

FREDRICKSON–ANDERSEN MODEL IN TWO DIMENSIONS *

IVAILO HARTARSKY¹, FABIO MARTINELLI² AND CRISTINA TONINELLI³

Abstract. The present article overviews recent mathematical advances on the Fredrickson–Andersen kinetically constrained spin model in two dimensions. It was introduced in physics as a toy model for recovering the glassy phenomenology in supercooled liquids close to the glass transition via dynamic constraints as opposed to static interactions.

Résumé. Dans cet article on passe en revue des avancées mathématiques récentes sur le modèle cinétiquement contraint de Fredrickson–Andersen en deux dimensions. Il fut introduit en physique comme modèle jouet pour reproduire la phénoménologie des liquides surfondus près de la transition vitreuse, moyennant une dynamique contrainte plutôt que des interactions statiques.

MSC2020: Primary 60K35; Secondary 82C22, 60J27, 60C05.

Keywords: Kinetically constrained models, interacting particle systems, sharp threshold, bootstrap percolation, Glauber dynamics, Poincaré inequality.

1. MODELS

The Fredrickson–Andersen model (FA) is a family of models named after their inventors [15]. They are paradigmatic examples within the larger class of *kinetically constrained models* (KCM), [44, 50]. Their purpose is to provide an accessible toy model featuring glassy behaviour only via dynamical facilitation (see [47]), while having a trivial stationary state. From the mathematical viewpoint these interacting particle systems (see [33]) are challenging to analyse as compared to close relatives such as the stochastic Ising model (see [34]) notably due to their lack of attractiveness. FA and, more generally, KCM have deep links to synchronous deterministic monotone processes known as bootstrap percolation (see [37]), which we will be led to discuss in detail. Building on bootstrap percolation knowledge and developing new tools for tackling the more intricate FA dynamics, significant rigorous progress has been made, often settling controversial nonrigorous predictions. Our goal is to account for the current state of the art, primarily focusing on the most recent advances [24, 25, 46] in addition to earlier results [6, 10, 36, 38, 40, 41]. Although most techniques carry over to higher dimensions, we focus on the two-dimensional setting for the sake of simplicity and concreteness.

* *This work is supported by ERC Starting Grant 680275 “MALIG.”*

¹ CEREMADE, CNRS, UMR 7534, Université Paris-Dauphine, PSL University, Place du Maréchal de Lattre de Tassigny, 75016 Paris, France

² Dipartimento di Matematica e Fisica, Università Roma Tre, Largo S.L. Murialdo, 00146, Roma, Italy

³ CEREMADE, CNRS, UMR 7534, Université Paris-Dauphine, PSL University, Place du Maréchal de Lattre de Tassigny, 75016 Paris, France

1.1. Bootstrap percolation

Let us start by introducing the j -neighbour bootstrap percolation model. Let $\Omega = \{0, 1\}^{\mathbb{Z}^2}$ and call a site $x \in \mathbb{Z}^2$ *infected* (or *empty*) for $\omega \in \Omega$ if $\omega_x = 0$ and *healthy* (or *filled*) otherwise. For fixed $0 < q < 1$, we denote by μ_q the product Bernoulli probability measure with parameter $1 - q$ under which each site is infected with probability q . When confusion does not arise, we write $\mu = \mu_q$. Given $j \in \{1, 2, 3, 4\}$, the j -neighbour bootstrap percolation model on \mathbb{Z}^2 is the monotone cellular automaton evolving as follows. Let $A_0 \subset \mathbb{Z}^2$ be the set of *initially infected sites* distributed according to μ . Then for any integer *time* $t \geq 0$ we recursively define

$$A_{t+1} = A_t \cup \{x \in \mathbb{Z}^2 : |N_x \cap A_t| \geq j\},$$

where N_x denotes the set of nearest neighbours of x in the usual graph structure of \mathbb{Z}^2 . In other words, a site becomes infected *forever* as soon as its constraint becomes satisfied, namely as soon as it has at least j already infected neighbours.

We denote by $[A] = \bigcup_{t \geq 0} A_t$ the *closure* of $A \subset \mathbb{Z}^2$ and define the *critical probability*

$$q_c = \inf \{q \in [0, 1] : \mu_q([A] = \mathbb{Z}^d) > 0\}. \tag{1}$$

Another key quantity for bootstrap percolation is the *infection time of the origin* defined as $\tau_0^{\text{BP}} = \inf\{t \geq 0 : 0 \in A_t\}$, so that $\tau_0^{\text{BP}} < \infty$ a.s. for $q < q_c$.

1.2. The Fredrickson–Andersen model

We next introduce the FA model, a natural stochastic counterpart of bootstrap percolation and our main focus.

For integers $1 \leq j \leq 4$, the *Fredrickson–Andersen j -spin facilitated model* (FA- j f) is the continuous time Markov process with state space $\Omega = \{0, 1\}^{\mathbb{Z}^2}$ constructed as follows. Each site is endowed with an independent Poisson clock with rate 1. At each clock ring the state of the site is updated to a Bernoulli random variable with parameter $1 - q$ subject to the crucial constraint that if the site has fewer than j infected (nearest) neighbours currently, then the update is rejected. We refer to updates occurring at sites with at least j infected neighbours at the time of the update as *legal*.

Remark 1.1. Contrary to bootstrap percolation, the FA- j f process is clearly non-monotone because of the possible recovery of infected sites with at least j infected neighbours. This feature is one of the major obstacles in the analysis of the process, along with the lack of attractiveness caused by the fact that more infections may make a healing update legal.

It is standard to show (see [33]) that the FA- j f process is well defined and it is reversible w.r.t. μ_q . For any function f depending on finitely many spins, its Dirichlet form reads

$$\mathcal{D}(f) = \sum_{x \in \mathbb{Z}^d} \mu \left(\mathbb{1}_{\{|y \in N_x : \eta_y = 0\}| \geq j} \cdot \text{Var}_x(f(\eta)) \right), \tag{2}$$

where the average μ is over $\eta \in \Omega$, η_y is the state of the spin at y and Var_x stands for the average over η_x , given the restriction of the configuration η to $\mathbb{Z}^2 \setminus \{x\}$. This enables the definition of a key characteristic timescale—the *spectral gap* or inverse *relaxation time*

$$(T_{\text{rel}})^{-1} = \inf_{f \neq \text{const}} \frac{\mathcal{D}(f)}{\text{Var}(f)}, \tag{3}$$

with $\text{Var}(f)$ the variance w.r.t. μ . Alternatively, like for bootstrap percolation, one may prefer to investigate the first infection time of the origin under the FA- j f dynamics

$$\tau_0 = \inf\{t \geq 0 : \eta_0(t) = 0\}.$$

A central goal is to precisely quantify the asymptotics as $q \rightarrow 0$ of T_{rel} and τ_0 for the stationary process with initial law μ_q . In fact, whenever possible, one would also like to treat non-equilibrium initial conditions, but this is mostly open, so unless otherwise stated, we consider the stationary process. Since it is clear that FA-3f and FA-4f have $T_{\text{rel}} = \infty$ and $\tau_0 = \infty$ with positive probability for any $q < 1$, we focus on $j \in \{1, 2\}$.

2. FA-1F

Initial results on FA-1f were obtained in [10], focusing on the relaxation time. There it was proved that T_{rel} is asymptotically q^{-2} up to a logarithmic factor, more precisely there exists a constant C s.t. for any $q \in (0, 1)$ it holds

$$C/q^2 \leq T_{\text{rel}} \leq \log(1/q)/(Cq^2) \quad (4)$$

As we will see, this reflects the fact that, to first order, as $q \rightarrow 0$ infections are typically isolated and perform random walks, jumping at rate of order $1/q$ to neighbouring positions. In order to equilibrate, these random walks need to cover a volume of the order of their inverse density, namely $1/q$, hence the intuition behind the $1/q^2$ scaling.

More rigorously, the lower bound of [46] consists in examining the number of connected clusters of infections truncated at distance $1/q$ from the origin as test function f in Eq. (3). Since infections are rare, they are mostly isolated and $\text{Var}(f)$ scales like $q \cdot q^{-2}$. Moreover, the number of clusters changes by at most 4 after a flip and only changes if the flip occurs at a site with two or more adjacent infections. Thus, contributions to the Dirichlet form of Eq. (2) only come from transitions with three infections at or around a given vertex, yielding $\mathcal{D}(f) \approx q^3 \cdot q^{-2}$. Hence, Eq. (3) gives $T_{\text{rel}} \geq 1/q^2$. A slightly more involved argument allows one to also deal with the expectation $\mathbb{E}_\mu(\tau_0)$ of τ_0 under the stationary process [45].

Rather than explaining the upper bound's original proof from [10], we will take a simpler but less direct route by defining a closely related model of *coalescing and branching simple exclusion process* (CBSEP) and then deducing the result on FA-1f. Essentially, FA-1f is CBSEP's evil twin lacking nice properties, but behaving exactly the same way.

2.1. An auxiliary model: CBSEP

Let $G = (V, E)$ be a connected graph. Minimum, maximum, and average degrees in G are denoted by d_{\min}, d_{\max} and d_{avg} , respectively. The degree of $x \in V$ is denoted by d_x . For any $\omega \in \Omega = \{0, 1\}^V$ and any vertex $x \in V$ we say that x is *filled* (resp. *empty*), or that there is a *particle* (resp. *hole*) at x , if $\omega_x = 1$ (resp. 0). We define $\Omega_+ = \Omega \setminus \{\mathbf{0}\}$ to be the event that there exists at least one particle. Similarly, for any edge $e = \{x, y\} \in E$ we refer to $(\omega_x, \omega_y) \in \{0, 1\}^{\{x, y\}}$ as the state of e in ω and write $E_e = \{\omega \in \Omega : \omega_x + \omega_y \neq 0\}$ for the event that e is not empty (its vertices are not both empty).

Given $p \in (0, 1)$, let $\pi = \bigotimes_{x \in V} \pi_x$ be the product Bernoulli measure in which each vertex is filled with probability p and let $\mu(\cdot) := \pi(\cdot | \Omega_+)$ (if G is infinite, then simply $\mu = \pi$). Given an edge $e = \{x, y\}$, we write $\pi_e := \pi_x \otimes \pi_y$ and $\lambda(p) := \pi(E_e) = p(2 - p)$.

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

$$\mathcal{D}^{\text{CBSEP}}(f) = \frac{1}{2} \sum_{\omega, \omega'} \mu(\omega) c(\omega, \omega') (f(\omega') - f(\omega))^2 = \sum_{e \in E} \mu(\mathbb{1}_{E_e} \text{Var}_e(f | E_e)). \quad (5)$$

Notice that the branching and coalescing moves of CBSEP are exactly the moves allowed in FA-1f, if we identify the particles of CBSEP with the infected sites of FA-1f. Moreover, the SEP move for the edge $\{x, y\}$ from $(1, 0)$ to $(0, 1)$ can be reconstructed using two consecutive FA-1f moves, the first one filling the hole at y and the second one emptying x . If we also take into account the rate for each move, we easily get the following comparison between the respective Dirichlet forms (see Eq. (2) and e.g. [32, Chapter 13.4]): there exists an absolute constant $c > 0$ such that for all $f : \Omega_+ \rightarrow \mathbb{R}$ it holds that

$$c^{-1} \mathcal{D}^{\text{FA-1f}}(f) \leq \mathcal{D}^{\text{CBSEP}}(f) \leq cd_{\max} p^{-1} \mathcal{D}^{\text{FA-1f}}(f), \tag{6}$$

setting the parameter q of FA-1f equal to the parameter p of CBSEP. In our application to FA-1f for $p \rightarrow 0$ only the upper bound, which we believe to be sharper, counts.

Although the two models are clearly closely related, we would like to emphasise that CBSEP has many advantages over FA-1f, making its study simpler. Most notably, CBSEP is *attractive* in the sense that there exists a grand-coupling (see e.g. [32]) which preserves the partial order on Ω given by $\omega \prec \omega'$ iff $\omega_x \leq \omega'_x$ for all $x \in V$. Furthermore, it is also natural to embed in CBSEP a continuous time random walk $(W_t)_{t \geq 0}$ on G such that CBSEP has a particle at W_t for all $t \geq 0$. The latter is a particularly fruitful feature, which is challenging to reproduce for FA-1f [6].

Thanks to Eqs. (3) and (6), in order to upper bound $T_{\text{rel}}^{\text{FA-1f}}$ and recover the result of [10], it suffices to prove the following.

Proposition 2.1 (Hartarsky, Martinelli, Toninelli [24]). *If $G = \mathbb{Z}^2$, then $T_{\text{rel}}^{\text{CBSEP}} \leq O(\log(1/p)/p)$.*

A proof was given in [24] and [25, Appendix B] up to minor modifications. In fact, much more is proved there. Namely, CBSEP on arbitrary graphs is treated, establishing often sharp bounds on T_{rel} , but, more importantly, also on its logarithmic Sobolev constant.¹ A corollary of such stronger results and Eq. (6) is the control of the mixing and L^2 -mixing times of FA-1f. This recovers, strengthens and generalises results of Pillai and Smith [40, 41] proved in a different and somewhat more involved way.

In addition, [24, 25] study a generalised version of CBSEP with general state spaces per site instead of $\{0, 1\}$. For this generalised model they establish appropriate mixing time bounds crucial for the results on FA-2f discussed in Section 4.

3. 2-NEIGHBOUR BOOTSTRAP PERCOLATION

We next turn our attention to 2-neighbour bootstrap percolation in two dimensions, which is a prerequisite for FA-2f. The 2-neighbour bootstrap percolation originates from [11] (see also [30, 42]). Initially it was believed that $q_c > 0$ based on simulations (see [2] and references therein) with estimated values in $(0.035, 0.17)$. However, it was proved soon after [52] that in fact $q_c = 0$. This was the first manifestation of what would grow to be called the *bootstrap percolation paradox* we will keep returning to. To give it in a somewhat simplistic sentence, it refers to the observation that predictions on bootstrap percolation based on simulations always fail, no matter how advanced rigorous results they take into account. An early discussion of this paradox concerning the above can be found in [53], while subsequent reassessments include [12, 20].

3.1. Coarse threshold

The first quantitative statement in the domain of bootstrap percolation, which naturally laid its foundations is due to Aizenman and Lebowitz [4] (for nonrigorous precursors see [31]). They proved that

$$\tau_0^{\text{BP}} = \exp(\Theta(1/q)) \tag{7}$$

w.h.p. We provide a sketch of the argument, as it introduces ingredients essential to us. The first thing to note about 2-neighbour bootstrap percolation is that the closure of any set of infections is the smallest (in terms of

¹This constant is defined like the spectral gap in Eq. (3) with $\text{Var}(f)$ replaced by the entropy $\mu(f^2 \log(f^2/\mu(f^2)))$.

inclusion) collection of rectangles (with sides parallel to the axes of the lattice) at graph distance at least 3 from each other containing the infections. Thus, the closure of any set can be determined via the following *rectangles process*. We start off with a collection of rectangles consisting of each of the initial infections. At each step we merge two of them at graph distance 2 or less, replacing them by the smallest rectangle containing their union. Repeating this until the process becomes stationary yields the collection of rectangles in the closure. A corollary of this process is the following fundamental lemma.

Lemma 3.1 (Aizenman–Lebowitz [4]). *We say that a rectangle R is internally filled (by the set A of initial infections), if $[A \cap R] = R$. If R is internally filled, then for every $k \leq \text{long}(R)$ there exists an internally filled rectangle $S \subset R$ such that $k \leq \text{long}(S) \leq 2k$, where $\text{long}(R)$ denotes the number of sites on the longer side of R .*

Clearly, $\tau_0^{\text{BP}} < \exp(c/q)$ implies that the origin belongs to an internally filled rectangle with long side at most $\exp(c/q)$ with c to be chosen appropriately later. Then Lemma 3.1 shows that within distance $\exp(c/q)$ of the origin there should be an internally filled rectangle R of long side of our choice up to a factor 2. The right side length to choose, which we refer to as *critical scale*, is $1/q$. Observing that such an internally filled rectangle cannot contain two consecutive healthy rows/columns, we get

$$\mu([A \cap R] = R) \leq \left(1 - (1 - q)^{2 \text{long}(R)}\right)^{\lfloor \text{long}(R)/2 \rfloor} = \exp(-\Theta(1/q)),$$

concluding the proof that $\tau_0^{\text{BP}} \geq \exp(\Omega(1/q))$ w.h.p. by the union bound on all possible positions of R , choosing c small enough.

A matching upper bound is guided by a similar idea (explaining the title ‘Metastability effects in bootstrap percolation’ of [4]). We first make sure to internally fill a square of (supercritical) side, say, q^{-3} and then this *critical droplet* is likely to grow and infect the entire grid at roughly linear speed. The internal filling can be directly forced starting from one infection and asking for it to find another one on its right and top side on each line as it progressively infects a growing square. This has probability

$$q \prod_{k=1}^{\infty} \left(1 - (1 - q)^k\right)^2 \approx \exp\left(2 \int_0^{\infty} \log(1 - e^{-qx}) \, dx\right) = \exp(-\Theta(1/q)) \quad (8)$$

and thus is likely to occur within distance $\exp(C/q)$ of the origin for C large enough. We may then ensure that with overwhelming probability every vertical or horizontal line of length q^{-3} at distance at most $\exp(C/q)$ from the origin contains an infection, so that the critical droplet does grow roughly linearly until it engulfs the origin after time $\exp(O(1/q))$, proving Eq. (7).

3.2. Sharp threshold

Naturally, following Eq. (7) the question of the day became determining the implicit constant. This came about in a breakthrough of Holroyd 15 years later [28], proving that w.h.p.

$$\tau_0^{\text{BP}} = \exp\left(\frac{\pi^2 + o(1)}{18q}\right). \quad (9)$$

We will prove stronger lower and upper bounds in the sequel, so it is useful to give an idea of the proof, which introduced several crucial techniques commonly used thereafter. As in the Aizenman–Lebowitz result, the main difficulty is controlling the probability of a rectangle of size roughly $1/q$ being internally filled. More precisely, Eq. (9) follows once we show that for R of size C/q for C large

$$\mu([A \cap R] = R) \approx e^{-\pi^2/(9q)}. \quad (10)$$




FIGURE 1. Recursive structure of droplets used in 2-neighbour bootstrap percolation [28] and FA-2f [25] respectively. Arrows indicate regions with no two consecutive healthy lines.

We will only discuss the easier lower bound in Eq. (10). Once again we start from a single infection and make it infect progressively larger rectangles. However, it grows by an amount larger than 1 in each direction before switching to the other (see Fig. 1). Namely, the right choice is to grow in steps of $1/\sqrt{q}$. The use of this is that we do not need an infection on every line, but on every second line. This is the origin of the constant $\pi^2/9$: it arises like the integral in Eq. (8), but for a function corresponding to the lack of two consecutive rows/columns of healthy sites. If one thinks about the two-term recurrence relation this function should come from (we only need to remember if an infection was found on the previous line or the one before it), it is not surprising that it appears as the root of a certain quadratic equation. The reader interested in the links of this function and its integral with integer partitions may consult [7, 29]. Actually, the sketch above is not quite the way the result is proved in [28], but anticipates [20] and [25] discussed below.

3.3. Speed of convergence

Equation (9) might as well have been the end of the story, had it not been a new manifestation of the bootstrap percolation paradox. Numerical estimates [1, 3, 39] of the constant $\pi^2/18$ above had yielded less than half the correct value. This naturally leads to the question of how fast the convergence in Eq. (9) is. For this reason, we quantify the error term in Eq. (9), again contradicting simulation predictions [48] (see also [21] for more) and showing that the convergence is very slow.

Theorem 3.2 (Gravner, Holroyd, Morris [20], Hartarsky, Morris [26]). *For 2-neighbour bootstrap percolation in two dimensions it holds w.h.p.*

$$\tau_0^{\text{BP}} = \exp\left(\frac{\pi^2 - \Theta(\sqrt{q})}{18q}\right).$$

The upper bound was established in [20] and is based on the mechanism presented for Eq. (9). Roughly speaking, the main difference, which is at the origin of the negative sign of the second term, is taking entropy into account. More precisely, rather than growing our squares in steps of $1/\sqrt{q}$, we allow the exact length of these increments to vary, while being of order $1/\sqrt{q}$. The entropy gained from this is sufficient to outweigh the energetic cost of deviating from a square shape.

The lower bound is significantly harder and is the subject of [26].

4. FA-2F

Moving on to FA-2f (again in two dimensions), the story is much shorter. Indeed, the analogues of all the results for 2-neighbour bootstrap—from the 1988 Aizenman–Lebowitz [4] one (Eq. (7)) to the recent Theorem 3.2—were not known before the recent contribution [25]. Our task is then to review the only two previous rigorous results [10, 36] and copious nonrigorous ones.

4.1. Background

As for bootstrap percolation, the initial expectation was that FA-2f would exhibit a nontrivial transition [15]. We should emphasise that here and in the other works to be quoted below, predictions were made, taking into account bootstrap percolation results already available. In particular, a transition was expected despite its absence in bootstrap percolation [16]. This was quickly dissipated by physicists [17, 43], though rigorous results in this direction came only two decades later [10] (see also [9]). Denoting the semigroup of FA-2f by $(P_t)_{t \geq 0}$ and its generator by \mathcal{L} , the *ergodicity critical parameter* is defined as

$$q_c = \inf \left\{ q > 0 : \forall f \in L^2(\mu), \lim_{t \rightarrow \infty} P_t f = \mu(f) \right\}.$$

It was proved in [10] that this transition coincides with the one of 2-neighbour bootstrap percolation (Eq. (1)), which is why we still denote it q_c . It also coincides with the more standard ergodic theory definition: for $q > q_c$ the eigenvalue 0 of \mathcal{L} is simple and, therefore, by the ergodic theorem we also have

$$q_c = \inf \{ q > 0 : \mathbb{P}_\mu(\tau_0 < \infty) = 1 \}.$$

The same paper also discarded the possibility that for FA-2f e.g. the tail $\mathbb{P}_\mu(\tau_0 \geq t)$ of the infection time would decay as a stretched exponential.² The pure exponential decay they established was quite unexpected as numerous nonrigorous works had exhibited evidence of stretching, though with various stretching exponents [5, 8, 13, 14, 16–19, 22, 43] according to [9, 44]. The exponential decay of the above quantity follows rather easily, once it is established that $T_{\text{rel}} < \infty$, though this had seemingly eluded physicists, who also had various predictions for the scaling of T_{rel} as $q \rightarrow 0$, as we will see.

The last results of [10] for FA-2f are the quantitative bounds on T_{rel}

$$\exp\left(\frac{\pi^2 - o(1)}{18q}\right) = \Omega(\mu(\tau_0^{\text{BP}})) \leq \mathbb{E}_\mu(\tau_0) \leq T_{\text{rel}}/q \leq \exp(O(1/q^5)), \quad (11)$$

in particular establishing that it is finite. The first two inequalities hold in great generality and are not hard. The upper bound is both harder and not useful to us, so we do not discuss it further. Unfortunately, Eq. (11) does not give the precise scaling of $\mathbb{E}_\mu(\tau_0)$. Therefore, discriminating between the conflicting expressions suggested by physicists [8, 14, 17–19, 39, 43, 49, 51] remained an open problem (e.g. [44] asked for settling this controversy). Progress in this direction was made recently in [36], improving the upper bound to

$$\exp\left(\frac{O(\log^2(1/q))}{q}\right), \quad (12)$$

much closer to the lower one, but still inconclusive. Indeed, by 2019, when [36] was published, several (different) predictions not only for the presence or absence of a logarithmic factor but also on the potential sharp constant, based on Eq. (9), had been accumulated in 35 years. The proof of [36] is again not very useful to us, so we do not discuss it.

Before settling the matter, let us explain the different predictions. The first one appeared in [39], where, based on numerical simulations, a faster than exponential divergence in $1/q$ was conjectured. The first to claim an exponential scaling $\exp(\Theta(1)/q)$ was Reiter [43]. He argued that the infection process of the origin is dominated by the motion of ‘macro-defects,’ i.e. rare regions having probability $\exp(-\Theta(1)/q)$ and size $q^{-\Theta(1)}$ that move at an exponentially small rate $\exp(-\Theta(1)/q)$. Later [51] refined the above picture. There it was argued that macro-defects should coincide with the critical droplets of 2-neighbour bootstrap percolation, having probability $\exp(-\pi^2/(9q))$ and that the time scale of the relaxation process inside a macro-defect should be $\exp(c/\sqrt{q})$,

²In [10], the first time when the origin changes state was considered, rather than the time when it becomes infected, but this is unimportant.

i.e. sub-dominant with respect to the inverse of their density, in sharp contrast with the prediction of [43]. Based on this and on the idea that macro-defects move diffusively, the relaxation time scale of FA-2f in $d = 2$ was conjectured to diverge as $\exp(\pi^2/(9q))$ [51, Section 6.3]. Yet, a different prediction was later made in [49] implying a different scaling of the form $\exp(2\pi^2/(9q))$.

4.2. Result

The main result of [25] shows that the scaling prediction of [43, 51] is correct, contrary to those of [39, 49]. Moreover, they show that the characteristic time scale of the relaxation process *inside* a macro-defect agrees with the prediction of [51] and disproves the one of [43].

Theorem 4.1 (Hartarsky, Martinelli, Toninelli [25]). *As $q \rightarrow 0$ the stationary FA-2f model on \mathbb{Z}^2 satisfies:*

$$\mathbb{E}_\mu(\tau_0) \geq \exp\left(\frac{\pi^2}{9q} (1 - \sqrt{q} \cdot O(1))\right), \tag{13}$$

$$\mathbb{E}_\mu(\tau_0) \leq \exp\left(\frac{\pi^2}{9q} \left(1 + \sqrt{q} \cdot (\log(1/q))^{O(1)}\right)\right). \tag{14}$$

Moreover, these also hold for τ_0 w.h.p.

Remark 4.2. Despite the resemblance, Theorem 4.1 is by no means a corollary of Theorem 3.2. While the lower bound Eq. (13) does indeed follow rather easily from Theorem 3.2 together with an improvement of the ‘automatic’ lower bound $\mathbb{E}_\mu(\tau_0) \geq \Omega(\mu(\tau_0^{\text{BP}}))$ from Eq. (11), the proof of Eq. (14) is much more involved. In particular, it requires guessing an efficient infection/healing mechanism to infect the origin, which has no counterpart in the monotone 2-neighbour bootstrap percolation model.

4.3. Behind Theorem 4.1: high-level ideas

The main intuition behind Theorem 4.1 is that for $q \rightarrow 0$ the relaxation to equilibrium of the stationary FA-2f process is dominated by the slow motion of unusually unlikely patches of infection, dubbed *mobile droplets* or just *droplets*. In analogy with the *critical droplets* of bootstrap percolation, mobile droplets have a linear size which is polynomially increasing in q (with some arbitrariness), i.e. they live on a much smaller scale than the metastable length scale $e^{\Theta(1/q)}$ arising in bootstrap percolation. One of the main requirements dictating the choice of the scale of mobile droplets is the requirement that the typical infection environment around a droplet is w.h.p. such that the droplet is able to move under the FA-2f dynamics in any direction. Within this scenario the main contribution to the infection time of the origin for the stationary FA-2f process should come from the time it takes for a droplet to reach the origin.

In order to translate the above intuition into a mathematically rigorous proof, one is faced with two different fundamental problems:

- (1) a precise, yet workable, definition of mobile droplets;
- (2) an efficient model for their ‘effective’ random evolution.

In [23, 35, 36] mobile droplets (dubbed ‘super-good’ regions there) have been defined rather rigidly as fully infected regions of suitable shape and size and their motion has been modelled as a *generalised FA-1f process* on \mathbb{Z} [35, Section 3.1]. In the latter process mobile droplets are freely created or destroyed with the correct heat-bath equilibrium rates but *only at locations which are adjacent to an already existing droplet*.

While rather powerful and robust, this solution has no chance to give the *exact* asymptotics of either (1), or (2) above. Indeed, a mobile droplet should be allowed to deform itself and move to a nearby position like an amoeba, by rearranging its infection using the FA-2f moves. This ‘amoeba motion’ between nearby locations should occur on a time scale much smaller than the global time scale necessary to bring a droplet from far away to the origin. In particular, it should not require to first create a new droplet from the initial one and only later destroy the original one (the main mechanism of the droplet dynamics under the generalised FA-1f process).

The work [25] offered a solution to (1) and (2) above which indeed leads to determining the exact asymptotics of the infection time. Concerning (1), their treatment consists of two steps. They first propose a sophisticated multiscale definition of mobile droplets which, in particular, introduces a crucial degree of *softness* in their microscopic infection's configuration. Namely, on each scale the lower scale droplet is allowed to have an arbitrary position within the higher scale one (see Fig. 1), instead of systematically being at the bottom-left corner. This allows for the droplet to move by rearranging its inside, which corresponds to changing the position of its core on each scale.

The second and much more technically involved step is developing the tools necessary to analyse the FA-2f dynamics inside a mobile droplet. In particular, [25] then proves two key features:

(1.a) to the leading order the probability ρ_D of mobile droplets matches Eq. (10):

$$\rho_D \geq \exp\left(-\frac{\pi^2}{9q} - \frac{O(\log^2(1/q))}{\sqrt{q}}\right),$$

(1.b) the ‘amoeba motion’ of mobile droplets to a nearby location occurs on time scale $\exp(O(\log(1/q)^3)/\sqrt{q})$ which is sub-leading w.r.t. the main time scale of the problem and only manifests in the second term of Eq. (14).

Property (1.a) follows rather easily essentially as explained for Eq. (10), while proving property (1.b) required several ideas. In particular, let us mention how the use of the standard technique of canonical paths is avoided. One proves a Poincaré inequality on each scale for the product measure μ on the volume of the droplet, conditioned to have the droplet structure present. Deducing the higher-scale inequality from the lower scale one, on the other hand, relies on a suitable adaptation the bisection technique from [10].

While properties (1.a) and (1.b) above are essential, they are not sufficient on their own for solving problem (2) above. That is where CBSEP's key role is to be played (recall Section 2.1). In [25] the random evolution of mobile droplets is treated as a generalised CBSEP at the level of Poincaré inequalities. More precisely, they considered a renormalisation, so that the state of a box is 1 if it contains a droplet and 0 otherwise. Thus, the parameter p of CBSEP corresponds to the density ρ_D of droplets. We next recall the relaxation time of CBSEP from Proposition 2.1, as well as the definitions Eqs. (3) and (5). Finally, we observe that the terms $\text{Var}_e(f|E_e)$ in Eq. (5) call out for a conditional Poincaré inequality for a droplet, as discussed in the context of (1.b). Putting these ingredients together, one does obtain the scaling of Theorem 4.1.

5. OPEN PROBLEMS

The results we accounted for in the previous sections concern the behaviour of the stationary process started from the equilibrium measure μ . A key issue, both from the mathematical and physical points of view, is to analyse the out of equilibrium dynamics, namely the evolution starting from an initial measure $\nu \neq \mu$. For ν a product Bernoulli measure of parameter $1 - q'$ with $q' \in [0, 1)$, we expect convergence to equilibrium to occur for FA-1f and FA-2f. Namely we expect

$$\lim_{t \rightarrow \infty} \nu(P_t f(\eta)) = \mu(f)$$

to hold for any local function f , where P_t denotes the semigroup of the process. Progress in this direction has been hindered by the lack of robust tools. Indeed, due to non-attractiveness, many powerful tools that have been developed to analyse relaxation to equilibrium for other interacting particle systems, such as coupling and censoring arguments, are unavailable here. Furthermore, due to the presence of constraints, a worst case analysis is not possible (if the initial configuration is completely healthy, no site can ever be infected and equilibrium is never attained). Similarly, the Sobolev and logarithmic Sobolev constants diverge in infinite volume, discarding classic arguments for proving relaxation to equilibrium based on hypercontractivity (see [27]).

The only results available in this out of equilibrium setting concern FA-1f and establish convergence to equilibrium only in the restricted parameter region $q > 1/2$ [6]. As for the speed of convergence, the result of [6] proves convergence at least as fast as a stretched exponential with exponent $1/2$. This was later improved

to exponential convergence in [38], albeit in an even more restricted parameter region (for q sufficiently close to 1). The reason why the above results hold only in a restricted region is that, in order to circumvent the problems stemming from non attractiveness, both [6, 38] compare FA-1f to an auxiliary dynamics. These dynamics have the advantage of being attractive, but the disadvantage of producing less infected sites and being efficient only if q is sufficiently high. For FA-2f even the regime q close to 1 is completely open due to the cooperative dynamics (no finite set of infections is able to move around on its own).

Let us conclude by mentioning two open questions for the stationary process started from μ .

- (i) For FA-1f a logarithmic gap remains to be filled between the upper and lower bounds to obtain the scaling of T_{rel} (see Eq. (4)).
- (ii) For FA-2f the second order term in the exponent remains to be determined to reach the precision of Theorem 3.2 for bootstrap percolation or even determine its sign.

REFERENCES

- [1] J. Adler, *Bootstrap percolation*, Phys. A **171** (1991), no. 3, 453–470 pp.
- [2] J. Adler and A. Aharony, *Diffusion percolation. I. Infinite time limit and bootstrap percolation*, J. Phys. A **21** (1988), no. 6, 1387–1404 pp. MR939745
- [3] J. Adler, D. Stauffer, and A. Aharony, *Comparison of bootstrap percolation models*, J. Phys. A **22** (1989), no. 7, L297–L301 pp.
- [4] M. Aizenman and J. L. Lebowitz, *Metastability effects in bootstrap percolation*, J. Phys. A **21** (1988), no. 19, 3801–3813 pp. MR968311
- [5] G. B. Alers, M. B. Weissman, A. Kinzig, and N. Israeloff, *Noise simulations of an Ising model of glassy kinetics*, Phys. Rev. B **36** (1987), no. 16, 8429–8434 pp.
- [6] 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
- [7] K. Bringmann and K. Mahlborg, *Improved bounds on metastability thresholds and probabilities for generalized bootstrap percolation*, Trans. Amer. Math. Soc. **364** (2012), no. 7, 3829–3859 pp. MR2901236
- [8] S. Butler and P. Harrowell, *The origin of glassy dynamics in the 2D facilitated kinetic Ising model*, J. Chem. Phys. **95** (1991), no. 6, 4454–4465 pp.
- [9] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Relaxation times of kinetically constrained spin models with glassy dynamics*, J. Stat. Mech. Theory Exp. **03** (2007), L03001, 8 pp.
- [10] 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
- [11] J. Chalupa, P. L. Leath, and G. R. Reich, *Bootstrap percolation on a Bethe lattice*, J. Phys. C **12** (1979), no. 1, L31–L35 pp.
- [12] P. de Gregorio, A. Lawlor, P. Bradley, and K. A. Dawson, *Clarification of the bootstrap percolation paradox*, Phys. Rev. Lett. **93** (2004), no. 2, 025501, 4 pp.
- [13] G. H. Fredrickson, *Linear and nonlinear experiments for a spin model with cooperative dynamics*, Ann. N. Y. Acad. Sci. **484** (1986), no. 1, 185–205 pp.
- [14] G. H. Fredrickson, *Recent developments in dynamical theories of the liquid-glass transition*, Annu. Rev. Phys. Chem. **39** (1988), no. 1, 149–180 pp.
- [15] G. H. Fredrickson and H. C. Andersen, *Kinetic Ising model of the glass transition*, Phys. Rev. Lett. **53** (1984), no. 13, 1244–1247 pp.
- [16] G. H. Fredrickson and H. C. Andersen, *Facilitated kinetic Ising models and the glass transition*, J. Chem. Phys. **83** (1985), no. 11, 5822–5831 pp.
- [17] G. H. Fredrickson and S. A. Brawer, *Monte Carlo investigation of a kinetic Ising model of the glass transition*, J. Chem. Phys. **84** (1986), no. 6, 3351–3366 pp.
- [18] I. S. Graham, L. Piché, and M. Grant, *Dynamics and effective thermodynamics of a model structural glass*, J. Phys.: Condens. Matter **5** (1993), no. 29, L349–L354 pp.
- [19] I. S. Graham, L. Piché, and M. Grant, *Model for dynamics of structural glasses*, Phys. Rev. E **55** (1997), no. 3, 2132–2144 pp.
- [20] J. Gravner and A. E. Holroyd, *Slow convergence in bootstrap percolation*, Ann. Appl. Probab. **18** (2008), no. 3, 909–928 pp. MR2418233
- [21] J. Gravner, A. E. Holroyd, and R. Morris, *A sharper threshold for bootstrap percolation in two dimensions*, Probab. Theory Related Fields **153** (2012), no. 1-2, 1–23 pp. MR2925568
- [22] P. Harrowell, *Visualizing the collective motions responsible for the α and β relaxations in a model glass*, Phys. Rev. E **48** (1993), no. 6, 4359–4363 pp.
- [23] I. Hartarsky, F. Martinelli, and C. Toninelli, *Universality for critical KCM: finite number of stable directions*, Ann. Probab. **49** (2021), no. 5, 2141–2174 pp. MR4317702

- [24] I. Hartarsky, F. Martinelli, and C. Toninelli, *Coalescing and branching simple symmetric exclusion process*, Ann. Appl. Probab. **32** (2022), no. 4, 2841–2859 pp.
- [25] I. Hartarsky, F. Martinelli, and C. Toninelli, *Sharp threshold for the FA-2f kinetically constrained model*, Probab. Theory Related Fields (To appear).
- [26] I. Hartarsky and R. Morris, *The second term for two-neighbour bootstrap percolation in two dimensions*, Trans. Amer. Math. Soc. **372** (2019), no. 9, 6465–6505 pp. MR4024528
- [27] R. Holley and D. Stroock, *Logarithmic Sobolev inequalities and stochastic Ising models*, J. Stat. Phys. **46** (1987), no. 5-6, 1159–1194 pp. MR893137
- [28] A. E. Holroyd, *Sharp metastability threshold for two-dimensional bootstrap percolation*, Probab. Theory Related Fields **125** (2003), no. 2, 195–224 pp. MR1961342
- [29] A. E. Holroyd, T. M. Liggett, and D. Romik, *Integrals, partitions, and cellular automata*, Trans. Amer. Math. Soc. **356** (2003), no. 8, 3349–3368 pp. MR2052953
- [30] P. M. Kogut and P. L. Leath, *Bootstrap percolation transitions on real lattices*, J. Phys. C **14** (1981), no. 22, 3187–3194 pp.
- [31] R. Lenormand and C. Zarcone, *Growth of clusters during imbibition in a network of capillaries*, Kinetics of aggregation and gelation, 1984, 177–180 pp.
- [32] 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
- [33] T. M. Liggett, *Interacting particle systems*, Classics in mathematics, Springer, Berlin, Heidelberg, 2005. Originally published by Springer, New York (1985). MR2108619
- [34] F. Martinelli, *Lectures on Glauber dynamics for discrete spin models*, Lectures on probability theory and statistics (Saint-Flour, 1997), 1999, 93–191 pp. MR1746301
- [35] F. Martinelli, R. Morris, and C. Toninelli, *Universality results for kinetically constrained spin models in two dimensions*, Comm. Math. Phys. **369** (2019), no. 2, 761–809 pp. MR3962008
- [36] 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
- [37] R. Morris, *Monotone cellular automata*, Surveys in combinatorics 2017, 2017, 312–371 pp. MR3728111
- [38] 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
- [39] H. Nakanishi and H. Takano, *Numerical study on the kinetic Ising model for glass transition*, Phys. Lett. A **115** (1986), no. 3, 117–121 pp.
- [40] 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
- [41] 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
- [42] M. Pollak and I. Riess, *Application of percolation theory to 2D-3D Heisenberg ferromagnets*, Phys. stat. sol. (b) **69** (1975), no. 1, K15–K18 pp.
- [43] J. Reiter, *Statics and dynamics of the two-spin-facilitated kinetic Ising model*, J. Chem. Phys. **95** (1991), no. 1, 544–554 pp.
- [44] F. Ritort and P. Sollich, *Glassy dynamics of kinetically constrained models*, Adv. Phys. **52** (2003), no. 4, 219–342 pp.
- [45] A. Shapira, *Bootstrap percolation and kinetically constrained models in homogeneous and random environments*, Ph.D. Thesis, Université Paris Diderot, Sorbonne Paris Cité, 2019.
- [46] A. Shapira, *A note on the spectral gap of the Fredrickson-Andersen one spin facilitated model*, J. Stat. Phys. **181** (2020), no. 6, 2346–2352 pp. MR4179809
- [47] T. Speck, *Dynamic facilitation theory: a statistical mechanics approach to dynamic arrest*, J. Stat. Mech. Theory Exp. **8** (2019), 084015, 14 pp. MR4031660
- [48] E. Teomy and Y. Shokef, *Finite-density effects in the Fredrickson-Andersen and Kob-Andersen kinetically-constrained models*, J. Chem. Phys. **141** (2014), no. 6, 064110, 7 pp.
- [49] E. Teomy and Y. Shokef, *Relation between structure of blocked clusters and relaxation dynamics in kinetically constrained models*, Phys. Rev. E **92** (2015), no. 3, 032133, 10 pp.
- [50] C. Toninelli, *Interacting particle systems with kinetic constraints*, In preparation.
- [51] C. Toninelli, G. Biroli, and D. S. Fisher, *Cooperative behavior of kinetically constrained lattice gas models of glassy dynamics*, J. Stat. Phys. **120** (2005), no. 1-2, 167–238 pp. MR2165529
- [52] A. C. D. van Enter, *Proof of Straley’s argument for bootstrap percolation*, J. Stat. Phys. **48** (1987), no. 3-4, 943–945 pp. MR914911
- [53] A. C. D. van Enter, J. Adler, and J. A. M. S. Duarte, *Finite-size effects for some bootstrap percolation models*, J. Stat. Phys. **60** (1990), no. 3-4, 323–332 pp. MR1069635