameter at most d is less than c^d for some constant c depending only on \mathcal{S} .*

Proof Note that a DYD or CDYD is uniquely determined by its rugged edge formed by its u_1 and u_2 -sides. However, this edge injectively defines an oriented percolation path with directions perpendicular to u_1 and u_2 on the lattice

$$\{x \in \mathbb{R}^2 : \exists x_1, x_2 \in \mathbb{Z}^2, \langle x, u_1 \rangle = \langle x_1, u_1 \rangle, \langle x, u_2 \rangle = \langle x_2, u_2 \rangle\}$$

(except its endpoints, which lie on similar lattices). Since the graph-length of this path is bounded by $O(d)$ and its endpoints are within distance d from a , the result follows. □

5.3 Span

We next introduce a procedure of merging DYD and CDYD. This will be used only for couples of intersecting ones, but can be defined regardless of whether they intersect. The operation is illustrated in Fig. 4.

Lemma 5.9 *For any two DYD, D_1 and D_2 , the minimal DYD containing $D_1 \cup D_2$ is well defined. We denote it by $D_1 \vee D_2$ and call it their span. The operation \vee is associative⁶ and commutative.*

Proof Let D_1 be defined by $Y^1 = \{y_i^1 : i \in I\}, x^1$ (see (8)) and similarly for D_2 . Let $x \in \mathbb{R}^2$ be the vector such that $\mathbb{H}_{v_i}(x^1) \cup \mathbb{H}_{v_i}(x^2) = \mathbb{H}_{v_i}(x)$ for $i \in \{1, 2\}$. Let Y be the set of $y_i \in Y^1 \cup Y^2$ such that for all $y_j \in Y^1 \cup Y^2$ with $y_i \neq y_j$ we have $\mathbb{H}_{u_1}(y_j) \cap \mathbb{H}_{u_2}(y_j) \not\supset \mathbb{H}_{u_1}(y_i) \cap \mathbb{H}_{u_2}(y_i)$. We denote by D the DYD defined by Y, x and claim that for any DYD $D' \supset D_1 \cup D_2$ we have $D' \supset D$, which is enough to conclude that $D = D_1 \vee D_2$ is well defined. Let D' be defined by Y', x' .

Note that for each $y_i \in Y$ (and in fact in $Y_1 \cup Y_2$) there is a sequence of points in D_1 or D_2 converging to y_i , so that (by extraction of a subsequence) there exists y'_j with $\mathbb{H}_{u_1}(y'_j) \cap \mathbb{H}_{u_2}(y'_j) \supset \mathbb{H}_{u_1}(y_i) \cap \mathbb{H}_{u_2}(y_i)$. Similarly, there is a sequence of points in D_1 or D_2 converging to the boundary of $\mathbb{H}_{v_1}(x)$, so that $\mathbb{H}_{v_1}(x') \supset \mathbb{H}_{v_1}(x)$ and similarly for v_2 . Thus, we do have $D' \supset D$.

⁶ Associativity was referred to as commutativity by previous authors [8].

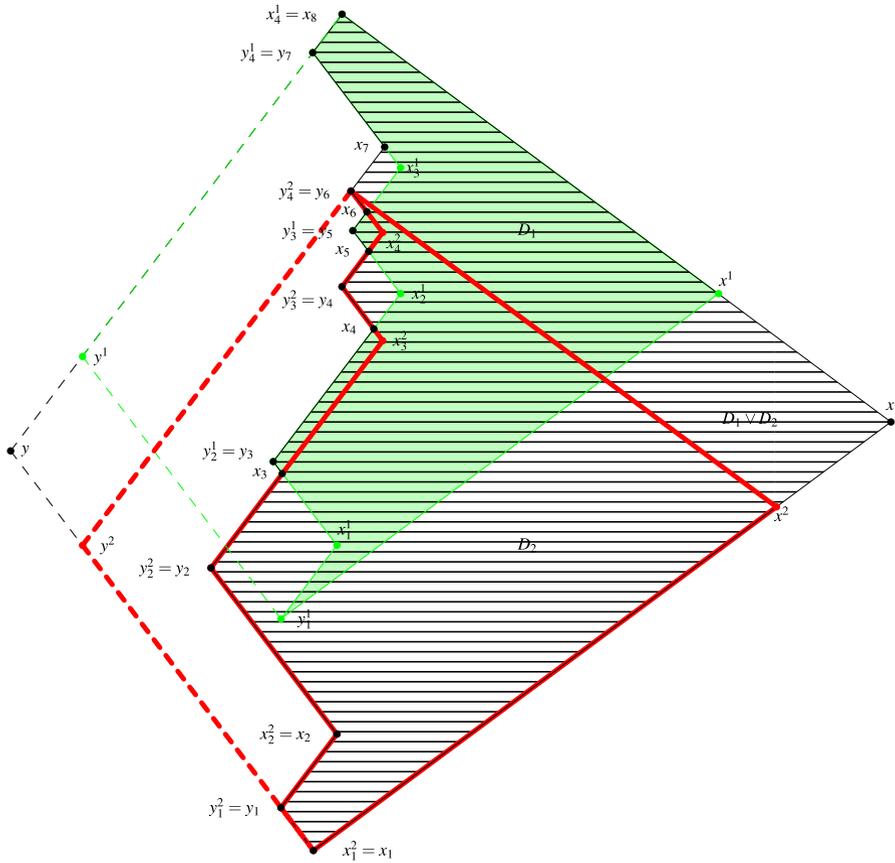


Fig. 4 The shaded region D_1 and thickened region D_2 are DYD. Their respective quadrilaterals $Q(D_i)$ are completed by dashed lines. Their span $D_1 \vee D_2$ is hatched and its quadrilateral $Q(D_1 \vee D_2)$ is also completed by dashed lines

Finally, the commutativity is obvious and the associativity follows from the characterisation of $D_1 \vee D_2$ as the minimal DYD containing both D_1 and D_2 . \square

We analogously define the *span* $D_1 \vee D_2$ of two CDYD D_1 and D_2 – the minimal CDYD containing both – and note that it coincides with their union (which is also commutative and associative). We also define the *span* $C \vee D$ of a DYD D and a CDYD C as the minimal CDYD containing $(C \cup D) \setminus \partial$, which coincides with $C \vee C(D)$. The proof that it is well defined is analogous to Lemma 5.9.

We have thus defined an associative and commutative binary operation \vee on all DYD and CDYD. Moreover, the idempotent unary operation $C(\cdot)$ is distributive with respect to \vee and $C(D_1) \vee D_2 = C(D_1 \vee D_2)$. Furthermore, the span of several DYD is the minimal DYD containing all of them, while the span of several DYD and at least one CDYD is the minimal CDYD containing all the corresponding CDYD.

5.4 Droplet algorithm and spanned droplets

A *droplet* is any DYD contained in Λ or CDYD. We are now ready to define our *droplet algorithm*, which takes as input a finite set $K \subset \Lambda \cap \mathbb{Z}^2$ of infections and outputs a set \mathcal{D} of disjoint connected droplets. It proceeds as follows.

- Form an initial collection of DYD \mathcal{D} consisting of $Q(C)$ for all clusters C of K .
If a DYD $D \in \mathcal{D}$ intersects ∂ , replace it by its CDYD, $C(D)$, to obtain a droplet.
- As long as it is possible, replace two intersecting droplets of \mathcal{D} by their span. If the span intersects ∂ , replace it by its CDYD to obtain a droplet.
- Output the collection \mathcal{D} obtained when all droplets are disjoint.

We similarly define the *modified droplet algorithm* by replacing $Q(C)$ by $Q'(C)$ and clusters by modified clusters above.

The output \mathcal{D} is clearly a collection of disjoint connected droplets. Indeed, by induction all x_i corners of droplets remain in Λ (see Fig. 4), so that DYD remain connected when replaced by CDYD.

Remark 5.10 From the results of Sect. 5.3 it is clear that the order of merging does not impact the output of the algorithm, which is thus well defined. It can also be expressed as the minimal collection of disjoint droplets containing the intersection with Λ of the original collection of quadrilaterals. This minimal collection is well defined. Consequently, the union of the output is increasing in the input.

Definition 5.11 (*Spanned droplets*) Let D be a droplet and $K \subset \mathbb{Z}^2$. We say that D is *spanned* for K with boundary ∂ if the output of the droplet algorithm for $K \cap D$ has a droplet containing D . We omit K and ∂ if they are clear from the context. Similarly, D is *modified spanned* if the output of the modified droplet algorithm for $K \cap D$ has a droplet containing D .

Note that, when seen as an event, a droplet being spanned is monotone. It is also clear that each droplet appearing in (the intermediate or final stages of) the droplet algorithm is spanned and similarly for the modified droplet algorithm. Indeed, the clusters responsible for creating a droplet in the course of the algorithm are contained in the droplet, so each of them is still a cluster of $K \cap D$ (recall that crumbs have diameter much smaller than C_3).

5.5 Properties of the algorithm

We next establish several properties of the algorithm. The approach is similar to the one of [7] with the notable exception of the key Closure Proposition 5.20. We start with the following purely geometric statement.

Lemma 5.12 (Subadditivity) *Let D_1 and D_2 be two DYD or CDYD with non-empty intersection. Then*

$$|D_1 \vee D_2| \leq |D_1| + |D_2|.$$

Furthermore, if D is a DYD intersecting ∂ , then $|C(D)| \leq |D|$.

Proof First assume that D_1 and D_2 are DYD. Since $|D| = |Q(D)|$ for any DYD D and $D_1 \vee D_2 \subset Q(Q(D_1) \vee Q(D_2))$, it suffices to prove the assertion for merging quadrilaterals instead of DYD. But in that case it is not hard to check directly and is a particular case of Lemma 15 of the first arXiv version of [8] (or Lemma 23 of the second version). Since similar (but actually slightly more involved) details were omitted in the proof of the corresponding Lemma 4.6 of [8] and differed to earlier versions, we will not go into useless detail here either. To give a sketch of a possible argument, one can check that for fixed shapes of $Q(D_1)$ and $Q(D_2)$ the maximal $Q(Q(D_1) \vee Q(D_2))$ is achieved when their intersection is reduced to a vertex. Yet, in those configurations one can obtain the v_1 and v_2 sides of $Q(Q(D_1) \vee Q(D_2))$ as the union of those of $Q(D_1)$ and translates of those of $Q(D_2)$ (see Fig. 4). This concludes the proof, as only v_1 and (possibly) v_2 sides contribute to $|\cdot|$ by Lemma 4.1.

Next assume that D_1 is a DYD and D_2 is a CDYD. Let $Y = \{y_i : i \in I\}$ be the set of vectors defining $C(D_1)$ and let $a \in D_1 \cap D_2$. Since $Y \subset \overline{D_1}$, we have that $d(y_i, a) \leq \text{diam}(D_1)$. It then easily follows that the CDYD defined by only one corner, y_i , which we denote $C(y_i)$, is within distance $O(\text{diam}(D_1))$ from $C(a)$. But then $C(D_1) = \bigcup_{i \in I} C(y_i)$ is within distance $O(\text{diam}(D_1))$ from $C(a)$. Thus, $|D_1 \vee D_2| \leq (\text{diam}(D_2) + O(\text{diam}(D_1)))/C_1 \leq |D_2| + |D_1|$, since $\text{diam}(D_1) = O(|D_1|)$ and all implicit constants depend only on \mathcal{S} and are thus much smaller than C_1 .

Next assume that D_1 and D_2 are CDYD. Then the statement is trivial, because $D_1 \vee D_2 = D_1 \cup D_2$, so $\text{diam}(D_1) + \text{diam}(D_2) \geq \text{diam}(D_1 \vee D_2)$ by the triangle inequality.

Finally, let D be a DYD intersecting ∂ . Then, $|C(Q(D))| \geq |C(D)|$ and $|Q(D)| = |D|$, so we may assume that $D = Q(D)$ and prove $|C(D)| \leq |D|$. But in this case it is easy to see that $\text{diam}(C(D)) = O(\text{diam}(D)) = O(|D|)$ with constants depending only on \mathcal{S} , which concludes the proof. \square

The subadditivity lemma will be used to prove the next two adaptations of classical results.

Lemma 5.13 (Aizenman-Lebowitz) *Let K be a finite set and let D be a spanned droplet with $|D| \geq C_4^2$. Then for all $C_4^2/C_1 \leq k \leq |D|/C_1$ there exists a connected spanned droplet D' with $k \leq |D'| \leq 2k$. The same statement holds for modified spanned droplets.*

Proof By Lemma 5.12 at each step of the droplet algorithm the largest size of a droplet appearing in the collection at most doubles. Initially the largest size is at most $C_1 C_4$ and in the end there is a (unique) droplet $D'' \supset D$, so that $|D''| \geq |D|/C_1 \geq C_4^2/C_1 > C_1 C_4$. Then there is a stage of the algorithm at which the maximal size of a droplet in \mathcal{D} is between k and $2k$, which is enough since all droplets appearing in the droplet algorithm are connected and spanned. The proof for modified spanned droplets is identical, using the modified droplet algorithm. \square

Lemma 5.14 (Extremal) *Let $K \subset \mathbb{Z}^2$ and let D be a droplet spanned for K . Then the total number of disjoint clusters for $K \cap D$ in D is at least $\text{diam}(D)/C_4^2$.*

Proof In this proof all clusters will be clusters for $K \cap D$. Assume that at the initial stage of the algorithm there are k clusters (not disjoint). One can then find k/C'_4 disjoint ones, since their diameter is at most C_3 . Furthermore, by Lemma 5.12 the total size of droplets in the collection \mathcal{D} is decreasing, so that $|D|/C_1 \leq |D'| \leq kC_1C_4$, where $D' \supset D$ is some droplet in the output of the algorithm. Indeed, $|Q(C)| \leq C_1C_4$ for all clusters C . This concludes the proof, since $|D| \geq \text{diam}(D)/C_1$ for all DYD and CDYD. \square

We next transform this extremal bound into an exponential decay of the probability that a droplet is spanned until saturation at the critical size. In the following lemma, we identify the configuration ω having law μ and the set of its zeroes.

Lemma 5.15 (Exponential decay) *Let D be a droplet with $|D| \leq 2/(C_5q^\alpha)$. Then*

$$\mu(D \text{ is spanned for } \omega) \leq \exp(-C_4|D|).$$

Proof Let D be a droplet with $|D| \leq 2/(C_5q^\alpha)$, so that $\text{diam}(D) = d \leq 2C_1/(C_5q^\alpha)$. By Lemma 5.14 if D is spanned for ω , it contains at least d/C_4^2 disjoint clusters for $\omega \cap D$, each one having diameter at most C_3 . Each non-boundary cluster has at least α sites by Observation 5.4, while boundary clusters are non-empty and located at distance at most C_2 from ∂ . Thus, we have the union bound

$$\begin{aligned} &\mu(D \text{ is spanned for } \omega) \\ &\leq \sum_{l=0}^{d/C_4^2} \binom{C_3^\alpha d^2}{l} \binom{C_3 d}{d/C_4^2 - l} q^{l\alpha + (d/C_4^2 - l)} \\ &\leq \sum_{l=d/(2C_4^2)}^{d/C_4^2} (C'_4 q^\alpha d^2/l)^l \cdot e^d + \sum_{l'=d/(2C_4^2)}^{d/C_4^2} (C'_4 q d/l')^{l'} \cdot e^d \\ &\leq \sum_{l=d/(2C_4^2)}^{d/C_4^2} \left(\frac{C'_4 e^{2C_4^2} q^\alpha}{1/(2C_4^2)} \cdot \frac{2C_1}{C_5 q^\alpha} \right)^l + \sum_{l'=d/(2C_4^2)}^{d/C_4^2} \left(2C_4^2 C'_4 e^{2C_4^2} q \right)^{l'} \\ &\leq \exp(-C_4 d), \end{aligned}$$

recalling that C_5 is sufficiently large depending on C_4, C'_4 and C_1 . \square

Our next aim is to prove that the closure of a set is contained in its droplet collection up to very local infections next to initial ones. To that end we will need some preliminary results, similar to those used by Bollobás et al. [7].

Observation 5.16 (Lemma 6.5 of [7]) *Let u be a rational non-semi-isolated stable direction. Let $K \subset \mathbb{Z}^2$ with $|K| < \alpha(u)$ (if $\alpha(u) = \infty$ the condition is that K is finite, but there is no a priori bound on its size). Then there exists a constant $C(\mathcal{U}, u, |K|)$ not depending on K such that $[K]_{\mathbb{H}_u}$ is within distance $C(\mathcal{U}, u, |K|)$ from K .*

Since we will require some improvements later, we spell out a proof of the above result for completeness (actually our proof is slightly different from the one in [7]).

Proof of Observation 5.16 We prove the statement by induction on $|K|$. For a $K = \{x\}$ this is easy, since if $\langle x, u \rangle$ is sufficiently large $[K]_{\mathbb{H}_u} = K$ and otherwise there is a single possible configuration for each value of $\langle x, u \rangle$ up to translation. Assume the result holds for $|K| < n$. If one can write $K = K_1 \sqcup K_2$ with $K_1, K_2 \neq \emptyset$ and $d(K_1, K_2) > 2C(\mathcal{U}, u, n - 1) + O(1)$, then $[K]_{\mathbb{H}_u} = [K_1]_{\mathbb{H}_u} \sqcup [K_2]_{\mathbb{H}_u}$, since $[K_1]_{\mathbb{H}_u}$ and $[K_2]_{\mathbb{H}_u}$ are at sufficiently large distance, hence no site can use both to become infected. Assume that, on the contrary, there are no large gaps between parts of K . There is a finite number of such K up to translation and for each of these $[K]$ is finite (e.g. since K is contained in a quadrilateral with sides perpendicular to \mathcal{S}), so within uniformly bounded distance from K . Therefore, if \mathbb{H}_u is sufficiently far from K , $[K]_{\mathbb{H}_u} = [K]$. Otherwise, there is a finite number of possible K up to translation perpendicular to u and for each of them $[K]_{\mathbb{H}_u}$ is finite, so that one can indeed find a finite uniform constant $C(\mathcal{U}, u, n)$ as claimed. \square

A quantitative version of this result was proved by Hartarsky and Mezei [21]. An easy corollary of Observation 5.16 is the fact that crumbs can only grow very locally (see Fig. 5a).

Corollary 5.17 *Let C_1 be sufficiently large depending on \mathcal{U} . Let $K \subset \mathbb{Z}^2$ with $|K| < \alpha$. Then $[K]$ is within distance $C_1/(6\alpha)$ from K . Also, for a (modified) crumb κ we have that $\text{diam}([\kappa]) \leq \alpha C_2$ and $[\kappa]$ is within distance C_1 from κ .*

Proof The first assertion follows from Observation 5.16, since if it were wrong, one could simply translate a set K sufficiently far from a half-plane yielding a contradiction with the observation.

Next consider a (modified) crumb κ and P_κ minimal with $|P_\kappa| < \alpha$ and $[P_\kappa] \supset \kappa$. Then $[\kappa] \subset [P_\kappa]$ is within distance $C_1/(6\alpha)$ from P_κ . If the sites of P_κ are not connected in the graph Γ'' on \mathbb{Z}^2 with connections at distance at most $C_1 + C_2$, then either κ is not connected in Γ or P_κ is not minimal, which are both contradictions. Similarly, if there is no site of κ at distance smaller than $C_1/(2\alpha)$ from a $C_1/(2\alpha)$ -connected component of P_κ , that component can be removed from P_κ , contradicting minimality. Hence, P_κ is within distance $C_1/2$ from κ . The result is then immediate, as $[\kappa]$ is within distance $C_1/2 + C_1/(6\alpha)$ from κ and its diameter is at most $C_1/(3\alpha) + \text{diam}(P_\kappa)$, while $\text{diam}(P_\kappa) \leq (\alpha - 1)(C_1 + C_2)$. \square

In order to treat infection at the concave corners of droplets we will need the following modification of Observation 5.16.

Corollary 5.18 *Let u_1 and u_2 be rational strongly stable directions such that $\mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}$ is stable for the bootstrap percolation dynamics i.e. $\nexists U \in \mathcal{U}, U \subset \mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}$. Let $K \subset \mathbb{Z}^2$ with $|K| \leq \alpha - 1$. Then $[K]_{\mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}}$ is within distance $C(\mathcal{U}, u_1, u_2)$ from K .*

Proof We apply a similar induction to the one in the proof of Observation 5.16. The only difference is that we can no longer use translation invariance. If $d(K, \mathbb{H}_{u_2}) >$

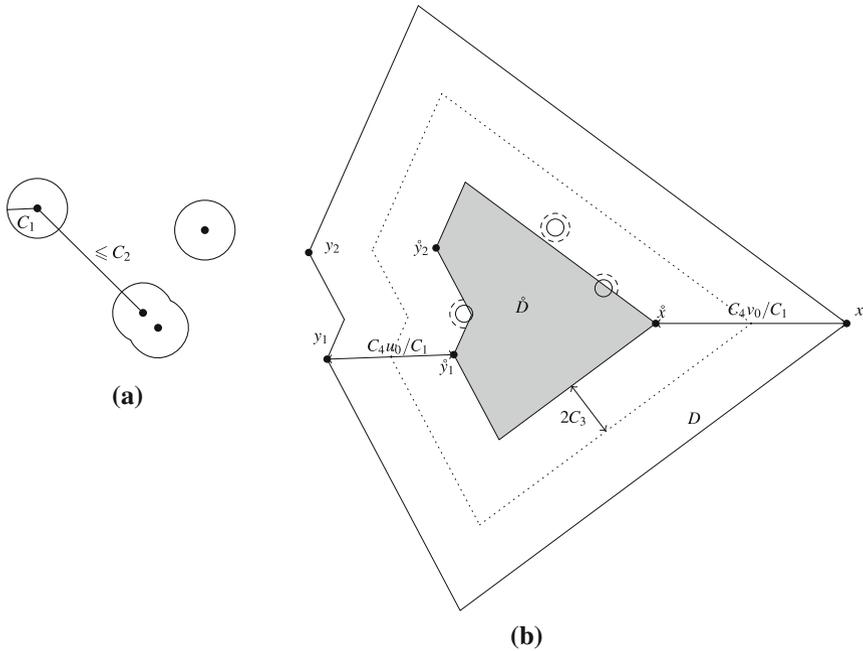


Fig. 5 Illustrations of Corollary 5.17, Lemma 5.19 and Proposition 5.20. **a** The dots represent the sites of a crumb. The (disconnected) circled shape bounds its closure. Note that crumbs may have gaps of size C_2 while the growth allowed is only $C_1 \ll C_2$. **b** The shaded region is the shrunken DYD \hat{D} of the largest DYD D . The solid circles represent crumbs and the dashed arcs are the bound for their growth provided by Lemma 5.19. The modified clusters of the closure are included in the dotted DYD

$C(\mathcal{U}, u_1, |K|) + O(1)$, by Observation 5.16, we have $[K]_{\mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}} = [K]_{\mathbb{H}_{u_1}}$ and similarly for u_1 and u_2 interchanged. We can thus assume that K is within distance $C'(\mathcal{U}, u_1, u_2)$ from the origin. But then

$$[K \cup \mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}] \subset \mathbb{H}_{u_1} \cup \mathbb{H}_{u_2} \cup \mathbb{H}_{u'}(C''(\mathcal{U}, u_1, u_2)u'),$$

where $u' = (u_1 + u_2)/2$, since the latter region is stable by the hypothesis on u_1, u_2 . □

We next transform these results for infinite regions into a result for droplets. It states that a crumb next to a droplet cannot grow significantly (see Fig. 5b).

Lemma 5.19 *Let C_1 be sufficiently large depending on \mathcal{U} and \mathcal{S} . Let D be a DYD at distance at least C_3 from ∂ or be a CDYD and let κ be a crumb. Then $[\kappa]_{D \cup \partial} = [\kappa]_D$ is within distance C_1 of κ .*

Proof Assume that D is a DYD at distance at least C_3 from ∂ . The proof of [7, Lemma 6.10] applies using (7), Observation 5.16, Corollary 5.18 and the arguments in the proof of Corollary 5.17 to give the result for $[\kappa]_D$, which is therefore at distance at least $C_2 - C_1$ from ∂ since $d(\kappa, \partial) \geq C_2$, so that in fact $[\kappa]_D = [\kappa]_{D \cup \partial}$.

Assume next that D is a CDYD. Then actually $D \cup \partial$ can be viewed as a DYD on the entire plane without boundary specified by an infinite number of vectors x_i , so that we are in the previous case. In order to avoid introducing the corresponding notion of infinite DYD, one can consider an increasing exhaustive sequence of DYD D_i converging to $D \cup \partial$ in the product topology and apply the previous result for $[\kappa]_{D_i}$, which will thereby apply to $D \cup \partial$. Finally, $[\kappa]_D = [\kappa]_{D \cup \partial}$ follows, since $d([\kappa]_{D \cup \partial}, \partial) \geq C_2 - C_1$. \square

The next proposition is key to making the output of the algorithm essentially invariant under the KCM dynamics without having to pay for the fact that the closure for the bootstrap percolation dynamics of infections at equilibrium is not at all at equilibrium itself. The proof is illustrated in Fig. 5b.

Proposition 5.20 (Closure) *Let K be a finite set and \mathcal{D}' be the collection of droplets given by the modified droplet algorithm with input $[K]_\partial$. Let \mathcal{D} be the output of the droplet algorithm for K . Then*

$$\forall D' \in \mathcal{D}' \exists D \in \mathcal{D}, D' \subset D.$$

Proof Let \mathcal{K} be the set of crumbs for K . Set $\kappa_0 = \bigcup_{\kappa \in \mathcal{K}} \kappa$. \square

Claim 1 For each crumb $\kappa \in \mathcal{K}$ its closure $[\kappa] = [\kappa]_\partial$ consists of at most $\alpha - 1$ modified crumbs of $[\kappa]$ all contained within distance C_1 from κ .

Proof of Claim 1 There exists a set P_κ as in Definition 5.1, such that $[P_\kappa] \supset \kappa$ and thus $[P_\kappa] \supset [\kappa]$, which proves that all connected components of $[\kappa]$ for Γ' are modified crumbs. The fact that $[\kappa]$ is within distance C_1 of κ (and thus at distance at least C'_2 from ∂) was proved in Corollary 5.17, which also shows that $[\kappa] = [\kappa]_\partial$, since κ is at distance more than C_2 from ∂ . \square

We can thus define $\mathcal{K}'(\kappa)$ to be the set of modified crumbs of $[\kappa]_\partial$, so that their union is disjoint and equal to $[\kappa]_\partial$. Moreover, crumbs in \mathcal{K} are at distance at least C_2 from each other, so for any two of them $\kappa_1 \neq \kappa_2$ we have that any $\kappa'_1 \in \mathcal{K}'(\kappa_1)$ and $\kappa'_2 \in \mathcal{K}'(\kappa_2)$ are at distance at least $C_2 - 2C_1 \gg C'_2$ and also at such distance from ∂ , so that $[\kappa_0]_\partial = \bigcup_{\kappa \in \mathcal{K}} [\kappa]_\partial$ has no modified cluster and consists of modified crumbs at distance at most C_1 from κ_0 .

For a droplet $D \in \mathcal{D}$ consider the set of vectors Y and x (x is absent for CDYD) defining it. Then define $\hat{Y} = Y + C_4 u_0 / C_1$ and $\hat{x} = x + C_4 v_0 / C_1$, where $u_0 \in \mathbb{R}^2$ is the vector such that $\langle u_0, u_1 \rangle = \langle u_0, u_2 \rangle = -1$ and v_0 is defined identically in terms of v_1 and v_2 . We denote by \hat{D} the droplet defined by \hat{Y} and \hat{x} and call it a *shrunk droplet*. Let $D_0 = \bigcup_{D \in \mathcal{D}} D$ and $\hat{D}_0 = \bigcup_{D \in \mathcal{D}} \hat{D}$. It is clear that \hat{D} is at distance at least C_4 / C_1 from $\Lambda \setminus D$ for all droplets D . In particular, all shrunk droplets are at distance at least C_4 / C_1 from each other and shrunk DYD are at distance at least C_4 / C_1 from ∂ , so that Lemma 5.19 applies to them and $[\hat{D}_0]_\partial = \hat{D}_0$.

Claim 2 $\hat{D}_0 \cup \kappa_0 \supset K$.

Proof of Claim 2 Note that it is enough to prove that the clusters of K are contained in \dot{D}_0 . Assume that there exists $a \in K \setminus \dot{D}_0$ and $a \in C$ for some cluster. Then, $Q(C) \cap \Lambda$ is contained in some $D \in \mathcal{D}$, which is defined by Y and x (x is absent for CDYD). Then since $a \notin \dot{D}$, either for all $\dot{y}_i \in \dot{Y}$ we have $a \notin \mathbb{H}_{u_1}(\dot{y}_i) \cap \mathbb{H}_{u_2}(\dot{y}_i)$ or $a \notin \mathbb{H}_{v_1}(\dot{x}) \cap \mathbb{H}_{v_2}(\dot{x})$. In the former case, $a - C_4 u_0 / C_1 \notin \mathbb{H}_{u_1}(y_i) \cap \mathbb{H}_{u_2}(y_i)$ for all $y_i \in Y$. However, $Q(C)$ contains the ball of radius C_4 centered at a and $\|u_0\| = O(1)$, so we get a contradiction. If $a \notin \mathbb{H}_{v_1}(\dot{x}) \cap \mathbb{H}_{v_2}(\dot{x})$, the first point on the segment from a to $a - C_4 v_0 / C_1$ that is not in D is in Λ and in $Q(C)$, hence a contradiction. \square

Claim 3 The set $[K]_\partial \setminus [\kappa_0]_\partial$ is within distance C_3 of \dot{D}_0 .

Proof of Claim 3 By Claim 2 we have $K_0 = \dot{D}_0 \cup \kappa_0 \supset K$. It then clearly suffices to prove that $[K_0]_\partial \setminus [\kappa_0]_\partial$ is within distance C_3 of \dot{D}_0 .

Consider a crumb $\kappa \in \mathcal{K}$ at distance at most C_2 from \dot{D}_0 , so at distance at most C_2 from a shrunken droplet \dot{D} and necessarily at distance at least $C_4/C_1 - C_2 - C_3$ from any other shrunken droplet and from ∂ if D is a DYD. By Lemma 5.19 $[\kappa]_{\dot{D}} = [\kappa]_{\dot{D} \cup \partial}$ is within distance C_1 of κ . Hence,

$$[K_0 \cup \partial] = \dot{D}_0 \cup \partial \cup [\kappa_0] \cup \bigcup_{\kappa, D} [\kappa]_{\dot{D}}, \tag{9}$$

where the last union is on couples (κ, D) as above. Indeed, all $[\kappa]_{\dot{D}}$ and $[\kappa]$ (for different κ) are at distance at least $C_2 - 2C_1$ from each other and from $\dot{D}_0 \setminus \dot{D}$ (by the reasoning above), so for each site of Λ the intersection of the ball of radius $O(1)$ centered at it with the set on the right-hand side of (9) coincides with the intersection with one of the sets $[\kappa \cup \dot{D}]$, $[\kappa]$ or $\dot{D}_0 \cup \partial$, which are all stable, so no infections occur, which proves (9).

The claim follows easily from (9), since for every couple κ, D the set $[\kappa]_{\dot{D}}$ is within distance C_1 of κ , which is itself at distance at most C_2 from \dot{D}_0 , and κ has diameter much smaller than C_3 by Corollary 5.17. \square

Let C' be a modified cluster of $[K]_\partial$ and assume for a contradiction that $C' \subset [\kappa_0]_\partial$. From Definition 5.3 we get that C' is also a modified cluster of $[\kappa_0]_\partial$, but this is a contradiction, since $[\kappa_0]_\partial$ only consists of modified crumbs.

Since any modified cluster C' of $[K]_\partial$ has diameter at most C_3 (by Definition 5.3) and intersects $[K]_\partial \setminus [\kappa_0]_\partial$, which is within distance C_3 of \dot{D}_0 by Claim 3, we get that C' is within distance $2C_3$ of \dot{D}_0 . Therefore, $\bigcup_{C' \in \mathcal{C}'([K]_\partial)} Q'(C') \subset D_0 \cup \partial$, where the union is over all modified clusters of $[K]_\partial$, since $\text{diam}(Q'(C')) \ll C_4/C_1 \leq d(\dot{D}_0, \Lambda \setminus D_0)$. As \mathcal{D} is the output of the droplet algorithm, D_0 is the union of disjoint DYD non-intersecting ∂ and CDYD, so it necessarily contains $\bigcup_{D' \in \mathcal{D}'} D'$ (see Remark 5.10), which concludes the proof.

Remark 5.21 It should be noted that the algorithm is more easily and naturally defined with no boundary, but that will not be sufficient for our purposes. However, this ‘free’ algorithm is trivially obtained as a specialisation of ours. It is also possible to deal with more general boundaries, with infinite input sets, as well as with droplets defined by more directions and possibly with several rugged sides.

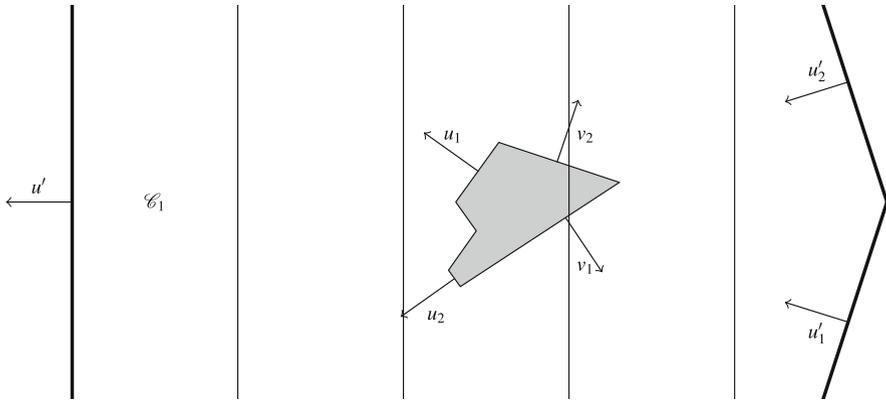


Fig. 6 The domain V is the thickened triangle, a portion of which is displayed. Solid lines separate columns \mathcal{C}_i . Inside the domain is drawn a DYD, which witnesses $\Phi(\omega)_3 = \uparrow$

6 Renormalised East dynamics

In this section we map the original dynamics into an East one and conclude the proof of our main result. In Sect. 6.1 we introduce the necessary notation for the relevant geometry. In Sect. 6.2 we consider a renormalised dynamics on the slices of Fig. 6 by algorithmically selecting certain modified spanned droplets of size $\Omega(1/q^\alpha)$. In Sect. 6.3 we further renormalise to recover an exact East dynamics where q is replaced by q_{eff} corresponding to the probability of spanning such a droplet. Finally, in Sect. 6.4 we prove Theorem 2.8 roughly as in [25].

6.1 Geometric setup

Let us start by defining the domain V we will work in, recalling the notation from Lemma 4.1. Roughly speaking, V is an isosceles triangle with height $e^{1/(C_5 q^\alpha)}$ directed by u' (see Fig. 6). It is divided into ‘columns’ \mathcal{C}_i perpendicular to u' of width roughly $1/q^\alpha$, so that the origin of \mathbb{Z}^2 is in the middle of the last column, close to the tip of V .

More formally, set $L = 1/(C_5 q^\alpha)$ and let ι be the smallest $x \geq 1$ such that the site $\frac{x}{2q^\alpha} u'$ is in \mathbb{Z}^2 , so that $\iota = 1 + O(q^\alpha)$. This way our columns will have width ι/q^α and be separated along rational lines. We define the domain

$$V = \mathbb{H}_{u'}(e^L u') \setminus \left(\mathbb{H}_{u'_2}(-\iota/(2q^\alpha) u') \cup \mathbb{H}_{u'_1}(-\iota/(2q^\alpha) u') \right).$$

Let us choose C_5 so that half the number of columns

$$N = e^L q^\alpha / (2\iota) + 1/4 = e^L q^\alpha (1/2 + O(q^\alpha))$$

is an integer. We then partition the domain $V = \bigcup_{i=1}^{2N} \mathcal{C}_i$ into columns with

$$\mathcal{C}_i = \{x \in V : e^L - \iota(i - 1)/q^\alpha > \langle x, u' \rangle \geq e^L - \iota i/q^\alpha\},$$

so that 0 is in the middle of \mathcal{C}_{2N} and $e^L u' \in \mathbb{Z}^2$. We shall refer to \mathcal{C}_i as the i -th column. Finally, define the half-plane containing \mathcal{C}_{i+1} , but not intersecting \mathcal{C}_i

$$\mathbb{H}_i = \mathbb{H}_{u'}((e^L - i/q^\alpha)u')$$

and the natural boundary for \mathcal{C}_i

$$\partial_i = \mathbb{H}_i \cup \bar{\partial},$$

obtained by considering $\mathcal{C}_j, j \geq i + 1$ as fully infected, where

$$\bar{\partial} = \mathbb{H}_{u'_2}(-\iota/(2q^\alpha)u') \cup \mathbb{H}_{u'_1}(-\iota/(2q^\alpha)u').$$

Note that these boundaries are of the form considered in Sect. 5.

6.2 Arrow variables

Let $\omega \in \Omega$. We will now define a collection of arrow variables which depend only on the restriction of ω to V . We naturally identify the restriction of ω to V with the subset of V where ω is 0 and we use the notation $\omega = \emptyset$ to indicate that all sites are filled (healthy) in V , namely $\omega_x = 1$ for all $x \in V$. Let $\omega^{(0)} = \omega \cap V$. We define the position of the first *up-arrow* as the smallest index $i_1(\omega) \in \{1, 2, \dots, 2N\}$ such that there is a modified spanned droplet of size at least L for $[\omega^{(0)}]_{\partial_{i_1(\omega)}}$ with boundary $\partial_{i_1(\omega)}$. If no such i_1 exists, we say that there are no up-arrows and set $i_1(\omega) = \infty$. We further denote $\omega^{(1)} = \omega^{(0)} \cap \mathbb{H}_{i_1(\omega)}$ as soon as $i_1(\omega) < \infty$, while otherwise $\omega^{(1)} = \emptyset$.

We define the set $I(\omega) = \{i_1(\omega), i_2(\omega), \dots\} \subset \{1, \dots, 2N\}$ containing the positions of up-arrows recursively as follows. If there are no up-arrows, then $I = \emptyset$. Otherwise, we set $I(\omega) = \{i_1(\omega)\} \cup I(\omega^{(1)})$ and $\omega^{(k)} = (\omega^{(k-1)})^{(1)}$, which defines $\omega^{(k)}$ for all k . Let us note that if $i_1(\omega) \neq \infty$, then $i_1(\omega) < i_1(\omega^{(1)})$, since by definition $[\omega^{(1)}]_{\partial_{i_1(\omega)}} = \emptyset$. Finally, we define $\Phi(\omega) \in \{\uparrow, \downarrow\}^{\{1, \dots, 2N\}}$ as

$$\Phi(\omega)_k = \begin{cases} \uparrow & \text{if } k \in I(\omega), \\ \downarrow & \text{otherwise.} \end{cases}$$

The next Lemma states that the probability to find at least one up-arrow decays as

$$q_{\text{eff}} = e^{-L}.$$

Lemma 6.1

$$\mu(i_1 < \infty) \leq q_{\text{eff}}.$$

Proof Fix $1 \leq i \leq 2N$ and consider the event $i_1 = i$. It is clearly included in the event E_i that there is a modified spanned droplet of size at least L for $[\omega^{(0)}]_{\partial_i}$ with boundary ∂_i . By Proposition 5.20 there is also a spanned droplet of size at least

L/C_1 for $\omega^{(0)} \setminus \partial_i$ with boundary ∂_i . By Lemma 5.13 this implies that there is also a spanned connected droplet of size between L/C_1^2 and $2L/C_1^2$. Then one can rewrite E_i as the union over all such droplets D of the event that D is spanned. Note that for each discretised DYD $D \cap \mathbb{Z}^2$ the event that there exists a spanned DYD D' with $D' \cap \mathbb{Z}^2 = D \cap \mathbb{Z}^2$ coincides with the event that a suitably chosen such D'_0 is spanned. Indeed, the intersection of two DYD is a DYD by (7) and the spanning of all D' depend only on the finite number of sites in $D \cap \mathbb{Z}^2$, so there is a finite number of possible events associated to different D' and one can consider the intersection of a D' defining each of these events. The same reasoning holds for CDYD and so for each discretised droplet $D \cap \mathbb{Z}^2$ one can bound the probability that there exists a spanned droplet with such discretisation using Lemma 5.15. Thus, by the union bound on discretised droplets counted in Observation 5.8, one obtains

$$\mu(E_i) \leq |V| \cdot e^L 2e^{-C_4 L/C_1^2} \leq q_{\text{eff}}/(2N),$$

which concludes the proof. □

We next consider the event of having at least n up-arrows

$$\mathcal{B}(n) = \{\omega \in \Omega : |I(\omega)| \geq n\}.$$

Corollary 6.2 *For any $1 \leq n \leq 2N$ we have*

$$\mu(\mathcal{B}(n)) \leq q_{\text{eff}}^n.$$

Proof We prove the statement by induction on n . The base, $n = 1$, is given by Lemma 6.1. For $n > 1$ we have

$$\begin{aligned} \mu(|I| \geq n) &= \sum_{i=1}^{2N} \mu(i_1(\omega) = i; |I(\omega \cap \mathbb{H}_i)| \geq n - 1) \\ &\leq \sum_{i=1}^{2N} \mu(i_1 = i) \mu(|I| \geq n - 1) \\ &\leq q_{\text{eff}}^n, \end{aligned}$$

where we used that the event $i_1 = i$ only depends on $\omega \setminus \mathbb{H}_i$ (i_1 is a stopping time for the filtration induced by the columns) and that the event $|I| \geq n - 1$ is increasing for the order defined by $\omega \leq \omega'$ when $\omega \subset \omega'$. □

We will now state a key deterministic property of the arrows under legal moves of the KCM dynamics.

Lemma 6.3 *Let $\omega \in \Omega$. Let $x \in \mathcal{C}_i$ be such that $\omega_x = 1$ and the constraint at x is satisfied by $\omega \cup \bar{\partial}$. Assume that $\Phi(\omega) \neq \Phi(\omega^x)$. Let $j = \max\{k : \Phi(\omega)_k \neq \Phi(\omega^x)_k\}$.*

Then

$$\begin{aligned} \Phi(\omega)_{[i-1,j]} &= (\uparrow, \downarrow, \uparrow, \downarrow, \uparrow, \dots) \\ \Phi(\omega^x)_{[i-1,j]} &= (\uparrow, \uparrow, \downarrow, \uparrow, \downarrow, \dots) \\ \Phi(\omega)_{[0,i-1]} &= \Phi(\omega^x)_{[0,i-1]} \end{aligned}$$

with the convention that $\Phi(\omega)_0 = \uparrow$ for all ω .

Proof We denote $\Phi := \Phi(\omega)$ and $\Phi' := \Phi(\omega^x)$. Clearly, $\Phi_{[0,i-1]} = \Phi'_{[0,i-1]}$, since those values do not depend on $\omega \cap \mathbb{H}_{i-1}$. \square

Claim 1 Let $k \geq i$. If $\Phi_k = \uparrow$, then $\Phi_{[k+1,2N]} \geq \Phi'_{[k+1,2N]}$ for the lexicographic order associated to $\uparrow < \downarrow$. If $\Phi'_k = \uparrow$, then $\Phi_{[k+1,2N]} \leq \Phi'_{[k+1,2N]}$.

Proof of Claim 1 The two assertions being analogous, we only prove the first one, so assume that $\Phi_k = \uparrow$. Let $j' = \min\{l > k : \Phi_l = \uparrow\}$. Then there is a modified spanned droplet of size at least L for $[\omega^{(0)} \cap \mathbb{H}_k]_{\partial_{j'}}$ with boundary $\partial_{j'}$. But this is also true for ω^x instead of ω , as they coincide in \mathbb{H}_k , and in particular the position of the first up-arrow of Φ' after k is at most j' . \square

Claim 2 Let $k \geq i - 1$ be such that $\Phi_k = \Phi'_k = \downarrow$. Then $k > j$ i.e. $\Phi_{[k,2N]} = \Phi'_{[k,2N]}$.

Proof of Claim 2 We can clearly assume that $k < 2N$. Further assume for a contradiction that $\Phi_{k+1} = \uparrow$ and $\Phi'_{k+1} = \downarrow$. Let $i' = \max\{l < k : \Phi_l = \uparrow\}$. Then there exists a modified spanned droplet D of size at least L for $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_{k+1}}$ with boundary ∂_{k+1} . By Lemma 5.13 we can assume that $L \leq |D| \leq C_1 L$. However, if $d(D, \mathcal{C}_{k+1}) > C_5$, then D is also modified spanned for $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_k}$ with boundary ∂_k , contradicting the definition of i' . Indeed, from the output of the modified droplet algorithm for $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_k} \cap D$ with boundary ∂_k we can create a collection $\hat{\mathcal{S}}$ of droplets for ∂_{k+1} by extending CDYD appropriately, thus $\hat{\mathcal{S}}$ contains $Q'(C') \setminus \partial_k = Q'(C') \setminus \partial_{k+1}$ for every modified cluster C' of $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_k} \cap D$ with boundary ∂_k . Moreover, the modified clusters of $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_{k+1}} \cap D$ with boundary ∂_{k+1} are contained in the modified clusters of $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_k} \cap D$ with boundary ∂_k , so $\hat{\mathcal{S}}$ contains the output of the modified droplet algorithm for $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_{k+1}} \cap D$ with boundary ∂_{k+1} by Remark 5.10, itself containing D .

Therefore, $d(D, \mathcal{C}_{k+1}) \leq C_5$. Moreover, D is not modified spanned for $[(\omega^x)^{(0)} \cap \mathbb{H}_{k-1}]_{\partial_{k+1}}$ with boundary ∂_{k+1} (otherwise $\Phi'_{[k,k+1]} \neq (\downarrow, \downarrow)$). Therefore, there exists a site $y \in D$ such that

$$y \in [\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_{k+1}} \setminus [(\omega^x)^{(0)} \cap \mathbb{H}_{k-1}]_{\partial_{k+1}}.$$

We consider two subcases. First assume that $d(x, \mathbb{R}^2 \setminus \mathbb{H}_{i-1}) \geq C_1$. Then, the constraint at x is satisfied by $(\omega \cap \mathbb{H}_{i-1}) \cup \bar{\partial}$, so $[\omega^{(0)} \cap \mathbb{H}_{k-1}]_{\partial_{k+1}} = [(\omega^x)^{(0)} \cap \mathbb{H}_{k-1}]_{\partial_{k+1}}$, and there is a path

$$P \subset [\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_{k+1}} \setminus [(\omega^x)^{(0)} \cap \mathbb{H}_{k-1}]_{\partial_{k+1}}$$

from $\mathbb{R}^2 \setminus \mathbb{H}_{k-1}$ to y such that each two consecutive sites are at distance at most $O(1)$. But $d(y, \mathbb{R}^2 \setminus \mathbb{H}_{k-1}) \geq \iota/q^\alpha - \text{diam}(D) - C_5 \geq C_2(L + 1)$, so one can find a subpath $P' \subset \mathcal{C}_k \cap P$ of diameter at least C_2L . Yet, it is clear that $P' \subset [\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_k}$ implies the existence of a modified spanned droplet of size larger than L with boundary ∂_k , so one would have an up-arrow of Φ in $[i' + 1, k]$ —a contradiction. If, on the contrary, $d(x, \mathbb{R}^2 \setminus \mathbb{H}_{i-1}) \leq C_1$, we can redo the same reasoning, but P needs to extend to either $\mathbb{R}^2 \setminus \mathbb{H}_{k-1}$ or x , both of which are sufficiently far from y .

Thus, $\Phi_{k+1} = \Phi'_{k+1}$, as the case $\Phi_{k+1} = \downarrow, \Phi'_{k+1} = \uparrow$ is treated identically. But then either both are \uparrow , in which case we are done by Claim 1 or both are \downarrow and we are done by induction. \square

It is easy to see that the only non-identical arrow sequences $\Phi_{[i-1, j]}$ and $\Phi'_{[i-1, j]}$ satisfying the two claims are $(\uparrow, \downarrow, \uparrow, \downarrow, \dots)$ and $(\uparrow, \uparrow, \downarrow, \uparrow, \dots)$ (in this order using that $\omega_x = 1$). Indeed, by Claims 1 and 2 $\Phi_k \neq \Phi'_k$ for all $i \leq k \leq j$, by Claim 1 one cannot have two consecutive up-arrows neither in Φ nor in Φ' in the interval $[i, j]$ and by Claim 2 $\Phi_{i-1} = \Phi'_{i-1} = \uparrow$.

6.3 Renormalised East dynamics

We partition $\{1, \dots, 2N\}$ into blocks $B_i = \{2i - 1, 2i\}$ for $1 \leq i \leq N$. Given $\omega \in \Omega$, we define $\eta(\omega) \in \{0, 1\}^{\{1, \dots, N\}}$ by

$$\eta(\omega)_i = \mathbb{1}_{\{\forall j \in B_i: \Phi(\omega)_j = \downarrow\}}$$

for all $i \in \{1, \dots, N\}$. Let

$$n = \lfloor L \rfloor = \left\lfloor \frac{1}{C_5 q^\alpha} \right\rfloor < \lfloor \log_2 N \rfloor.$$

Recall the definition of legal paths, Definition 2.4. Given an event $\mathcal{E} \subset \Omega$ and a legal path $\gamma = (\omega_{(0)}, \dots, \omega_{(k)})$ we will say that $\gamma \cap \mathcal{E} = \emptyset$ if $\omega_{(i)} \notin \mathcal{E}$ for all $i \in \{0, \dots, k\}$. Also, given $\omega \in \Omega$ and $\mathcal{A} \subset \Omega$, we say that γ connects ω to \mathcal{A} if $\omega_{(0)} = \omega$ and $\omega_{(k)} \in \mathcal{A}$. Recall that $\mathcal{B}(n) \subset \Omega$ is the set of configurations with at least n up-arrows. The following is a straightforward but important corollary of Lemma 6.3.

Corollary 6.4 *For any legal path $(\omega_{(0)}, \dots, \omega_{(k)})$, the path $(\eta(\omega_{(0)}), \dots, \eta(\omega_{(k)}))$ is legal for the East model on $\{1, \dots, N\}$ defined by fixing $\eta_0 = 0$.*

Proof By Lemma 6.3 $\eta(\omega_{(j)}) \neq \eta(\omega_{(j+1)})$ implies that $\Phi(\omega_{(j)})$ and $\Phi(\omega_{(j+1)})$ only differ on an alternating chain of arrows ending in some B_i , preceded by \uparrow . Then clearly $\eta(\omega_{(j)})_l = \eta(\omega_{(j+1)})_l$ for all $l \neq i$ and $\eta(\omega_{(j)})_{i-1} = 0$. \square

Let Ω_\downarrow and Ω_\uparrow^{2N} be respectively the set of configurations which do not have up-arrows, and the set of configurations with an up-arrow in the $2N$ -th column, namely

$$\begin{aligned} \Omega_\downarrow &= \{\omega \in \Omega : \Phi(\omega) = (\downarrow, \dots, \downarrow)\}, \\ \Omega_\uparrow^{2N} &= \{\omega \in \Omega : \Phi(\omega)_{2N} = \uparrow\}. \end{aligned}$$

Combining the last corollary with Proposition 2.7, we obtain the most important input for the proof of the main result.

Corollary 6.5 *For any $\omega \in \Omega_\downarrow$ there does not exist a legal path γ with $\gamma \cap \mathcal{B}(n+1) = \emptyset$ connecting ω to Ω_\uparrow^{2N} .*

6.4 Proof of Theorem 2.8

To prove Theorem 2.8 it is sufficient to prove the lower bound for the mean infection time and use the following inequality (see [10, Theorem 4.4] and also [26, Section 2.2])

$$T_{\text{rel}} \geq q \mathbb{E}(\tau_0). \tag{10}$$

However, it is instructive to construct at this stage a test function that directly gives the desired lower bound on T_{rel} without going through the comparison with the mean infection time. Indeed, the mechanism will appear more clearly this way.

Proof of Theorem 2.8 for T_{rel} We define the event

$$\tilde{\mathcal{A}} = \{ \omega \in \Omega : \exists \text{ a legal path } \gamma \text{ with } \gamma \cap \mathcal{B}(n) = \emptyset \text{ connecting } \omega \cup (\mathbb{Z}^2 \setminus V) \text{ to } \Omega_\downarrow \}$$

and the test function $f : \Omega \rightarrow \{0, 1\}$

$$f = \mathbb{1}_{\tilde{\mathcal{A}}}.$$

Then, by Definition 2.5 we get

$$T_{\text{rel}} \geq \frac{\mu(\tilde{\mathcal{A}})(1 - \mu(\tilde{\mathcal{A}}))}{\mathcal{D}(f)}, \tag{11}$$

where the Dirichlet form $\mathcal{D}(f)$ is defined in (4).

Lemma 6.6 (Bounds on $\mu(\tilde{\mathcal{A}})$)

$$\mu(\tilde{\mathcal{A}}) \left(1 - \mu(\tilde{\mathcal{A}}) \right) \geq \exp \left(\frac{\log q}{C_4 q^\alpha} \right).$$

Proof By Lemma 6.1 we have

$$\mu(\tilde{\mathcal{A}}) \geq \mu(\Omega_\downarrow) \geq 1 - q_{\text{eff}} \geq 1/2.$$

On the other hand,

$$1 - \mu(\tilde{\mathcal{A}}) \geq \mu(\Omega_\uparrow^{2N}) \geq q^{C_1 L} \geq \exp(C_1 \log q / (C_5 q^\alpha)),$$

where we used Corollary 6.5 for the first inequality together with the fact that if $(\omega_{(0)}, \dots, \omega_{(k)})$ is a legal path, then $(\omega_{(k)}, \dots, \omega_{(0)})$ is one as well, and for the second

inequality we notice that for the $2N$ -th arrow to be up it is sufficient to have an empty segment of length C_1L in \mathcal{C}_{2N} . \square

Lemma 6.7 (Estimate of the Dirichlet form) $\mathcal{D}(f) \leq \exp(-1/(C_5^3q^{2\alpha}))$.

Proof Using the fact that $f(\omega)$ depends only on the values of ω in V , we get

$$\begin{aligned} \mathcal{D}(f) &= \sum_{x \in V} \mu(c_x \text{Var}_x(f)) \\ &= q(1 - q) \sum_{x \in V} \mu \left(c_x \mathbb{1}_{\{\omega \in \mathcal{A}, \omega^x \notin \mathcal{A}\}} + c_x \mathbb{1}_{\{\omega \notin \mathcal{A}, \omega^x \in \mathcal{A}\}} \right) \tag{12} \\ &\leq |V| \mu(\mathcal{B}(n - 1)), \end{aligned}$$

since, by Lemma 6.3 $\|I(\omega) - I(\omega^x)\| \leq 1$ when $c_x = 1$, so the indicators both imply $\omega \in \mathcal{B}(n - 1)$. Indeed, $\omega \in \mathcal{A}$ implies the existence of a legal path γ from Ω_\downarrow to $\omega \cup (\mathbb{Z}^2 \setminus V)$ with each configuration not in $\mathcal{B}(n)$. Since $c_x = 1$, the path $\tilde{\gamma}$ obtained by adding the transition from $\omega \cup (\mathbb{Z}^2 \setminus V)$ to $\omega^x \cup (\mathbb{Z}^2 \setminus V)$ is also legal, thus the hypothesis $\omega^x \notin \mathcal{A}$ is not satisfied unless $\omega^x \in \mathcal{B}(n)$ (and similarly for $\omega \notin \mathcal{A}, \omega^x \in \mathcal{A}$). Thus, the result follows by using Corollary 6.2. \square

Then the lower bound for T_{rel} of Theorem 2.8 follows from (11), Lemma 6.6 and Lemma 6.7.

The above proof, together with the matching upper bound of Theorem 2(a) of [26] indicate that the bottleneck dominating the time scales is indeed the creation of $\Theta(\log(1/q_{\text{eff}}))$ simultaneous droplets of probability q_{eff} .

Proof of Theorem 2.8 for $\mathbb{E}(\tau_0)$ The proof of the lower bound for the infection time follows a similar route, with some complications due to the fact that we have to identify a (sufficiently likely) initial set starting from which we have to go through the bottleneck configurations before infecting the origin.

By [25, Corollary 3.4], to prove the desired lower bound on $\mathbb{E}(\tau_0)$ it suffices to construct a local function $\phi = \phi_q$ such that

- (i) $\mu(\phi^2) = 1$,
- (ii) $\frac{\mu(\phi^4)}{\mathcal{D}(\phi)} \geq \exp(1/(C_5^4q^{2\alpha}))$,
- (iii) $\phi(\omega) = 0$ if $\omega_0 = 0$.

Inspired by [25] we let

$$\Omega_g = \Omega_\downarrow \cap \{\omega \in \Omega : \omega_{\Lambda_0} = 1\}$$

where $\Lambda_0 = \{x \in \mathbb{Z}^2 : d(x, 0) \leq 1/(4q^\alpha)\} \subset \mathcal{C}_{2N}$ and

$$\mathcal{A} = \{\omega \in \Omega : \exists \text{ a legal path } \gamma \text{ with } \gamma \cap \mathcal{B}(n) = \emptyset \text{ connecting } \omega \cup (\mathbb{Z}^2 \setminus V) \text{ to } \Omega_g\}.$$

Then we set

$$\phi(\cdot) = \mathbb{1}_{\mathcal{A}}(\cdot) / \mu(\mathcal{A})^{1/2}. \tag{13}$$

We are now left with proving that this function satisfies (i)–(iii) above.

Property (i) follows immediately from (13). In order to verify (ii) we start by establishing a lower bound on $\mu(\mathcal{A})$. By definition it holds that

$$\mu(\mathcal{A}) \geq \mu(\Omega_g) \geq \mu(\omega_{\Lambda_0} = 1)\mu(\Omega_{\downarrow}) \geq e^{-O(1)/q^{2\alpha-1}}(1 - q_{\text{eff}}) = e^{-O(1)/q^{2\alpha-1}}, \tag{14}$$

where we used Harris’ inequality [18] ($\{\omega_{\Lambda_0} = 1\}$ and Ω_{\downarrow} are increasing events if we consider that $\omega \leq \omega'$ when $\omega_x \leq \omega'_x$ for all $x \in \mathbb{Z}^2$), Lemma 6.1 and $|\Lambda_0| = O(1/q^{2\alpha})$.

Furthermore, one can repeat the proof of Lemma 6.7 to obtain

$$\mathcal{D}(\phi) \leq e^{-1/(C_3^3 q^{2\alpha})}. \tag{15}$$

Thus, recalling (14), Property (ii) holds.

We are therefore only left with proving the next lemma establishing Property (iii), completing the proof of Theorem 2.8.

Lemma 6.8 *Let ω be such that $\omega_0 = 0$. Then any legal path connecting Ω_g to ω intersects $\mathcal{B}(n)$.*

As in the lower bound on $1 - \mu(\tilde{\mathcal{A}})$ for T_{rel} , the proof relies on Corollary 6.5, but an additional complication arises due to the fact that emptying the origin does not a priori require creating a critical droplet nearby.

Proof of Lemma 6.8 Suppose for a contradiction that there exists a configuration ω with $\omega_0 = 0$, a configuration $\omega_{(0)} \in \Omega_g$ and a legal path $\gamma = (\omega_{(0)}, \dots, \omega_{(k)})$ with $\omega_{(k)} = \omega$ and $\omega_{(j)} \notin \mathcal{B}(n)$ for all $j \in \{0, \dots, k\}$. Assuming without loss of generality that $\omega_{(j)} \neq \omega_{(j-1)}$ for all j , let x_j be such that $\omega_{(j)} = (\omega_{(j-1)})^{x_j}$. Consider the path $\tilde{\gamma} = (\tilde{\omega}_{(0)}, \dots, \tilde{\omega}_{(k)})$ obtained by performing the same updates as for γ except for flips in the column \mathcal{C}_{2N} , which are performed only if they correspond to emptying sites. More precisely, we let $\tilde{\omega}_{(0)} = \omega_{(0)}$ and

$$\tilde{\omega}_{(j)} = \begin{cases} (\tilde{\omega}_{(j-1)})^{x_j} & \text{if } x_j \notin \mathcal{C}_{2N} \text{ or } (\tilde{\omega}_{(j-1)})_{x_j} = 1, \\ \tilde{\omega}_{(j-1)} & \text{otherwise.} \end{cases}$$

It is not difficult to verify by induction that $\tilde{\gamma}$ is also a legal path with $\tilde{\omega}_{(j)} \leq \omega_{(j)}$ for all j (where $\omega \leq \omega'$ when $\omega_x \leq \omega'_x$ for all $x \in \mathbb{Z}^2$) and that $\tilde{\omega}_{(j)}$ and $\omega_{(j)}$ coincide outside of \mathcal{C}_{2N} . Then $(\tilde{\omega}_{(k)})_0 \leq (\omega_{(k)})_0 = 0$ and by definition $(\tilde{\omega}_{(0)})_{\Lambda_0} = 1$. Therefore, since inside \mathcal{C}_{2N} each site that has been emptied in γ is also empty in $\tilde{\omega}_{(k)}$, we conclude that necessarily $\tilde{\omega}_{(k)} \cap \mathcal{C}_{2N}$ contains a (modified) spanned droplet of size $1/(4C_1 q^\alpha) > L$ with boundary $\partial_{2N} = \bar{\partial}$. Indeed, there is a path of sites x with steps of size $O(1)$ from $\mathbb{Z}^2 \setminus \Lambda_0$ to 0 such that $(\tilde{\omega}_{(k)})_x = 0$. This means that $\tilde{\omega}_{(k)} \in \Omega_{\uparrow}^{2N}$. Furthermore, for all j we have $\Phi(\tilde{\omega}_{(j)})_{[1, 2N-1]} = \Phi(\omega_{(j)})_{[1, 2N-1]}$, as those do not depend on the sites in \mathcal{C}_{2N} . Thus, using Corollary 6.5, together with the facts that $\tilde{\omega}_{(0)} \in \Omega_g \subset \Omega_{\downarrow}$, $\tilde{\omega}_{(k)} \in \Omega_{\uparrow}^{2N}$ and $\tilde{\gamma} \cap \mathcal{B}(n+1) = \emptyset$, we reach a contradiction. \square

7 Open problems

With Theorem 2.8 the scaling of the infection time is determined up to a polylogarithmic factor. The next natural question is to pursue determining this factor in the spirit of the refined universality result of [7]. For the moment there is only one critical model with infinitely many stable directions for which this is known — the Duarte model [25]. In that case the corrective factor is $\Theta((\log q)^4)$. However, for bootstrap percolation there are already two different possible behaviours of this factor depending on whether the model is balanced or unbalanced (see Definition 2.3). Based on this one could expect the following.

Conjecture 7.1 Let \mathcal{U} be a critical update family with an infinite number of stable directions.

– If \mathcal{U} is balanced, then

$$\mathbb{E}(\tau_0) = \exp\left(\frac{\Theta(1)}{q^{2\alpha}}\right).$$

– If \mathcal{U} is unbalanced, then

$$\mathbb{E}(\tau_0) = \exp\left(\frac{\Theta((\log q)^4)}{q^{2\alpha}}\right).$$

The same asymptotics hold for T_{rel} .

In other words we expect the lower bound of Theorem 2.8 to be sharp for balanced models, while the upper bound of [26, Theorem 2(a)] to be sharp for unbalanced ones. The balanced case is not hard and only requires an improvement of the approach of [26]. It will be treated in a future work, since it shares none of the techniques discussed here. In the unbalanced case the $(\log q)^4$ should arise as the square of the $(\log q)^2$ factor for bootstrap percolation, itself caused by the one-dimensional geometry and larger size of critical droplets. This is indeed what happens for the Duarte model [25], an example of unbalanced critical constraint.

Acknowledgements We wish to thank Fabio Martinelli for numerous and enlightening discussions and the Departement of Mathematics and Physics of University Roma Tre for its kind hospitality.

References

1. Aizenman, M., Lebowitz, J.L.: Metastability effects in bootstrap percolation. *J. Phys. A* **21**(19), 3801–3813 (1988). <http://stacks.iop.org/0305-4470/21/3801>
2. Aldous, D., Diaconis, P.: The asymmetric one-dimensional constrained Ising model: rigorous results. *J. Stat. Phys.* **107**(5–6), 945–975 (2002). <https://doi.org/10.1023/A:1015170205728>
3. Bakry, D.: Functional inequalities for Markov semigroups. In: *Probability measures on groups: recent directions and trends*, pp. 91–147. Tata Inst. Fund. Res., Mumbai (2006)
4. Balister, P., Bollobás, B., Przykucki, M., Smith, P.: Subcritical \mathcal{U} -bootstrap percolation models have non-trivial phase transitions. *Trans. Amer. Math. Soc.* **368**(10), 7385–7411 (2016). <https://doi.org/10.1090/tran/6586>

5. Berthier, L., Biroli, G.: Theoretical perspective on the glass transition and amorphous materials. *Rev. Mod. Phys.* **83**, 587–645 (2011). <https://doi.org/10.1103/RevModPhys.83.587>
6. Bollobás, B., Duminil-Copin, H., Morris, R., Smith, P.: The sharp threshold for the Duarte model. *Ann. Probab.* **45**(6B), 4222–4272 (2017). <https://doi.org/10.1214/16-AOP1163>
7. Bollobás, B., Duminil-Copin, H., Morris, R., Smith, P.: Universality of two-dimensional critical cellular automata. *Proc. Lond. Math. Soc.* (to appear)
8. Bollobás, B., Smith, P., Uzzell, A.: Monotone cellular automata in a random environment. *Combin. Probab. Comput.* **24**(4), 687–722 (2015). <https://doi.org/10.1017/S0963548315000012>
9. Cancrini, N., Martinelli, F., Roberto, C., Toninelli, C.: Kinetically constrained spin models. *Probab. Theory Related Fields* **140**(3–4), 459–504 (2008). <https://doi.org/10.1007/s00440-007-0072-3>
10. Cancrini, N., Martinelli, F., Roberto, C., Toninelli, C.: Facilitated spin models: recent and new results. In: *Methods of contemporary mathematical statistical physics, Lecture Notes in Math.*, vol. 1970, pp. 307–340. Springer, Berlin (2009)
11. Cancrini, N., Martinelli, F., Schonmann, R., Toninelli, C.: Facilitated oriented spin models: some non equilibrium results. *J. Stat. Phys.* **138**(6), 1109–1123 (2010). <https://doi.org/10.1007/s10955-010-9923-x>
12. Chleboun, P., Faggionato, A., Martinelli, F.: Time scale separation and dynamic heterogeneity in the low temperature East model. *Commun. Math. Phys.* **328**(3), 955–993 (2014). <https://doi.org/10.1007/s00220-014-1985-1>
13. Chleboun, P., Faggionato, A., Martinelli, F.: Relaxation to equilibrium of generalized East processes on \mathbb{Z}^d : renormalization group analysis and energy-entropy competition. *Ann. Probab.* **44**(3), 1817–1863 (2016). <https://doi.org/10.1214/15-AOP1011>
14. Chung, F., Diaconis, P., Graham, R.: Combinatorics for the East model. *Adv. Appl. Math.* **27**(1), 192–206 (2001). <https://doi.org/10.1006/aama.2001.0728>
15. Faggionato, A., Martinelli, F., Roberto, C., Toninelli, C.: The East model: recent results and new progresses. *Markov Process. Related Fields* **19**(3), 407–452 (2013)
16. Fredrickson, G.H., Andersen, H.C.: Kinetic Ising model of the glass transition. *Phys. Rev. Lett.* **53**, 1244–1247 (1984). <https://doi.org/10.1103/PhysRevLett.53.1244>
17. Garrahan, P., Sollich, P., Toninelli, C.: Kinetically constrained models. In: Berthier, L., Biroli, G., Bouchaud, J.P., Cipelletti, L., van Saarloos, W. (eds.) *Dynamical heterogeneities in Glasses, colloids and granular media and jamming transitions*, International series of monographs on physics 150, pp. 341–369. Oxford University Press (2011). <https://hal.archives-ouvertes.fr/hal-00563884>
18. Harris, T.E.: A lower bound for the critical probability in a certain percolation process. *Proc. Camb. Philos. Soc.* **56**, 13–20 (1960)
19. Hartarsky, I.: \mathcal{Z} -bootstrap percolation: critical probability, exponential decay and applications. arXiv e-prints [arXiv:1806.11405](https://arxiv.org/abs/1806.11405) (2018)
20. Hartarsky, I., Martinelli, F., Toninelli, C.: Universality for critical KCM: finite number of stable directions. arXiv e-prints [arXiv:1910.06782](https://arxiv.org/abs/1910.06782) (2019)
21. Hartarsky, I., Mezei, T.: Complexity of 2D bootstrap percolation difficulty: Algorithm and NP-hardness. arXiv e-prints [arXiv:1809.01525](https://arxiv.org/abs/1809.01525) (2018)
22. Holroyd, A.E.: Sharp metastability threshold for two-dimensional bootstrap percolation. *Probab. Theory Related Fields* **125**(2), 195–224 (2003). <https://doi.org/10.1007/s00440-002-0239-x>
23. Jäckle, J., Eisinger, S.: A hierarchically constrained kinetic Ising model. *Z. Phys. B Con. Mat.* **84**(1), 115–124 (1991). <https://doi.org/10.1007/BF01453764>
24. Liggett, T.M.: *Interacting particle systems*. Classics in Mathematics. Springer-Verlag, Berlin (2005). <https://doi.org/10.1007/b138374>. Reprint of the 1985 original
25. Maréché, L., Martinelli, F., Toninelli, C.: Exact asymptotics for Duarte and supercritical rooted kinetically constrained models. *Ann. Probab.* (to appear)
26. Martinelli, F., Morris, R., Toninelli, C.: Universality results for kinetically constrained spin models in two dimensions. *Commun. Math. Phys.* **369**(2), 761–809 (2019). <https://doi.org/10.1007/s00220-018-3280-z>
27. Martinelli, F., Toninelli, C.: Towards a universality picture for the relaxation to equilibrium of kinetically constrained models. *Ann. Probab.* **47**(1), 324–361 (2019). <https://doi.org/10.1214/18-AOP1262>
28. Morris, R.: Bootstrap percolation, and other automata. *Eur. J. Combin.* **66**, 250–263 (2017). <https://doi.org/10.1016/j.ejc.2017.06.024>
29. Morris, R.: Monotone cellular automata. In: *Surveys in combinatorics 2017*, London Math. Soc. Lecture Note Ser., vol. 440, pp. 312–371. Cambridge Univ. Press, Cambridge (2017)

30. Mountford, T.S.: Critical length for semi-oriented bootstrap percolation. *Stoch. Process. Appl.* **56**(2), 185–205 (1995). [https://doi.org/10.1016/0304-4149\(94\)00061-W](https://doi.org/10.1016/0304-4149(94)00061-W)
31. Ritort, F., Sollich, P.: Glassy dynamics of kinetically constrained models. *Adv. Phys.* **52**(4), 219–342 (2003). <https://doi.org/10.1080/0001873031000093582>
32. Sollich, P., Evans, M.R.: Glassy time-scale divergence and anomalous coarsening in a kinetically constrained spin chain. *Phys. Rev. Lett.* **83**, 3238–3241 (1999). <https://doi.org/10.1103/PhysRevLett.83.3238>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Ivailo Hartarsky^{1,2}  · Laure Maréché³ · Cristina Toninelli² 

✉ Ivailo Hartarsky
hartarsky@ceremade.dauphine.fr

Laure Maréché
laure.mareche@epfl.ch

Cristina Toninelli
toninelli@ceremade.dauphine.fr

¹ DMA UMR 8553, École Normale Supérieure, PSL University, CNRS, 45 rue d'Ulm, 75005 Paris, France

² CEREMADE UMR 7534, Université Paris-Dauphine, PSL University, CNRS, Place du Maréchal de Lattre de Tassigny, 75775 Paris Cedex 16, France

³ LPSM UMR 8001, Université Paris Diderot, Sorbonne Paris Cité, CNRS, 75013 Paris, France