Kinetically constrained models out of equilibrium

Ivailo Hartarsky and Fabio Lucio Toninelli

Technische Universität Wien, Institut für Stochastik und Wirtschaftsmathematik, Wiedner Hauptstraße 8-10, A-1040, Vienna, Austria {ivailo.hartarsky,fabio.toninelli}@tuwien.ac.at

March 11, 2024

Abstract

We study the full class of kinetically constrained models in arbitrary dimension and out of equilibrium, in the regime where the density q of facilitating sites in the equilibrium measure (but not necessarily in the initial measure) is close to 1. For these models, we establish exponential convergence to equilibrium in infinite volume and linear time precutoff in finite volume with appropriate boundary condition. Our results are the first out-of-equilibrium results that hold for any model in the so-called critical class, which is covered in its entirety by our treatment. It includes e.g. the Fredrickson–Andersen 2-spin facilitated model, in which a site is updated only when at least two neighbouring sites are in the facilitating state. In addition, these results generalise, unify and sometimes simplify several previous works in the field. As byproduct, we recover and generalise exponential tails for the connected component of the origin in the upper invariant trajectory of perturbed cellular automata and in the set of eventually infected sites in subcritical bootstrap percolation models. Our approach goes through the study of cooperative contact processes, last passage percolation, Toom contours, as well as a very convenient coupling between contact processes and kinetically constrained models.

MSC2020: 60K35; 82C22
Keywords: kinetically constrained models, contact processes, perturbed cellular automata, bootstrap percolation, convergence to equilibrium, precutoff

1 Introduction

Kinetically constrained models (KCM) were introduced in order to study the liquid-glass transition [23] (see [13,25,47] for reviews). They are Markov processes featuring a parameter $q \in [0,1]$ tuning the density of facilitating sites. Purposefully, KCM are reversible w.r.t. the product Bernoulli measure with parameter q. Indeed, it has been proposed that the real-world liquid-glass transition has a purely dynamical origin, that is not reflected in the equilibrium measure. KCM were introduced precisely with the purpose of testing whether glassy behaviour could be explained purely in dynamical terms. Despite having a trivial stationary measure, the degenerate rates of KCM make them very hard to tackle mathematically, by making their dynamics non-attractive, cooperative, heterogeneous, admitting multiple invariant measures, not satisfying coercive inequalities and sometimes featuring ergodicity breaking phase transitions.

In recent years a detailed understanding of KCM at equilibrium has been achieved, especially in two dimensions (see [31, Chapter 1] for an overview). However, from the physical perspective it is essential to understand their behaviour out of equilibrium, typically after a quench from one temperature to a different one. Rigorous results in this direction are rather limited and will be the subject of a detailed account in Section 4. Suffice it to say that with the exception of [18], all out-of-equilibrium results pertain to the class of so-called supercritical modes, for which a finite patch of facilitating sites can trigger relaxation. On the other hand, [18] treats subcritical models with an orientation. The main goal of the present work is to deal with critical models. However, our methods work in the greatest possible generality, so we will also cover the other universality classes. Our main results are: the proof that the mixing time of the process in a box of side n and boundary condition entirely composed of facilitating sites is of order n (see Theorem 3.1) and the proof of exponential convergence to equilibrium for the infinite-volume dynamics started from a Bernoulli initial condition in the ergodic regime (see Theorem 3.3). A comparison with earlier results is given in Section 4.

The main novelties of the present work are as follows. Firstly, we set up a general scheme for proving exponential decay for the size of connected clusters of objects in dependent settings (see Section 9) and showcase two applications of independent interest going far beyond the needs of our main results. We further devise a simple and very robust technique for coupling interacting particle systems, even in the absence of attractiveness (see Section 7). Moreover, as opposed to previous studies, we develop tools to tackle cooperative models (such that no finite set of facilitating sites is able to propagate), including several useful renormalisation techniques (see Sections 6, 8 and 11). In particular, we provide the first out-of-equilibrium results for any kinetically constrained model in the *critical universality class*.

A major limitation for the study of KCM out of equilibrium is that, like ours, most results work in the perturbative regime of high density of facilitating sites for the equilibrium measure and not necessarily the initial one. Exceptions to this are [18], heavily relying on orientation, and results on the East model [7, 15, 15, 17, 19, 21, 22, 24, 40], where a model-specific miracle greatly simplifies the problem.[1] Our study will also be restricted to the perturbative regime for the equilibrium measure, but not the initial one. Nevertheless, it is our hope that, once a robust renormalisation scheme is found for controlling KCM in the non-perturbative regime in terms of the perturbative one, our tools will enable the treatment of the full class of KCM out-of-equilibrium.

2 Models

In this section we define our models of interest, KCM, along with several other models, which will play an auxiliary role in the proofs. The reader eager to see the statements of the results in Section 3 will only need Sections 2.1 and 2.2.

2.1 Update families

Let $\|\cdot\|$, $\langle\cdot,\cdot\rangle$ and $d(\cdot,\cdot)$ denote the Euclidean norm, scalar product and distance respectively. An *update family* $\mathcal{U}$ is a finite non-empty family of finite non-empty subsets of $\mathbb{Z}^d\setminus\{0\}$ called *update rules*. We refer to unit vectors

$$u\in S^{d-1}=\left\{v\in\mathbb{R}^d:\|v\|=1\right\}$$

as *directions*. We denote by $\mathbb{H}_u=\{x\in\mathbb{Z}^d:\langle x,u\rangle<0\}$. A direction u 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. An update family is *not trivial subcritical* if it has an unstable direction and *trivial subcritical* otherwise.

While we will not need to distinguish between the other universality classes, it is useful to introduce them for the sake of discussing previous results. We say that an update family is *supercritical*, if there is an open

[1] Some of the perturbative results on the Fredrickson–Andersen 1-spin facilitated model in one dimension [8, 9, 20] concern a neighbourhood of $q=1$ with noticeable length. When applied to this model, our treatment also yields reasonable quantitative bounds.

hemisphere consisting only of unstable directions. We say that it is *subcritical* (trivial or non trivial), if every hemisphere contains an open set of stable directions. Finally, the update family is *critical* if it is neither supercritical nor subcritical. We call update families such that there exists $u \in S^{d-1}$ such that $U \subset \mathbb{H}_u$ for all $U \in \mathcal{U}$ *oriented*.

Let us introduce an illustrative example corresponding to the classical *Fredrickson-Andersen j-spin facilitated model* (FA-jf). Its update family $\mathcal{U} = \{X \subset \{e_1, \dots, e_d, -e_1, \dots, -e_d\}, |X| = j\}$ is given by all j-element subsets of the $2d$ nearest neighbours of the origin. One can check that for $j = 1$ this family is supercritical, for $j \in \{2, \dots, d\}$ it is critical, while for $j > d$ it is trivial subcritical. The reader may keep in mind the case $j = d$, which is also the most interesting one, throughout the paper. Indeed, all difficulties we face are present for this model.

In the rest of the work, adopting the language of bootstrap percolation and of the contact process, we say that site $x \in \mathbb{Z}^d$ is "infected" in the configuration η, if its state at x, η_x, is 1.

2.2 Kinetically constrained models

The $\mathcal{U}$-KCM is a continuous time Markov process with state space $\Omega = \{0,1\}^{\mathbb{Z}^d}$ defined by the following graphical construction (see e.g. [38, Section III.6] for background). To each $x \in \mathbb{Z}^d$ we attach an independent Poisson point process P_x on $[0, \infty)$ of intensity 1 and uniform random variables $(\Upsilon_x(t))_{t \in P_x}$ on $[0,1]$, which are independent and independent of those for other sites. The model has a further parameter $q \in [0,1]$, which we call the *equilibrium density*. We denote by $\eta_x(t)$ the state of site $x \in \mathbb{Z}^d$ in the configuration at time $t \in [0, \infty)$. Define the *constraint* at $x \in \mathbb{Z}^d$ for a configuration $\eta \in \Omega$ by

$$c_x(\eta) = \mathbb{1}_{\exists U \in \mathcal{U}, \forall u \in U, \eta_{x+u}=1}. \tag{2.1}$$

We have

$$\eta_x(t) = \begin{cases} \mathbb{1}_{\Upsilon_x(t) \leq q} & t \in P_x \text{ and } c_x(\eta(t-)) = 1, \\ \eta_x(t-) & \text{otherwise.} \end{cases} \tag{2.2}$$

Before moving on, let us also give an informal but more intuitive description of the $\mathcal{U}$-KCM with parameter q. Each site $x \in \mathbb{Z}^d$ is equipped with a clock which rings at exponentially distributed intervals of time of mean 1 (P_x is the set of clock ring times). When it rings, we verify whether the constraint is satisfied, that is, if there is a completely infected update rule around x. If this is the case, we replace the state of x in the configuration by an independent Bernoulli random variable with parameter q. Note that in the above definition, we have coupled $\mathcal{U}$-KCM for all values of q and all initial conditions using the same clock rings and defining the Bernoulli random variables, using the uniform ones, $\Upsilon_x(t)$.

We remark, that $c_x(\eta)$ of Eq. (2.1) does not depend on η_x, so $c_x(\eta(t-)) = c_x(\eta(t))$ for all $t \in P_x$ and $x \in \mathbb{Z}^d$. It is easy to see that this implies that the $\mathcal{U}$-KCM is reversible w.r.t. the product Bernoulli measure $\mu_q = \mathrm{Ber}(q)^{\otimes \mathbb{Z}^d}$ (see e.g. [38, Section IV.2] for background). We emphasise that $\mathcal{U}$-KCM are not attractive, i.e. the natural stochastic order is not preserved by the dynamics (see [38, Section III.2] for background). We refer to sites in state 1 as *infected* and sites in state 0 as *healthy*. Hence, the constraint asks for the presence of suitably arranged infections around the site we are trying to update.

One defines $q_{\mathrm{c}}^{\mathrm{KCM}}$ as the infimum of all $q \in [0, 1]$ such that 0 is a simple eigenvalue of the generator of the $\mathcal{U}$-KCM with parameter q. That is, $q_{\mathrm{c}}^{\mathrm{KCM}}$ is the critical parameter for ergodicity. It is known by [12, Proposition 2.5] and [2, Corollary 1.6 and Theorem 7.1] that $q_{\mathrm{c}}^{\mathrm{KCM}} > 0$ if and only if $\mathcal{U}$ is subcritical and $q_{\mathrm{c}}^{\mathrm{KCM}} = 1$ if and only if $\mathcal{U}$ is trivial subcritical. We will also need the critical parameter of the spectral gap of the generator of the $\mathcal{U}$-KCM (see [12, Section 2] for background, but understanding this definition is not essential to the present work):

$$\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}} = \inf \{q > 0 : \mathrm{gap} > 0\}.$$

It is believed that $\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}} = q_{\mathrm{c}}^{\mathrm{KCM}}$ for all update families, but this has only been shown in some cases (see Remark 3.4).

Fix $\Lambda \subset \mathbb{Z}^d$. For any $\omega \in \Omega$ we denote by $\omega_\Lambda \in \Omega_\Lambda = \{0, 1\}^\Lambda$ the restriction of ω to Λ. We define the $\mathcal{U}$-KCM η on Λ with boundary condition $\tau \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ by setting the configuration equal to τ outside Λ at all times. We denote the fully infected (resp. healthy) configuration by $\mathbf{1}$ (resp. $\mathbf{0}$).

We next introduce the *mixing time* of the $\mathcal{U}$-KCM on a finite set Λ with some boundary condition $\tau \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$ (see [36] for background). Given $\delta \in (0, 1)$, we define

$$t_{\mathrm{mix}}(\delta) = \inf \left\{ t \geq 0 : \max_{\rho \in \Omega_\Lambda} d_{\mathrm{TV}} \left(\mathbb{P}\left(\eta^\rho(t) \in \cdot\right), \mu_q \right) \leq \delta \right\} \in (0, \infty], \tag{2.3}$$

where $d_{\mathrm{TV}}(\mu, \nu) = \sup_A (\mu(A) - \nu(A))$ with the supremum running over all events A for the arbitrary probability measures μ and ν and η^ρ is the $\mathcal{U}$-KCM with initial condition ρ.

2.3 Contact processes

The $\mathcal{U}$-contact process (CP) is defined by the same graphical construction as the $\mathcal{U}$-KCM and the same constraint as in Eq. (2.1). However, we set

$$\zeta_x(t) = \begin{cases} 1 & t \in P_x, c_x(\zeta(t-)) = 1, \Upsilon_x(t) \le q, \\ 0 & t \in P_x, \Upsilon_x(t) > q, \\ \zeta_x(t-) & \text{otherwise}, \end{cases} \tag{2.4}$$

instead of Eq. (2.2). That is, the constraint $c_x(\zeta(t-)) = 1$ is no longer required to be satisfied in order to update the configuration at site x to the value 0. We define the $\mathcal{U}$-CP in finite volume with a boundary condition analogously to what was done for KCM in Section 2.2. We emphasise, that we use the same Poisson processes (clock rings) and uniform random variables, so that now $\mathcal{U}$-KCM and $\mathcal{U}$-CP for all initial conditions and parameters $q \in [0,1]$ are coupled on the same probability space.

Contrary to KCM, CP are attractive (see [38, Section III.2] for background), we may therefore define its upper invariant measure $\bar{\nu}$ as the $t \to \infty$ limit in law of $\zeta^{\mathbf{1}}(t)$, the $\mathcal{U}$-CP with initial condition $\mathbf{1}$. We then define the critical point

$$q_{\mathrm{c}}^{\mathrm{CP}} = \inf \{q > 0 : \bar{\nu} \neq \delta_{\mathbf{0}}\}. \tag{2.5}$$

In other words, $q_{\mathrm{c}}^{\mathrm{CP}}$ is the critical parameter, above which $\mathcal{U}$-CP has multiple invariant measures.

Remark 2.1. It is known that $q_{\mathrm{c}}^{\mathrm{CP}} < 1$ if and only if $\mathcal{U}$ is not trivial subcritical [27, Corollary 18.3.2] (the easier "only if" direction is contained in the proof of [27], but also follows from [2, Lemma 7.3]). On the other hand, a classical comparison with a branching process shows that $q_{\mathrm{c}}^{\mathrm{CP}} > 0$ for any $\mathcal{U}$ (see e.g. [37, Section I.1]).

We also extend Eq. (2.3) to $\mathcal{U}$-CP without change.

2.4 Cellular automata with death

A *cellular automaton* (CA or ϕ-CA when we want to emphasise the dependence on the map ϕ) is specified by a map $\phi : \Omega \to \{0,1\}$ depending only on finitely many sites. Given an initial condition $\omega(0) \in \Omega$, we inductively define for all $t \ge 1, x \in \mathbb{Z}^d$

$$\omega_x(t) = \phi\left(\omega_{\cdot - x}(t-1)\right).$$

In other words, the map ϕ is applied at each site simultaneously in a translation invariant way. The CA is said to be an *eroder*, if for all finite $A \subset \mathbb{Z}^d$ there exists $T(A) < \infty$ such that $\omega(0) = \mathbb{1}_{\mathbb{Z}^d \setminus A}$ implies $\omega(T(A)) = \mathbf{1}$, that is, finite sets of 0s become extinct after a finite time. The CA is *attractive*, if ϕ is non-decreasing for the pointwise partial order on Ω. That is, if $\omega_x \geq \omega'_x$ for some $\omega, \omega' \in \Omega$ and all $x \in \mathbb{Z}^d$, then $\phi(\omega) \geq \phi(\omega')$.

Given a cellular automaton ϕ, we further consider its version with *death* as follows. For $x \in \mathbb{Z}^d$ and $t \in \mathbb{N}$, let $\xi_{x,t}$ be i.i.d. Bernoulli variables with parameter $\delta \in [0,1]$. Then the automaton with map ϕ and δ death is defined by

$$\omega_x(t) = \begin{cases} \phi(\omega_{\cdot - x}(t-1)) & \xi_{x,t} = 0, \\ 0 & \xi_{x,t} = 1, \end{cases} \tag{2.6}$$

for all $t \geq 1$ starting from a given initial condition $\omega(0) \in \Omega$. That is, at each space-time point, we apply the map ϕ, as in the ϕ-CA, with probability $1 - \delta$ and we change the state to healthy with probability δ.

For attractive ϕ, we further denote by $\bar{\nu}$ the *upper invariant measure*—the limit of the law of the cellular automaton with δ death and initial condition $\mathbf{1}$, $\omega^{\mathbf{1}}(t)$, as $t \to \infty$. One then defines its stability threshold

$$\delta_{\mathrm{c}} = \sup \left\{ \delta \in [0,1] : \bar{\nu} \neq \delta_{\mathbf{0}} \right\},$$

where $\delta_{\mathbf{0}}$ is the Dirac measure on $\mathbf{0}$. This threshold captures the point up to which death is not strong enough to extinguish infection, if we start from the completely infected state. It is a classical result of Toom [50] that $\delta_{\mathrm{c}} > 0$ if and only if the ϕ-CA is an eroder.

2.5 Bootstrap percolation

The $\mathcal{U}$-bootstrap percolation (BP) is the particular CA whose map ϕ is defined via

$$\phi(\omega) = \max(\omega_0, c_x(\omega)) \tag{2.7}$$

for any $\omega \in \Omega$ with c_x from Eq. (2.1). In other words, infected sites remain infected, while healthy ones become infected once their constraint is satisfied (i.e. there are enough infections around in the sense of $\mathcal{U}$). One commonly considers

$$q_{\mathrm{c}}^{\mathrm{BP}} = \inf \left\{ q \in [0,1] : \lim_{t \to \infty} \mathbb{P}\left(\omega_0^{\mu_q}(t) = 0\right) = 0 \right\},$$

where ω^{μ_q} is the $\mathcal{U}$-BP with initial condition distributed according to μ_q (there is no other randomness involved). This threshold reflects at which point initial infections with density q are sufficient to almost surely infect the entire lattice. It is known [12, Proposition 2.5] that in fact for the same update family $\mathcal{U}$ we have $q_c^{\mathrm{KCM}} = q_c^{\mathrm{BP}}$.

Since $\mathcal{U}$-BP is a CA, one may consider its version with death. That is, for $\delta \geq 0$, $\mathcal{U}$-BP with δ death is the ϕ-CA with δ death with ϕ from Eq. (2.7). Recalling Section 2.4, this means that at each time step infected vertices become healthy with probability δ and stay infected with probability $1-\delta$; healthy vertices whose constraint is satisfied become infected with probability $1-\delta$; healthy sites whose constraint is not satisfied remain healthy with probability 1.

2.6 Last passage percolation

Given an update rule $U \subset \mathbb{H}_u$ for some $u \in S^{d-1}$, we define the U-last passage percolation (LPP) on $\Lambda = \{1, \dots, n\}^d$ as follows. Endow each $x \in \Lambda$ with an i.i.d. exponentially distributed random variable $T(x)$ with mean 1. For each $x \notin \Lambda$ set $s_x = 0$. For every $x \in \Lambda$ inductively define the *U-LPP time of x* by

$$s_x = T(x) + \max_{y \in U} s_{x+y}. \tag{2.8}$$

Indeed, this is possible, because U is contained in an open half-plane. Another way to view U-LPP is the following. Sites in Λ are initially healthy and those in $\mathbb{Z}^d \setminus \Lambda$ are infected. When all neighbours of a site (in the sense of U) are infected, it becomes infected at rate 1 and never heals afterwards.

With this representation in mind, it is not hard to check that the set of vertices $x \in \mathbb{Z}^d$ where $s_x \leq t$ for the U-LPP coincides with the set of infected vertices at time t in the configuration of the $\{U\}$-KCM in the box Λ, with entirely infected boundary condition, healthy initial condition and parameter $q = 1$. Note also that if $U = \{0, -1\}^d \setminus \{0\}$, the U-LPP on Λ coincides with the standard $\{-e_1, \dots, -e_d\}$-LPP on Λ. Indeed, e.g. $s_{x-e_1-e_2} \leq s_{x-e_1} \leq s_x$ for any $x \in \Lambda$ for $\{-e_1, \dots, -e_d\}$-LPP, so the maximum over U in Eq. (2.8) coincides with the maximum over $\{-e_1, \dots, -e_d\} \subset U$.

3 Results

We are now ready to state our main results. The first one concerns the mixing time of KCM or CP in finite volume with infected boundary.

Theorem 3.1 (Linear mixing). *Let $\mathcal{U}$ be an update family, which is not trivial subcritical. There exists $\varepsilon = \varepsilon(\mathcal{U}) > 0$ such that for all $q \in [1-\varepsilon, 1]$ the $\mathcal{U}$-KCM on $\{1, \dots, n\}^d$ with infected boundary condition exhibits precutoff in* *linear time: there exists $C = C(\mathcal{U}) > 0$ such that for all $\delta \in (0,1)$ and n large enough depending on δ,*

$$n/C \leq t_{\mathrm{mix}}(\delta) \leq Cn. \tag{3.1}$$

The same holds for the $\mathcal{U}$-CP.

Remark 3.2. Let us note that all the conditions above are essential. Indeed, trivial subcritical models are simply not ergodic, as they admit finite healthy regions which cannot change [2, Lemma 7.3]. Moreover, there exist non-trivial subcritical models with $q_{\mathrm{c}}^{\mathrm{KCM}}$ arbitrarily close to 1 [30, Proposition 7.1], so one cannot hope ε to be independent of $\mathcal{U}$. Finally, one cannot change the boundary condition to healthy or periodic (restricting to the ergodic component), since it is known that even 2-neighbour bootstrap percolation with these boundary conditions may have quadratic infection time [5], and the $\mathcal{U}$-BP infection time is a lower bound on the $\mathcal{U}$-KCM one (see [42, Lemma 4.3]).

Our second main result establishes that, in infinite volume, $\mathcal{U}$-KCM converge exponentially fast to their equilibrium measure.

Theorem 3.3 (Exponential convergence)**.** *Let $\mathcal{U}$ be any update family and $\alpha > 0$. Then, there exist $\varepsilon \in (0,1)$ and $c > 0$, such that for any $p \in [\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}}+\alpha, 1]$ and $q \in [1-\varepsilon, 1]$ the following holds. Let $(\eta^{\mu_p}(t))_{t\geq 0}$ be the infinite volume $\mathcal{U}$-KCM with initial distribution[2] μ_p and parameter q. Then for all local functions $f : \Omega \to \mathbb{R}$ and $t \geq 0$*

$$|\mathbb{E}\left[f(\eta^{\mu_p}(t))\right] - \mu_q(f)| \leq e^{-ct}\|f\|_\infty \cdot |\operatorname{supp} f|/c, \tag{3.2}$$

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

As in Remark 3.2, one cannot hope for ε independent of $\mathcal{U}$. However, one should expect both Theorems 3.1 and 3.3 to hold for any $q > q_{\mathrm{c}}^{\mathrm{KCM}} + \alpha$.

Remark 3.4. Since $\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}}$ appears directly in Theorem 3.3, let us mention that for supercritical and critical models it is known that $\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}} = q_{\mathrm{c}}^{\mathrm{KCM}} = 0$ [3, 30].[3] For trivial subcritical update families we have $\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}} = q_{\mathrm{c}}^{\mathrm{KCM}} = 1$

[2]The initial condition is assumed to be product mostly for simplicity. It will be clear from the proof that e.g. any initial condition stochastically dominating μ_p would do.

[3] For supercritical and critical models [3] gives a stretched exponential decay of the tail of the infection time of the origin for any $q > 0$, so $\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}} = 0$ for non-subcritical update families by [30, Theorem 3.7]. Note that, while [30] is formulated in two dimensions, the parts we will use do generalise rather straightforwardly to higher dimensions.

[2, 12], so that Theorem 3.3 is empty, as it should. For subcritical non-trivial families it is known that $0 < q_\mathrm{c}^\mathrm{KCM} \le \tilde{q}_\mathrm{c}^\mathrm{KCM} < 1$ [2, 12, 30] and the second inequality is believed to be an equality, but this is an important open problem. Equality has been shown for oriented update families [32] and the general case was reduced to a related question involving only $\mathcal{U}$-BP in [30]. Furthermore, Theorem 3.3 cannot hold for $p < q_\mathrm{c}^\mathrm{KCM} = q_\mathrm{c}^\mathrm{BP}$, since then a.s. there are sites whose state remains 0 forever. Hence, if the conjecture $q_\mathrm{c}^\mathrm{KCM} = \tilde{q}_\mathrm{c}^\mathrm{KCM}$ holds, the range of values for p in the theorem is the best possible.

Finally, let us mention that in the course of the proof of our main results we will derive consequences on exponential decay in space-time for eroder attractive cellular automata with death, and on space exponential decay for the set of sites eventually reaching state 1 in $\mathcal{U}$-BP for subcritical $\mathcal{U}$ with initial condition μ_q. The reader interested in these developments (Corollaries 9.5 and 9.7) can directly refer to Section 9, which can be read independently of the remainder of the paper.

4 Background

Before turning to the proof of Theorems 3.1 and 3.3, let us discuss previous work on KCM out of equilibrium. We start by mentioning that results of a different kind, concerning graphs whose size diverges jointly with the parameter q tending to 0, can be found in [15, 33, 44, 45], while large deviations in trajectory space have been studied in [10].

Along the lines of our work, much more has been done, mostly for supercritical models, especially the $\{\{-e_1\}, \dots, \{-e_d\}\}$-KCM called the East model and FA-1f, that is, the $\{\{-e_1\}, \dots, \{-e_d\}, \{e_1\}, \dots, \{e_d\}\}$-KCM. In all cases roughly the same route has been followed, to the extent possible, along the following steps in that order, each one relying on the previous one.

Step 1. Theorem 3.3, possibly with a weaker stretched exponential decay.

Step 2. Positive speed of the infection front and the corresponding precutoff, that is, Theorem 3.1.

Step 3. Ergodicity of the process seen from the front and law of large numbers for the front position.

Step 4. CLT for the front position and cutoff.

Step 1 was performed first for the East model in $d = 1$ in [14]. Like all results on this model, that relies on orientation and further favourable properties, allowing the results to hold for all $q > 0$. Certain qualitative convergence results for a model closely related to FA-1f were obtained in [48] (also see there references therein). For FA-1f, Step 1 was done with stretched exponential decay for FA-1f for $q > 1/2$ in [8] (pure exponential in $d = 1$). For the East model in $d > 1$ stretched exponential decay was proved in [16]. Convergence for FA-1f was improved to pure exponential in [43] for q large enough. For the East model in $d > 1$ the same was done in [40]. The exponential decay was then generalised to all supercritical models in any dimension for q large enough in [41], thus including the ones of [8, 43], as well as [14, 16, 40] up to the restriction on q. The most general of the above results, [41], is contained in Theorem 3.3. In summary, before the present work, Theorem 3.3 was known only for supercritical KCM.

Turning to Step 2, Theorem 3.1 was proved for the East model for any $q > 0$ in [16]. We believe that even for FA-1f in dimension $d > 1$ Theorem 3.1 is new. However, there is another important work in the direction of Theorem 3.1. Namely, in [18] this result was proved with a weaker upper bound of order $n \log n$, assuming that the update family is oriented. It should be noted that orientation rules out the possibility for the model to be critical and is a very convenient feature for the analysis, as we will see. On the other hand, the approach of [18] has the major advantage of working for any $q > q_{\mathrm{c}}^{\mathrm{KCM}}$ owing to [32, Theorem 1.6] and [30, Theorem 3.7].

Moving on to Step 3, in $d = 1$ it has been established that the front has a well-defined speed and that the law of the configuration behind the front converges to a limit for large times. This is done for the East model in [7] for any $q > 0$ and in [9] for FA-1f for q large enough. In [17, 19] the East model was studied in $d > 1$ with the aim to examine the limit shape of the set of updated sites starting from a single infection. Results are still far from establishing that such a limit shape actually exists, but some control on the speed of the front in different directions is obtained.

Finally, Step 4 was achieved in $d = 1$ for the East model in [24] (see [21, 22] for further results about the $d = 1$ East process out of equilibrium) and for FA-1f at q close to 1 in [20]. This was also obtained for the East model in higher dimensions for a particular domain and boundary condition in [19]. We should also mention that in [35, Theorem 2.4] Step 4 was performed for FA-2f in $d = 2$ in the somewhat degenerate case $q = 1$, which also coincides with the zero-temperature Ising model with appropriate external field. This can also be viewed as a continuous time version of 2-neighbour BP or a non-oriented LPP.

While it would be extremely interesting to see Step 3 and Step 4 established for general update families, this seems rather remote, given that even the d-dimensional East model for $d > 1$ has not been handled at that level yet. As it is the case for the 1-dimensional East model and FA-1f, we expect that our Theorems 3.1 and 3.3 and the tools developed to prove them will play an important role in attacking these questions.

5 Outline of the proof

Let us start by sketching the proof of Theorem 3.3, which is slightly simpler than the one of Theorem 3.1. The proof is composed of several steps corresponding to Sections 6 to 9, which are put together in Section 10.

The first step (see Section 6) consists in "warming up" the initial condition. That is, we improve our initial condition μ_p with $p > \tilde{q}_{\mathrm{c}}^{\mathrm{KCM}}$ to one with high density of infections. This is achieved via a renormalisation drawing on [30]. It roughly says that, since $p > \tilde{q}_{\mathrm{c}}^{\mathrm{KCM}}$, the probability that a site does not become infected within time t in the $\mathcal{U}$-BP dynamics with the same initial condition decays exponentially with t. Since the parameter q of our $\mathcal{U}$-KCM is close to 1, the same holds for it up to a large enough time, since the $\mathcal{U}$-KCM essentially reduces to the $\mathcal{U}$-BP in the absence of recovery events. Hence, looking at a renormalised lattice, we may assume that the initial condition of our dynamics is product with high density of infections.

The second step (see Section 7) is to reduce the study of the $\mathcal{U}$-KCM to the CP with update family consisting of a single rule U_0 that is oriented. While there is a standard monotone comparison (see Claim 7.2) that guarantees that all infections of the $\{U_0\}$-CP are infections in the $\mathcal{U}$-KCM, we need to go further. Namely, we establish that studying a certain set that depends not only on the configuration, but also on the history of the CP, we are able to deduce that the $\mathcal{U}$-KCM not only has lots of infections, but has actually coupled for all initial conditions larger than the one of the CP.

Hence, we have reduced our problem to one about the $\{U_0\}$-CP with parameter q_0 close to 1 and initial condition close to $\mathbf{1}$. In doing so, we have lost the reversibility and the product invariant measure of the KCM, but we have gained attractiveness and orientation for the CP. Nevertheless, we are not done yet, because the $\{U_0\}$-CP is by far not as simple as the classical CP, as its dynamics is still cooperative, because U_0 may contain more than one vertex.

The third step (see Section 8) is a further renormalisation generalising and somewhat simplifying the one of [41], itself stemming from [43].

It transforms the $\{U_0\}$-CP into a $\mathcal{U}_0$-BP with (little) death, where $\mathcal{U}_0 = \{\{0,-1\}^d \setminus \{0\}\}$. To achieve this, we tessellate space-time into large boxes of carefully chosen geometry. Roughly speaking, we ensure that if all neighbours of a box in the directions given by $\mathcal{U}_0$ are fully infected, then infection propagates with high probability to the box of interest. Moreover, an infected box remains such at the next time step with high probability. If either of these high probability events fails, we view that as a death in the renormalised BP process. We have thus made our process even simpler, as it now evolves in discrete time and no longer depends on the original update family $\mathcal{U}$, while remaining cooperative.

The fourth step (see Section 9) is to show that the $\mathcal{U}_0$-BP is exponentially unlikely to have large space-time clusters of healthy sites. This can be traced back through the previous steps to the CP and then the KCM to yield Theorem 3.3 (see Section 10). In fact, we prove this in general for any CA with death which is an eroder (recall Section 2.4), by developing a novel and very general scheme for leveraging exponential bounds on the probability of occurrence of objects rooted at a given point to exponential bounds on clusters of such objects and applying this to Toom contours [50], as recently revisited in [49]. Alternatively, one could employ a multi-scale renormalisation argument, as in [2] transported to the setting of CA with death via [32], but this would degrade Theorem 3.3 to stretched exponential convergence at best.

Finally, we turn to the proof of Theorem 3.1, which still relies on all of the above, but requires a substitute for the initial "warming up" step above, since we need to deal with arbitrary initial conditions, including $\mathbf{0}$. This alternative first step (see Section 11) consists in yet another renormalisation, this time from the $\{U_0\}$-CP to the standard LPP. We show that after the LPP time at a renormalised site, the corresponding box is coupled for the CP. To do this, somewhat surprisingly, we look at times when sites become healthy in the CP. Using the orientation of the $\{U_0\}$-CP, we have that, once all sites on whose state the constraint at a given site v depends on are coupled, it remains to wait for a single update at v to the state 0, in order to couple the state at v, too. Although these updates are rare (q is close to 1), it still takes a time of order 1 to couple the entire box corresponding to a renormalised site, given that the ones it depends on (in the LPP sense) are already coupled. Thus, to ensure the CP is coupled on a box Λ of size n, it suffices to wait until the LPP on a (renormalised) box of size of order n reaches all sites. The latter is known to be linear from [28].

Hence, after a linear time the CP has reached its equilibrium distribution in the box with $\mathbf{1}$ boundary condition. By attractiveness of CP, this distribution stochastically dominates the restriction to the box Λ of the infinite volume upper stationary measure. Moreover, the CP lower bounds the KCM with any initial condition, so, after this burn-in time, we may perform the same procedure as for the proof of Theorem 3.3. Namely, we exploit the relation between $\{U_0\}$-CP and $\mathcal{U}$-KCM, renormalise the former to $\mathcal{U}_0$-BP with death and, finally, use that the latter has exponentially small probability to have large healthy space-time clusters. Using this exponential bound, we get that it suffices to wait for a time of order $\log n$ to ensure that the $\mathcal{U}$-KCM has coupled with high probability, once we have waited for the initial linear burn-in time needed for the LPP to reach all sites in the renormalised version of Λ.

6 Warming up the initial condition

In this section we start by showing that if the initial condition is μ_p, with $1-p<1-\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}}$ possibly much larger than the equilibrium density $1-q$ of healthy sites, after a sufficiently large but finite time the law of the process dominates a renormalised Bernoulli measure with *large* infection density, but still not the equilibrium one.

Let us note that the vectors v_i' in the next lemma are arbitrary at this point. The reader is encouraged to think of them as the canonical basis of $\mathbb{R}^d$, while a more convenient choice will appear in Section 8.

Lemma 6.1. *Let $\mathcal{U}$ be an update family, $\alpha>0$, $\varepsilon_0>0$ and let $v_1',\dots,v_d'\in\mathbb{Z}^d$ be linearly independent. There exists $R_0\in\{1,2,\dots\}$ such that for any $R\ge R_0$ there exists T_0 such that for any $T\ge T_0$ there exists $\varepsilon_1>0$ such that for any $q\in[1-\varepsilon_1,1]$ the following holds. Set*

$$\hat{B}=\sum_{i=1}^{d}(v_i'[0,R))=\left\{\sum_{i=1}^{d}a_iv_i':(a_i)_{i=1}^{d}\in[0,R)^d\right\}\subset\mathbb{R}^d$$

and $\hat{B}_x=\hat{B}+\sum_{i=1}^{d}Rx_iv_i'$ for every $x\in\mathbb{Z}^d$. We can couple all $\mathcal{U}$-KCM $\eta^{\mu_{p'}}$ with parameter q and initial conditions $\mu_{p'}$ for $p'\ge\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}}+\alpha$ together with $\xi\sim\mu_{1-\varepsilon_0}$ so that for every $x\in\mathbb{Z}^d$,

$$\xi_x=1\Rightarrow\ \forall p'\ge\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}}+\alpha,\eta_{\hat{B}_x}^{\mu_{p'}}(T)=\mathbf{1}_{\hat{B}_x}.$$

Proof. Fix $\alpha>0$ and set $p=\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}}+\alpha$. From [30, Theorems 3.5 and 3.7] we have that there exists $c=c(\alpha)\in(0,\infty)$ such that the $\mathcal{U}$-BP (without death) ω satisfies $\mathbb{P}(\omega_0^{\mu_p}(t)=0)\le e^{-ct}$ for every integer $t\ge0$. Hence, for any fixed $\varepsilon' > 0$, we can choose $R > 0$ large enough depending on c and ε' such that $\sqrt{R} \in \mathbb{N}$ and

$$\mathbb{P}\left(\omega_{\hat{B}}^{\mu_p}\left(\sqrt{R}\right) = \mathbf{1}_{\hat{B}}\right) \geq 1 - \varepsilon'.$$

Since R is large enough, the above event only depends on $\omega^{\mu_p}(0)$ restricted to $\bigcup_{z\in\{-1,0,1\}^d} \hat{B}_z$.

Recall that P_x is the Poisson process of clock ring times associated to $x \in \mathbb{Z}^d$ from Section 2.2 used to couple the $\mathcal{U}$-KCM with all initial conditions. We further assume initial conditions distributed according to $\mu_{p'}$ for $p' \geq p$ to be coupled in a monotone way. Let us now choose $T > 0$ large enough depending on R in such a way that with probability at least $1 - \varepsilon'$ for each $i \in \{0, \dots, \sqrt{R} - 1\}$ and $x \in \bigcup_{z\in\{-1,0,1\}^d} \hat{B}_z$ we have $P_x \cap (iT/\sqrt{R}, (i+1)T/\sqrt{R}) \neq \varnothing$. That is, in each of these $\sqrt{R}$ intervals of time the clock of each site rings. Finally, if q is close enough to 1 depending also on T, we get that with probability at least $1-\varepsilon'$, we have $\Upsilon_x(t) \leq q$ for all $t \in P_x \cap [0, T)$ and $x \in \bigcup_{z\in\{-1,0,1\}^d} \hat{B}_z$. That is, at each clock ring, we attempt to infect the corresponding site.

We claim that if all three events above occur for some $x \in \mathbb{Z}^d$, that is, $\omega_{\hat{B}_x}^{\mu_p}(\sqrt{R}) = \mathbf{1}_{\hat{B}_x}$, $P_y \cap (iT/\sqrt{R}, (i+1)T/\sqrt{R}) \neq \varnothing$ and $\Upsilon_y(t) \leq q$, for all $y \in \bigcup_{z\in x+\{-1,0,1\}^d} \hat{B}_z$, $i \in \{0, \dots, \sqrt{R} - 1\}$ and $t \in P_y \cap [0, T)$, then

$$\eta_{\hat{B}_x}^{\mu_{p'}}(T) = \mathbf{1}_{\hat{B}_x} \text{ for any } p' \geq p. \tag{6.1}$$

Since $\omega_{\hat{B}_x}^{\mu_p}(\sqrt{R}) = \mathbf{1}_{\hat{B}_x}$ and R is large (hence much larger than $\sqrt{R}$), there exists X with $\hat{B}_x \subset X \subset \bigcup_{z\in\{-1,0,1\}^d} \hat{B}_{x+z}$ such that the $\mathcal{U}$-BP process $\omega^{\mu_p, X}$ on X with boundary condition $\mathbf{0}_{\mathbb{Z}^d\setminus X}$ and initial state $\omega_X^{\mu_p}(0)$ satisfies $\omega^{\mu_p,X}(\sqrt{R}) = \mathbf{1}_X$. [4] Then we can prove by induction on $i \in \{0, \dots, \sqrt{R}\}$ that for all $y \in X$ we have $\eta_y^{\mu_{p'}}(iT/\sqrt{R}) \geq \omega_y^{\mu_p,X}(i)$.

The base is the monotone coupling of the initial conditions. For the induction step, observe that by assumption, for the KCM, no attempt is made to change any state to 0, but at least one attempt is made to change each site to 1. Since the constraint c_y is non-decreasing in the configuration and the $\eta_X^{\mu_{p'}}$ process is non-decreasing in time, if for some $y \in X$ and $i \in \{0, \dots, \sqrt{R} - 1\}$ we have $c_y(\omega^{\mu_p,X}(i)) = 1$, then for any $t \in$

[4] To see this, one may remove the infections at distance more than $C\sqrt{R}$ from $\hat{B}_x$ for some large constant $C > 0$, since they do not reach X by time $\sqrt{R}$, and take X to be the set of all sites infected by the remaining ones up to time $\sqrt{R}$.

$(iT/\sqrt{R},(i+1)T/\sqrt{R})$ we have $c_y(\eta^{\mu_{p'}}(t)) = 1$. Applying this to some $t \in P_y \cap (iT/\sqrt{R},(i+1)T/\sqrt{R}) \neq \varnothing$, we obtain the induction step:

$$\begin{aligned}\eta_y^{\mu_{p'}}\left((i+1)T/\sqrt{R}\right) &\geq \eta_y^{\mu_{p'}}(t) = c_y\left(\eta^{\mu_{p'}}(t)\right) \vee \eta_y^{\mu_{p'}}(t-)\\ &\geq c_y\left(\omega^{\mu_p,X}(i)\right) \vee \eta_y^{\mu_{p'}}\left(iT/\sqrt{R}\right)\\ &\geq c_y\left(\omega^{\mu_p,X}(i)\right) \vee \omega_y^{\mu_p,X}(i) = \omega^{\mu_p,X}(i+1).\end{aligned}$$

Hence, the claimed Eq. (6.1) follows by taking $i = \sqrt{R}$.

Thus, for every $x \in \mathbb{Z}^d$, on an event E_x of probability at least $1 - 3\varepsilon'$, the $\mathcal{U}$-KCM with initial condition $\mu_{p'}$ coupled in a monotone way satisfy Eq. (6.1). Moreover, by construction the events E_x are 1-dependent in terms of x, so by the Liggett–Schonmann–Stacey Theorem [39], the set of x such that E_x is realised stochastically dominates an i.i.d. configuration with parameter at least $1 - \varepsilon_0$ such that $\varepsilon_0 \to 0$ if $\varepsilon' \to 0$. □

7 Coupling KCM and CP

In this section we examine the coupling between KCM and CP. Let $\mathcal{U}$ be an update family which is not trivial subcritical and let

$$\|\mathcal{U}\| = \max_{U \in \mathcal{U}, x \in U} \|x\|. \tag{7.1}$$

Fix $U_0 \in \mathcal{U}$ such that $U_0 \subset \mathbb{H}_u$ for some $u \in S^{d-1}$ (this is possible since the update family $\mathcal{U}$ is assumed not to be a trivial subcritical one). Let us fix a domain $\Lambda \subset \mathbb{Z}^d$ and boundary condition $\tau \in \Omega_{\mathbb{Z}^d \setminus \Lambda}$. Fix two parameters $0 \leq q_0 \leq q \leq 1$. Fix an initial condition $\xi \in \Omega_\Lambda$ and denote by ζ the $\{U_0\}$-CP on Λ with boundary condition τ, initial condition ξ and parameter q_0. Recall from Sections 2.2 and 2.3 that KCM and CP for all update families, initial conditions, domains, boundary conditions and parameters are coupled on the same probability space using the same clock rings (P_x) and the same unifrom random variables $(\Upsilon_x(t))_{t \in P_x}$.

We next consider a set which will contain the discrepancies between $\mathcal{U}$-KCM with different initial conditions, based on the the trajectory of the $\{U_0\}$-CP. It can be seen as an analogue of second class particles for the exclusion process [37, Section III.1], the envelope probabilistic cellular automaton [11, Section 4.2] and is similar to the idea of [26, Section 1.3]. We define the set $O_t \subset \{x \in \Lambda : \zeta_x(t) = 0\}$ of *orange* healthy sites to be the càdlàg process with jumps at clock ring times $\bigcup_{x \in \Lambda} P_x$ defined as follows.

We first set $O_0 = \{x \in \Lambda : \xi_x = 0\}$, so that all healthy sites are initially orange. Then, for each $x \in \Lambda$ and $t \in P_x$, we set

$$O_t = \begin{cases} O_{t-} \setminus \{x\} & \zeta_x(t) = 1, \\ O_{t-} \cup \{x\} & \zeta_x(t) = 0, \exists y \in O_{t-}, d(x,y) \le \|\mathcal{U}\|, \\ O_{t-} & \zeta_x(t) = 0, \forall y \in O_{t-}, d(x,y) > \|\mathcal{U}\|. \end{cases} \tag{7.2}$$

In words, orange sites appear when a site becomes healthy close to an orange site, but they disappear whenever a site becomes infected.

It is clear that $x \in O_t$ implies $\zeta_x(t) = 0$, so that O_t is indeed a subset of the healthy sites in the $\{U_0\}$-CP ζ. Indeed, by construction, $\bigcup_{x\in\Lambda} P_x$ are the only times when the $\{U_0\}$-CP ζ may change.

Lemma 7.1. *Consider the $\{U_0\}$-CP ζ on Λ with boundary condition τ, initial condition $\xi \in \Omega_\Lambda$ and parameter q_0. Also consider the $\mathcal{U}$-KCM $\eta^{\mathbf{1}}$ and $\eta^{\xi'}$ with parameter $q \ge q_0$ on Λ with boundary condition τ and initial conditions $\mathbf{1}$ and $\xi' \in \Omega_\Lambda$ respectively, for some $\xi' \ge \xi$. Then almost surely, we have that*

$$\left\{x \in \Lambda : \eta^{\mathbf{1}}_x(t) \neq \eta^{\xi'}_x(t)\right\} \subset O_t \tag{7.3}$$

for any $t \ge 0$. In particular, if $O_t = \varnothing$, then $\eta^{\mathbf{1}}(t') = \eta^{\xi'}(t')$ for all $t' \ge t$ and $\xi' \ge \xi$. The same holds, if we replace the $\mathcal{U}$-KCM by the $\mathcal{U}$-CP.

Before proving the lemma, let us prove the following standard fact.

Claim 7.2. *Consider the $\mathcal{U}$-KCM $\eta^{\xi'}$ and $\{U_0\}$-CP ζ as in Lemma 7.1. Then almost surely, for all $t \ge 0$ we have*

$$\zeta_x(t) \le \eta^{\xi'}_x(t) \text{ for all } x \in \Lambda. \tag{7.4}$$

Proof. We first consider the case Λ finite, so that we can proceed by induction on the clock rings $\bigcup_{x\in\Lambda} P_x$. Equation (7.4) holds at $t = 0$ since $\xi' \ge \xi$. Assume Eq. (7.4) holds for all $t' < t \in P_x$ for some $x \in \Lambda$. If $\Upsilon_x(t) > q_0$, Eq. (7.4) clearly remains true, since $\zeta_x(t) = 0$. On the other hand, if $\Upsilon_x(t) \le q_0 \le q$ and $\zeta_x(t) = 1$, we have two possibilities. If $\zeta_x(t-) = 1$, we are done by using Eq. (7.4) for $t-$. Instead, if $\zeta_x(t-) = 0$, then necessarily $\zeta_{x+U_0}(t-) = \zeta_{x+U_0}(t) = \mathbf{1}$ (recall Eq. (2.4)). But then Eq. (7.4) for $t-$ implies that the constraint c_x (recall Eq. (2.1)) is satisfied in $\eta^{\xi'}(t)$, so that $\eta^{\xi'}_x(t) = 1$ (recall Eq. (2.2)), concluding the proof of the claim for finite Λ.

Next assume Λ is infinite. It follows from the fact that interactions have finite range, that for any $x \in \mathbb{Z}^d$ and $t \ge 0$, the states $\zeta_x(t')$ and $\eta^{\xi'}_x(t')$ for all

$\xi' \in \Omega_\Lambda$ and $t' \in [0,t]$ coincide with those obtained by the same clock rings and uniform random variables on a finite domain Λ' with boundary condition $\mathbf{1}_{\mathbb{Z}^d\setminus\Lambda'}$ with Λ' depending on x, t and the clock rings. This standard fact can be traced back to [29] (also see e.g. [37, Section I.1]). Thus, we can apply the result for finite Λ' to obtain the one for infinite Λ. $\square$

Proof of Lemma 7.1. As in the proof of Claim 9.12, we may assume that Λ is finite and proceed by induction on the clock rings. Since $\xi' \geq \xi$ and initially all healthy sites in ξ are orange, Eq. (7.3) holds at $t = 0$. Fix $x \in \Lambda$ and $t \in P_x$ and assume that Eq. (7.3) holds for any $t' < t$. Further assume for a contradiction that $\eta^{\mathbf{1}}_x(t) \neq \eta^{\xi'}_x(t)$, but $x \notin O_t$. We consider several cases.
Case 1. Assume $\zeta_x(t) = 1$. Then by Claim 7.2 $\eta^{\mathbf{1}}_x(t) = \eta^{\xi'}_x(t) = 1$, so Eq. (7.3) holds, since it holds for $t' < t$ and the only possible change in the left and right hand sides of Eq. (7.3) is at x. For the $\mathcal{U}$-CP instead of the $\mathcal{U}$-KCM, Claim 7.2 is a direct consequence of attractiveness, so the same reasoning applies.
Case 2. Assume $\zeta_x(t) = 0$ and there exists $y \in O_{t-}$ such that $d(x,y) \leq \|\mathcal{U}\|$. Then by definition $x \in O_t$. Moreover, $O_t \setminus \{x\} = O_{t-} \setminus \{x\}$ and

$$\left\{z \in \Lambda \setminus \{x\} : \eta^{\mathbf{1}}_z(t) \neq \eta^{\xi'}_x(t)\right\} = \left\{z \in \Lambda \setminus \{x\} : \eta^{\mathbf{1}}_z(t-) \neq \eta^{\xi'}_x(t-)\right\},$$

so Eq. (7.3) at $t-$ concludes the proof.
Case 3. Assume $\zeta_x(t) = 0$ and there does not exist $y \in O_{t-}$ such that $d(x,y) \leq \|\mathcal{U}\|$. Then from Eq. (7.3) at $t-$ the $\eta^{\mathbf{1}}$ and $\eta^{\xi'}$ processes coincide in the neighbourhood of x, so they also coincide after the attempted update at x, regardless whether it is successful or not. $\square$

8 Space-time renormalisation of CP to BP with death

Recall from Section 7 that $U_0 \in \mathcal{U}$ is such that $U_0 \subset \mathbb{H}_u$ for some $u \in S^{d-1}$. In this section we perform a simple renormalisation of the $\{U_0\}$-CP to the $\mathcal{U}_0$-BP with death (recall from Sections 2.4 and 2.5 that BP is a CA, so we may consider its version with death), where

$$\mathcal{U}_0 = \left\{\{0,-1\}^d \setminus \{0\}\right\}. \tag{8.1}$$

The renormalisation is similar to the ones used in [40, 41, 43], where one obtains oriented percolation as a result of the renormalisation. We start by fixing the relevant geometry.

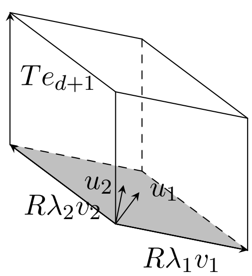

(a) The box B defined in Eq. (8.2) with its base $\hat{B}$, containing the vectors $u_1, \ldots, u_d$, shaded.

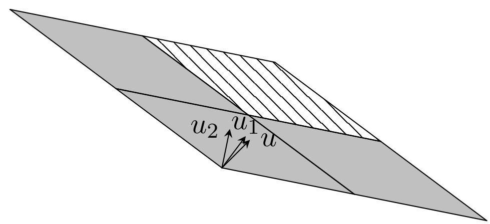

(b) The 2^d box bases $\hat{B}_{x+y}$ for $y \in \{0, -1\}^d$. The Poisson process points p_z occur in the hatched base $\hat{B}_x$ in the order of increasing $\langle z, u \rangle$, as indicated by the hatching direction.

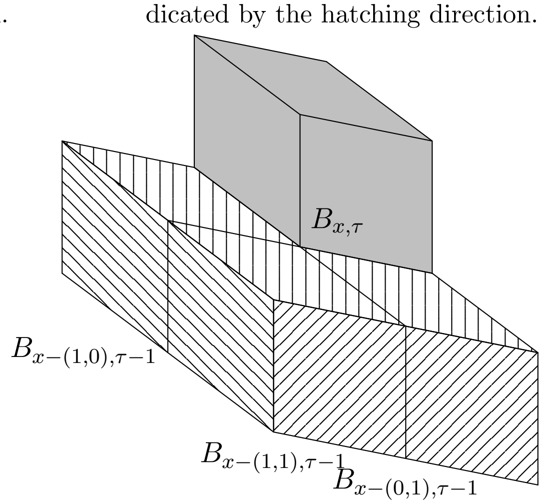

(c) If the shaded box $B_{x,\tau}$ is good, we are able to propagate the infection from the three hatched boxes to $B_{x,\tau}$.

Figure 1: Illustration of the renormalisation of Section 8 in the case $d = 2$.

Fix linearly independent directions $(u_i)_{i=1}^d$ such that:

- for all $i \in \{1, \dots, d\}$, $\lambda'_i u_i \in \mathbb{Z}^d$ for some $\lambda'_i > 0$;
- the u_i are sufficiently close to u so that $U_0 \subset \mathbb{H}_{u_i}$.

The latter condition can be guaranteed thanks to the fact that the $\mathbb{H}_u$ is defined as the open half-space (recall Section 2.1). For each i let us decompose, via the Gram–Schmidt algorithm, $u_i = v_i + u'_i$ with $\langle v_i, u_j \rangle = 0$ for all $j \neq i$ and $u'_i \in \operatorname{span}(\{u_j : j \neq i\})$, that is, the linear span of the remaining vectors. Note that each v_i can be computed via its own Gram–Schmidt process, rather than the entire family, and these vectors are not necessarily orthogonal (see Fig. 1a). By examining the Gram–Schmidt algorithm, one can check that there also exist $\lambda_i > 0$ such that $\lambda_i v_i \in \mathbb{Z}^d$. In what follows we consider time as the $d+1$-th coordinate and we abusively identify $u_i \in S^{d-1}$ with $(u_i, 0) \in S^d$ for all i to lighten notation and similarly we identify v_i with $(v_i, 0)$.

Define the *space-time box* (see Fig. 1a)

$$\begin{aligned} B &= \sum_{i=1}^{d} (\lambda_i v_i [0, R)) + [0, T) e_{d+1} \\ &= \left\{ a \in \mathbb{R}^d : \forall i \in \{1, \dots, d\}, \frac{\langle a, u_i \rangle}{\lambda_i \|v_i\|^2 R} \in [0, 1) \right\} \times [0, T) \subset \mathbb{R}^{d+1} \end{aligned} \tag{8.2}$$

for an integer constant $R > 0$ chosen sufficiently large depending on U_0 and all u_i, v_i and another constant $T > 0$ chosen also sufficiently large depending on R. We refer to $\hat{B} = \sum_{i=1}^d (\lambda_i v_i [0, R)) \subset \mathbb{R}^d$ as the *base* of B and to $\hat{H} = [0, T)$ as its *height*. For any $(x, \tau) = (x_1, \dots, x_d, \tau) \in \mathbb{Z}^{d+1}$ we set $B_{x,\tau} = B + \sum_{i=1}^d R \lambda_i x_i v_i + T \tau e_{d+1}$, which we view as our renormalised space-time points. We similarly define $\hat{B}_x$ and $\hat{H}_\tau$ as the corresponding space and time projections. This notation is consistent with that of Section 6 for $v'_i = \lambda_i v_i$.

For each $(x, \tau) \in \mathbb{Z}^{d+1}$ we define the event $E_{x,\tau}$ that the point (x, τ) is *good* if the following two conditions are satisfied:

- $\Upsilon_z(t) \leq q_0$ for all $z \in \bigcup_{y \in \{0,-1\}^d} \hat{B}_{x+y}$ and $t \in (\hat{H}_\tau \cup \hat{H}_{\tau-1}) \cap P_z$, where we recall from Section 2.2 that P_z is the Poisson process associated to vertex z;
- For each $z \in \hat{B}_x$ there exists $p_z \in P_z \cap \hat{H}_{\tau-1}$ and $\langle z, u \rangle > \langle z', u \rangle$ implies $p_z > p_{z'}$.

The box is called *bad* if it is not good. In words, a good space-time box has a suitable sequence of clock rings and all updates in the box and its $\mathcal{U}_0$-neighbours attempt to infect the corresponding sites. Notice that $E_{x,\tau}$ only depends on the Poisson process points $P_z \cap \{t\}$ and the uniform random variables $\Upsilon_z(t)$ with $(z,t) \in \bigcup_{y\in\{0,-1\}^d} B_{x+y,\tau} \cup B_{x+y,\tau-1}$,

The following lemma provides the desired coupling between the $\{U_0\}$-CP with parameter q_0, denoted ζ, and a $\mathcal{U}_0$-BP with death, denoted $\widetilde{\omega}$. In fact, the $\widetilde{\omega}$ process will not quite be a $\mathcal{U}_0$-BP with death, but is defined in the same way, given the Bernoulli variables $\xi_{x,t}$ from Eq. (2.6), which will however not be i.i.d. (the tilde is there to remind us of this difference). Nevertheless, we will somewhat abusively refer to it as a $\mathcal{U}_0$-BP with death despite this.

Lemma 8.1. *Consider the $\mathcal{U}_0$-BP $\widetilde{\omega}$ with death on $\mathbb{Z}^d$ with death marks $\xi_{x,\tau} = 1$ at points $(x,\tau) \in \mathbb{Z}^d \times \mathbb{N}$ such that $B_{x,\tau}$ is bad. Further let its initial condition be given by the configuration where infected sites x are exactly the sites $x \in \mathbb{Z}^d$ such that:*

- *$\hat{B}_x$ is fully in state 1 in the initial condition $\zeta(0)$ of the $\{U_0\}$-CP with parameter q_0;*
- *$\Upsilon_z(t) \le q_0$ for all $z \in \hat{B}_x$ and $t \in \hat{H}_0 \cap P_z$.*

Then for all $(x,\tau) \in \mathbb{Z}^d \times \mathbb{N}$ and $(y,t) \in B_{x,\tau}$, if $\zeta_y(t) = 0$, then $\widetilde{\omega}_x(\tau) = 0$.

Proof. We prove the statement by induction on $\tau \ge 0$. The base case follows from the definition of the initial condition of $\widetilde{\omega}$. Therefore, let $\tau \ge 1$ and assume for a contradiction that $\zeta_y(t) = 0$, $\widetilde{\omega}_x(\tau) = 1$ for some $(y,t) \in B_{x,\tau}$. If $\widetilde{\omega}_x(\tau-1) = 1$, then by induction $\zeta_y((\tau-1)T) = 1$ and by the first condition for $B_{x,\tau}$ being good (which holds, because otherwise $\widetilde{\omega}_x(\tau) = 0$, since bad boxes are deaths) this remains true until time t, leading to a contradiction. Thus, we may assume that $\widetilde{\omega}_x(\tau-1) = 0$, so that (by the definition Eq. (2.7) of BP) $\widetilde{\omega}_{x+z}(\tau-1) = 1$ for all $z \in \{0,-1\}^d \setminus \{0\}$ and $\zeta_w((\tau-1)T) = 1$ for all $w \in \hat{B}_{x+z}$ by induction hypothesis. Since the box $B_{x,\tau}$ is good, there has been no attempt to put 0 at y after time τT in the contact process, so it suffices to show that $\zeta_y(\tau T) = 1$ to reach a contradiction.

We claim that by our choice of geometry (see Fig. 1),

$$U_0 + \hat{B}_x \subset \bigcup_{z\in\{0,-1\}^d} \hat{B}_{x+z}. \tag{8.3}$$

To see this, notice that

$$\hat{B}_x = \left\{ a \in \mathbb{R}^d : \forall i \in \{1,\dots,d\}, \frac{\langle a, u_i\rangle}{\lambda_i \|v_i\|^2 R} \in [x_i, x_i+1) \right\},$$

$$\bigcup_{z\in\{0,-1\}^d} \hat{B}_{x+z} = \left\{ a \in \mathbb{R}^d : \forall i \in \{1,\dots,d\}, \frac{\langle a, u_i\rangle}{\lambda_i \|v_i\|^2 R} \in [x_i-1, x_i+1) \right\}.$$

Therefore, for any $a \in \hat{B}_x$, $b \in a + U_0$ and $i \in \{1,\dots,d\}$ we have

$$\lambda_i \|v_i\|^2 R(x_i - 1) \le \langle a, u_i\rangle - \|\mathcal{U}\| \le \langle b, u_i\rangle < \langle a, u_i\rangle < \lambda_i \|v_i\|^2 R(x_i+1),$$

if we choose R large enough so that $\lambda_i\|v_i\|^2 R > \|\mathcal{U}\|$ for every i. Hence, $b \in \bigcup_{z\in\{0,-1\}^d} \hat{B}_{x+z}$ as claimed.

Finally, observe that by the second condition for $B_{x,\tau}$ being good, there has been a sequence of attempts at times $p_a \in \hat{H}_{\tau-1}$ to put 1 at each site $a \in \hat{B}_x$. Since the sequence is in the order of increasing scalar product with u and $a + U_0$ is contained in $(a + \mathbb{H}_u) \cap \bigcup_{z\in\{0,-1\}^d} \hat{B}_{x+z}$ (again, provided that R is large enough), the constraint $c_a(\zeta(p_a))$ is fulfilled for each of them, so we are done. □

The following corollary will be more convenient for our purposes.

Corollary 8.2. *Fix U_0 as above and $\delta > 0$. There exist $\varepsilon_0 > 0$ and $R_1 > 0$ such that for any $R \ge R_1$ there exists T_1 such that for any $T \ge T_1$ there exists $\varepsilon_2 > 0$ such that for any $q_0 \in [1-\varepsilon_2, 1]$ the following holds. Consider the $\{U_0\}$-CP ζ with parameter q_0 and initial condition given by $\zeta_y(0) = \xi_x$ for all $y \in \hat{B}_x$ and $x \in \mathbb{Z}^d$, where $\xi \sim \mu_{1-\varepsilon_0}$. Then the trajectory at times $\tau \ge 1$ of the $\mathcal{U}_0$-BP with δ death and initial condition* **1** *is stochastically dominated by the process given by*

$$\hat{\omega}_x(\tau) = \mathbb{1}_{\forall (y,t)\in B_{x,\tau-1}, \zeta_y(t)=1}.$$

Proof. Fix $\varepsilon_0 > 0$ small enough depending on δ. Taking T large after R and then $\varepsilon_2 > 0$ small, it is clear that the probability of a box being good can be made larger than $1-\varepsilon_0$. Moreover, good boxes together with the initial condition $\widetilde{\omega}(0)$ from Lemma 8.1 form a percolation with bounded range of dependence. Therefore, by the Liggett–Schonmann–Stacey Theorem [39] it stochastically dominates an independent Bernoulli field with parameter $1-\delta$, provided ε_0 is small enough. By attractiveness of $\mathcal{U}_0$-BP, this together with Lemma 8.1 completes the proof. Indeed, we simply observed that the first step of $\mathcal{U}_0$-BP with δ death and initial condition **1** has distribution $\mu_{1-\delta}$. □

9 Exponential decay

In this section we establish several exponential decay properties. It can be viewed independently of the rest of the paper and will entail results of independent interest. We therefore adopt a rather abstract and general framework, to keep the approach as flexible as possible. It will be convenient to work in $\mathbb{Z}^D$ with $D \geq 1$, which will play the role of $d+1$ in our original setting.

9.1 Decorated set systems

Definition 9.1 (k-connectivity)**.** Fix a positive real k. We say that a set $X \subset \mathbb{R}^D$ is *k-connected*, if for every $x, y \in X$ there exists a sequence $x_0 = x, x_1, \ldots, x_m = y$ of distinct elements of X such that $d(x_i, x_{i+1}) \leq k$ for all $i \leq m-1$. We call such a sequence a *k-connected path* with *endpoints* x and y.

Definition 9.2 (Decorated set system)**.** For any set $Z \subset \mathbb{Z}^d$ we fix an arbitrary set Γ_Z of *possible decorations*. We allow some Γ_Z to be empty, making the corresponding sets Z impossible to decorate. A *decorated set* is a pair (Z, γ) with $Z \subset \mathbb{Z}^D$ nonempty and bounded and $\gamma \in \Gamma_Z$. Two decorated sets (Z_1, γ_1) and (Z_2, γ_2) are called *disjoint*, if $Z_1 \cap Z_2 = \varnothing$. A *decorated set system* is a probability measure $\mathbb{P}$ and a function E that associates to each decorated set (Z, γ) an event $E(Z, \gamma)$ in such a way that for any finite set of disjoint decorated sets $(Z_i, \gamma_i)_{i \in I}$ we have

$$\mathbb{P}\left(\bigcap_{i \in I} E(Z_i, \gamma_i)\right) \leq \prod_{i \in I} \mathbb{P}\left(E(Z_i, \gamma_i)\right). \tag{9.1}$$

For $x \in \mathbb{Z}^D$, we denote by

$$E(x) = \bigcup_{(Z,\gamma): x \in Z \subset \mathbb{Z}^D, \gamma \in \Gamma_Z} E(Z, \gamma)$$

the event that there exists a decorated set containing x whose event occurs.

In all applications of this construction below, we will actually have equality in Eq. (9.1) but we work under this more general condition, as the proof of Proposition 9.3 works exactly the same with inequality or with equality. In what follows for any $X \subset \mathbb{R}^D$ we denote $\operatorname{diam}(X) = \sup_{x,y \in X}(d(x, y))$ with the convention $\operatorname{diam}(\varnothing) = -\infty$.

Proposition 9.3. *Consider a decorated set system. Assume that for some $C > 1$ and $\epsilon > 0$ small enough depending on C the following hold:*

1. *For all decorated sets (Z, γ) we have* $\mathrm{diam}(Z) \leq C|Z|$.
2. *For every $z \in \mathbb{Z}^D$ the number of decorated sets (Z, γ) such that $|Z| = m$ and $z \in Z$ is at most C^m.*
3. *For all decorated sets (Z, γ) we have $\mathbb{P}(E(Z, \gamma)) \leq \epsilon^{|Z|/C}$.*

Fix $n, k \geq 1$ and $x \in \mathbb{Z}^D$. Let $\mathcal{E}(x, n, k)$ denote the event that there exist $y \in \mathbb{Z}^D$ with $d(x, y) \geq n$ and a k-connected path P with endpoints x and y such that $E(p)$ occurs for each $p \in P$. Then

$$\mathbb{P}(\mathcal{E}(x, n, k)) \leq \epsilon^{n/(7C+7k)^2}.$$

Note that the second condition implies that Γ_Z is finite for every finite Z. Proposition 9.3 will be proved in Section 9.3, but before that, let us provide a few applications to make the abstract setting more concrete.

9.2 Applications

While it is not hard to imagine examples of decorated set systems satisfying the conditions of Proposition 9.3, let us give a more explicit toy example to get used to the notion before turning to more interesting applications based on Toom contours and variants thereof.

Example 9.4. Consider a field of i.i.d. Bernoulli random variables ξ_x for $x \in \mathbb{Z}^D$. For each finite non-empty $Z \subset \mathbb{Z}^D$ the set of decorations Γ_Z is empty if Z is not a 1-connected path and a singleton otherwise. The event $E(Z, \gamma)$ corresponds to $\bigcap_{x \in Z}\{\xi_x = 1\}$. Then Eq. (9.1) is satisfied by independence. Condition 1 of Proposition 9.3 follows by 1-connectedness. Condition 2 holds, because one can encode a 1-connected path by the sequence of its increments, so the number of paths containing $z \in \mathbb{Z}^D$ of cardinality n is at most $n(2D)^n$. Condition 3 is also verified, since $\mathbb{P}(E(Z, \gamma)) = (\mathbb{P}(\xi_0 = 1))^{|Z|}$. Hence, if the Bernoulli variables have a sufficiently small parameter, Proposition 9.3 yields an exponentially small bound on the probability that one can find a k-connected path starting at 0 such that each of its points belongs to a 1-connected path with all Bernoulli variables equal to 1. Thus, in this case this degenerates into looking for a k-connected path in the set of sites with Bernoulli variable equal to 1, so the conclusion of the proposition is a classical fact.

The next corollary will be used to control the $\mathcal{U}_0$-BP with death we recovered in Corollary 8.2. We formulate it more generally for cellular automata with death.

Corollary 9.5. *Let k be a positive real number and let the map ϕ define a CA which is attractive and an eroder, as defined in Section 2.4. Then there exist $c > 0$ and $\delta_0 > 0$ such that for all $\delta \in (0, \delta_0]$ the ϕ-CA $\omega^{\bar{\nu}}$ with death parameter δ and with initial condition given by its upper invariant measure $\bar{\nu}$ satisfies the following. Denote by $A_{x,t}$ the maximal k-connected component containing $(x, t) \in \mathbb{Z}^{d+1}$ with $\omega_a^{\bar{\nu}}(s) = 0$ for all $(a, s) \in A_{x,t}$. It holds that*

$$\mathbb{P}\left(\operatorname{diam}(A_{x,t}) \geq \ell\right) \leq \delta^{c(\ell+1)}.$$

Remark 9.6. Let us note that Corollary 9.5 provides a much more straightforward and general proof of several of the main results of [46] (see Theorems 5, 6 and 7 there).

Proof of Corollary 9.5. Corollary 9.5 follows directly from Proposition 9.3, applied to the decorated set system given by Toom contours [49, 50] and their presence. Since these notions are rather technical to define, while the details of the definition are irrelevant, let us instead highlight the high level viewpoint, referring to [49] for more details (the facts we will need were actually already known since [50]).

A Toom contour rooted at $z \in \mathbb{Z}^{d+1}$ consists of a finite set $Z \subset \mathbb{Z}^{d+1}$ (with $z \in Z$) equipped with a complicated decoration taking the form of a connected coloured oriented multigraph with vertex set Z, such that the endpoints of each edge are at most at some bounded mutual distance (depending on the support of ϕ). Connectedness readily implies condition 1.

In each Toom contour one can identify a set $Z_* \subset Z$ of sinks. A contour is said to be present if $\xi_{x,t} = 1$ (recall Section 2.4) for each sink (x, t), so that disjoint decorated sets occur independently. In particular, Toom contours form a decorated set system and their probability of occurrence is $\delta^{|Z_*|}$. A key and highly non-trivial fact [49, Theorem 7] is that if $\omega_x^{\bar{\nu}}(t) = 0$, then some non-empty finite Toom contour rooted at (x, t) occurs. Moreover, [49, Lemma 13] ensures that the number of edges (and, therefore, the number of vertices by connectedness) of a Toom contour is at most a constant multiple of $|Z_*|$, thus proving condition 3 . Finally, [49, Lemma 14] shows that the number of Toom contours containing a given point and with N edges is at most exponential in N and, therefore, in the number of vertices by the previous result, so condition 2 is also satisfied. $\square$

The next corollary will not be used in the proof of our main results, but we include it, since it is of independent interest and follows analogously.

Corollary 9.7. *Consider $\mathcal{U}'$-BP ω with subcritical $\mathcal{U}'$. Fix $p > 0$ and let C_0 be the k-connected component containing the origin in $\{x \in \mathbb{Z}^d : \lim_{t\to\infty} \omega_x^{\mu_p}(t) = 1\}$. Then there exists $c = c(\mathcal{U}', k) > 0$ such that for all $p > 0$ small enough and $\ell \geq 0$ we have*

$$\mathbb{P}(\mathrm{diam}(C_0) \geq \ell) \leq p^{c(\ell+1)}.$$

We recall that non-subcritical models are exactly those with $q_{\mathrm{c}}^{\mathrm{BP}} = 0$ [2, 3], so $C_0 = \mathbb{Z}^d$ almost surely for all $p > 0$ and it is meaningless to consider them in the above sense.

Remark 9.8. Corollary 9.7 provides a positive answer to [4, Question 12] in the perturbative regime. As pointed out in [30, Section 7.1.2], the stated exponential decay cannot hold for all $p < q_{\mathrm{c}}^{\mathrm{BP}}$ in general. Nevertheless, in the spirit of [1] one could expect that it does hold up to a different critical threshold $p_{\mathrm{c}}(k) \leq q_{\mathrm{c}}^{\mathrm{BP}}$, past which the diameter is a.s. infinite. On a different note, Corollary 9.7 was proved by more classical means in [6, Theorem 4.2] for *trivial* subcritical update families in two dimensions for $k = 1$. Thus, Corollary 9.7 vastly generalises this result and solves [6, Problem 6.1].

Proof of Corollary 9.7. Corollary 9.7 follows from Proposition 9.3 essentially along the same lines as Corollary 9.5 follows from Proposition 9.3, but using a different decorated set system. Namely, we consider the space embeddings of shattered contours, which are the central object of study in [34][5], similarly to Toom contours in [49]. Again, the definition of these objects, which are projections of equivalence classes of Toom contours, is rather technical and unimportant for us, so we refer to [34] for those details. Instead, let us indicate that the presence of a finite non-empty shattered contour rooted at $x \in \mathbb{Z}^d$ is implied by $\lim_{t\to\infty} \omega_x^{\mu_p}(t) = 1$ [34, Corollary 4.3]; their numbers of edges and vertices are bounded by a constant multiple of the number of sinks [34, Lemma 5.2]; the number of shattered Toom contours rooted at a given point is at most exponentially large in the number of sinks [34, Lemma 5.3]. □

Thus, our only remaining task in this section is to prove Proposition 9.3. Before doing so, let us mention a natural question closely related to Corollary 9.5.

[5]Note that in [34] the roles of 0 and 1 are exchanged with respect to the present work.

Question 9.9. *Fix $\varepsilon > 0$. Is it true that for the $\mathcal{U}_0$-BP with δ death, with $\delta > 0$ small enough, the upper invariant measure stochastically dominates $\mu_{1-\varepsilon}$? More generally, is this true for any attractive eroder?*

9.3 Proof of Proposition 9.3

We fix a decorated set system and C, ϵ, n, k, and x as in Proposition 9.3. Since our sets can be quite fuzzy we begin by regularising them.

Definition 9.10. Given a finite non-empty $Z \subset \mathbb{Z}^d$, we set

$$\bar{Z} := \left\{ x \in \mathbb{Z}^d : d(x, Z) \leq 3(1 + \mathrm{diam}(Z)) \right\}.$$

Let us fix a k-connected path $P = (p_0 = x, p_1, \dots, p_l = y)$ with $d(x, y) \geq n$. We further fix decorated sets (Z_p, γ_p) with $p \in Z_p$ for each $p \in P$. We next run the following algorithm.

Algorithm 9.11. *Define $i_0 = 0$, $I_0 = \{0\}$ and $X_0 = \bar{Z}_{p_0}$ and initialise $t = 0$. While $P \not\subset X_t$, repeat the following, then return (I_t, X_t, t). Increment t by setting $t := t + 1$. Set $i_t = \min\{j \leq l : p_j \notin X_{t-1}\}$. Let $J_t = \{j \in I_{t-1} : Z_{p_j} \cap Z_{p_{i_t}} \neq \varnothing\}$. Set $I_t = \{i_t\} \cup (I_{t-1} \setminus J_t)$. Set $X_t = \bigcup_{j \in I_t} \bar{Z}_{p_j}$.*

By definition, if the algorithm terminates and outputs (I_t, X_t, t), then $P \subset X_t$. Moreover, by induction we have that for all $t' \leq t$ and $a, b \in I_{t'}$ with $a \neq b$, it holds that

$$Z_{p_a} \cap Z_{p_b} = \varnothing, \tag{9.2}$$

using the definition of $J_{t'}$. To see that Algorithm 9.11 terminates, it suffices to see that $X_{t'} \cap P$ is strictly increasing in t', since P is finite (on the other hand, we note that $I_{t'}$ is not necessarily monotone in t'). In order to prove this, we first show that $X_{t'} \cap P$ is non-decreasing and then exhibit an element which is in $X_{t'} \cap P$, but not in $X_{t'-1} \cap P$.

Claim 9.12. *For any $t' \in \{1, \dots, t\}$ we have $X_{t'} \cap P \supset X_{t'-1} \cap P$.*

Proof. By the definitions of $X_{t'}$ and $I_{t'}$, we have that

$$X_{t'-1} \setminus X_{t'} \subset \bigcup_{j \in J_{t'}} \bar{Z}_{p_j} \setminus \bar{Z}_{p_{i_{t'}}}. \tag{9.3}$$

Thus, it remains to show that $\bar{Z}_{p_{i_{t'}}} \supset \bar{Z}_{p_j}$ for all $j \in J_{t'}$.

Fix $j \in J_{t'} \subset I_{t'-1}$, so $Z_{p_j} \cap Z_{p_{i_{t'}}} \neq \varnothing$ by the definition of $J_{t'}$. Then

$$\mathrm{diam}\left(Z_{p_{i_{t'}}}\right) \geq d\left(p_{i_{t'}}, p_j\right) - \mathrm{diam}\left(Z_{p_j}\right) \geq d\left(p_{i_{t'}}, p_j\right)/3 + 2 + \mathrm{diam}\left(Z_{p_j}\right),$$

where we first used the triangle inequality, then the fact that $p_{i_{t'}} \notin \bar{Z}_{p_j}$ by definition of $i_{t'}$ and $X_{t'-1}$. But then $\bar{Z}_{p_{i_{t'}}} \supset \bar{Z}_{p_j}$ by the triangle inequality and we are done. $\square$

We next claim that for any $t' \in \{1, \ldots, t\}$ we have

$$p_{i_{t'}} \in X_{t'} \setminus X_{t'-1} = \bar{Z}_{p_{i_{t'}}} \setminus \bigcup_{j \in I_{t'-1}} \bar{Z}_{p_j}. \tag{9.4}$$

Indeed, if $p_{i_{t'}} \in \bar{Z}_{p_j}$ for some $j \in I_{t'-1}$, that would imply $p_{i_{t'}} \in X_{t'-1}$ by the definition of $X_{t'-1}$, but this contradicts the definition of $i_{t'}$.

Combining Claim 9.12 and Eq. (9.4), we get that Algorithm 9.11 does terminate, so I_t, X_t, t are well defined.

Claim 9.13. *X_t is k-connected.*

Proof. We prove by induction that $X_{t'}$ is k-connected for all $t' \in \{0, \ldots, t\}$. The base follows since $X_0 = \bar{Z}_{p_0}$ is k-connected, by Definitions 9.1 and 9.10 and $k \geq 1$. By Eq. (9.4), $X_{t'} \setminus X_{t'-1} \subset \bar{Z}_{p_{i_{t'}}} \subset X_{t'}$, the last inclusion using the definition of $X_{t'}$ and $I_{t'}$. Moreover, by the proof of Claim 9.12 $X_{t'} \supset X_{t'-1}$. Thus, since $\bar{Z}_{p_{i_{t'}}}$ is k-connected, it remains to show that $d(X_{t'-1}, \bar{Z}_{p_{i_{t'}}}) \leq k$. In fact, the stronger statement $d(p_{i_{t'}}, X_{t'-1}) \leq k$ holds, because the definition of $i_{t'}$ gives $p_{i_{t'}-1} \in X_{t'-1}$, and $i_{t'} \neq 0$ for $t' > 0$, since $X_{t'-1} \supset X_0 = \bar{Z}_{p_0} \ni p_0$ and because consecutive p_j are at distance at most k. $\square$

Definition 9.14 (Chain)**.** A *chain starting at x of length at least n* is a sequence of disjoint decorated sets $(V_j, \gamma_j)_{j=1}^m$ such that $d(\bar{V}_j, \bar{V}_{j+1}) \leq k$ for all $j \leq m-1$, $x \in \bar{V}_1$ and there exists $y \in \bar{V}_m$, such that $d(x, y) \geq n$.

By Claim 9.13 and Eq. (9.2) we can extract from I_t a sequence $i'_1, \ldots, i'_m$ such that $(V_j, \gamma_j)_{j=1}^m$ is a chain starting at x of length at least n, where $V_j = Z_{p_{i'_j}}$ and $\gamma_j = \gamma_{p_{i'_j}}$ for $j \in \{1, \ldots, m\}$. Thus, recalling $\mathcal{E}(x, n, k)$ from Proposition 9.3, we have proved the following.

Lemma 9.15. *If $\mathcal{E}(x, n, k)$ occurs, then there exists a chain $(V_j, \gamma_j)_{j=1}^m$ starting at x of length at least n such that $\bigcap_{j=1}^m E(V_j, \gamma_j)$ occurs.*

We are now ready to conclude the proof of Proposition 9.3 by a union bound over all such chains, since their decorated sets are disjoint, so the events involved are negatively correlated.

Proof of Proposition 9.3. By Lemma 9.15 and Definitions 9.2 and 9.14,

$$\mathbb{P}\left(\mathcal{E}(x,n,k)\right) \leq \sum_{(V_j,\gamma_j)_{j=1}^m} \prod_{j=1}^{m} \epsilon^{|V_j|/C}, \tag{9.5}$$

where the sum is over all chains starting at x of length at least n. Observe that by condition 1 of Proposition 9.3, we have

$$n \leq \sum_{i=1}^{m} \left(k + \operatorname{diam}\left(\bar{V}_i\right)\right) \leq mk + 6m + 7\sum_{i=1}^{m} \operatorname{diam}\left(V_i\right) \leq 7(C+k)\sum_{i=1}^{m} |V_i|. \tag{9.6}$$

Further note that for any $i \in \{1, \dots, m-1\}$ the distance between an arbitrarily chosen point in V_i and one in V_{i+1} is at most

$$k + \operatorname{diam}\left(\bar{V}_i\right) + \operatorname{diam}\left(\bar{V}_{i+1}\right) \leq k + 12 + 14C\max(|V_i|, |V_{i+1}|).$$

Therefore, the number of ways to fix the positions of one distinguished point π_i in each V_i with $\pi_1 = x$ is at most

$$\prod_{i=1}^{m-1} \left(1 + 2(k + 12 + 14C\max(|V_i|, |V_{i+1}|))\right)^d \leq \prod_{i=1}^{m} 30(k + C|V_i|)^{2d}. \tag{9.7}$$

Combining Eqs. (9.5) to (9.7) with condition 2 of Proposition 9.3, the probability we seek to bound in Proposition 9.3 is at most

$$\begin{aligned}&\sum_{\substack{m,n_1,\dots,n_m \geq 1\\ \sum_{i=1}^m n_i \geq n/(7(C+k))}} \prod_{i=1}^{m} C^{n_i} 30(k + Cn_i)^{2d} \epsilon^{n_i/C}\\ &\qquad \leq 30 \sum_{N=\lceil n/(7C+7k)\rceil}^{\infty} 2^N C^N (C+k)^{2dN} \epsilon^{N/C} \leq 30\epsilon^{n/(2C(7C+7k))}. \end{aligned} \tag{9.8}$$

For ϵ small enough Eq. (9.8) is clearly at most $\epsilon^{n/(7C+7k)^2}$ as desired, completing the proof of Proposition 9.3. $\square$

10 Assembling Theorem 3.3

In this section we assemble the results of Sections 6 to 9 in order to prove Theorem 3.3. Therefore, let us fix an update family $\mathcal{U}$ which is not trivial subcritical (otherwise $\tilde{q}_{\mathrm{c}}^{\mathrm{KCM}} = 1$ and there is nothing to prove) and $U_0 \in \mathcal{U}$

such that $U_0 \subset \mathbb{H}_u$ for some $u \in S^{d-1}$. Let $v_i' = v_i\lambda_i$ for all $i \in \{1,\dots,d\}$, where v_i, λ_i are chosen as in Section 8. Let δ_0 be as in Corollary 9.5 for $k = \sqrt{d+1}$ and ϕ be the map corresponding to $\mathcal{U}_0$-BP (recall Eq. (8.1)), which is clearly an eroder attractive cellular automaton. Then let ε_0 be as in Corollary 8.2, setting $\delta = \delta_0$. Fix $\alpha \in (0, 1-\tilde{q}_c^{\mathrm{KCM}})$. Let $R = \max(R_0, R_1)$, where R_0 is as in Lemma 6.1 and R_1 is as in Corollary 8.2. Let $T = \max(T_0, T_1)$, where T_0 is as in Lemma 6.1 and T_1 is as in Corollary 8.2. Let ε_1 be as in Lemma 6.1. Finally, let $p \in [\tilde{q}_c^{\mathrm{KCM}} + \alpha, 1]$ and $\varepsilon = \min(\varepsilon_1, \varepsilon_2, 1 - \tilde{q}_c^{\mathrm{KCM}} - \alpha)$, with ε_2 from Corollary 8.2, and $q \in [1-\varepsilon, 1]$.

Let η^{μ_p} denote the $\mathcal{U}$-KCM on $\mathbb{Z}^d$ with parameter q and initial condition with law μ_p and η^{μ_q} be the stationary $\mathcal{U}$-KCM with the same parameter. The initial conditions are coupled so that $\eta_x^{\mu_q}(0) \le \eta_x^{\mu_p}(0)$ for all $x \in \mathbb{Z}^d$, if $q \le p$ and $\eta_x^{\mu_q}(0) \ge \eta_x^{\mu_p}(0)$ for all $x \in \mathbb{Z}^d$, if $p \le q$. Note, however, that, due to the non-attractiveness of KCM, this inequality need not be preserved by the dynamics.

Since η^{μ_q} is stationary, we get that

$$|\mathbb{E}\left[f(\eta^{\mu_p}(t)) - \mu_q(f)\right]| \le 2\|f\|_\infty \cdot \mathbb{P}\left(\eta_S^{\mu_p}(t) \neq \eta_S^{\mu_q}(t)\right), \tag{10.1}$$

where S denotes the support of the local function f. By a union bound over the sites of S and translation invariance, for all $t \ge 0$

$$\mathbb{P}\left(\eta_S^{\mu_p}(t) \neq \eta_S^{\mu_q}(t)\right) \le |S| \cdot \mathbb{P}\left(\eta_0^{\mu_p}(t) \neq \eta_0^{\mu_q}(t)\right). \tag{10.2}$$

We start by running the two KCM up to time T. Then Lemma 6.1 gives that $\min(\eta_y^{\mu_p}(T), \eta_y^{\mu_q}(T)) \ge \xi_x$ for all $x \in \mathbb{Z}^d$, $y \in \hat{B}_x$, where $\xi \sim \mu_{1-\varepsilon_0}$ is suitably coupled with the two KCM.

Let ζ be the $\{U_0\}$-CP on $\mathbb{Z}^d$ starting at time T with parameter $q_0 = 1-\varepsilon_2 \le q$ and initial condition given by $\zeta_y(T) = \xi_x$ for all $y \in \hat{B}_x$ and $x \in \mathbb{Z}^d$. Further recall its orange healthy sites O_t from Section 7, which we now define w.r.t. the initial time T instead of 0. By the Markov property and Lemma 7.1 applied once to $\xi' = \eta^{\mu_p}(T)$ and once to $\xi' = \eta^{\mu_q}(T)$ gives that for all $t \ge T$

$$\mathbb{P}\left(\eta_0^{\mu_p}(t) \neq \eta_0^{\mu_q}(t)\right) \le \mathbb{P}(0 \in O_t). \tag{10.3}$$

Recall Eqs. (7.1) and (7.2) and Definition 9.1. Observe that by construction $0 \in O_t$ implies not only that $\zeta_0(t) = 0$, but also that there exists a $\|\mathcal{U}\|$-connected set $K \subset \mathbb{Z}^d \times [T, \infty)$ in space-time which contains $(0,t)$, intersects the hyperplane $\mathbb{Z}^d \times \{T\}$ and satisfies $\zeta_x(\theta) = 0$ for all $(x,\theta) \in K$. Indeed, whenever a site is added to O_t, it has to be at distance at most $\|\mathcal{U}\|$

from an orange site. Let us call E_t the event that $(0,t)$ belongs to a $\|\mathcal{U}\|$-connected component of space-time points $(x,\theta)\in\mathbb{Z}^d\times[0,t]$ with $\zeta_x(\theta)=0$, that intersects $\mathbb{Z}^d\times\{T\}$. Then we just showed that

$$\mathbb{P}(0\in O_t)\le\mathbb{P}(E_t). \tag{10.4}$$

We are then ready to apply Corollary 8.2. To that end, let ω be the $\mathcal{U}_0$-BP with δ_0-death and initial condition $\mathbf{1}$ (recall Eq. (8.1)). Corollary 8.2 gives that we can couple ω and ζ in such a way that for all $x\in\mathbb{Z}^d$ and $\tau\ge1$, $\omega_x(\tau)=1$ implies that $\zeta_y(t)=1$ for all $(y,t)\in B_{x,\tau}$.

Then for $t\ge T$ the event E_t implies that there is a $\sqrt{d+1}$-connected component of space-time points $(x,\tau)\in\mathbb{Z}^{d+1}$ such that $\omega_x(\tau)=0$ containing both $(0,\lfloor t/T\rfloor)$ and a point in $\mathbb{Z}^d\times\{1\}$. Let us denote this event by $F_{\lfloor t/T\rfloor-1}$, so that that for any $\tau\ge0$

$$\mathbb{P}(E_t)\le\mathbb{P}\left(F_{\lfloor t/T\rfloor-1}\right). \tag{10.5}$$

Since the $\mathcal{U}_0$-BP with δ_0 death is attractive and $\mathbf{1}$ is the maximal initial condition, denoting by $F_\tau^{\bar\nu}$ the event F_τ with ω replaced by the process $\omega^{\bar\nu}$ with initial condition distributed according to its upper invariant measure, we get that for some $c>0$ and any $\tau\ge0$

$$\mathbb{P}(F_\tau)\le\mathbb{P}\left(F_\tau^{\bar\nu}\right)\le\delta_0^{c(\tau+1)}, \tag{10.6}$$

where we applied Corollary 9.5. Note that c is now an absolute constant, as it no longer depends on $\mathcal{U}$, but only on $\mathcal{U}_0$, which is fixed by Eq. (8.1).

Putting Eqs. (10.4) to (10.6) together, we get that for all $t\ge T$

$$\mathbb{P}(0\in O_t)\le\delta_0^{c\lfloor t/T\rfloor}. \tag{10.7}$$

Further recalling Eqs. (10.1) to (10.3), this yields

$$|\mathbb{E}\left[f(\eta^{\mu_p}(t))-\mu_q(f)\right]|\le2\|f\|_\infty|S|\delta_0^{c\lfloor t/T\rfloor},$$

completing the proof of Theorem 3.3.

11 Renormalisation of CP to LPP

In this section we perform a renormalisation somewhat similar to the one from Section 8. Recall $\mathcal{U}_0$ from Eq. (8.1) and U_0 from Section 7 and set $\hat U_0=\{0,-1\}^d\setminus\{0\}$, so that $\mathcal{U}_0=\{\hat U_0\}$. We seek to control the $\{U_0\}$-CP on the box $\Lambda = \{1, \dots, n\}^d$ with parameter $q_0 \in [0, 1)$ in terms of the $\hat{U}_0$-LPP slowed down by some factor. As noted in Section 2.6, the $\hat{U}_0$-LPP coincides with the usual $\{-e_1, \dots, -e_d\}$-LPP, so we will be able to exploit the fact that this model is known to propagate ballistically.

Recall the boxes $\hat{B}_x$ from Section 8 (also recall Fig. 1) and the direction $u \in S^{d-1}$ such that $U_0 \subset \mathbb{H}_u$. For each $x \in \mathbb{Z}^d$ we will define its *passage time* $t_x \in [0, \infty)$. Let

$$\hat{\Lambda} = \left\{ x \in \mathbb{Z}^d : \hat{B}_x \cap \Lambda \neq \varnothing \right\} \tag{11.1}$$

denote the renormalised version of Λ. We start by recalling a result of [28].

Proposition 11.1. *Denote by $(s_x)_{x \in \mathbb{Z}^d}$ the passage times of the standard $\{-e_1, \dots, -e_d\}$-LPP in the set $\hat{\Lambda}$ defined in Eq. (11.1), with $s_x = 0$ if $x \notin \hat{\Lambda}$. Then, there exists $C = C(U_0) < \infty$ such that for every $\delta > 0$, for n large enough (depending on δ) we have*

$$\mathbb{P}\left(\max_{x \in \mathbb{Z}^d} s_x \geq nC \right) \leq \delta. \tag{11.2}$$

Remark 11.2. The proposition is, essentially, a byproduct of the main result of [28]. In that paper, the authors upper bound the mixing time of a discrete-time Markov chain on $d-$dimensional discrete monotone sets. That dynamics depends on a bias parameter λ, that is assumed to be sufficiently large, depending only on the dimension d. In the limit $\lambda \to \infty$, the dynamics reduces to (a discrete-time version of) standard $d-$dimensional LPP. For the reader's convenience, we give a streamlined proof of Proposition 11.1 in the context of (continuous-time) LPP and along the way we slightly improve the main statement of [28] (which, translated into our language, would give Eq. (11.2) with nC replaced by $nC \log(1/\delta)$.)

Proof of Proposition 11.1. First of all, using monotonicity of LPP, we can replace the set $\hat{\Lambda}$ by a cube $\{1, \dots, \ell\}^d$ that contains it. Using the fact that $\hat{\Lambda}$ has diameter upper bounded by n times a U_0-dependent constant, we assume henceforth that $\hat{\Lambda}$ is a cube with $\ell = O(n)$.

We say that a subset $\sigma \subset \hat{\Lambda}$ is a *monotone set* if the conditions that $x \in \sigma$ and $y \preceq x$ (that is, $y \leq x$ componentwise) imply that $y \in \sigma$. We define a continuous-time Markov chain on the collection Σ of monotone sets; the state of the chain at time t is denoted $\sigma(t)$, and we let $\sigma(0) := \varnothing$. Each $x \in \hat{\Lambda}$ has an independent Poisson clock of rate 1. When the clock at x rings, if $x \notin \sigma(t-)$ and if $\sigma(t-) \cup \{x\} \in \Sigma$, then we let $\sigma(t) = \sigma(t-) \cup \{x\}$. Otherwise, nothing happens. It is easy to check that this dynamics is equivalent to standard $\{-e_1,\ldots,-e_d\}$-LPP in the sense that the process $(\sigma(t))_{t\geq 0}$ has the same law as the process $(\{x\in\hat\Lambda : s_x\leq t\})_{t\geq 0}$. In particular, Eq. (11.2) is equivalent to proving that

$$\mathbb{P}\left(\sigma(nC)=\hat\Lambda\right)>1-\delta. \tag{11.3}$$

The trivial measure concentrated on the absorbing state $\sigma=\hat\Lambda$ is stationary for the monotone set dynamics, because $\sigma(t)\supset\sigma(s)$ for $t>s$. For this reason, Eq. (11.3) is equivalent to the fact that the mixing time of the monotone set dynamics satisfies $t_{\mathrm{mix}}(\delta)\leq Cn$, for n large enough. The idea of [28] is to use path coupling with an exponential metric. That is, fix $\gamma>0$ and let $\underline{1}=(1,\ldots,1)\in\mathbb{Z}^d$. Given $\sigma,\sigma'\in\Sigma$ that differ by a single vertex, say, $\sigma=\sigma'\cup\{x\}$, define $d_\gamma(\sigma,\sigma')=e^{-\gamma\langle x,\underline{1}\rangle}$, and extend d_γ to be a distance on the whole state space Σ. The path coupling method (see [28] and [36, Ch. 14] in the discrete time setting, but the method is analogous in continuous time) consists in proving that there exists some $\alpha>0$ such that for each pair (σ,σ') differing by a single vertex there exists a coupling of the processes $(\sigma(t),\sigma'(t))$ with initial conditions σ,σ' such that

$$\frac{d}{dt}\mathbb{E}\left[d_\gamma(\sigma(t),\sigma'(t))\right]\Big|_{t=0}\leq-\alpha d_\gamma(\sigma,\sigma'). \tag{11.4}$$

If this is the case, then

$$t_{\mathrm{mix}}(\delta)\leq\frac{1}{\alpha}\log(\mathrm{diam}_\gamma(\Sigma)/\delta), \tag{11.5}$$

with $\mathrm{diam}_\gamma(\Sigma)$ the diameter of Σ with respect to the metric d_γ. In the case of the exponential metric above, $\log\mathrm{diam}_\gamma(\Sigma)=O(\ell)$ and, recalling that $\ell=O(n)$, Eq. (11.5) implies that $t_{\mathrm{mix}}(\delta)\leq Cn$, for n large enough, as desired. The computation of the time derivative in Eq. (11.4) is very similar to the computation of the discrete-time derivative in [28, Sec. 4], so we do not repeat it: one finds that there is exactly one update (the one at x) that decreases the distance by $d_\gamma(\sigma,\sigma')$, and at most d updates (at vertices neighbouring x) that increase the distance by $e^{-\gamma}d_\gamma(\sigma,\sigma')$. Altogether,

$$\frac{d}{dt}\mathbb{E}\left[d_\gamma\left(\sigma(t),\sigma'(t)\right)\right]\Big|_{t=0}\leq-\left(1-de^{-\gamma}\right)d_\gamma(\sigma,\sigma'), \tag{11.6}$$

which gives Eq. (11.4) with $\alpha>0$ for γ large enough. This concludes the proof. $\square$

Recall the clock rings $(P_x)_{x\in\mathbb{Z}^d}$ and uniform random variables $(\Upsilon_x(t))_{t\in P_x}$ from Section 2.2. We set $t_x = 0$ for all $x \in \mathbb{Z}^d \setminus \hat{\Lambda}$. For $x \in \hat{\Lambda}$, we define $t_x \geq 0$ by induction as follows, assuming that t_{x+y} is already defined for all $y \in \hat{U}_0$. The stopping time t_x is the first time $t > \tilde{t}_x := \max_{y\in\hat{U}_0} t_{x+y}$ such that for every $z \in \hat{B}_x \cap \Lambda$ there exists $p_z \in P_z \cap [\tilde{t}_x, t]$ such that:

- $\Upsilon_z(p_z) > q_0$,
- the collection $(p_z)_{z\in\hat{B}_x}$ satisfies that $\langle z, u\rangle > \langle z', u\rangle$ implies $p_z > p_{z'}$.

In words, once the boxes $\hat{B}_{x+y}$ for $y \in \hat{U}_0$ have been treated, we require the occurrence of a sequence of clock rings corresponding to healing in the $\{U_0\}$-CP and occurring in the order of increasing scalar product with the direction u.

The use of these passage times is clear in the following lemma, where, as usual, ζ denotes the $\{U_0\}$-CP on Λ with parameter q_0 and boundary condition $\mathbf{1}_{\mathbb{Z}^d\setminus\Lambda}$.

Lemma 11.3. *For any $x \in \mathbb{Z}^d$, $z \in \hat{B}_x$ and $t \geq t_x$ we have $\zeta_z^{\mathbf{1}}(t) = \zeta_z^{\mathbf{0}}(t)$.*

Proof. If $z \in \hat{B}_x \setminus \Lambda$, the claim is trivial because $\zeta_z^{\mathbf{1}}(t) = \zeta_z^{\mathbf{0}}(t) = 1$ for all t. Otherwise, we proceed by induction and we assume that the claim has been proven for all $z' \in \hat{B}_y$, $y \in x + \hat{U}_0$, and we want to prove it for $z \in \hat{B}_x$. By the induction hypothesis, the two processes are perfectly coupled for $t \geq \tilde{t}_x$ in the boxes $\hat{B}_y$ for $y \in x + \hat{U}_0$.

For $z, z' \in \mathbb{Z}^d$ we write $z \prec z'$ if $\langle z, u\rangle < \langle z', u\rangle$. Within the box $\hat{B}_x$, we proceed by induction with respect to this partial order to show that $\zeta_z^{\mathbf{1}}(t) = \zeta_z^{\mathbf{0}}(t)$ for all $t \geq p_z$, where $(p_z)_{z\in\hat{B}_x}$ are provided by the definition of t_x. Assume that this is proved for all $z' \prec z \in \hat{B}_x$. Then $\zeta_z^{\mathbf{1}}(p_z) = \zeta_z^{\mathbf{0}}(p_z) = 0$, since $\Upsilon_z(p_z) > q_0$ (recall Eq. (2.4)). Moreover, using Eq. (8.3) and the fact that $U_0 \subset \mathbb{H}_u$, we obtain

$$z + U_0 \subset \bigcup_{y\in\{0,-1\}^d} \hat{B}_{x+y} \cap \left\{z' \in \mathbb{Z}^d : z' \prec z\right\}.$$

Since the r.h.s. above is coupled at all times $t \geq p_z$ by induction hypothesis, the definition of the $\{U_0\}$-CP (see Eqs. (2.1) and (2.4)) gives that $\zeta_z^{\mathbf{1}}(t) = \zeta_z^{\mathbf{0}}(t)$ for all $t \geq p_z$, as desired. □

Combining Proposition 11.1 and Lemma 11.3, we can now prove the following.

Proposition 11.4. *Let $q_0 \in [0,1)$. There exists $c_0 = c_0(U_0, q_0) < \infty$ such that for every $\delta > 0$ and n large enough (depending on δ), the mixing time of the $\{U_0\}$-CP with parameter q_0 (in the box Λ, with $\mathbf{1}$ boundary condition) satisfies*

$$t_{\mathrm{mix}}(\delta) \le c_0 n. \tag{11.7}$$

Note that in Proposition 11.4 we do not need q_0 to be large. However, for small q_0 the CP in infinite volume has trivial upper invariant measure and the mixing time in the box should be logarithmic in n.

Proof. It is not hard to check that $t_x - \tilde{t}_x$ is stochastically dominated by the sum of $|\hat{B}|$ independent exponential random variables of parameter $1 - q_0$. On the other hand, the size of $\hat{B}$ depends only on the set U_0. Since the sum of N exponential random variables is stochastically dominated by a single exponential random variable with suitably large (N-dependent) expectation, we have that $t_x - \tilde{t}_x$ is stochastically dominated by an exponential random variable of parameter depending on q_0 and U_0. Since LPP is monotone in the quantities $t_x - \tilde{t}_x$, which are clearly independent for different x, we get that t_x is smaller than the passage time of a $\hat{U}_0$-LPP slowed down by a factor $\theta(q_0, U_0) \in [1, \infty)$, which we noted in Section 2.6 to coincide with the standard $\{-e_1, \dots, -e_d\}$-LPP slowed down by the same factor.

By Lemma 11.3 and attractiveness of CP we have that $\max_{x \in \mathbb{Z}^d} t_x$ is an upper bound on the coupling time of the $\{U_0\}$-CP on Λ with boundary condition $\mathbf{1}$. The proof is concluded by Proposition 11.1 and Eq. (2.3), as in Eq. (10.1). $\square$

12 Assembling Theorem 3.1

Proof of the lower bound in Eq. (3.1). Let

$$\Lambda_\ell = \{\lfloor n/2 \rfloor - \ell, \dots, \lfloor n/2 \rfloor + \ell\}^d.$$

One can choose ℓ large enough (depending on δ and q, but not on n) so that the probability, under the stationary measure for the process in the box $\Lambda = \{1, \dots, n\}^d$ with boundary condition $\mathbf{1}$, that Λ_ℓ is in state $\mathbf{0}$ is smaller than $\delta/2$. For the $\mathcal{U}-$KCM, this is obvious, since the stationary measure is a Bernoulli product measure with parameter $q > 0$. For the $\mathcal{U}-$CP, this follows from the fact that the stationary measure is stochastically larger than the restriction to Λ of the upper invariant measure $\bar{\nu}$ on $\mathbb{Z}^d$ (recall Section 2.3). The latter is ergodic (see [38, Theorem III.2.3.(f)]) and non-trivial under the assumption that $q > q_c^{\mathrm{CP}}$ (recall Eq. (2.5)). But $q > q_c^{\mathrm{CP}}$, since q is sufficiently close to 1 and $q_c^{CP} < 1$ for $\mathcal{U}$ which is not trivial subcritical (recall Remark 2.1).

On the other hand, starting the dynamics (either KCM or CP) from the $\mathbf{0}$ configuration in Λ with $\mathbf{1}$ boundary condition, with high probability it takes a time at least cn (with c independent of δ) before the state in Λ_ℓ changes. This follows from the definition of the dynamics, the fact that the box Λ_ℓ is at distance at least $n/4$ from the boundary of Λ and from standard estimates on first passage percolation (see e.g. [37, Section I.1]). $\square$

Proof of the upper bound in Eq. (3.1). We need to show that, for every $\delta > 0$ and n large enough (depending on δ), the $\mathcal{U}$-KCM in Λ with parameter q and boundary condition $\mathbf{1}$, started from the initial state $\mathbf{1}$ has coupled at time cn with the one started from an arbitrary configuration $\xi' \in \Omega_\Lambda$, with probability at least $1-\delta$. The proof for the $\mathcal{U}$-CP is identical and therefore omitted.

From Proposition 11.4 we know that (up to a total variation error $\delta/2$) for any $t \geq T_0 := c_0 n$ the $\{U_0\}$-CP with parameter $q_0 \leq q$ large enough, in Λ with boundary condition $\mathbf{1}$ and any initial condition, has coupled. Namely, denoting these processes by $\zeta^{\xi'',\Lambda}$ for initial conditions $\xi'' \in \Omega_\Lambda$, we get

$$\mathbb{P}\left(\forall t \geq T_0, \zeta^{\mathbf{0},\Lambda}(t) = \zeta^{\mathbf{1},\Lambda}(t)\right) \geq 1 - \delta/2.$$

By attractiveness we have that at any time $\zeta^{\mathbf{1},\Lambda}$ dominates the restriction to Λ of the infinite volume $\{U_0\}$-CP $\zeta^{\mathbf{1}}$ with initial condition $\mathbf{1}$ and parameter q_0. Thus,

$$\mathbb{P}\left(\forall t \geq T_0, \forall x \in \Lambda, \zeta^{\mathbf{1}}_x(t) \leq \zeta^{\mathbf{0},\Lambda}_x(t)\right) \geq 1 - \delta/2. \tag{12.1}$$

For all $t \geq T_0$, let $O_t^{\mathbf{1}}$ denote the set of orange healthy sites for $\zeta^{\mathbf{1}}$ and let $O_t^{\mathbf{0},\Lambda}$ denote the one of $\zeta^{\mathbf{0},\Lambda}$, where both are initialised at time T_0 to be equal to the set of all healthy sites and then defined via Eq. (7.2). Using Eq. (7.2), is not hard to check by induction on the number of clock rings in Λ in the time interval $[T_0, t]$, that, if the event in the left hand side of Eq. (12.1) occurs, then for any $t \geq T_0$ we have $O_t^{\mathbf{1}} \cap \Lambda \supset O_t^{\mathbf{0},\Lambda}$. Hence,

$$\mathbb{P}\left(\forall t \geq T_0, O_t^{\mathbf{1}} \cap \Lambda \supset O_t^{\mathbf{0},\Lambda}\right) \geq 1 - \delta/2.$$

Moreover, applying Lemma 7.1 (and Claim 7.2 up to time T_0), a union bound and translation invariance, we have that for $t \geq T_0$

$$\mathbb{P}\left(\eta^{\mathbf{1}}(t) \neq \eta^{\xi'}(t)\right) \leq \mathbb{P}\left(O_t^{\mathbf{0},\Lambda} \neq \varnothing\right) \leq \delta/2 + |\Lambda| \cdot \mathbb{P}\left(0 \in O_t^{\mathbf{1}}\right). \tag{12.2}$$

The proof of Theorem 3.1 is therefore concluded, using monotonicity once again together with Eq. (10.7), taking e.g. $t = T_0 + \sqrt{n}$, since none of the other quantities depend on n. Indeed, in Section 10, Eq. (10.7) was valid for the $\{U_0\}$-CP starting from an initial configuration renormalising to $\xi \sim \mu_{1-\varepsilon_0}$ with ε_0 small enough, while here we simply have $\varepsilon_0 = 0$. □

Acknowledgements

This work was supported by the Austrian Science Fund (FWF): P35428-N. We thank Paul Chleboun, Damiano de Gaspari, Markus Heydenreich, Laure Marêché, Rob Morris, Réka Szabó and Cristina Toninelli for enlightening discussions and to the anonymous referees for valuable comments.

References

[1] C. Alves, G. Amir, R. Baldasso, and A. Teixeira, *Percolation phase transition on planar spin systems*, Ann. Inst. Henri Poincaré Probab. Stat. (To appear).

[2] P. Balister, B. Bollobás, R. Morris, and P. Smith, *Subcritical monotone cellular automata*, Random Structures Algorithms **64** (2024), no. 1, 38–61 pp. MR4672996

[3] P. Balister, B. Bollobás, R. Morris, and P. Smith, *The critical length for growing a droplet*, Mem. Amer. Math. Soc. (To appear).

[4] P. Balister, B. Bollobás, M. Przykucki, and P. Smith, *Subcritical $\mathcal{U}$-bootstrap percolation models have non-trivial phase transitions*, Trans. Amer. Math. Soc. **368** (2016), no. 10, 7385–7411 pp. MR3471095

[5] F. Benevides and M. Przykucki, *Maximum percolation time in two-dimensional bootstrap percolation*, SIAM J. Discrete Math. **29** (2015), no. 1, 224–251 pp. MR3304258

[6] D. Blanquicett, *Anisotropic bootstrap percolation in three dimensions*, Ann. Probab. **48** (2020), no. 5, 2591–2614 pp. MR4152652

[7] O. Blondel, *Front progression in the East model*, Stochastic Process. Appl. **123** (2013), no. 9, 3430–3465 pp. MR3071385

[8] O. Blondel, N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Fredrickson-Andersen one spin facilitated model out of equilibrium*, Markov Process. Related Fields **19** (2013), no. 3, 383–406 pp. MR3156958

[9] O. Blondel, A. Deshayes, and C. Toninelli, *Front evolution of the Fredrickson-Andersen one spin facilitated model*, Electron. J. Probab. **24** (2019), Paper No. 1, 32 pp. MR3903501

[10] T. Bodineau and C. Toninelli, *Activity phase transition for constrained dynamics*, Comm. Math. Phys. **311** (2012), no. 2, 357–396 pp. MR2902193

[11] A. Bušić, J. Mairesse, and I. Marcovici, *Probabilistic cellular automata, invariant measures, and perfect sampling*, Adv. in Appl. Probab. **45** (2013), no. 4, 960–980 pp. MR3161292

[12] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Kinetically constrained spin models*, Probab. Theory Related Fields **140** (2008), no. 3-4, 459–504 pp. MR2365481

[13] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Facilitated spin models: recent and new results*, Methods of contemporary mathematical statistical physics, 2009, 307–340 pp. MR2581609

[14] N. Cancrini, F. Martinelli, R. Schonmann, and C. Toninelli, *Facilitated oriented spin models: some non equilibrium results*, J. Stat. Phys. **138** (2010), no. 6, 1109–1123 pp. MR2601425

[15] P. Chleboun, A. Faggionato, and F. Martinelli, *Time scale separation and dynamic heterogeneity in the low temperature East model*, Comm. Math. Phys. **328** (2014), no. 3, 955–993 pp. MR3201217

[16] P. Chleboun, A. Faggionato, and F. Martinelli, *Mixing time and local exponential ergodicity of the East-like process in* $\mathbb{Z}^d$, Ann. Fac. Sci. Toulouse Math **24** (2015), no. 4, 717–743 pp. MR3434253

[17] P. Chleboun, A. Faggionato, and F. Martinelli, *Relaxation to equilibrium of generalized East processes on* $\mathbb{Z}^d$*: renormalization group analysis and energy-entropy competition*, Ann. Probab. **44** (2016), no. 3, 1817–1863 pp. MR3502595

[18] P. Chleboun and F. Martinelli, *Mixing time bounds for oriented kinetically constrained spin models*, Electron. Commun. Probab. **18** (2013), Paper No. 60, 9 pp. MR3084571

[19] Y. Couzinié and F. Martinelli, *On a front evolution problem for the multidimensional East model*, Electron. J. Probab. **27** (2022). MR4505380

[20] A. Ertul, *Cutoff for the Fredrickson-Andersen one spin facilitated model*, ALEA Lat. Am. J. Probab. Math. Stat. **19** (2022), no. 1, 511–536 pp. MR4394307

[21] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *Aging through hierarchical coalescence in the East model*, Comm. Math. Phys. **309** (2012), no. 2, 459–495 pp. MR2864800

[22] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *The East model: recent results and new progresses*, Markov Process. Related Fields **19** (2013), no. 3, 407–452 pp. MR3156959

[23] G. H. Fredrickson and H. C. Andersen, *Kinetic Ising model of the glass transition*, Phys. Rev. Lett. **53** (1984), no. 13, 1244–1247 pp.

[24] S. Ganguly, E. Lubetzky, and F. Martinelli, *Cutoff for the East process*, Comm. Math. Phys. **335** (2015), no. 3, 1287–1322 pp. MR3320314

[25] P. Garrahan, P. Sollich, and C. Toninelli, *Kinetically constrained models*, Dynamical heterogeneities in glasses, colloids and granular media and jamming transitions, 2011, 341–369 pp.

[26] M. Gottschau, M. Heydenreich, K. Matzke, and C. Toninelli, *Phase transition for a non-attractive infection process in heterogeneous environment*, Markov Process. Related Fields **24** (2018), no. 1, 39–56 pp. MR3793942

[27] L. F. Gray, *Toom's stability theorem in continuous time*, Perplexing problems in probability, 1999, 331–353 pp. MR1703139

[28] S. Greenberg, A. Pascoe, and D. Randall, *Sampling biased lattice configurations using exponential metrics*, Proc. Twent. Annu. ACM-SIAM Symp. Discrete Algorithms, 2009, 76–85 pp. MR2809307

[29] T. E. Harris, *Additive set-valued Markov processes and graphical methods*, Ann Probab. **6** (1978), no. 3, 355–378 pp. MR488377

[30] I. Hartarsky, *$\mathcal{U}$-bootstrap percolation: critical probability, exponential decay and applications*, Ann. Inst. Henri Poincaré Probab. Stat. **57** (2021), no. 3, 1255–1280 pp. MR4291442

[31] I. Hartarsky, *Bootstrap percolation and kinetically constrained models: two-dimensional universality and beyond*, Ph.D. Thesis, Université Paris Dauphine, PSL University, 2022. Available at `https://tel.archives-ouvertes.fr/tel-03527333`.

[32] I. Hartarsky, *Bootstrap percolation, probabilistic cellular automata and sharpness*, J. Stat. Phys. **187** (2022), no. 3, Article No. 21, 17 pp. MR4408459

[33] I. Hartarsky, F. Martinelli, and C. Toninelli, *Coalescing and branching simple symmetric exclusion process*, Ann. Appl. Probab. **32** (2022), no. 4, 2841–2859 pp. MR4474521

[34] I. Hartarsky and R. Szabó, *Subcritical bootstrap percolation via Toom contours*, Electron. Commun. Probab. **27** (2022), Paper No. 55, 13 pp. MR4510850

[35] H. Lacoin, *The scaling limit for zero-temperature planar Ising droplets: with and without magnetic fields*, Topics in percolative and disordered systems, 2014, 85–120 pp. MR3229287

[36] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, Providence, RI, 2009. With a chapter by J. G. Propp and D. B. Wilson. MR2466937

[37] T. M. Liggett, *Stochastic interacting systems: contact, voter and exclusion processes*, Grundlehren der mathematischen Wissenschaften, vol. 324, Springer, Berlin, Heidelberg, 1999. MR1717346

[38] T. M. Liggett, *Interacting particle systems*, Classics in mathematics, Springer, Berlin, Heidelberg, 2005. Originally published by Springer, New York (1985). MR2108619

[39] T. M. Liggett, R. H. Schonmann, and A. M. Stacey, *Domination by product measures*, Ann. Probab. **25** (1997), no. 1, 71–95 pp. MR1428500

[40] L. Marêché, *Exponential convergence to equilibrium for the d-dimensional East model*, Electron. Commun. Probab. **24** (2019), Paper No. 55, 10 pp. MR4003129

[41] L. Marêché, *Exponential convergence to equilibrium in supercritical kinetically constrained models at high temperature*, Markov Process. Related Fields **26** (2020), no. 5, 861–884 pp. MR4246011

[42] F. Martinelli and C. Toninelli, *Towards a universality picture for the relaxation to equilibrium of kinetically constrained models*, Ann. Probab. **47** (2019), no. 1, 324–361 pp. MR3909971

[43] T. Mountford and G. Valle, *Exponential convergence for the Fredrickson-Andersen one-spin facilitated model*, J. Theoret. Probab. **32** (2019), no. 1, 282–302 pp. MR3908915

[44] N. S. Pillai and A. Smith, *Mixing times for a constrained Ising process on the torus at low density*, Ann. Probab. **45** (2017), no. 2, 1003–1070 pp. MR3630292

[45] N. S. Pillai and A. Smith, *Mixing times for a constrained Ising process on the two-dimensional torus at low density*, Ann. Inst. Henri Poincaré Probab. Stat. **55** (2019), no. 3, 1649–1678 pp. MR4010947

[46] L. Poncelet, *Phase transitions in probabilistic cellular automata*, Ph.D. Thesis, Université Catholique de Louvain, 2013.

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

[48] A. Sudbury, *The convergence of the biased annihilating branching process and the double-flipping process in* $\mathbb{Z}^d$, Stochastic Process. Appl. **68** (1997), no. 2, 255–264 pp. MR1454835

[49] J. M. Swart, R. Szabó, and C. Toninelli, *Peierls bounds from Toom contours*, arXiv e-prints (2022), available at arXiv:2202.10999.

[50] A. L. Toom, *Stable and attractive trajectories in multicomponent systems*, Multicomponent random systems, 1980, 549–575 pp. MR599548