

Monotone cellular automata

Robert Morris

Abstract

Cellular automata are interacting particle systems whose update rules are local and homogeneous. Since their introduction by von Neumann almost 50 years ago, many particular such systems have been investigated, but no general theory has been developed for their study, and for many simple examples surprisingly little is known. Understanding the rules that govern their typical global behaviour is an important and challenging problem in statistical physics, probability theory and combinatorics.

In this survey we will consider the behaviour of a particular (large) family of *monotone* cellular automata – those which can naturally be embedded in d -dimensional space – with random initial conditions. For example, in the case where a site updates (from inactive to active) if at least r of its neighbours are already active, these models are known as *bootstrap percolation*, and have been extensively studied for various specific underlying graphs.

Our aims are threefold: to provide a relatively gentle introduction to some of the key techniques in the area; to describe some dramatic recent progress relating to an extremely general class of models, which were introduced by Bollobás, Smith and Uzzell; and to discuss applications of these techniques to models in statistical physics, such as the Glauber dynamics of the Ising model, the abelian sandpile, and kinetically constrained spin models.

1 Introduction

Consider the following deterministic process on a graph G : at each time step, a healthy vertex becomes *infected* if at least r of its neighbours are already infected, while infected vertices stay infected forever. Thus, writing A_t for the set of infected vertices at time t , we have

$$A_{t+1} = A_t \cup \{v \in V(G) : |N(v) \cap A_t| \geq r\}$$

for each $t \geq 0$. This process, which is called *r -neighbour bootstrap percolation*, was introduced in 1979 by Chalupa, Leath and Reich [23], who were motivated by the Glauber dynamics of the Ising model on \mathbb{Z}^d (see Section 5). In this context, the following basic question is of fundamental importance: What is the typical behaviour of the bootstrap process when the initial set of infected vertices is chosen randomly?

In order to state this question more precisely, let us write $A = A_0$ for the set of vertices infected at time $t = 0$, and define \mathbb{P}_p to be the probability distribution obtained by placing each vertex of G in A independently at random with probability p . (We say that a subset $A \subseteq V(G)$ is *p-random* if it is chosen according to \mathbb{P}_p .) The *closure* of A is denoted

$$[A] := \bigcup_{t \geq 0} A_t$$

and we say that the set A *percolates* if $[A] = V(G)$. The *critical probability* for r -neighbour bootstrap percolation on G is then defined to be

$$p_c(G, r) := \inf \left\{ p : \mathbb{P}_p([A] = V(G)) \geq 1/2 \right\}.$$

We remark that the event $\{[A] = V(G)\}$ typically has a sharp threshold¹, and so the choice of the constant $1/2$ is not important.

The r -neighbour bootstrap process has been studied on a wide range of graphs, including trees [23, 10, 15], the Erdős-Rényi random graph [42] and the random regular graph [11]. In this survey, however, we will restrict our attention to a particular class of processes: roughly speaking, those that can naturally be embedded in d -dimensional Euclidean space. As we will see, this leads to an extremely rich and complex class of models, but we will nevertheless be able to make significant progress towards understanding the behaviour of such processes.

The structure of this survey is as follows. As a warm-up, we will begin (in Sections 2–4) by discussing in some detail an important specific class of models: the two-neighbour process on the two-dimensional torus \mathbb{Z}_n^2 , and its (significantly more challenging) generalization, the r -neighbour process on \mathbb{Z}_n^d . We will also (in Section 5) give a couple of applications of these results in a more complex setting: the zero-temperature Glauber dynamics of the Ising model on \mathbb{Z}^d .

After this extended introduction, we will be ready to expand our horizons and consider more general classes of processes. We will begin (in Section 6) by discussing a few instructive examples, in order to build up our intuition; we will then proceed (in Sections 7 and 8) to the main topic of this survey: a series of recent breakthroughs that provide a very precise characterization of monotone cellular automata in two dimensions. Finally, in Sections 9 and 10, we will discuss possible future directions, including a conjectured characterization in higher dimensions, and some related models including nucleation and growth, kinetically constrained spin models, and the abelian sandpile.

¹In the sense that if $p < (1 + o(1))p_c(G, r)$ then $\mathbb{P}_p([A] = V(G)) = o(1)$, whereas if $p > (1 + o(1))p_c(G, r)$ then $\mathbb{P}_p([A] = V(G)) = 1 - o(1)$.

2 The two-neighbour model on \mathbb{Z}_n^2

How many *randomly chosen* infected vertices of the graph \mathbb{Z}_n^2 are needed to (typically) result in the eventual infection of the entire torus? On the one hand, it is easy to see that there exists a set A of size $n - 1$ such that $[A] = \mathbb{Z}_n^2$ (for example, one could take all but one element of a diagonal line); on the other, if the elements of A are typically far apart, it is hard to see how they could help one another.

The first step in developing our intuition is the following theorem, first proved by van Enter [29] in 1987.

Theorem 2.1 $p_c(\mathbb{Z}_n^2, 2) \rightarrow 0$ as $n \rightarrow \infty$.

Proof The proof has two steps: we first find a large infected square, and then (using sprinkling to maintain independence) show that this square is likely to grow and infect the entire torus. To be precise, for each constant $\varepsilon > 0$ we will show that, if $p = 2\varepsilon$, then $\mathbb{P}_p([A] = \mathbb{Z}_n^2) \rightarrow 1$ as $n \rightarrow \infty$. Indeed, let A_1 and A_2 be (independent) ε -random subsets of \mathbb{Z}_n^2 , so that (via a trivial coupling) we have $A \supseteq A_1 \cup A_2$, and set $k = \log \log n$. We will prove the following two statements:

1. With high probability A_1 contains a square R of side-length k .
2. With high probability $[A_2 \cup R] = \mathbb{Z}_n^2$.

Indeed, the first claim follows since each $k \times k$ square lies in A_1 with probability $\varepsilon^{k^2} = n^{-o(1)}$, for disjoint squares the corresponding events are independent, and there exists a collection of $n^{2-o(1)}$ disjoint $k \times k$ squares in \mathbb{Z}_n^2 . To see the second claim, let us suppose (without loss of generality) that $R = [k]^2$, and observe that if $[m]^2 \subseteq [A_2 \cup R]$ and A_2 contains at least one element of each of the sets $\{(m+1, y) : y \in [m]\}$ and $\{(x, m+1) : x \in [m]\}$ then $[m+1]^2 \subseteq [A_2 \cup R]$. It follows that

$$\mathbb{P}([A_2 \cup R] = \mathbb{Z}_n^2) \geq 1 - 2 \sum_{m=k}^{n-1} (1 - \varepsilon)^m \rightarrow 1$$

as $n \rightarrow \infty$, as claimed. \square

The following year, Aizenman and Lebowitz [1] proved the following fundamental theorem, which determines the *threshold* for the event that A percolates.²

²In fact, Aizenman and Lebowitz determined the threshold for the two-neighbour model on \mathbb{Z}_n^d for all $d \geq 2$, but for simplicity we shall work in two dimensions. We remark that a second proof of this theorem was obtained independently several years later by Balogh and Pete [11].

Theorem 2.2

$$p_c(\mathbb{Z}_n^2, 2) = \Theta\left(\frac{1}{\log n}\right).$$

The proof of the upper bound in Theorem 2.2 is essentially just a refinement of the proof of Theorem 2.1, the key observation being that in the first step we can replace the event $R \subseteq A_1$ by the (much more likely) event that R is *internally filled* by A_1 , that is, if $R \subseteq [A_1 \cap R]$. Indeed, if $R = [k]^2$ and A_1 is p -random, then the probability of this event can be bounded by

$$\mathbb{P}(R \subseteq [A_1 \cap R]) \geq \prod_{m=1}^{k-1} (1 - (1-p)^m)^2 \geq e^{-C/p} \quad (2.1)$$

for some large constant C , since if each of the sets $\{(m+1, y) : y \in [m]\}$ and $\{(x, m+1) : x \in [m]\}$ contains an element of A_1 then $R \subseteq [A_1 \cap R]$, and these sets are all disjoint. (The second inequality can be proved, for example, by noting that if $m = o(1/p)$ then $1 - (1-p)^m \sim pm$, if $m = \Theta(1/p)$ then $1 - (1-p)^m = \Theta(1)$, and if $m \gg 1/p$ then $(1-p)^m \leq e^{-pm}$.)

Now, suppose that $p \geq C/\log n$, and let $k = (\log n)^3$ and A_1 be a p -random set. By applying (2.1) to (roughly) $n^2/k^2 \gg n$ disjoint squares, it follows that with high probability there exists a $k \times k$ square R that is internally filled by A_1 . As in the proof of Theorem 2.1, if A_2 is also a p -random set then

$$\mathbb{P}([A_2 \cup R] = \mathbb{Z}_n^2) \geq 1 - 2 \sum_{m=k}^{n-1} (1-p)^m \geq 1 - 2ne^{-pk} \rightarrow 1$$

as $n \rightarrow \infty$, so $p_c \leq 2C/\log n$, as required.

For the lower bound, the key tool is the following deterministic lemma, which allows us to identify the ‘‘bottleneck’’ that prevents A from percolating. For technical reasons (see the proofs of the lemmas below), it will be slightly simpler to work on the grid $[n]^2$; however, the proof can be easily extended to the torus \mathbb{Z}_n^2 without any additional ideas. Let $\phi(R)$ denote the semi-perimeter of a rectangle $R \subseteq [n]^2$.

Lemma 2.3 (The Aizenman–Lebowitz lemma) *Let R be a rectangle that is internally filled by a set $A \subseteq [n]^2$. For every $1 \leq k \leq \phi(R)$, there exists an internally filled rectangle $S \subseteq R$ with $k \leq \phi(S) < 2k$.*

We remark that variants of this lemma will play a crucial role in almost all of the proofs discussed below. The lemma is a straightforward consequence of the following process, which is simply a convenient reordering of the order in which we infect the sites of $[A] \setminus A$.

Definition 2.4 (The rectangles process) Let $A = \{x_1, \dots, x_m\} \subseteq [n]^2$, and consider the collection $\{R_1, \dots, R_m\}$, where $R_j = \{x_j\}$ for each $j \in [m]$. Now repeat the following steps until STOP:

1. If there exist two rectangles R_i and R_j in the current collection at distance at most two from one another, then choose such a pair, remove them from the collection, and replace them by $[R_i \cup R_j]$.
2. If there do not exist such a pair of rectangles, then STOP.

Before going any further, let's record a couple of simple but important observations for future reference. We leave the (easy) proofs to the reader.

Lemma 2.5 *Let R and S be rectangles in $[n]^2$ at distance at most two from one another. Then $[R \cup S]$ is a rectangle, and moreover*

$$\phi([R \cup S]) \leq \phi(R) + \phi(S).$$

Moreover, if R and S are internally filled, then $[R \cup S]$ is internally filled.

Note that, by Lemma 2.5, every rectangle that appears in the rectangles process applied to A is internally filled by A .

For each $A \subseteq [n]^2$, let us write $\langle A \rangle = \{R_1, \dots, R_k\}$ for the collection of rectangles obtained by applying the rectangles process to A ; we call $\langle A \rangle$ the *span* of A . The next lemma is also just an observation.

Lemma 2.6 *For each $A \subseteq [n]^2$, we have*

$$[A] = \bigcup_{S \in \langle A \rangle} S.$$

In particular, if R is internally filled by A , then $\langle A \cap R \rangle = \{R\}$.

Proof Each rectangle that appears in the rectangles process applied to A is internally filled by A , and so each $S \in \langle A \rangle$ is clearly contained in $[A]$. On the other hand, suppose (for a contradiction) that there exists a vertex of $[A]$ that is not contained in some rectangle in $\langle A \rangle$, and consider the first such vertex to be infected. This vertex (v , say) has at least two previously infected neighbours, which (by minimality) must be contained in (one or two) rectangles in $\langle A \rangle$. But these cannot be different rectangles, since the distance between them is at most two, and if they are the same rectangle then they also contain v , which is the desired contradiction.

Finally, note that, since the rectangles in $\langle A \rangle$ are at pairwise distance at least three, if they cover R then they must consist of a single rectangle. \square

We can now easily deduce Lemma 2.3.

Proof of Lemma 2.3 If R is internally filled by $A \subseteq [n]^2$, then we have $R \in \langle A \cap R \rangle$, by Lemma 2.6. Let S be the first rectangle that appears in the rectangles process with $\phi(S) \geq k$. Note that S is internally filled by A , and that $S = [R_i \cup R_j]$ for some rectangles with $\phi(R_i), \phi(R_j) < k$. Hence, by Lemma 2.5, we have $\phi(S) < 2k$, as required. \square

The final tool we need to prove the lower bound in Theorem 2.2 is the following lemma, which also follows easily from the rectangles process.

Lemma 2.7 *Let $R \subseteq [n]^2$ be a rectangle. If R is internally filled by A , then*

$$|A \cap R| \geq \frac{\phi(R)}{2}.$$

Proof We use induction on $\phi(R)$, noting that if $|R| = 1$ then $\phi(R) = 2$, so the base case holds trivially. To prove the induction step, we will show that there exist two *disjointly* internally filled rectangles S_1 and S_2 such that

$$\phi(S_1), \phi(S_2) < \phi(R) \quad \text{and} \quad [S_1 \cup S_2] = R.$$

That is, there exist disjoint sets $A_1, A_2 \subseteq A$ with $S_1 = [A_1]$ and $S_2 = [A_2]$.

To prove this, observe first that at every step of the rectangles process all of the rectangles in the current collection are disjointly internally filled. (Proof: this is true at the beginning, and is maintained by uniting two rectangles.) Thus, if we run the rectangles process for the set $A \cap R$, stopping just before the rectangle R appears for the first time, we obtain two rectangles S_1 and S_2 as required.

Now, by Lemma 2.5 and the induction hypothesis, it follows that

$$|A \cap R| \geq |A_1| + |A_2| \geq \frac{\phi(S_1)}{2} + \frac{\phi(S_2)}{2} \geq \frac{\phi(R)}{2},$$

as required. \square

This is in fact not the simplest proof of Lemma 2.7, but it has the significant advantage of being fairly “robust”; in particular, it motivates the strategy we will use later in more complex contexts.

Exercise: Find an alternative (one line!) proof of Lemma 2.7.

We are now ready to prove Theorem 2.2.

Proof of Theorem 2.2 The proof of the upper bound was outlined above, so we will prove only the lower bound. Suppose that $A \subseteq [n]^2$ is such that $[A] = [n]^2$. By Lemma 2.3, there exists an internally filled rectangle R with

$$\log n \leq \phi(R) \leq 2 \log n.$$

Let X denote the random variable that counts the number of such rectangles when A is a p -random subset of $[n]^2$, and observe that

$$\mathbb{E}[X] \leq \sum_{k=\log n}^{2 \log n} n^3 \binom{k^2}{k/2} p^{k/2} \leq \sum_{k=\log n}^{2 \log n} n^3 (2ekp)^{k/2} \leq \frac{1}{n} \quad (2.2)$$

if $1/p > e^{12} \log n$, by Lemma 2.7, since there are at most n^3 rectangles in $[n]^2$ with semi-perimeter k , and each has area at most k^2 . It follows from Markov's inequality that a p -random set $A \subseteq [n]^2$ percolates with probability at most $1/n$, and hence

$$p_c([n]^2, 2) \geq \frac{1}{e^{12} \log n},$$

as required. \square

The proof of Aizenman and Lebowitz can be extended (in a relatively straightforward manner) to determine the threshold for the two-neighbour model in d dimensions (i.e., on the grid $[n]^d$ or the torus \mathbb{Z}_n^d). We leave this generalization as an exercise for the reader (see also Section 3).

Exercise: Prove that

$$p_c(\mathbb{Z}_n^d, 2) = \Theta\left(\frac{1}{\log n}\right)^{d-1}$$

for every $d \geq 2$.

2.1 A sharp threshold for the two-neighbour process on \mathbb{Z}_n^2

The next major breakthrough in the study of the two-neighbour process was made by Holroyd [39] in 2003, who determined the following sharp threshold in two dimensions.

Theorem 2.8

$$p_c(\mathbb{Z}_n^2, 2) = \left(\frac{\pi^2}{18} + o(1)\right) \frac{1}{\log n}.$$

The main technical innovation in Holroyd's proof, which has had an extremely significant impact on the subsequent development of the area, is the concept of a *hierarchy*. Roughly speaking, this is obtained by recording a (cleverly-chosen) subset of the rectangles that appear in the rectangles process. This subset has the following properties:

- (a) There are at most $n^{o(1)}$ possible hierarchies of a rectangle R of semi-perimeter $O(\log n)$ (this is useful as we will take a union bound over hierarchies when bounding the probability that R is internally filled).
- (b) Each hierarchy is unlikely to occur (at least as unlikely as that corresponding to the simplest possible kind of growth, in which a single droplet finds at each step a single infected site on its boundary).

To be slightly more precise, whenever two 'large' rectangles combine we will record the event, and if a droplet grows for some time without meeting another 'large' rectangle, then we will periodically record its latest size and position. Much more precisely, we have the following technical definition.

Definition 2.9 Let R be a rectangle. A *hierarchy* \mathcal{H} for R consists of a directed rooted tree $G_{\mathcal{H}}$, with all of its edges directed downwards (i.e., away from the root v_0), and for each vertex $v \in V(G_{\mathcal{H}})$ an associated rectangle D_v , such that the following conditions are satisfied:

- (i) the root vertex is associated to R , i.e., $D_{v_0} = R$;
- (ii) each vertex has out-degree at most 2;
- (iii) if $v \in N(u)$ then $D_v \subseteq D_u$;
- (iv) if $N(u) = \{v, w\}$ then $D_u = [D_v \cup D_w]$,

where $N(u)$ denotes the out-neighbourhood of u in $G_{\mathcal{H}}$.

Given $s, t > 0$, we say that \mathcal{H} is (s, t) -good if it satisfies the following conditions for each $u \in V(G_{\mathcal{H}})$:

- (v) u is a leaf (i.e, $|N(u)| = 0$) if and only if $\phi(D_u) \leq s$;
- (vi) if $N(u) = \{v\}$ and $|N(v)| = 1$ then

$$t \leq \phi(D_u) - \phi(D_v) \leq 2t;$$

- (vii) if $N(u) = \{v\}$ and $|N(v)| \neq 1$ then $\phi(D_u) - \phi(D_v) \leq 2t$;
- (viii) if $N(u) = \{v, w\}$ then $\phi(D_u) - \phi(D_v) \geq t$.

Finally, \mathcal{H} is *satisfied* by A if the following events all occur *disjointly*:

- (ix) if v is a leaf then D_v is internally filled by A ;
- (x) if $N(u) = \{v\}$ then $D_u = [D_v \cup (D_u \cap A)]$.

That is, there exists a family of disjoint subsets of A , one for each of these events, such that each guarantees that the corresponding event occurs.

Let us denote by $\mathcal{H}_R(s, t)$ the family of (s, t) -good hierarchies for a rectangle R . The following lemma is a straightforward consequence of the rectangles process (as described above), so we will leave the details of the proof to the reader.

Lemma 2.10 *Let R be a rectangle that is internally filled by $A \subseteq \mathbb{Z}_n^2$. Then, for every $s \geq t \geq 1$, there exists an (s, t) -good hierarchy $\mathcal{H} \in \mathcal{H}_R(s, t)$ that is satisfied by A .*

In order to deduce a bound on the probability that a rectangle is internally filled, we will need the following fundamental inequality of van den Berg and Kesten [12]. Given two increasing³ events $E, F \subseteq \{0, 1\}^{\mathbb{Z}_n^2}$, we write $E \circ F$ for the event that E and F occur disjointly.

Lemma 2.11 (The van den Berg-Kesten inequality) *If E and F are increasing events then*

$$\mathbb{P}_p(E \circ F) \leq \mathbb{P}_p(E) \cdot \mathbb{P}_p(F).$$

We remark that the assumption in Lemma 2.11 that E and F are increasing is in fact not necessary for the conclusion to hold, as was conjectured by van den Berg and Kesten [12], and later proved by Reimer [53].

Let us write $L(\mathcal{H})$ for the set of leaves of $G_{\mathcal{H}}$ (we will refer to the rectangles corresponding to these vertices as the *seeds* of \mathcal{H}), and $\prod_{u \rightarrow v}$ for the product over all pairs $\{u, v\} \subseteq V(G_{\mathcal{H}})$ such that $N(u) = \{v\}$ (we will refer to these as *steps*). Also, given rectangles $R \subseteq R'$, define

$$I(R) = \{[A \cap R] = R\} \quad \text{and} \quad \Delta(R, R') := \{R' = [R \cup (R' \cap A)]\},$$

and note that these events correspond to R being internally filled by A , and R' being internally filled by $A \cup R$, respectively, cf. conditions (ix) and (x) of Definition 2.9.

We are now ready to state and prove Holroyd's fundamental bound on the probability that a rectangle is internally filled by a p -random set A .

³An event $E \subseteq \{0, 1\}^{\mathbb{Z}_n^2}$ is *increasing* if $A \in E$ and $A \subseteq B$ implies $B \in E$.

Lemma 2.12 *Let $R \subseteq \mathbb{Z}_n^2$ be a rectangle. For every $s \geq t \geq 1$, we have*

$$\mathbb{P}_p(I(R)) \leq \sum_{\mathcal{H} \in \mathcal{H}_R(s,t)} \left(\prod_{u \in L(\mathcal{H})} \mathbb{P}_p(I(D_u)) \right) \left(\prod_{u \rightarrow v} \mathbb{P}_p(\Delta(D_v, D_u)) \right). \quad (2.3)$$

Proof If R is internally filled by A then, by Lemma 2.10, there must exist an (s, t) -good hierarchy $\mathcal{H} \in \mathcal{H}_R(s, t)$ that is satisfied by A . By Definition 2.9, if \mathcal{H} is satisfied by A then the events $I(D_u)$ (for each $u \in L(\mathcal{H})$) and $\Delta(D_v, D_u)$ (for each pair $\{u, v\}$ such that $N(u) = \{v\}$) all occur disjointly. The lemma now follows by applying the van den Berg-Kesten inequality to this family of events. \square

We will apply Lemma 2.12 with $s = \varepsilon \log n$ and $t = \varepsilon^2 \log n$, for some sufficiently small constant $\varepsilon > 0$. This choice will allow us to obtain sufficiently strong bounds on the probabilities of the events $I(D_u)$ and $\Delta(D_v, D_u)$ on the right-hand side of (2.3), whilst keeping the total number of hierarchies relatively small. Indeed, we have the following bound on the size of $\mathcal{H}_R(s, t)$.

Lemma 2.13 *If $\phi(R) = O(\log n)$ and $t = \Omega(\log n)$, then*

$$|\mathcal{H}_R(s, t)| = (\log n)^{O(1)}.$$

Proof This follows since if $\mathcal{H} \in \mathcal{H}_R(s, t)$ then the height of the tree $G_{\mathcal{H}}$ is bounded, and hence $|V(G_{\mathcal{H}})| = O(1)$. Since each rectangle in the set $\{R_u : u \in V(G_{\mathcal{H}})\}$ is contained in R , it follows that we have only $(\log n)^{O(1)}$ choices for each, and so the claimed bound follows. \square

We also easily obtain the following bound on the probability that a seed is internally filled, cf. (2.2).

Lemma 2.14 *If R is a rectangle and $k = \lceil \phi(R)/2 \rceil$, then*

$$\mathbb{P}_p(I(R)) \leq (ekp)^k.$$

Proof By Lemma 2.7, if R is internally filled by A then $|A \cap R| \geq k$. Since the area of R is at most k^2 , it follows that

$$\mathbb{P}_p(I(R)) \leq \binom{k^2}{k} p^k \leq (ekp)^k,$$

as claimed. \square

It follows immediately from Lemmas 2.12, 2.13 and 2.14 that the expected number of satisfied hierarchies in $\mathcal{H}_R(s, t)$ such that the sum of the semi-perimeters of the seeds is at least $\delta \log n$ is at most

$$(\log n)^{O(1)} (ep \cdot \varepsilon \log n)^{\delta \log n} \leq \frac{1}{n^3}$$

if $p \leq 1/\log n$ and $\varepsilon = \varepsilon(\delta) > 0$ is sufficiently small. Since there are only $n^{2+o(1)}$ rectangles in \mathbb{Z}_n^2 of semi-perimeter $O(\log n)$, this will be sufficient for our planned application of Markov's inequality.

We may therefore consider from now on only hierarchies with relatively few seeds; however, since doing so is rather more difficult (and technical), we will give only a brief sketch of the necessary ideas. The basic idea is to compare the steps of \mathcal{H} to those in an 'ideal' hierarchy with one large seed of semi-perimeter

$$z(\mathcal{H}) := \sum_{u \in L(\mathcal{H})} \phi(D_u).$$

Since growth becomes easier as the droplet gets larger, it is natural to expect (and possible to prove) that the product

$$\prod_{u \rightarrow v} \mathbb{P}_p(\Delta(D_v, D_u))$$

will be (asymptotically) minimized (over those $\mathcal{H} \in \mathcal{H}_R(s, t)$ with a given value of $z(\mathcal{H})$) by taking a hierarchy of this form. Moreover, the 'easiest' way for such a droplet to grow is 'along the diagonal', i.e., as an approximate square. The basic fact that allows one to show this can be paraphrased as follows:

"A rectangle is crossed by A if and only if it has no double gap."

More precisely, let us write ∂S for the set of sites directly to the left of a rectangle S , and say that S is crossed by A if $S \subseteq [\partial S \cup (A \cap S)]$. Then S is crossed by A if and only if there does not exist a 'double gap', i.e., a consecutive pair of columns of S neither of which contains an element of A . We leave the proof of the following lemma as an exercise for the reader.

Lemma 2.15 *Let S be an $a \times b$ rectangle. Then*

$$\beta(u)^a \leq \mathbb{P}_p(S \text{ is crossed by } A) \leq \beta(u)^{a-1},$$

where $u = 1 - (1 - p)^b$ and $\beta(u) = (u + \sqrt{u(4 - 3u)})/2$.

Moreover, the choice of s and t (with $t = \varepsilon s$) allows one to approximate the event $\Delta(D_v, D_u)$ by the intersection of four ‘crossing’ events (to the right, left, up and down), since the ‘corner regions’ are relatively small, and therefore are unlikely to contain many elements of A .

It follows (after a fair amount of work) that one can bound the probability that a hierarchy $\mathcal{H} \in \mathcal{H}_R(s, t)$ is satisfied by A by

$$\prod_{u \rightarrow v} \mathbb{P}_p(\Delta(D_v, D_u)) \leq \prod_{b=z(\mathcal{H})}^{\phi(R)/2} \beta(u(b))^{2-\delta},$$

where $u(b) = 1 - (1-p)^b$ and $\delta > 0$ is an arbitrarily small⁴ constant. After some (non-trivial) integration (see [39, Section 2], and also [41]), this implies, for every $\gamma > 0$, the bound

$$\mathbb{P}_p(\mathcal{H} \text{ is satisfied by } A) \leq \exp\left(-\frac{\pi^2 - \gamma}{9p}\right)$$

for every $\mathcal{H} \in \mathcal{H}_R(s, t)$ with $z(\mathcal{H}) \leq \delta \log n$, if $\phi(R) \geq C \log n$ for some sufficiently large constant $C = C(\gamma)$, and if $\delta = \delta(\gamma) > 0$ was chosen sufficiently small. Thus, if

$$p \leq \left(\frac{\pi^2}{18} - \gamma\right) \frac{1}{\log n},$$

then we obtain

$$\mathbb{P}_p(I(R)) \leq (\log n)^{O(1)} \exp\left(-\frac{\pi^2 - \gamma}{9p}\right) \leq n^{-2-\gamma}. \quad (2.4)$$

Finally, applying Lemma 2.3 (cf. the proof of Theorem 2.2), if $[A] = \mathbb{Z}_n^2$ then there exists an internally filled rectangle R with

$$C \log n \leq \phi(R) \leq 2C \log n,$$

where $C = C(\gamma)$ is the sufficiently large constant chosen above. By (2.4) (and Markov’s inequality) this has probability at most $n^{-\gamma+o(1)}$, and hence

$$p_c(\mathbb{Z}_n^2, 2) \geq \left(\frac{\pi^2}{18} - \gamma\right) \frac{1}{\log n},$$

for every $\gamma > 0$, as required. This completes our sketch of Holroyd’s proof of the lower bound in Theorem 2.8; we leave the upper bound to the reader.

⁴To prove this for a given δ we must take $\varepsilon = \varepsilon(\delta) > 0$ sufficiently small.

Exercise: Use Lemma 2.15 and the fact (see [39, Section 2]) that

$$\int_0^\infty -\log(\beta(1 - e^{-z})) \, dz = \frac{\pi^2}{18}$$

to prove that

$$p_c(\mathbb{Z}_n^2, 2) \leq \left(\frac{\pi^2}{18} + o(1) \right) \frac{1}{\log n}$$

as $n \rightarrow \infty$.

Much more precise bounds on $p_c(\mathbb{Z}_n^2, 2)$ have since been obtained by Gravner and Holroyd [35], Gravner, Holroyd and Morris [36], and Morris [46], culminating in the following (unpublished) theorem.

Theorem 2.16

$$p_c(\mathbb{Z}_n^2, 2) = \frac{\pi^2}{18 \log n} - \frac{\Theta(1)}{(\log n)^{3/2}}.$$

The proof of Theorem 2.16 requires a significantly more complicated notion of hierarchy than that described above (requiring the use of Reimer's theorem instead of the van den Berg–Kesten inequality), as well as extremely precise bounds on the probability various types of 'steps'. We refer the reader to [46] for (most of) the details.

3 The r -neighbour model on \mathbb{Z}_n^d

The methods of the previous section can be generalized without too much difficulty to the setting of \mathbb{Z}_n^d when $r = 2$. However, when $r \geq 3$ things fall apart very quickly, as the closed sets are no longer (unions of) rectangles. The first breakthrough was made by Schonmann [54] in 1992, who proved the following generalization of van Enter's theorem.

Theorem 3.1 *If $d \geq r \geq 1$, then*

$$p_c(\mathbb{Z}_n^d, r) \rightarrow 0$$

as $n \rightarrow \infty$.

Note that if $r > d$ then $p_c(\mathbb{Z}_n^d, r) = 1 - o(1)$, since a translate of $\{1, 2\}^d$ that contains no element of A will remain uninfected forever (cf. the discussion of 'subcritical update rules' in Sections 7 and 9, below).

Since $p_c(\mathbb{Z}_n^d, r)$ is an increasing function of r , we may assume in the proof of Theorem 3.1 that $r = d$. The key observation is that the d -neighbour process on the side of an infected droplet $D = \{1, \dots, m\}^d$ is ‘dominated’ by a $(d-1)$ -neighbour process on $\{1, \dots, m\}^{d-1}$. To be more precise, suppose that D is internally filled, and note that each element of the set $S = \{m+1\} \times \{1, \dots, m\}^{d-1}$ has exactly one neighbour in D , and thus requires only $d-1$ additional infected neighbours inside S in order to become infected. Hence, if we can obtain a sufficiently strong lower bound on the probability of percolation in the $(d-1)$ -neighbour process on $\{1, \dots, m\}^{d-1}$ for all sufficiently large $m \in \mathbb{N}$, then we will be able to deduce a lower bound on the probability of percolation in the d -neighbour process on \mathbb{Z}_n^d . In other words, we will use induction on d .

Let us write $I_r(R) = \{R \subseteq [A \cap R]\}$ for the event that a set of vertices $R \subseteq \mathbb{Z}_n^d$ is internally filled in the r -neighbour process on \mathbb{Z}_n^d . The induction hypothesis is as follows.

Lemma 3.2 *For each $p > 0$ and $d \geq 1$, there exists $\varepsilon = \varepsilon(d, p) > 0$ such that the following holds. If $R = \{1, \dots, m\}^d$, then*

$$\mathbb{P}_p(I_d(R)) \geq 1 - e^{-\varepsilon m}$$

for all sufficiently large $m \in \mathbb{N}$.

Proof When $d = 1$ the result holds with $\varepsilon(1, p) = p$, since R is internally filled with probability $1 - (1-p)^m$. So let $d \geq 2$ and assume that the claim holds for all smaller values of d . The first step is to show that

$$\mathbb{P}_p(I_d(R)) \rightarrow 1$$

as $m \rightarrow \infty$. To do so, as in the proof of Theorem 2.1, we consider two independent $(p/2)$ -random subsets A_1 and A_2 of R , find a large droplet $D \subseteq A_1$, and use the induction hypothesis to show that $[D \cup A_2] = R$ with high probability. Indeed, if $D = \{1, \dots, m_0\}^d$ for some sufficiently large m_0 then, by the induction hypothesis, the probability that it fails to grow one step in each direction decreases exponentially in its side-length. Thus

$$\mathbb{P}([D \cup A_2] = R) \geq 1 - 2d \sum_{k=m_0}^{m-1} e^{-\varepsilon k} \rightarrow 1$$

as $m_0 \rightarrow \infty$, as claimed.

Now, let $\delta > 0$ be sufficiently small, and fix k such that if $S = \{1, \dots, k\}^d$ then $\mathbb{P}(I_d(S)) \geq 1 - \delta$. Now, by tiling R with copies of S ,

it follows from standard results in percolation theory that with probability at least $1 - e^{-\varepsilon m}$, all components of $R \setminus [A \cap R]$ have diameter at most $m/4$. However, it is easy to see that any such component must in fact be empty, and so the lemma (and hence also the theorem) follows. \square

The proof above can be modified to give the following upper bound:

$$p_c(\mathbb{Z}_n^d, r) \leq \left(\frac{\Theta(1)}{\log_{(r-1)} n} \right)^{d-r+1}, \tag{3.1}$$

where $\log_{(r)} n = \log(\log_{(r-1)} n)$ denotes an r -times iterated logarithm. In a major breakthrough, a matching lower bound was obtained by Cerf and Cirillo [19] (in the case $d = r = 3$), and by Cerf and Manzo [20] (for all $d \geq r \geq 3$). As noted above, the case $r = 2$ was resolved by Aizenman and Lebowitz [1], while the case $r = 1$ is trivial.

Theorem 3.3 *If $d \geq r \geq 1$, then*

$$p_c(\mathbb{Z}_n^d, r) = \left(\frac{\Theta(1)}{\log_{(r-1)} n} \right)^{d-r+1}.$$

The basic idea behind the proof of Theorem 3.3 is as follows: there must exist an internally filled component of diameter about $\log n$, and this implies that a box⁵ of side-length about $\log n$ must be ‘internally spanned’. To be precise, we make the following important definition.

Definition 3.4 A box $R \subseteq \mathbb{Z}_n^d$ is *internally spanned* by A if there exists a connected set $S \subseteq [A \cap R]$ such that R is the smallest box containing S .

Note that if S is a connected subset of \mathbb{Z}_n^d such that $[A \cap S] = S$, then the smallest box containing S is internally spanned by A . This simple observation allows one to prove the following analogue of the Aizenman-Lebowitz lemma (Lemma 2.3) when $r \geq 3$. Let us write $\text{diam}(S)$ for the L^∞ -diameter of a connected set $S \subseteq \mathbb{Z}_n^d$, i.e., the maximum side-length of the smallest box containing S .

Lemma 3.5 *Let R be a box in \mathbb{Z}_n^d , and suppose that R is internally spanned. For every $1 \leq k \leq \text{diam}(R)$, there exists an internally spanned box $Q \subseteq R$ with $k \leq \text{diam}(Q) \leq 2k$.*

To prove this lemma, we replace the rectangles process of Section 2 by the following ‘components process’. Since we will later need to use this process in other settings, we define it on a general (finite) graph G .

⁵That is, a set of the form $[a_1, b_1] \times \dots \times [a_d, b_d]$.

Definition 3.6 (The components process) Given a finite graph G , let $A = \{x_1, \dots, x_m\} \subseteq V(G)$, and set $\mathcal{S} := \{S_1, \dots, S_m\}$, where $S_j = \{x_j\}$ for each $j \in [m]$. Now repeat the following steps until STOP:

1. If there exist a family of $2 \leq s \leq r$ sets $\{S_{i_1}, \dots, S_{i_s}\} \subseteq \mathcal{S}$ such that⁶

$$[S_{i_1} \cup \dots \cup S_{i_s}]$$

is connected in G , then choose a minimal such family, remove them from \mathcal{S} , and replace them by $[S_{i_1} \cup \dots \cup S_{i_s}]$.

2. If there does not exist such a family of sets in \mathcal{S} , then STOP.

Observe that all members of \mathcal{S} (at all points of the process) are closed under the (r -neighbour) bootstrap process, and that

$$A \subseteq V(\mathcal{S}) := \bigcup_{S \in \mathcal{S}} S \subseteq [A].$$

Therefore, if $V(\mathcal{S})$ is not equal to $[A]$ then it is not closed, i.e., there must exist a vertex v with at least r neighbours in $V(\mathcal{S})$. Since the closure of (the union of) the corresponding elements of \mathcal{S} is connected (via v), it follows that the process does not stop until $V(\mathcal{S}) = [A]$.

Now, for the graph $G = \mathbb{Z}_n^d$, it is moreover straightforward to show that if $[S_{i_1} \cup \dots \cup S_{i_s}]$ is connected (and minimal) then

$$\text{diam}([S_{i_1} \cup \dots \cup S_{i_s}]) \leq 2 \cdot \max_{i \in [s]} \{\text{diam}(S_i)\} + 2,$$

since boxes are closed under the r -neighbour bootstrap process on \mathbb{Z}_n^d . Lemma 3.5 now follows easily (cf. the proof of Lemma 2.3), so we leave the details as an exercise for the reader.

Exercise: Write down a complete proof of Lemma 3.5.

In order to prove Theorem 3.3, it remains to bound the probability that a box of side-length k is internally spanned by a p -random set A . For simplicity, we will only sketch the proof in the case $d = r = 3$ (the general case is similar, but the details are somewhat more complicated). Let us denote by $I^\times(R)$ the event that R is internally spanned by A .

Lemma 3.7 *There exists $c > 0$ such that the following holds if*

$$p \leq \frac{c}{\log \log n}.$$

⁶The closure here refers to the r -neighbour bootstrap process on G .

Fix $\delta > 0$, and let $n \in \mathbb{N}$ be sufficiently large. If $R \subseteq \mathbb{Z}_n^3$ is a box of diameter $k \leq \log n$, then

$$\mathbb{P}_p(I^\times(R)) \leq \delta^k.$$

Proof The key idea is to partition R into $k \times k \times 2$ ‘slices’, where we assume (for simplicity) that $R = \{1, \dots, k\}^3$. Each vertex has at most one neighbour in a different slice, and we may therefore couple the 3-neighbour process on R with $k/2$ independent 2-neighbour processes on the slices. Since

$$p \leq \frac{c}{\log \log n} \leq \frac{c}{\log k},$$

it follows from the proof of Aizenman and Lebowitz (see Section 2) that the closure in a given slice contains a component of size larger than $\log k$ with probability at most $1/k$. Moreover, the same method implies that the expected number of vertices v that are contained in the same internally filled component of a given vertex u (only counting those components of diameter at most $\log k$) is $o(1)$. More precisely, if we define, for vertices u and v in a slice X ,

$$u \sim v \Leftrightarrow \text{there exists a component } Y \subseteq X, \text{ with } u, v \in Y \text{ and} \\ \text{diam}(Y) \leq \log k, \text{ such that } Y \text{ is internally spanned} \\ \text{by } A \text{ in the 2-neighbour bootstrap process on } X,$$

then the expected size of $\{v \in X : u \sim v\}$ is $o(1)$ for each $v \in X$.

Now, in order to bound the probability that R is internally spanned, we first condition on the subset Z of slices with a component of size larger than $\log k$. We then count the expected number of ‘minimal’ paths between each consecutive pair in Z , where each step of the path corresponds either to moving between vertices u and v in the same slice with $u \sim v$, or between neighbouring vertices in adjacent slices. By minimality, different steps in the same slice occur disjointly, and so we may use the van den Berg–Kesten inequality (Lemma 2.11) to bound the probability of each such path. The result now follows by some simple counting. \square

For the details, we refer the reader to the original papers [19, 20], or to the more recent applications of the Cerf–Cirillo method in [40, 8, 6].

Exercise: Expand the sketch proof above into a full proof of Lemma 3.7.

It is now easy to deduce Theorem 3.3 (in the case $d = r = 3$) from Lemmas 3.5 and 3.7, cf. the proof of Theorem 2.2.

Proof of Theorem 3.3 The proof of the upper bound was outlined above, so we will prove only the lower bound. If $[A] = \mathbb{Z}_n^3$, then by Lemma 3.5 there exists a box $Q \subseteq \mathbb{Z}_n^3$, with

$$\frac{\log n}{2} \leq \text{diam}(Q) \leq \log n$$

that is internally spanned by A . By Lemma 3.7, the expected number of such boxes is

$$n^{3+o(1)} \delta^{\log n/2} \rightarrow 0$$

as $n \rightarrow \infty$, if $\delta = e^{-8}$, say. Hence

$$p_c(\mathbb{Z}_n^3, 3) \geq \frac{c}{\log \log n},$$

where $c > 0$ is the constant in Lemma 3.7, as required. \square

3.1 A sharp threshold for the r -neighbour process on \mathbb{Z}_n^d

Despite the breakthroughs of Cerf and Cirillo [19] and Holroyd [39], the problem of determining a sharp threshold for $p_c(\mathbb{Z}_n^d, r)$ remained open until a few years ago, when it was finally resolved by Balogh, Bollobás, Duminil-Copin and Morris [6, 8].

Theorem 3.8 *For every $d \geq r \geq 2$, there exists a constant $\lambda(d, r) > 0$ such that*

$$p_c(\mathbb{Z}_n^d, r) = \left(\frac{\lambda(d, r) + o(1)}{\log_{(r-1)} n} \right)^{d-r+1}$$

as $n \rightarrow \infty$.

The constant $\lambda(d, r)$ is defined as follows. For each $k \in \mathbb{N}$, let⁷

$$\beta_k(u) = \frac{1}{2} \left(1 - (1-u)^k + \left(1 + (4u-2)(1-u)^k + (1-u)^{2k} \right)^{1/2} \right),$$

and set $g_k(z) = -\log(\beta_k(1 - e^{-z}))$. Then

$$\lambda(d, r) = \int_0^\infty g_{r-1}(z^{d-r+1}) dz.$$

This definition, while a little complicated, has the following nice properties:

$$\lambda(d, 2) = \frac{d-1}{2} + o(1) \quad \text{and} \quad \lambda(d, d) = \left(\frac{\pi^2}{6} + o(1) \right) \frac{1}{d}$$

⁷Note that $\beta_k(u)^2 = (1 - (1-u)^k)\beta_k(u) + u(1-u)^k$.

as $d \rightarrow \infty$.

The proof of Theorem 3.8 is complicated, but the basic idea (at least when $d = r = 3$) is quite simple: one would like to apply the Cerf–Cirillo method to slices whose width is a large constant C , instead of two. For any fixed C , the two sides of the slice (which are assumed to get help from the adjacent slices) will be able to cooperate enough to affect the constant, but as $C \rightarrow \infty$ the effect of this cooperation becomes arbitrarily small. The main technical difficulty is then to extend Holroyd’s method to the (significantly) more complicated context of a slice; we remark that proving an analogue of Lemma 2.15 is particularly challenging, and the tools that were developed in [6, 8] in order to resolve this problem have proven extremely useful in other contexts, see Section 8. For the details of the proof, we refer the interested reader to [6].

4 Bootstrap in high dimensions

In Sections 2 and 3 we studied the r -neighbour bootstrap process on \mathbb{Z}_n^d with d and r fixed, and n sufficiently large. It is natural to ask how large n needs to be for these results to hold, and (similarly) to ask for bounds on $p_c(\mathbb{Z}_n^d, r)$ when n and r are allowed to be arbitrary functions of $d \rightarrow \infty$. This question was first asked by Balogh and Bollobás [5], who determined $p_c(Q_d, 2)$ up to a constant factor, where $Q_d = \mathbb{Z}_2^d$ is the d -dimensional hypercube. The sharp threshold was determined several years later by Balogh, Bollobás and Morris [9], who proved that

$$\left(1 + \frac{\log d}{\sqrt{d}}\right) \frac{16\lambda}{d^2} 2^{-2\sqrt{d}} \leq p_c(Q_d, 2) \leq \left(1 + \frac{5(\log d)^2}{\sqrt{d}}\right) \frac{16\lambda}{d^2} 2^{-2\sqrt{d}}$$

for all sufficiently large $d \in \mathbb{N}$, where $\lambda \approx 1.17$ is the smallest positive root of the equation

$$\sum_{k=0}^{\infty} \frac{(-1)^k \lambda^k}{2^{k^2-k} k!} = 0.$$

They also obtained the following sharp threshold whenever $d \gg \log n \gg 1$.

Theorem 4.1 *If $d \gg \log n \gg 1$, then*

$$p_c(\mathbb{Z}_n^d, 2) = \frac{4\lambda + o(1)}{d^2} 2^{-\sqrt{d \log_2 n}}$$

as $d \rightarrow \infty$.

The proof of Theorem 4.1 is based on a very precise enumeration of the number of close to minimum-size percolating sets in a subcube of

‘critical’ size. This is proved using induction on the dimension, using (a higher dimensional analogue of) the rectangles process. Since the analysis is rather involved, we refer the interested reader to [9] for the details.

Theorems 3.8 and 4.1 determine $p_c(\mathbb{Z}_n^d, 2)$ up to a factor of $1 + o(1)$ when $d = O(1)$ and $d \gg \log n$, but no bounds are known for intermediate values of d .

Open problem: Determine $p_c(\mathbb{Z}_n^d, 2)$ when $1 \ll d = O(\log n)$.

It seems likely that the method described in Section 2 could be extended to give reasonable bounds for a relatively slowly-growing function (perhaps $d = \log \log n$), but new ideas are likely to be necessary in order to deal with larger d , and in particular the range $d = \Theta(\log n)$.

As when $d = O(1)$, the r -neighbour bootstrap process becomes much more complicated when $r \geq 3$, and very little is known about $p_c(Q_d, r)$ when $r \geq 3$ is fixed $d \rightarrow \infty$. In an important breakthrough, however, the *extremal* problem was recently resolved by Morrison and Noel [48]. To be precise, given a graph G and $r \in \mathbb{N}$, define

$$m(G, r) = \min \{|A| : A \subseteq V(G) \text{ percolates in the } r\text{-neighbour process}\}.$$

Morrison and Noel proved the following theorem, which confirms (in a strong form) a conjecture of Balogh and Bollobás [5].

Theorem 4.2 For each $r \in \mathbb{N}$,

$$m(Q_d, r) = \frac{d^{r-1}}{r!} + \Theta(d^{r-2}).$$

This breakthrough provides hope that it might be possible to make progress on the following conjecture of Balogh, Bollobás and Morris [9].

Conjecture 4.3 For each $r \in \mathbb{N}$,

$$p_c(Q_d, r) = \exp\left(-\Theta(d^{1/2^{r-1}})\right).$$

To motivate this conjecture, let us give a brief sketch of the proof of the upper bound, which proceeds by considering a growing subcube of infected vertices. The key lemma states that

$$\mathbb{P}_p(I_r(Q_\ell)) \geq \exp\left(-(\log 1/p)^{2^{r-1}}\right)$$

for every $\ell \in \mathbb{N}$, and the proof proceeds by induction on r , noting that if Q_ℓ is already infected, then on a neighbouring subcube we may couple

with an $(r - 1)$ -neighbour process. After roughly $(\log 1/p)^{2^{r-2}}$ steps, each of which adds one dimension to the growing cube, we pass the ‘bottleneck’ and growth becomes easier. The proof now proceeds as usual: we find a large internally filled subcube somewhere in Q_d , and (using sprinkling) show it is likely to grow to infect the entire hypercube.

4.1 Large d and r

For the d -neighbour model, Balogh, Bollobás and Morris [7] proved the following theorem.

Theorem 4.4 *If $d \gg (\log \log n)^2 \log \log \log n$, then*

$$p_c(\mathbb{Z}_n^d, d) \rightarrow \frac{1}{2}$$

as $d \rightarrow \infty$.

The proof of Theorem 4.4 is completely different from any of those described above. Note that if $p = 1/2 + \varepsilon$ then almost all vertices are infected in the first step, whereas if $p = 1/2 - \varepsilon$ then almost no vertices are infected. For the upper bound, the idea is to show that if a vertex v is not infected after $k \approx \sqrt{d/\log d}$ steps, then there exists a set X of about $(cd/k)^k$ vertices (for some constant $c > 0$) at distance k from v , none of which was infected in the first step. Since for vertices $x, y \in X$ at distance at least three, the events $x \notin A_1$ and $y \notin A_1$ are independent, it is straightforward to bound the probability that such a set exists, and hence the expected number of vertices that are not in A_k .

The proof of the lower bound is more interesting (and the key idea will also be needed in the next section). The main difficulty is that the process can continue for a long time, and this creates long-distance dependencies that complicate the calculations. The solution is to introduce a modified ‘more generous’ process that infects vertices with slightly fewer than d infected neighbours in the first few steps: if a vertex is not infected during the early ‘generous’ part of the algorithm, it needs to gain many new infected neighbours in order to be infected later. This allows one to show that, if a vertex v is infected after k steps of the generous process, there must be a set of about $(c'd/k)^k$ vertices at distance k from v , all of which were infected in the first step. See [7, Section 6] for the details.

Combining Theorems 3.3 and 4.4, it follows that (roughly speaking):

$$p_c(\mathbb{Z}_n^d, d) \rightarrow \begin{cases} 1/2 & \text{if } n \ll 2^{2^{\sqrt{d}}} \\ 0 & \text{if } n \gg 2^{2^{\dots^2}}, \text{ a tower of 2s of height } d. \end{cases}$$

It would be very interesting to improve these bounds, and it seems likely that a significant breakthrough will be necessary in order to describe the transition between the two regimes.

Open problem: Determine $p_c(\mathbb{Z}_n^d, d)$ in the range $2^{2^{\sqrt{d}}} \leq n \leq 2^{2^{\dots^2}}$.

The lower bound seems ‘softer’ to us, and we expect the threshold to lie closer to the upper bound. More precisely, we suspect that the function

$$f(d) := \min \{n \in \mathbb{N} : p_c(\mathbb{Z}_n^d, d) \leq 1/4\}$$

grows like a tower of 2s of height $\Theta(d)$.

5 The Glauber dynamics of the Ising model on \mathbb{Z}^d

Before introducing the general model that is the main topic of this survey, let us conclude our gentle introduction to the area by describing an application of the results presented above to the Ising model on \mathbb{Z}^d . This is an extremely well-studied model of ferromagnetism (see, for example, [43, 52] and the references therein), and provided the original motivation for the introduction of bootstrap percolation in [23].

Definition 5.1 The zero-temperature Glauber dynamics of the Ising model are defined as follows:

- (a) At each time $t \geq 0$, the system is in a state $\sigma(t) \in \{+, -\}^{\mathbb{Z}^d}$.
- (b) Each vertex $x \in \mathbb{Z}^d$ has an independent exponential clock, which rings randomly at rate 1.
- (c) If the clock at vertex x rings at time t , then its state $\sigma_x(t) \in \{+, -\}$ resets according to the following rule:
 - if x has at least $d + 1$ neighbours in state $+$, then $\sigma_x(t) := +$.
 - if x has at least $d + 1$ neighbours in state $-$, then $\sigma_x(t) := -$.
 - if x has exactly d neighbours in each state, then $\sigma_x(t)$ is chosen uniformly at random, independently of all other events.

For the sake of brevity, in this section we will refer to this process as simply ‘the Glauber dynamics on \mathbb{Z}^d ’.

We are interested in the long-term behaviour of the dynamics, starting from a randomly chosen initial state. In particular, does the system ‘fixate’, or do some vertices change state infinitely many times?

In order to state this question more precisely, let us write \mathbb{P}_p^+ for the probability distribution obtained by setting $\sigma_x(0) = +$ independently at random with probability p for each $x \in \mathbb{Z}^d$. We say that a vertex $x \in \mathbb{Z}^d$ *fixates* if its state changes only finitely many times, and we say that the dynamics *fixates* if every vertex fixates. (If the state of every vertex is eventually $+$, then we moreover say that the dynamics *fixates at $+$* .) The critical probability for the zero-temperature Glauber dynamics of the Ising model on \mathbb{Z}^d is then defined as follows:

$$p_c^{\text{Is}}(\mathbb{Z}^d) := \inf \left\{ p : \mathbb{P}_p^+ (\text{the dynamics on } \mathbb{Z}^d \text{ fixates at } +) = 1 \right\}.$$

In one dimension, the critical probability was determined by Arratia [2] in 1983, who proved (a significant generalisation of) a conjecture of Erdős and Ney [32] about annihilating random walks on \mathbb{Z} . His result implies that, for every $p \in (0, 1)$, every vertex almost surely changes state infinitely many times in the dynamics on \mathbb{Z} , and hence that $p_c^{\text{Is}}(\mathbb{Z}) = 1$.

In two or more dimensions, on the other hand, the problem of determining $p_c^{\text{Is}}(\mathbb{Z}^d)$ is wide open. The following (possibly folklore) conjecture is the main topic of this section.

Conjecture 5.2

$$p_c^{\text{Is}}(\mathbb{Z}^d) = \frac{1}{2}$$

for every $d \geq 2$.

An important breakthrough on Conjecture 5.2 was achieved by Fontes, Schonmann and Sidoravicius [34], who proved the following theorem.

Theorem 5.3

$$p_c^{\text{Is}}(\mathbb{Z}^d) < 1$$

for every $d \geq 2$.

The proof of Theorem 5.3 (which we will sketch below) uses some of the results of Aizenman and Lebowitz [1] described in Section 2. More recently, Morris [45] used the main result of [34], together with the techniques introduced by Balogh, Bollobás and Morris [7] in their work on majority bootstrap percolation on the hypercube, in order to prove the following strengthening in high dimensions.

Theorem 5.4

$$p_c^{\text{Is}}(\mathbb{Z}^d) \rightarrow \frac{1}{2}$$

as $d \rightarrow \infty$.

The proof of Fontes, Schonmann and Sidoravicius [34] uses multi-scale analysis, which (roughly speaking) means using induction to control the process inside an increasingly-large sequence of boxes. Indeed, consider two scales, n and N , with $N = n^C$ for some large constant C , and tile \mathbb{Z}^d with copies of $[n]^d$ and $[N]^d$. Suppose that we have already found a coupling of the original dynamics with a dynamics that is ‘more generous’⁸ to $-$, such that, in the coupled dynamics:

- (a) Each tile (i.e., copy of $[n]^d$ in our tiling) is entirely $+$ at time t with probability at least $1 - q$, for suitable t and q .
- (b) For non-adjacent tiles (i.e., copies of $[n]^d$ in our tiling that do not share a corner), the events described in (a) are independent.

We then aim to find a similar coupling at scale N , except replacing t and q by some (much larger) t' and (much smaller) q' . The key observation is that the evolution of the $-$ spins in $[N]^d$ can be coupled with two-neighbour bootstrap percolation on $[N/n]^d$, and if $q \ll (\log(N/n))^{-d+1}$, then, by the results of Aizenman and Lebowitz [1] (cf. the proof of Theorem 2.2), the closure of this bootstrap process contains a component (of $[n]^d$ -tiles) of size larger than $\log n$ with probability at most $1/n$. Now, each such component is surrounded (and hence attacked) by $+$ vertices, and it can be shown that the probability that it survives until time n^{2d} (say) is extremely small. Moreover, since information only propagates at rate 1, it is extremely unlikely that non-adjacent $[N]^d$ -tiles interact before this time, if $N \geq n^{3d}$ (say). Using these facts, it is straightforward to construct the required coupling, with $t' = t + n^{2d}$ and $q' = 1/n$.

The proof of Theorem 5.4 is similar, except we replace the first step of the argument above (which holds trivially if p is sufficiently close to 1) by a more careful analysis, assuming that the dimension d is sufficiently large. More precisely, we tile \mathbb{Z}^d with copies of $[n]^d$, where $n = 2^d$, and couple the dynamics in each tile with a more generous process which satisfies conditions (a) and (b) above. The coupling is constructed in several steps, but the key idea is to run the more generous bootstrap process used to prove Theorem 4.4 for a *bounded* number of steps on the set of vertices in $[n]^d$ that were initially in state $-$. With (extremely) high probability this process reaches a set X that is closed under majority bootstrap percolation, and (crucially) there are no long-range dependencies between the events $\{x \in X\}$. This allows one to couple the state at time d with a distribution in which each vertex is in state $-$ with super-polynomially

⁸That is, at every time t , the set of vertices in state $-$ in the original dynamics is contained in the the set of vertices in state $-$ in the coupled dynamics.

small probability (in d), and vertices at distance at least 20 have independent states. A similar idea then allows one to turn the remaining $-$ vertices to $+$ by time d^5 (with extremely high probability), and since 2^d is much larger than d^5 , it is extremely unlikely that non-adjacent $[n]^d$ -tiles interact before this time. Inserting this argument into the proof of Fontes, Schonmann and Sidoravicius, we obtain Theorem 5.4.

5.1 The symmetric setting, $p = 1/2$

Recall that Arratia [2] proved that on \mathbb{Z} , every vertex almost surely changes state infinitely many times for every $p \in (0, 1)$. Nanda, Newman and Stein [50] proved that this is also the case on \mathbb{Z}^2 when $p = 1/2$.

Theorem 5.5 *If $p = 1/2$ then, in Glauber dynamics on \mathbb{Z}^2 , almost surely every vertex changes state an infinite number of times.*

Proof It follows by a standard ergodic theory argument that a given row or column contains a site which fixates at $+$ with probability either 0 or 1. If this probability is zero then we are done (by symmetry between $+$ and $-$, and between rows and columns), so let us assume that it occurs almost surely. This implies that there exists an axis-parallel rectangle R such that three of its corners fixate in alternating states. In other words, there exist x, y and z , with x and y in the same column, and x and z in the same row, such that x fixates at $+$, whereas y and z fixate at $-$.

Now, at any time t after x, y and z have fixated, the boundary (in the dual lattice) between $+$ and $-$ must connect two adjacent sides of R , either side of one of x, y or z . But then that vertex will change state in $[t, t + 1]$ with positive probability, which is a contradiction. \square

The problem appears to be much harder in higher dimensions.

Open problem: If $p = 1/2$ and $d \geq 3$ then, in Glauber dynamics on \mathbb{Z}^d , does every vertex almost surely change state an infinite number of times?

6 A few other specific bootstrap processes on \mathbb{Z}^2

Before defining the general model of Bollobás, Smith and Uzzell [17], let us provide a gentle warm-up (and some motivation) by considering a few specific bootstrap-like processes in two dimensions. The only difference between the processes discussed below and the two-neighbour model discussed in Section 2 will be the *update rule*: in each case, we need to define under what conditions a vertex is infected in step $t + 1$.

6.1 Modified bootstrap percolation

The simplest variant of the two-neighbour model is a process known as *modified bootstrap percolation*, whose update rule is as follows: a vertex x is infected once it has at least one already-infected horizontal neighbour, and at least one already-infected vertical neighbour. In other words, $x \in A_{t+1}$ if either $x \in A_t$, or $x + X \subseteq A_t$ for some $X \in \mathcal{M}^{(2)}$, where

$$\mathcal{M}^{(2)} = \left\{ \{(0, 1), (1, 0)\}, \{(1, 0), (0, -1)\}, \{(0, -1), (-1, 0)\}, \{(-1, 0), (0, 1)\} \right\}.$$

In this model the closed sets are still unions of rectangles, and the proofs of Aizenman and Lebowitz [1] and Holroyd [39] go through with only minor modifications. In fact, the natural generalization of the modified model to \mathbb{Z}_n^d (in which a vertex is infected if it has at least one infected neighbour in each dimension) is significantly simpler to study than the d -neighbour model when $d \geq 3$. Indeed, Holroyd [40] was able to determine the sharp threshold in d dimensions, some years before the corresponding result was obtained for the r -neighbour model, by combining the technique of [39] with the method of Cerf and Cirillo [19].

Theorem 6.1

$$p_c(\mathbb{Z}_n^d, \mathcal{M}^{(d)}) = \left(\frac{\pi^2}{6} + o(1) \right) \frac{1}{\log_{(d-1)} n}.$$

Note that here we write $p_c(\mathbb{Z}_n^d, \mathcal{M}^{(d)})$ for the critical probability of the modified model in d dimensions, see Definition 7.1, below.⁹

Exercise: Verify that the proof of Theorem 2.2 also implies that

$$p_c(\mathbb{Z}_n^2, \mathcal{M}^{(2)}) = \frac{\Theta(1)}{\log n}.$$

6.2 Anisotropic bootstrap percolation

A significantly more challenging variant, which exhibits rather different behaviour, is the following so-called *anisotropic* model, which was first studied by Gravner and Griffeath [37]: The update rule in this setting is as follows: a vertex x is infected once there are at least three infected vertices in the set

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

⁹To be precise, $\mathcal{M}^{(d)}$ denotes the family of sets that contain exactly one element of the pair $\{e_i, -e_i\}$ for each $i \in [d]$, and no other elements, where $e_i = (0, \dots, 0, 1, 0, \dots, 0)$ is the i th basis vector.

The threshold for this model was determined by van Enter and Hulshof [31], and the sharp threshold by Duminil-Copin and van Enter [28], who proved the following theorem.

Theorem 6.2

$$p_c(\mathbb{Z}_n^2, \mathcal{A}) = \left(\frac{1}{12} + o(1) \right) \frac{(\log \log n)^2}{\log n}.$$

Here we write $p_c(\mathbb{Z}_n^2, \mathcal{A})$ for the critical probability in the anisotropic model.¹⁰ Note that in the anisotropic model the critical probability has increased by a factor of $(\log \log n)^2$ (compared with the two-neighbour model); this is because the typical growth of a droplet in the anisotropic model is (in a certain sense) *asymptotically one-dimensional*, unlike in the case of the two-neighbour model. This is an important point, so let us spend a little time discussing the proof of the following weaker result which, as mentioned above, was proved by van Enter and Hulshof [31]:

$$p_c(\mathbb{Z}_n^2, \mathcal{A}) = \Theta\left(\frac{(\log \log n)^2}{\log n}\right). \quad (6.1)$$

Let us begin with the upper bound. Let R_0 be a rectangle of height $a = p^{-1} \log(1/p)$ and width $b = p^{-2}$, and note that if A contains the two left-most columns of R_0 , and at least one vertex in every column, then it is internally filled by A . This has probability at least

$$p^{2a}(1 - (1 - p)^a)^b \geq p^{2a}(1 - p)^b \geq p^{2a}e^{-2/p}.$$

Next, let $R_1 \supseteq R_0$ be a rectangle of height $5a$ and width b , and note that if A contains two adjacent vertices of each row of R_1 , then $[R_0 \cup (A \cap R_1)] = R_1$. This has probability at least

$$(1 - (1 - p^2)^{b/2})^{5a} \geq e^{-10a}.$$

Now, let $R_2 \supseteq R_1$ be a rectangle of height $5a$ and width b^2 , and note that if A contains a vertex of each column of R_2 , then $[R_1 \cup (A \cap R_2)] = R_2$. This has probability at least

$$(1 - (1 - p)^{5a})^{b^2} \geq (1 - p^5)^{b^2} \geq \frac{1}{2}.$$

It follows that if

$$p \geq \frac{2(\log \log n)^2}{\log n}$$

¹⁰This matches the notation introduced in the next section if \mathcal{A} denotes the collection of subsets of $\{(-2, 0), (-1, 0), (0, 1), (0, -1), (1, 0), (2, 0)\}$ of size three.

then with high probability there exist at least

$$\frac{n^2}{5ab^2} \cdot p^{2a} e^{-2/p-10a-1} \geq n^2 \exp\left(-\frac{3}{p}\left(\log\frac{1}{p}\right)^2\right) \gg 1$$

internally filled translates of R_2 in \mathbb{Z}_n^2 . Finally, if every row of $b^2 = p^{-4}$ consecutive vertices contains at least two adjacent elements of A , and every column of \mathbb{Z}_n^2 intersects A , then any internally filled translate of R_2 grows to infect the whole of \mathbb{Z}_n^2 . Since both of these events occur with high probability, this proves that $p_c(\mathbb{Z}_n^2, \mathcal{A}) = O((\log \log n)^2 / \log n)$.

The alert reader will have noticed several points in the above argument that were not sharp, and indeed a more careful (and somewhat more complicated) proof along the same lines gives the upper bound in Theorem 6.2; we leave the details as an exercise for the reader.

Exercise: By modifying the argument above, show that

$$p_c(\mathbb{Z}_n^2, \mathcal{A}) \leq \left(\frac{1}{12} + o(1)\right) \frac{(\log \log n)^2}{\log n}.$$

To prove the lower bound in (6.1), we apply the components process (see Definition 3.6) with G the graph on \mathbb{Z}_n^2 with edge set

$$E(G) = \{xy : |x_1 - y_1| + 2 \cdot |x_2 - y_2| \leq 2\}.$$

Using this process, one can show that if A percolates, then there must exist an internally spanned¹¹ rectangle R with either

(a) height at most $\delta p^{-1} \log(1/p)$ and width at least $p^{-3/2}$, or

(b) height at least $\delta p^{-1} \log(1/p)$ and width at most $3p^{-3/2}$,

for some (small) $\delta > 0$. (Indeed, simply run the components process until we find an internally spanned rectangle with width between $p^{-3/2}$ and $3p^{-3/2}$.) Now, if R is internally spanned, then every three consecutive columns of R must contain at least one element of A , and every pair of consecutive rows of R must contain at least two elements of A that are within distance two in G . Thus, setting $a = \delta p^{-1} \log(1/p)$ and $b = p^{-3/2}$, the probability that R is internally spanned is at most

$$\max\left\{(1 - (1-p)^{3a})^{b/3}, (30p^2b)^{a/2}\right\} \leq \exp\left(-\frac{\delta}{5p}\left(\log\frac{1}{p}\right)^2\right).$$

¹¹For the anisotropic model, we say that a box $R \subseteq \mathbb{Z}_n^d$ is *internally spanned* by A if there exists a set $S \subseteq [A \cap R]$ that is connected in the graph G , such that R is the smallest box containing S , cf. Definition 3.4.

Since there are at most n^3 rectangles R in \mathbb{Z}_n^2 that satisfy either (a) or (b) above, it now follows, by Markov's inequality, that A percolates with probability at most

$$n^3 \exp\left(-\frac{\delta}{5p}\left(\log\frac{1}{p}\right)^2\right) \leq \frac{1}{n}$$

if $p \leq \delta^2(\log \log n)^2 / \log n$, which implies the lower bound in (6.1).

The lower bound in Theorem 6.2 is significantly more difficult, and follows by adapting the technique of Holroyd [39] to the anisotropic setting. We refer the interested reader to [28, 31] for the details.

Finally, we note that the threshold for a class of anisotropic models in three dimensions was determined by van Enter and Fey [30]. More precisely, they considered the r -neighbour process with (a, b, c) -neighbourhood

$$\{(x, 0, 0) : \pm x \in [a]\} \cup \{(0, y, 0) : \pm y \in [b]\} \cup \{(0, 0, z) : \pm z \in [c]\}$$

and $r = a + b + c$. They showed that if $a = b \leq c$, then

$$p_c(\mathbb{Z}_n^3, \mathcal{A}(a, b, c)) = \Theta\left(\frac{1}{\log \log n}\right)^{1/a},$$

whereas if $a < b \leq c$, then

$$p_c(\mathbb{Z}_n^3, \mathcal{A}(a, b, c)) = \Theta\left(\frac{(\log \log \log n)^2}{\log \log n}\right)^{1/a},$$

where we write $\mathcal{A}(a, b, c)$ for the update family of the model. It would be interesting to determine a sharp threshold for this class of models, and more generally for all $c + 1 \leq r \leq a + b + c$.

6.3 The Duarte model

Our final example is the simplest example of a class of models that cause the greatest difficulty in the general setting of the next two sections: those with 'drift'. Its update rule is as follows: a vertex x is infected once there are at least two infected vertices in the set

$$x + \{(-1, 0), (0, 1), (0, -1)\}.$$

This model was first introduced by Duarte [26] in 1988, and first studied rigorously a few years later by Mountford [49], who determined the threshold for percolation. The following sharp threshold was determined only very recently, by Bollobás, Duminil-Copin, Morris and Smith [14].

Theorem 6.3

$$p_c(\mathbb{Z}_n^2, \mathcal{D}) = \left(\frac{1}{8} + o(1)\right) \frac{(\log \log n)^2}{\log n}.$$

Here we write $p_c(\mathbb{Z}_n^2, \mathcal{D})$ for the critical probability in the Duarte model.¹² The threshold is therefore the same (up to a constant factor) as in the anisotropic model, and the typical growth of a droplet in the anisotropic model is again asymptotically one-dimensional. However, in the Duarte model there is a significant difference: growth only occurs in one direction ('to the right'), and a typical droplet is not a rectangle, but a non-polygonal convex shape that is taller at the right end than the left.

To illustrate these points, let us briefly discuss how to modify the proof of (6.1) above to deduce the following weaker bounds, which were first proved by Mountford [49]:

$$p_c(\mathbb{Z}_n^2, \mathcal{D}) = \Theta\left(\frac{(\log \log n)^2}{\log n}\right). \quad (6.2)$$

The proof of the upper bound is very similar to that given above, the main difference being that for the droplet to grow vertically, it is not sufficient that A contains two adjacent vertices of each row. Instead, we need to find a sequence of (single) elements of A , one in each pair of consecutive rows, which form an 'increasing' sequence (meaning their x -coordinates are an increasing function of their y -coordinates). We will typically be able to find such a sequence of length roughly p times the current length of the droplet, and we thus obtain a vertical line of this height in $[A]$. This vertical line can then be used to continue growth to the right, which can then be used to grow upwards, and so on.

To be more precise, let R_0 be a rectangle of height $a = 4p^{-1} \log(1/p)$ and width $b = p^{-5}$, and note that if A contains the left-most column of R_0 , and at least one vertex in every column, then it is internally filled by A . This has probability at least

$$p^a(1 - (1 - p)^a)^b \geq p^a(1 - p^a)^b \geq p^a e^{-2/p}.$$

Next, let $R_1 \supseteq R_0$ be a rectangle of height p^{-3} and width b , and note that if A contains an increasing sequence of vertices as described above, then the right-hand side of R_1 is contained in $[R_0 \cup (A \cap R_1)]$. This occurs with high probability, since $pb \gg p^{-3}$. Finally, if every row or columns of p^{-3} consecutive vertices intersects A , then any translate of the right-hand

¹²This matches the notation introduced in the next section if \mathcal{D} denotes the collection of subsets of $\{(-1, 0), (0, 1), (0, -1)\}$ of size two.

side of R_1 grows to infect the whole of \mathbb{Z}_n^2 . Since this occurs with high probability, it is easy to deduce that $p_c(\mathbb{Z}_n^2, \mathcal{D}) = O((\log \log n)^2 / \log n)$.

Once again, a more careful (and somewhat more complicated) proof along the same lines gives the upper bound in Theorem 6.3; we again leave the details as an exercise for the reader.

Exercise: By modifying the argument above, show that

$$p_c(\mathbb{Z}_n^2, \mathcal{D}) \leq \left(\frac{1}{8} + o(1)\right) \frac{(\log \log n)^2}{\log n}.$$

The proof of the lower bound in (6.2) is significantly more complicated than that of the lower bound in (6.1); we will try to explain the new difficulties that arise, and give an idea of how they may be overcome.

The first problem we run into – the fact that the interaction between vertices is no longer symmetric, cf. Definition 3.6 – is easy to overcome in this particular setting: we simply apply a slight modification of the components process in which we take the closure in the Duarte model, but define connectivity according to the graph G on \mathbb{Z}_n^2 with edge set

$$E(G) = \{xy : |x_1 - y_1| \leq 1 \text{ and } |x_1 - y_1| + |x_2 - y_2| \leq 2\}.$$

Using this process, one can show (as before) that if A percolates, then there must exist an internally spanned rectangle R with either

- (a) height at most $\delta p^{-1} \log(1/p)$ and width at least $p^{-4/3}$, or
- (b) height at least $\delta p^{-1} \log(1/p)$ and width at most $3p^{-4/3}$,

for some (small) $\delta > 0$. In the first case we are done, since every column of an internally spanned rectangle must intersect A , and so a rectangle satisfying condition (a) is internally spanned with probability at most

$$(1 - (1 - p)^a)^b \leq \exp(-p^{-4/3+2\delta}),$$

where $a = \delta p^{-1} \log(1/p)$ and $b = p^{-4/3}$. However, it is not so easy to bound the probability that a rectangle R satisfying condition (b) is internally spanned, as even if $A \cap R$ does not contain a large ‘increasing’ sequence as in the proof of the upper bound, it can still be internally spanned (for example, via the cooperation of many small internally spanned droplets).

This obstacle is overcome in different ways in the papers [13, 14, 49]; we will describe the method of [13], specialised to the setting of the Duarte model, where most of the difficulties encountered for general update rules do not occur. We first partition the droplet into strips of height $p^{-2/3}$,

and observe that each of these must either be ‘half-crossed’ from above, or from below. Next, we partition (most of) each half-strip into constant width strips that are rotated slightly, so that the right-hand end of each is about $p^{-1/2}$ higher than its left-hand end. The crucial observation is now that each constant width strip must either contain a ‘cluster’ of elements of A , close enough together that they can cooperate, or must intersect a large internally spanned ‘saver’ droplet, and moreover all of these events occur disjointly. Since sites in a cluster must lie within distance $p^{-1/2}$ of one another, this allows us to prove a bound on the probability that R satisfying condition (b) is internally spanned of the form

$$p^{\delta a} \leq \exp\left(-\frac{\delta^2}{p}\left(\log\frac{1}{p}\right)^2\right),$$

where $a = \delta p^{-1} \log(1/p)$. The claimed lower bound in (6.2) now follows as in Section 6.2.

The proof of the sharp threshold for the Duarte model requires two important additional innovations, which in turn create a large number of additional technical difficulties. These are the ‘method of iterated hierarchies’, and the use of non-polygonal droplets. The former technique was introduced in [13] in order to determine the threshold for general ‘critical’ update families (see Sections 7 and 8, below), and involves applying Holroyd’s method of hierarchies (see Section 2.1) at various different scales. We deduce sharp bounds at a given scale by using (even sharper) bounds at a smaller scale in order to control the probability that a saver droplet is internally spanned. At sufficiently small scales, it is relatively easy to obtain extremely strong bounds using (a suitably modified version of) the method of Aizenman and Lebowitz, see Section 2. However, one cannot obtain sharp bounds using rectangular droplets; instead, one needs to define a droplet as follows.

Definition 6.4 Given $\varepsilon > 0$ and $p > 0$, a *Duarte region* $D^* \subseteq \mathbb{R}^2$ is a set of the form

$$D^* = (a, b) + \{(x, y) \in \mathbb{R}^2 : 0 \leq x \leq w, |y| \leq f(x)\}, \quad (6.3)$$

for some $a, b, w \in \mathbb{R}$, where $f: [0, \infty) \rightarrow [0, \infty)$ is the function

$$f(x) := \frac{1}{p} \log\left(1 + \frac{\varepsilon^2 p x}{\log 1/p}\right).$$

A *Duarte droplet* $D \subseteq \mathbb{Z}^2$ is the intersection of a Duarte region with \mathbb{Z}^2 .

With this crucial definition in hand, it is possible (though rather non-trivial) to carry out the proof outlined above; see [14] for the details.

7 General cellular automata: the BSU model

We now make a significant transition: from considering specific update rules one at a time, to dealing with large families of models simultaneously. As a consequence, many of the definitions become somewhat more technical, and the processes can be more difficult to visualise; nevertheless, the approach we take is surprisingly simple and natural.

Let us begin by defining a large class of d -dimensional monotone cellular automata, introduced by Bollobás, Smith and Uzzell [17].

Definition 7.1 (The \mathcal{U} -bootstrap process) Let $\mathcal{U} = \{X_1, \dots, X_m\}$ be an arbitrary finite collection of finite subsets of $\mathbb{Z}^d \setminus \{\mathbf{0}\}$. The \mathcal{U} -bootstrap process on the d -dimensional torus \mathbb{Z}_n^d is defined as follows: given a set $A \subseteq \mathbb{Z}_n^d$ of initially *infected* sites, set $A_0 = A$, and define, for each $t \geq 0$,

$$A_{t+1} = A_t \cup \{x \in \mathbb{Z}_n^d : x + X \subseteq A_t \text{ for some } X \in \mathcal{U}\}.$$

The \mathcal{U} -bootstrap process on \mathbb{Z}^d is defined in an analogous way.

Let $[A]_{\mathcal{U}} = \bigcup_{t \geq 0} A_t$ denote the *closure* of A under the \mathcal{U} -bootstrap process, and define the *critical probability* of the \mathcal{U} -bootstrap process on \mathbb{Z}_n^d to be

$$p_c(\mathbb{Z}_n^d, \mathcal{U}) := \inf \left\{ p : \mathbb{P}_p([A]_{\mathcal{U}} = \mathbb{Z}_n^d) \geq 1/2 \right\}.$$

We will refer to \mathcal{U} as the *update family* of the process. Note that each of the processes considered above is a \mathcal{U} -bootstrap process for some update family \mathcal{U} . For example, the r -neighbour model on \mathbb{Z}_n^d is obtained by setting $\mathcal{U} = \mathcal{N}_r^d$, the family of r -subsets of the $2d$ nearest neighbours of the origin in \mathbb{Z}^d , and the Duarte model is obtained by setting $\mathcal{U} = \mathcal{D}$, the collection of subsets of $\{(-1, 0), (0, 1), (0, -1)\}$ of size two.

One of the key insights of Bollobás, Smith and Uzzell [17] was that the typical global behaviour of a two-dimensional \mathcal{U} -bootstrap process acting on random initial sets should be determined by the action of the process on discrete half-spaces. (The same also turns out to be true in higher dimensions, but there the situation is significantly more complicated.) For the rest of this section, let us therefore restrict ourselves to \mathbb{Z}^2 .

Definition 7.2 (The set of stable directions) For each $u \in S^1$, let

$$\mathbb{H}_u := \{x \in \mathbb{Z}^2 : \langle x, u \rangle < 0\}$$

denote the discrete half-plane whose boundary is perpendicular to u .

We say that u is a *stable direction* if

$$[\mathbb{H}_u]_{\mathcal{U}} = \mathbb{H}_u$$

and we denote by $\mathcal{S} = \mathcal{S}(\mathcal{U}) \subseteq S^1$ the collection of stable directions.

It is easy to determine the stable set of an update family: simply remove from S^1 the (possibly empty) open interval of directions that are destabilised by X , for each $X \in \mathcal{U}$. In particular, note that $\mathcal{S}(\mathcal{N}_1^2)$ is empty and $\mathcal{S}(\mathcal{N}_3^2) = S^1$, while $\mathcal{S}(\mathcal{N}_2^2)$ consists of four isolated points, and $\mathcal{S}(\mathcal{D})$ consists of a closed semicircle and an isolated point. We remark that $\mathcal{S}(\mathcal{U})$ always consists of a finite collection of closed intervals, some of which may be isolated points, all of whose endpoints are rational.

Let \mathcal{C} denote the collection of open semicircles in S^1 . The following key definition is due to Bollobás, Smith and Uzzell [17].

Definition 7.3 (Critical and super/subcritical update families) We say that a two-dimensional update family \mathcal{U} is:

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

Note that this is a partition of the two-dimensional update families.

The importance of Definition 7.3 is demonstrated by the following theorem. Parts (a) and (b) were proved by Bollobás, Smith and Uzzell [17], and part (c) was proved by Balister, Bollobás, Przykucki and Smith [4].

Theorem 7.4 *Let \mathcal{U} be a two-dimensional update family.*

- (a) *if \mathcal{U} is supercritical then $p_c(\mathbb{Z}_n^2, \mathcal{U}) = n^{-\Theta(1)}$.*
- (b) *if \mathcal{U} is critical then $p_c(\mathbb{Z}_n^2, \mathcal{U}) = (\log n)^{-\Theta(1)}$.*
- (c) *if \mathcal{U} is subcritical then $\liminf_{n \rightarrow \infty} p_c(\mathbb{Z}_n^2, \mathcal{U}) > 0$.*

It is perhaps difficult to convey to the reader how surprising it is that such a simple and beautiful characterization could be proved in such extraordinary generality. The proof of Theorem 7.4 is quite complicated, and so we will give only a brief outline of some of the key ideas.

7.1 Supercritical update families

The only supercritical model we have come across in the sections above is the 1-neighbour process, which is somewhat trivial. This could easily lead one to suspect that all supercritical models are ‘easy’, but this would be mistaken: they exhibit an *extremely* rich variety of complex behaviours, and even proving the relatively weak bound claimed in Theorem 7.4(a) will require some important new ideas. Nevertheless, this is easiest part of the proof, and will serve as a useful warm-up.

Recall that if \mathcal{U} is supercritical, then there exists an open semicircle C in S^1 that contains no stable direction. Choose such a semicircle with rational endpoints, and let u^+ denote its midpoint. For simplicity we will rotate \mathbb{Z}_n^2 so that $u^+ = (1, 0)$, i.e., u^+ points to the right. We claim that there exists a constant $k = k(\mathcal{U})$ such that the closure of the continuous square¹³ $R = [-k, k]^2$ contains all vertices in the (minimal) horizontal strip containing R . If $p \geq n^{-1/2k^2}$ (say), then with high probability A contains such a square in every translate of this strip, so this will be sufficient to prove Theorem 7.4(a).

To prove the claim, it will suffice to show that the column L directly to the right of R lies in the closure of R . Since u^+ is unstable, it follows that there exists a rule $X_1 \in \mathcal{U}$ that is entirely contained in \mathbb{H}_{u^+} , and this allows us to infect all but a constant (depending on the size of $|X_1|$) number of vertices of L , which lie at the top and bottom. Infecting the last few vertices is rather more complicated, and requires the entire semicircle C to be unstable; roughly speaking, the idea is as follows. First, using the rule X_1 we can construct a ‘pyramid’ of infected sites on the right-hand side of R . The sides of this pyramid correspond to the endpoints of the open interval that is destabilised by X_1 , and so (since C is entirely unstable) there must exist another rule $X_2 \in \mathcal{U}$ that destabilises the sides of the pyramid. This gives us another pyramid with slightly steeper sides, and by repeating this process sufficiently many times, we eventually infect the whole of L .

In order to make the above sketch precise we will need an important idea from [17], which also turns out to be extremely useful when proving upper bounds for critical update families. Define a set

$$\mathcal{Q} := \bigcup_{X \in \mathcal{U}} \bigcup_{x \in X} \{u \in S^1 : \langle u, x \rangle = 0\} \tag{7.1}$$

of *quasi-stable* directions by taking the two unit vectors u and $-u$ perpendicular to x (considered as a vector) for every site $x \in X$ and every

¹³By this, we mean the set of vertices of the rotated \mathbb{Z}_n^2 that lie in $[-k, k]^2$.

update rule $X \in \mathcal{U}$. The following lemma allows one to control the growth of droplets whose sides are all perpendicular to members of \mathcal{Q} .

Lemma 7.5 *If u, v are consecutive elements of the set $\mathcal{S}(\mathcal{U}) \cup \mathcal{Q}$, then there exists an update rule X such that*

$$X \subseteq (\mathbb{H}_u \cup \ell_u) \cap (\mathbb{H}_v \cup \ell_v),$$

where $\ell_u = \{x \in \mathbb{Z}^2 : \langle x, u \rangle = 0\}$.

Once \mathcal{Q} has been defined as above, the proof of Lemma 7.5 is actually quite straightforward (see Figure 1)¹⁴. Indeed, let $w \in S^1$ be a direction between u and v , and note that since w is not stable, there exists an update rule $X \subseteq \mathbb{H}_w$. Suppose the conclusion of the lemma fails; then without loss of generality there exists $x \in X$ such that $\langle x, v \rangle < 0$ and $\langle x, u \rangle > 0$. But this implies that there exists $w' \in S^1$ perpendicular to x with w' between u and v , contradicting the construction of \mathcal{Q} .

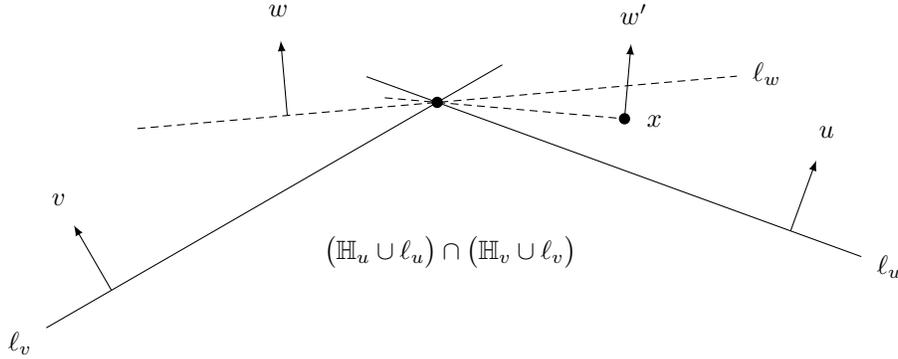


Figure 1: The proof of Lemma 7.5.

Now, it follows from Lemma 7.5 that a ‘frontier’ consisting of (reasonably long) intervals perpendicular to the elements of $\mathcal{Q} \cap C$ can be extended (under the action of \mathcal{U}) by one line in each direction, and hence L is contained in the closure of the rectangle R , as claimed. We refer the reader to [17, Sections 5 and 7] for the details.

¹⁴All figures used in this paper are by Paul Smith, and are reproduced here with his kind permission. We remark that Figure 1 originally appeared as Figure 4 in [17].

7.2 Critical update families: the upper bound

The upper bound in Theorem 7.4(b) again uses quasi-stable directions, but also requires one or two extra important ideas. Let us begin by stating the key lemma in the proof of the upper bound for critical families.

Lemma 7.6 *Let $u \in S^1$ be an isolated point of $\mathcal{S}(\mathcal{U})$. Then there exists a finite set $Z \subseteq \ell_u$ such that $\ell_u \subseteq [\mathbb{H}_u \cup Z]_{\mathcal{U}}$.*

Proof It is sufficient to show that a sufficiently long interval $Z \subseteq \ell_u$ grows by one in both directions. To see that it grows to the right, consider a direction v a little to the right of u , and observe that v is unstable if it is chosen sufficiently close to u . Now let $X \in \mathcal{U}$ be a rule which destabilises v , and observe that if v is chosen sufficiently close to u then $X \setminus \mathbb{H}_u$ is contained in Z . Therefore Z grows to the right, and by symmetry also to the left, as required. \square

The proof is now almost identical to that outlined in Section 7.1, the only difference being that the rectangle R must have size at least p^{-a} , where a is sufficiently large so that for every $u \in \mathcal{S}(\mathcal{U}) \cap C$, there exists a set Z as in Lemma 7.6 of size at most $a - 2$. Indeed, once our growing droplet D is this large, we are very likely to find a translate of Z on the side of D at each step, and this allows one to show that D grows to fill an entire horizontal strip (with very high probability), as before.

7.3 Critical update families: the lower bound

The lower bound in Theorem 7.4(b) is significantly more difficult to prove than the upper bounds sketched above, and required the introduction of another important concept: the *covering algorithm*. To define this algorithm, observe first that if \mathcal{U} is a critical two-dimensional update family, then there must exist a set \mathcal{T} of three or four stable directions which intersects every open semicircle in S^1 . We will control the process using \mathcal{T} -droplets, defined as follows:

Definition 7.7 A \mathcal{T} -droplet is a non-empty set of the form

$$D = \bigcap_{u \in \mathcal{T}} (\mathbb{H}_u + a_u)$$

for some collection $\{a_u \in \mathbb{Z}^2 : u \in \mathcal{T}\}$.

The covering algorithm replaces the rectangles process of Section 2, with \mathcal{T} -droplets playing the role of rectangles. The first step is to choose a

sufficiently large constant $\kappa = \kappa(\mathcal{U})$, fix a \mathcal{T} -droplet \hat{D} of diameter roughly κ , and place a copy of \hat{D} (arbitrarily) on each element of A . Now, at each step, if two \mathcal{T} -droplets in the current collection are within distance κ^2 of one another, then remove them from the collection, and replace them by the smallest \mathcal{T} -droplet containing both. If κ is chosen sufficiently large, then one can prove that the final collection of \mathcal{T} -droplets cover $[A]_{\mathcal{U}}$.

If a \mathcal{T} -droplet occurs at some point in the covering algorithm, then let us say that it is *covered* by A . Also, let us write $\text{diam}(D)$ for the diameter of a \mathcal{T} -droplet D , i.e., the maximum distance between two points in D . Using the covering algorithm, one can prove the following analogues of Lemmas 2.3 and 2.7; we leave the details to the reader.¹⁵

Lemma 7.8 (Aizenman–Lebowitz lemma for covered droplets) *If $[A]_{\mathcal{U}} = \mathbb{Z}_n^2$, then for every $\kappa^3 \leq k \leq n$, there exists a covered \mathcal{T} -droplet $D \subseteq \mathbb{Z}_n^2$ with $k \leq \text{diam}(D) \leq 3k$.*

Lemma 7.9 (Extremal lemma for covered droplets) *There exists a constant $\varepsilon = \varepsilon(\mathcal{U}) > 0$ such that*

$$|D \cap A| \geq \varepsilon \cdot \text{diam}(D)$$

for every covered \mathcal{T} -droplet D .

Once we have these two lemmas, it is relatively straightforward to deduce the claimed lower bound, using the method of Aizenman and Lebowitz [1].

Proof of the lower bound in Theorem 7.4(b) Suppose that $A \subseteq \mathbb{Z}_n^2$ is such that $[A]_{\mathcal{U}} = \mathbb{Z}_n^2$. By Lemma 7.8, there exists a covered \mathcal{T} -droplet D with

$$\frac{\log n}{\varepsilon} \leq \text{diam}(D) \leq \frac{3 \log n}{\varepsilon}.$$

Let X denote the random variable that counts the number of such \mathcal{T} -droplets when A is a p -random subset of \mathbb{Z}_n^2 , and observe that

$$\mathbb{E}[X] \leq \sum_{k=\varepsilon^{-1} \log n}^{3\varepsilon^{-1} \log n} n^3 \binom{k^2}{\varepsilon k} p^{\varepsilon k} \leq \sum_{k=\varepsilon^{-1} \log n}^{3\varepsilon^{-1} \log n} n^3 \left(\frac{ekp}{\varepsilon} \right)^{\varepsilon k} \leq \frac{1}{n}$$

if $p < \varepsilon^3 / \log n$, by Lemma 7.9, since the number of \mathcal{T} -droplets in \mathbb{Z}_n^2 with semi-perimeter k is $n^{2+o(1)}$, and each has area at most k^2 . It follows

¹⁵Note that in both lemmas we implicitly assume that \mathcal{U} is a critical two-dimensional update family, so that there exists a set $\mathcal{T} \subseteq \mathcal{S}(\mathcal{U})$ and (sufficiently large) constant $\kappa > 0$ as described above.

from Markov's inequality that a p -random set $A \subseteq \mathbb{Z}_n^2$ percolates with probability at most $1/n$, and hence

$$p_c(\mathbb{Z}_n^2, \mathcal{U}) \geq \frac{\varepsilon^3}{\log n},$$

as required. \square

7.4 Subcritical update families

When there is no 'easy' direction in which to grow, the behaviour of the \mathcal{U} -bootstrap process changes dramatically, and the proof of Theorem 7.4(c) more closely resembles that of Theorem 5.3 than the lower bounds on $p_c(\mathbb{Z}_n^2, \mathcal{U})$ that we have seen so far. Indeed, subcritical processes behave much more like classical models from percolation theory than bootstrap processes, and the main result of Balister, Bollobás, Przykucki and Smith [4] is actually the following theorem.

Theorem 7.10 *Let \mathcal{U} be a subcritical two-dimensional update family. If $p > 0$ is sufficiently small, and A is a p -random subset of \mathbb{Z}^2 , then almost surely $[A]_{\mathcal{U}}$ contains no infinite component.*

If \mathcal{U} is a subcritical two-dimensional update family, then there must exist a set \mathcal{T} consisting of three stable *intervals* in S^1 , such that \mathcal{T} intersects every open semicircle in S^1 (cf. Section 7.3). The proof of Theorem 7.10 is via multi-scale analysis, using \mathcal{T} -droplets to control the process.

The basic idea is to partition \mathbb{Z}^2 into squares of side-length $n(i)$ for each $i \in \mathbb{N}$, where $n(1) < n(2) < \dots$ is a suitable increasing sequence of 'scales', and to define (at each scale) a notion of a 'bad' square, so that good squares at scale $i + 1$ contain only small, well-separated clusters of bad squares at scale i . (When $i = 1$, a square is good if and only if it contains no element of A .) The key step is then to cover each cluster of bad squares with a \mathcal{T} -droplet whose sides are sufficiently far from any other bad square. This is done by first placing a triangle (whose sides correspond to points in the interior of \mathcal{T}) on top of the cluster, and then 'locally adjusting' the sides of the triangle so as to avoid bad squares that happen to be a little too close. (The corners of the initial triangle, the definition of a bad square, and the scales $n(i)$, all need to be chosen rather carefully for this to be possible.) We refer the reader to [4] for the details.

Finally, we remark that the authors of [4] posed a number of extremely interesting (and challenging) open problems; see [4, Section 7].

8 Critical update rules in two dimensions

We saw in the previous section that two-dimensional bootstrap processes have poly-logarithmic thresholds if, and only if, they are critical. These are therefore (in some sense) the ‘correct’ generalization of the classical two-neighbour process, and of the anisotropic and Duarte models discussed in Section 6. In this section, we will see how to determine the threshold for *every* process in this family.

To be precise, we will define a parameter $\alpha = \alpha(\mathcal{U}) \in \mathbb{N}$, and a bipartition of the set of critical update families into ‘balanced’ and ‘unbalanced’ families, such that the following ‘universality’ theorem of Bollobás, Duminil-Copin, Morris and Smith [13] holds.

Theorem 8.1 *Let \mathcal{U} be a critical two-dimensional update family.*

(a) *If \mathcal{U} is balanced, then*

$$p_c(\mathbb{Z}_n^2, \mathcal{U}) = \Theta\left(\frac{1}{\log n}\right)^{1/\alpha}.$$

(b) *If \mathcal{U} is unbalanced, then*

$$p_c(\mathbb{Z}_n^2, \mathcal{U}) = \Theta\left(\frac{(\log \log n)^2}{\log n}\right)^{1/\alpha}.$$

Roughly speaking, the parameter α is determined by the ‘difficulty’ of growth in the ‘easiest’ direction, where the difficulty of a stable direction $u \in \mathcal{S}(\mathcal{U})$ is the number of ‘nearby’ infected sites needed for \mathbb{H}_u to infect the line ℓ_u , and the difficulty of growth in direction u is the maximum difficulty of a stable direction in the open semicircle centred at u . Even more roughly speaking, a critical two-dimensional update family \mathcal{U} is ‘balanced’ if growth under the \mathcal{U} -bootstrap process is completely two-dimensional (like the two-neighbour process) and unbalanced if it is asymptotically one-dimensional (like the anisotropic and Duarte models).

To define these concepts precisely, we will need some extra terminology. Let $\mathcal{Q}_1 \subseteq S^1$ denote the set of rational directions on the circle, and for each $u \in \mathcal{Q}_1$, let ℓ_u^+ (resp. ℓ_u^-) be the (infinite) subset of the line ℓ_u consisting of the sites to the right (resp. left) of the origin as one looks in the direction of u . Now, let $\alpha_{\mathcal{U}}^+(u)$ (resp. $\alpha_{\mathcal{U}}^-(u)$) denote the minimum (possibly infinite) cardinality of a set $Z \subseteq \mathbb{Z}^2$ such that the set $[\mathbb{H}_u \cup Z]_{\mathcal{U}}$ contains infinitely many sites of ℓ_u^+ (resp. ℓ_u^-).

Definition 8.2 Given $u \in \mathcal{Q}_1$, the *difficulty* of u (with respect to \mathcal{U}) is¹⁶

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

We define the *difficulty* of \mathcal{U} to be

$$\alpha := \min_{C \in \mathcal{C}} \max_{u \in C} \alpha(u), \quad (8.1)$$

where (as before) \mathcal{C} denotes the collection of open semicircles of S^1 .

We remark that $\alpha(u) = 0$ if and only if u is unstable, and that

$$\alpha(u) < \infty \quad \Leftrightarrow \quad u \text{ is unstable or isolated in } \mathcal{S}(\mathcal{U}).$$

Indeed, it follows from Lemma 7.6 that if u is an isolated point of $\mathcal{S}(\mathcal{U})$ then $\alpha(u) < \infty$. It is moreover not hard to show that if u is not an isolated point of $\mathcal{S}(\mathcal{U})$ then $\alpha(u) = \infty$ (consider a stable direction v a little to one side of u , and observe that $\mathbb{H}_u \cup \mathbb{H}_v$ is stable).

It follows (cf. Lemma 7.6) that a direction u has finite difficulty if, and only if, there exists a finite set Z of sites that, such that $\ell_u \subseteq [\mathbb{H}_u \cup Z]_{\mathcal{U}}$. Moreover, the difficulty is at least k if every such Z contains at least k infected sites within a bounded distance of one another (see Lemma 8.4, below). If the open semicircle centred at u contains no direction of difficulty greater than k , then it is possible for a ‘critical droplet’ of infected sites to grow in direction u without ever finding more than k infected sites close together.

We can now define what it means for an update family to be balanced.

Definition 8.3 A critical update family \mathcal{U} is *balanced* if there exists a closed semicircle C such that $\alpha(u) \leq \alpha$ for all $u \in C$. It is said to be *unbalanced* otherwise.

As noted above, it turns out that growth under the action of balanced critical families is completely two-dimensional, while that for unbalanced critical families is asymptotically one-dimensional. Despite this fact, it turns out that analyzing the \mathcal{U} -bootstrap process when \mathcal{U} is unbalanced is significantly more difficult than when it is balanced; the most problematic class of all are the (Duarte-like) ‘drift’ models, see below.

¹⁶In order to slightly simplify the notation, and since the update family \mathcal{U} will always be clear from the context, we will drop the dependence of the difficulty on \mathcal{U} .

8.1 Upper bounds

The upper bounds in Theorem 8.1 follow from a more careful application of the method of quasi-stable directions (see Sections 7.1 and 7.2), combined (in the case of unbalanced families) with the method used in Section 6.3 to bound the critical probability of the Duarte model from above. The main new tool is the following sharp version of Lemma 7.6.

Lemma 8.4 *For every $u \in S^1$, there exists a set $Z \subseteq \mathbb{Z}^2$ of size $\alpha(u)$ such that*

$$\ell_u \subseteq [\mathbb{H}_u \cup (Z + a_1 + k_1 b) \cup \dots \cup (Z + a_r + k_r b)]_{\mathcal{U}}$$

for some $r \in \mathbb{N}$ and $a_1, \dots, a_r, b \in \ell_u$, and every $k_1, \dots, k_r \in \mathbb{Z}$.

In words, this says that a bounded number of voracious sets are sufficient (together with \mathbb{H}_u) to infect ℓ_u , and moreover we may choose any suitable translation of each voracious set. To prove Lemma 8.4, suppose (without loss of generality) that $\alpha_{\mathcal{U}}^+(u) = \alpha(u)$, and observe that (by definition) there exists a set $Z \subseteq \mathbb{Z}^2$ of size $\alpha(u)$ such that $[\mathbb{H}_u \cup Z]_{\mathcal{U}}$ contains infinitely many sites of ℓ_u^+ . One next needs to show (exercise) that, given such a set Z , there exists an infinite arithmetic progression in $[\mathbb{H}_u \cup Z]_{\mathcal{U}}$, and hence deduce that a bounded collection of translates of Z (as in the statement of the lemma) are sufficient to infect all of ℓ_u^+ . Finally, assuming that $\alpha(u) < \infty$ (otherwise the lemma is trivial), it follows from Lemma 7.6 that $[\mathbb{H}_u \cup \ell_u^+]_{\mathcal{U}} = \ell_u$. We leave the details to the reader.

In order to deduce the claimed upper bounds, we again use the set \mathcal{Q} of quasi-stable directions defined in (7.1). Suppose first that \mathcal{U} is balanced, and observe that there exists a closed arc C in S^1 , strictly containing a closed semicircle, such that $\alpha(u) \leq \alpha$ for every $u \in C$. We set

$$\mathcal{T} := (\mathcal{S}(\mathcal{U}) \cup \mathcal{Q} \cup \{w_1, w_2\}) \cap C$$

where w_1 and w_2 are the endpoints of C , and consider \mathcal{T} -droplets.

By Lemma 8.4, a \mathcal{T} -droplet D grows by one step in direction $u \in \mathcal{T} \setminus \{w_1, w_2\}$ (see Figure 2) with probability at least

$$(1 - (1 - p^\alpha)^{\Omega(m)})^{O(1)},$$

where m is the length of the side of D corresponding to u . The proof now follows the usual argument, noting that the ‘critical scale’ of our droplet (above which growth becomes likely) is therefore of order $p^{-\alpha}$.

If \mathcal{U} is unbalanced, on the other hand, then we replace the closed arc C above by a closed semicircle, and repeat the proof, except starting with a rectangular droplet of height $O(p^{-\alpha} \log(1/p))$, and alternately growing

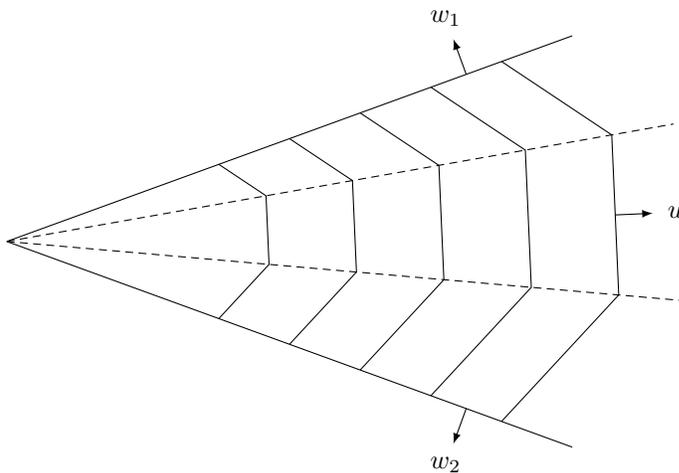


Figure 2: The growth of a droplet when the update family is balanced.

in the ‘easy’ and ‘hard’ directions, cf. Section 6.3. We refer the reader to [13, Section 5] for the details of the proof.

Exercise: Expand the above sketch to a complete proof of the upper bound in Theorem 8.1.

8.2 Balanced update families

The proof of the lower bound for balanced update families is again not all that much more difficult than the lower bound of Bollobás, Smith and Uzzell [17]. The key new idea is to modify the covering algorithm (see Section 7.3) by only covering ‘clusters’ of α nearby sites. The crucial observation is that, by the definition of α , the remaining sites do not contribute significantly to the growth of a droplet.

Observe first that, by (8.1), if \mathcal{U} is a critical two-dimensional update family then there must exist a set \mathcal{T} of three or four stable directions, each of difficulty at least α , which intersects every open semicircle in S^1 . Choose a sufficiently large constant $\kappa = \kappa(\mathcal{U})$, and say that two vertices are *strongly connected* if the distance between them is at most κ . Now, define an α -cluster to be any strongly connected set of α vertices.

The α -covering algorithm is similar to the covering algorithm, but instead of placing a copy of the \mathcal{T} -droplet \hat{D} on each element of A , we place one only on each α -cluster in A .

Definition 8.5 (The α -covering algorithm) Given $K \subseteq A$, choose a maximal collection \mathcal{B} of disjoint α -clusters in K , and a collection \mathcal{D} of copies of a fixed, sufficiently large \mathcal{T} -droplet \hat{D} , such that the elements of \mathcal{D} cover the elements of \mathcal{B} . Now repeat the following steps until STOP:

1. If there are two droplets D and D' in the current collection and an $x \in \mathbb{Z}^2$ such that the set

$$D \cup D' \cup (x + \hat{D})$$

is strongly connected, then remove them from the collection, and replace them by the smallest \mathcal{T} -droplet containing $D \cup D'$.

2. If there do not exist such a pair of droplets, then STOP.

We call the final collection of \mathcal{T} -droplets an α -cover of K .

We will also say that a droplet D is α -covered if the single droplet $\mathcal{D} = \{D\}$ is an α -cover of $D \cap A$, i.e., a possible output of the α -covering algorithm. The following two key lemmas follow easily from the algorithm, as in Section 7.3; we again leave the details to the reader.

Lemma 8.6 (Aizenman–Lebowitz lemma for α -covered droplets) *Let D be an α -covered droplet. Then for every $1 \ll k \leq \text{diam}(D)$, there exists an α -covered droplet $D' \subseteq D$ such that $k \leq \text{diam}(D') \leq 3k$.*

Lemma 8.7 (Extremal lemma for α -covered droplets) *There exists a constant $\varepsilon = \varepsilon(\mathcal{U}) > 0$ such that every α -covered \mathcal{T} -droplet D contains at least $\varepsilon \cdot \text{diam}(D)$ disjoint α -clusters of elements of A .*

It remains to show that an α -cover \mathcal{D} of a set K is a reasonable approximation of the closure $[K]_{\mathcal{U}}$. The basic idea is simple: since all α -clusters are contained in some droplet of \mathcal{D} , the remaining ‘dust’ of $K \setminus (D_1 \cup \dots \cup D_k)$ should contribute only ‘locally’ to the set of eventually infected sites. The following lemma makes this notion precise.

Lemma 8.8 *There exists a constant $\rho = \rho(\mathcal{U}) > 0$ such that if \mathcal{D} is an α -cover of a set K , then every vertex in $[K]_{\mathcal{U}}$ is either contained in an element of \mathcal{D} , or lies within distance ρ of some element of K .*

To prove the lemma, we partition the ‘dust’

$$Y := K \setminus X, \quad \text{where} \quad X := \bigcup_{D \in \mathcal{D}} D,$$

into a collection Y_1, \dots, Y_s of maximal strongly connected components. We then show that, since each strongly connected component has size at most $\alpha - 1$, and each element of \mathcal{T} has difficulty at least α , each Y_i can (with the help of X) only infect a bounded number of extra vertices. Since κ was chosen sufficiently large, this means that different strongly connected components are too far apart to cooperate, and the lemma follows. We refer the reader to [13, Section 6.1] for the details.

Proof of the lower bound in Theorem 8.1(a) Suppose that $A \subseteq \mathbb{Z}_n^2$ is such that $[A]_{\mathcal{U}} = \mathbb{Z}_n^2$. By Lemmas 8.6 and 8.8, there exists an α -covered \mathcal{T} -droplet D with

$$\log n \leq \text{diam}(D) \leq 3 \log n.$$

Let X denote the random variable that counts the number of such \mathcal{T} -droplets when A is a p -random subset of \mathbb{Z}_n^2 , and observe that there are $n^{2+o(1)}$ \mathcal{T} -droplets in \mathbb{Z}_n^2 with semi-perimeter k , and each contains at most Ck^2 (not necessarily infected) α -clusters for some constant $C = C(\mathcal{U}) > 0$. Thus, by Lemma 8.7, we have

$$\mathbb{E}[X] \leq \sum_{k=\log n}^{3 \log n} n^3 \binom{Ck^2}{\varepsilon k} p^{\varepsilon \alpha k} \leq \sum_{k=\log n}^{3 \log n} n^3 (C^2 k p^\alpha)^{\varepsilon k} \leq \frac{1}{n}$$

if $p^\alpha \log n$ is sufficiently small. It follows from Markov's inequality that a p -random set $A \subseteq \mathbb{Z}_n^2$ percolates with probability at most $1/n$, and hence

$$p_c(\mathbb{Z}_n^2, \mathcal{U}) = \Omega\left(\frac{1}{\log n}\right)^{1/\alpha},$$

as required. □

For a particular class of balanced models, a sharp threshold is known. Given a two-dimensional update family \mathcal{U} , let us say that \mathcal{U} is a *threshold rule* if it consist of the the r -subsets of some neighbourhood $N \subseteq \mathbb{Z}^2$ of the origin. Moreover, let us say that N is a symmetric star-neighbourhood if $x \in N$ implies that $-x \in N$, and moreover that every vertex of \mathbb{Z}^2 on the straight line between x and $-x$ is in N . The following theorem was proved by Duminil-Copin and Holroyd [27].

Theorem 8.9 *Let \mathcal{U} be a balanced critical two-dimensional update family. If \mathcal{U} is a threshold rule with a symmetric star-neighbourhood, then*

$$p_c(\mathbb{Z}_n^2, \mathcal{U}) = \left(\frac{\lambda + o(1)}{\log n}\right)^{1/\alpha}$$

for some constant $\lambda = \lambda(\mathcal{U}) > 0$.

The constant $\lambda(\mathcal{U})$ can be computed as the solution of a variational problem, see [27, Section 4] for the details. It is a very interesting (and likely very challenging) open problem to generalise Theorem 8.9 to non-symmetric and unbalanced update families.

8.3 Unbalanced update families

Finally we arrive at the real challenge: proving the lower bound in Theorem 8.1 for unbalanced update families. The proof is significantly more difficult than any of those described above, and we will only be able to give a very rough sketch of the main ideas.

The first key observation is that if \mathcal{U} is an unbalanced critical two-dimensional update family, then there exist a pair of opposite directions $\{u, -u\}$ that both have difficulty at least $\alpha + 1$. We set

$$\mathcal{T} = \{u, -u, v, v'\},$$

where v and v' are stable directions in opposite semicircles with endpoints $\{u, -u\}$, each of difficulty at least α . Observe that the pair $\{u, -u\}$ exists by Definition 8.3, and the pair $\{v, v'\}$ exists by (8.1). Let us rotate our perspective so that u is vertical, and write $h(D)$ and $w(D)$ for (respectively) the height and width of an \mathcal{T} -droplet D .

As in the balanced case, we will need a suitable variant of the rectangles process; however, instead of the α -covering algorithm (which fails to capture the non-linear geometry of unbalanced models), we use the following variant of the components process introduced in Section 3. As above, we say that two vertices are strongly connected if the distance between them is at most κ , where $\kappa = \kappa(\mathcal{U})$ is a sufficiently large constant.

Definition 8.10 (The spanning algorithm) Let $K = \{x_1, \dots, x_m\}$, and set $\mathcal{S} := \{S_1, \dots, S_m\}$, where $S_j = \{x_j\}$ for each $j \in [m]$. Now repeat the following steps until STOP:

1. If there exist a pair of sets $\{S_1, S_2\} \subseteq \mathcal{S}$ such that

$$[S_1]_{\mathcal{U}} \cup [S_2]_{\mathcal{U}}$$

is strongly connected in G , then remove S_1 and S_2 from \mathcal{S} , and replace them by $S_1 \cup S_2$.

2. If there do not exist such a family of sets in \mathcal{S} , then STOP.

The *span* of K is defined to be

$$\langle K \rangle = \{D(S) : S \in \mathcal{S}\},$$

where $D(S)$ denotes the smallest \mathcal{T} -droplet containing S , and \mathcal{S} is the final collection of sets in the spanning algorithm. We would like to say that a \mathcal{T} -droplet D is internally spanned by A if $D \in \langle D \cap A \rangle$. This turns out to be true if we modify Definition 3.4 as follows.

Definition 8.11 A \mathcal{T} -droplet D is *internally spanned* by A if there exists a strongly connected set $S \subseteq [D \cap A]_{\mathcal{U}}$ such that $D = D(S)$. We write $I^\times(D)$ for the event that D is internally spanned.

Fix a small constant $\varepsilon = \varepsilon(\mathcal{U}) > 0$. We say that a \mathcal{T} -droplet is *critical* if either of the following conditions hold:

- (a) $w(D) \leq p^{-\alpha-1/5}$ and $\varepsilon p^{-\alpha} \log(1/p) \leq h(D) \leq 3\varepsilon p^{-\alpha} \log(1/p)$.
- (b) $p^{-\alpha-1/5} \leq w(D) \leq 3p^{-\alpha-1/5}$ and $h(D) \leq \varepsilon p^{-\alpha} \log(1/p)$.

The spanning algorithm allows us to prove the following Aizenman–Lebowitz lemma for internally spanned \mathcal{T} -droplets.

Lemma 8.12 (Aizenman–Lebowitz lemma for critical \mathcal{T} -droplets)
If $[A]_{\mathcal{U}} = \mathbb{Z}_n^2$, then there exists an internally spanned critical \mathcal{T} -droplet.

Note that the reason we need both types of critical droplet is that the spanning algorithm cannot control the width and the height of the critical droplet simultaneously. We choose the constant ε so that a critical \mathcal{T} -droplet is not overwhelmingly likely to grow sideways (that is, perpendicular to u); in particular, if the probability of growing one step is roughly $1 - p^\varepsilon$, then the width $p^{-\alpha-1/5}$ is large enough so that the probability that a critical droplet of type (b) is internally spanned should be sufficiently small to compensate for the number of choices for the initial rectangle. Unfortunately, however, if \mathcal{U} is not symmetric then there is no easy way to prove this, and we need to use the method of hierarchies. In fact, when $\alpha > 1$ the situation is much worse, and we need a more complicated method, which we call the *method of iterated hierarchies*.

In order to motivate this method, let us first see why a straightforward application of the method of hierarchies cannot work. The first key reason is that the extremal lemma we obtain via the spanning algorithm is much less powerful than that given by the α -covering algorithm.

Lemma 8.13 (Extremal lemma for internally spanned \mathcal{T} -droplets)
There exists a constant $\varepsilon = \varepsilon(\mathcal{U}) > 0$ such that every internally spanned \mathcal{T} -droplet D contains at least $\varepsilon \cdot \max\{h(D), w(D)\}$ elements of A .

This lemma only implies a non-trivial bound on the probability that D is internally spanned if D has either height or width at most p^{-1} , whereas our critical droplets have height and width larger than $p^{-\alpha}$. The second main obstacle is that when bounding the probability of a ‘step’ of the hierarchy, we need to bound the probability of the existence of an internally spanned ‘saver’ droplet (cf. Section 6.3). However, if the seeds and steps both have size only p^{-1} then the number of possible hierarchies will be much too large for our application of the union bound.

Thus, in order to apply the method of hierarchies, we first need to give strong bounds on the probability that a seed or saver droplet of height and width roughly $p^{-\alpha}$ is internally spanned. We prove this in stages, building up from the result for much smaller droplets (which we prove using Lemma 8.13). More precisely, we use the following induction hypothesis:

Definition 8.14 For each $\beta_1, \beta_2 \in \mathbb{N}$, we say that $\text{IH}(\beta_1, \beta_2)$ holds if the following statement is true for some (small, fixed) constant $\eta > 0$:

There exists $\delta = \delta(\beta_1 + \beta_2) > 0$ such that

$$\mathbb{P}_p(I^\times(D)) \leq p^{\delta \max\{w(D), h(D)\}}$$

for every \mathcal{T} -droplet D such that

$$w(D) \leq p^{-\beta_1(1-2\eta)-\eta} \quad \text{and} \quad h(D) \leq p^{-\beta_2(1-2\eta)-\eta}. \quad (8.2)$$

The bounds we need on the probability that a seed or a saver droplet is internally spanned are given by the following lemma.

Lemma 8.15 *The assertions $\text{IH}(\alpha + 1, \alpha)$ and $\text{IH}(\alpha, \alpha + 1)$ both hold.*

We prove Lemma 8.15 by induction on $\beta_1 + \beta_2$. The base case, $\beta_1 = \beta_2 = 1$, follows easily from Lemma 8.13; indeed, the probability that D contains at least $k := \varepsilon \cdot \max\{h(D), w(D)\}$ elements of A is at most

$$\binom{O(k^2)}{k} p^k \leq O(pk)^k \leq p^{\eta k/2}$$

if $k \leq p^{-1+\eta}$, so $\text{IH}(1, 1)$ holds, as claimed.

The specific induction statements that were proved in [13] are:

$$\begin{aligned} \text{IH}(\beta, \beta) &\Rightarrow \text{IH}(\beta + 1, \beta) && \text{for all } 1 \leq \beta \leq \alpha; \\ \text{IH}(\beta, \beta) &\Rightarrow \text{IH}(\beta, \beta + 1) && \text{for all } 1 \leq \beta \leq \alpha; \\ (\text{IH}(\beta + 1, \beta) \wedge \text{IH}(\beta, \beta + 1)) &\Rightarrow \text{IH}(\beta + 1, \beta + 1) && \text{for all } 1 \leq \beta \leq \alpha - 1. \end{aligned}$$

The proof of these statements depends on whether, for the \mathcal{T} -droplet D in question, $h(D) \geq w(D)$ or vice-versa. This is because we have a pair $\{u, -u\}$ of opposite stable directions, so we may partition the droplet D into many smaller sub-droplets of the same width, and bound the probability that each is vertically crossed (possibly with help from above and below) independently, since these events depend on disjoint sets of infected sites. This turns out to be a good idea if $h(D) \geq w(D)$; otherwise, we need to use the spanning algorithm to construct an (s, t) -good hierarchy for an internally spanned \mathcal{T} -droplet D (cf. Section 2.1) with $s = t \approx p^{-\beta(1-2\eta)-\eta}$. Note that in this latter case, we can use the induction hypothesis to bound the probability that a seed is internally spanned.

In either case, it remains to bound the probability that it is possible to ‘cross’ a parallelogram of sites from one side to the other with ‘help’ from one of the sides in the form of an infected half-plane. We obtain bounds for the probabilities of crossings by showing that, to a certain level of precision, the most likely way these events could occur is via the droplet (or half-plane) advancing row-by-row, rather than via the merging of many smaller droplets. The hardest part of the proof is controlling vertical crossings in the case of models ‘with drift’ (that is, when one of u and $-u$ has infinite difficulty). As in Section 6.3, the trick is to rotate ones view slightly (by an angle of $p^{1-\eta}$), and show that each constant width (rotated) strip must either contain a reasonably large ‘cluster’ of elements of A , or must intersect a large internally spanned ‘saver’ droplet, and moreover all of these events occur disjointly. To prove this, we need introduce yet another algorithm for approximating the closure of a set of sites (the ‘iceberg algorithm’). Since the details are rather complicated, we refer the interested reader to [13, Sections 6.3 and 8.3].

9 A general conjecture in higher dimensions

Perhaps the most important open problem on monotone cellular automata is to extend Theorem 7.4 to higher dimensions. In this section we will state a conjecture as to the correct form of this generalization.

Fix an integer $d \geq 2$ and let \mathcal{U} be a d -dimensional update family. Let

$$\mathbb{H}_u^d := \{x \in \mathbb{Z}^d : \langle x, u \rangle < 0\}$$

denote the discrete d -dimensional half-space with normal $u \in S^{d-1}$, and define the set of stable directions to be

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

Let μ denote Lebesgue measure on S^{d-1} (this is usually called ‘spherical measure’), and let \mathcal{C}^d denote the collection of open hemispheres in S^{d-1} .

Definition 9.1 We say that a d -dimensional update family is:

- *supercritical* if there exists $C \in \mathcal{C}^d$ that is disjoint from \mathcal{S} ,
- *critical* if there exists $C \in \mathcal{C}^d$ such that $\mu(C \cap \mathcal{S}) = 0$, and every $C \in \mathcal{C}^d$ has non-empty intersection with \mathcal{S} ,
- *subcritical* if $\mu(C \cap \mathcal{S}) > 0$ for every $C \in \mathcal{C}^d$.

Note that, as in two dimensions, this trichotomy depends only on the stable set \mathcal{S} . The following conjecture was made by Bollobás, Dumitriu-Copin, Morris and Smith [13].

Conjecture 9.2 *Let \mathcal{U} be a d -dimensional update family.*

(a) *If \mathcal{U} is supercritical then $p_c(\mathbb{Z}_n^d, \mathcal{U}) = n^{-\Theta(1)}$.*

(b) *If \mathcal{U} is critical then there exists $r = r(\mathcal{U}) \in \{2, \dots, d\}$ such that*

$$p_c(\mathbb{Z}_n^d, \mathcal{U}) = \left(\frac{1}{\log_{(r-1)} n} \right)^{\Theta(1)}, \quad (9.1)$$

where $\log_{(r-1)}$ denotes an $(r-1)$ -times iterated logarithm.

(c) *If \mathcal{U} is subcritical then $\liminf_{n \rightarrow \infty} p_c(\mathbb{Z}_n^d, \mathcal{U}) > 0$.*

The intuition behind this conjecture is as follows. Each direction $u \in S^{d-1}$ should have a ‘difficulty’ $r(u) \in \{0, \dots, d\}$ that depends on the intersection of \mathcal{S} with its neighbourhood; this should correspond to the r in (9.1) when applied to the $(d-1)$ -dimensional process induced on the side of the half-space \mathbb{H}_u^d (if the direction is unstable then $r = 0$, and if the induced process is super-/subcritical then $r \in \{1, d\}$). Now, if there exists an open hemisphere in which all directions have difficulty at most $r-1$, then the difficulty of \mathcal{U} should be at most r (cf. the proof of (3.1)); if not, then it should be at least $r+1$ (cf. the proof of Lemma 3.7).

10 Some related models

We finish this survey by briefly discussing some applications (and potential future applications) of the results and techniques described above to three more complicated models: kinetically constrained spin models, a model of nucleation and growth, and the abelian sandpile.

10.1 Kinetically constrained spin models

Suppose we modify the update rule in \mathcal{U} -bootstrap percolation so that the process is no longer monotone, i.e., so that infected sites may become uninfected? For example, if we randomly (and independently) resample the state (infected or uninfected) of a vertex v whenever $v + X$ is entirely infected for some $X \in \mathcal{U}$, we obtain an important class of models of the ‘liquid-glass transition’ known as ‘kinetically constrained spin models’. The following general family of models was introduced by Cancrini, Martinelli, Roberto and Toninelli [18].

Definition 10.1 Let \mathcal{U} be a finite collection of finite subsets of $\mathbb{Z}^d \setminus \{\mathbf{0}\}$, and let $p \in (0, 1)$. The \mathcal{U} -kinetically constrained spin model on \mathbb{Z}^d with density p is defined as follows:

- (a) Each vertex has an independent exponential clock which rings randomly at rate 1.
- (b) If the clock at vertex v rings at (continuous) time $t \geq 0$, and $v + X \subseteq A_t$ for some $X \in \mathcal{U}$, where $A_t \subseteq \mathbb{Z}^d$ is the set of infected vertices at time t , then v becomes infected with probability p , and healthy with probability $1 - p$, independently of all other events.

Note that if the set of infected vertices at time $t = 0$ percolates in the \mathcal{U} -bootstrap process, then the set of infected vertices at all later times also has this property. The model can therefore be thought of as a biased random walk on the family of percolating sets. Let us choose the infected sites at time $t = 0$ to be p -random; the distribution is therefore also given by \mathbb{P}_p at every later time t , though the distributions at different times are (of course) not independent of one another.

Let us write

$$\tau(\mathbb{Z}^d, \mathcal{U}) := \inf \{t \geq 0 : \mathbf{0} \in A_t\},$$

for the (random) time at which the origin is first infected. It was pointed out in [18] that the time taken for the \mathcal{U} -bootstrap process to infect the origin provides a lower bound on $\tau(\mathbb{Z}^d, \mathcal{U})$. In particular, the following theorem is an immediate consequence of Theorem 8.1.

Theorem 10.2 *Let \mathcal{U} be a critical two-dimensional update family. There exists a constant $c = c(\mathcal{U}) > 0$ such that the following holds with high probability as $p \rightarrow 0$.*

- (a) *If \mathcal{U} is balanced, then*

$$\tau(\mathbb{Z}^2, \mathcal{U}) \geq \exp\left(cp^{-\alpha}\right).$$

(b) If \mathcal{U} is unbalanced, then

$$\tau(\mathbb{Z}^2, \mathcal{U}) \geq \exp\left(cp^{-\alpha}(\log(1/p))^2\right).$$

Cancrini, Martinelli, Roberto and Toninelli [18] introduced a powerful and general analytic technique that allows one to prove upper bounds on $\tau(\mathbb{Z}^d, \mathcal{U})$ for many specific update families. There is some hope (see [44, 47]) that their ideas, combined with the techniques introduced in [13, 17], might be sufficient to prove an almost-matching upper bound for a large collection of critical two-dimensional update families.

10.2 Nucleation and growth

We next discuss a model that was first studied by Dehghanpour and Schonmann [24, 25], who used it to study the metastable behavior of the kinetic Ising model on \mathbb{Z}^2 with a small magnetic field and vanishing temperature. (Their results were recently generalized to higher dimensions by Cerf and Manzo [21, 22].) The following general (monotone) version of this model was recently introduced in [16].

Definition 10.3 Let $\mathcal{U}_1, \dots, \mathcal{U}_r$ be d -dimensional update families, and let $1 \leq k_1(n) \leq \dots \leq k_r(n) = n$ be functions. In the $(\mathcal{U}_1, \dots, \mathcal{U}_r)$ -*nucleation and growth model*, each vertex $v \in \mathbb{Z}^d$ is initially uninfected, but becomes (permanently) infected at rate $k_\ell(n)/n$ at each time $t \geq 0$, where

$$\ell = \ell(v, t) := \max \{j \in [r] : v + X \subseteq A_t \text{ for some } X \in \mathcal{U}_j\},$$

where A_t is the set of infected sites at time t , and at rate $1/n$ otherwise.

Similarly to the previous subsection, let us write

$$\tau(\mathbb{Z}^d; \mathcal{U}_1, \dots, \mathcal{U}_r) := \inf \{t \geq 0 : \mathbf{0} \in A_t\},$$

for the (random) time at which the origin is first infected.

Dehghanpour and Schonmann [24] studied the case $d = r = 2$, with $\mathcal{U}_1 = \mathcal{N}_1^2$ and $\mathcal{U}_2 = \mathcal{N}_2^2$ (recall that \mathcal{N}_r^d denotes the family of r -subsets of the $2d$ nearest neighbours of the origin in \mathbb{Z}^d), proving that

$$\tau(\mathbb{Z}^2; \mathcal{N}_1^2, \mathcal{N}_2^2) = \begin{cases} \left(\frac{n}{k}\right)^{1+o(1)} & \text{if } k \leq \sqrt{n} \\ \left(\frac{n^2}{k}\right)^{1/3+o(1)} & \text{if } k \geq \sqrt{n} \end{cases} \quad (10.1)$$

with high probability as $n \rightarrow \infty$. The following more precise bounds were proved recently by Bollobás, Griffiths, Morris, Rolla and Smith [16].

Theorem 10.4 *With high probability as $n \rightarrow \infty$, the following hold.*

(a) *If $k \ll \log n$ then*

$$\tau(\mathbb{Z}^2; \mathcal{N}_1^2, \mathcal{N}_2^2) = \left(\frac{\pi^2}{18} + o(1) \right) \frac{n}{\log n}.$$

(b) *If $\log n \ll k \ll \sqrt{n}(\log n)^2$ then*

$$\tau(\mathbb{Z}^2; \mathcal{N}_1^2, \mathcal{N}_2^2) = \left(\frac{1}{4} + o(1) \right) \frac{n}{k} \log \left(\frac{k}{\log n} \right).$$

(c) *If $\sqrt{n}(\log n)^2 \ll k \ll n$ then*

$$\tau(\mathbb{Z}^2; \mathcal{N}_1^2, \mathcal{N}_2^2) = \Theta \left(\frac{n^2}{k \log(n/k)} \right)^{1/3}.$$

It would be very interesting to generalize these results to other update families. In particular, we expect the following problem to already be very challenging.

Open problem: Determine $\tau(\mathbb{Z}^2; \mathcal{U}_1, \mathcal{U}_2)$ up to a constant factor whenever \mathcal{U}_1 is supercritical and \mathcal{U}_2 is critical.

Cerf and Manzo [21] determined $\tau(\mathbb{Z}^d; \mathcal{N}_1^d, \dots, \mathcal{N}_d^d)$ (with high probability) up to a factor of $1 + o(1)$ in the exponent, cf. (10.1). Generalizing their result to arbitrary collections of d -dimensional update families is an important (and likely extremely difficult) problem.

10.3 The abelian sandpile

The final model we would like to briefly discuss was first introduced almost 30 years ago by Bak, Tang and Wiesenfeld [3] as an example of so-called ‘self-organised criticality’. In the model, grains of sand are placed on each vertex of a graph G ; if there are at least $d(v)$ particles on vertex v , then the pile ‘topples’ by sending one grain to each of its neighbours. The process is ‘abelian’ in the sense that the order of toppling does not affect the final configuration. We are interested in this model on the graph \mathbb{Z}^d , and with a random initial configuration.

Definition 10.5 Given a function $A: \mathbb{Z}^d \rightarrow \mathbb{Z}$, the *abelian sandpile* on \mathbb{Z}^d with initial state A is defined as follows:

(a) At time $t = 0$, place $A(v)$ grains of sand on v for each $v \in \mathbb{Z}^d$.

- (b) For each $t \in \mathbb{N}$, do the following at time t : for each vertex $v \in \mathbb{Z}^d$ with at least $2d$ grains of sand at time $t - 1$, remove $2d$ grains from v and place one at each nearest neighbour of v .

Let us say that A *percolates on* \mathbb{Z}^d if every site topples infinitely many times in the abelian sandpile on \mathbb{Z}^d starting from A , and define

$$\lambda_c^S(\mathbb{Z}^d) := \inf \left\{ \lambda : \mathbb{P}_\lambda(A \text{ percolates on } \mathbb{Z}^d) = 1 \right\},$$

where \mathbb{P}_λ denotes the probability measure in which each $A(v)$ is chosen according to an independent Poisson random variable of mean λ .

The problem of determining $\lambda_c^S(\mathbb{Z}^d)$ was introduced over 15 years ago by Dickman, Muñoz, Vespagnani and Zapperi [55], but the best known bounds (see e.g. [33]) are only

$$d \leq \lambda_c^S(\mathbb{Z}^d) \leq 2d - 1.$$

Note that there exists an initial distribution with density $2d - 1$ which does not percolate (take $A(v)$ to be constant), and distributions with density arbitrarily close to d which do percolate (take $A(v) = 2d$ with probability ε , and $A(v) = d$ otherwise, and apply Theorem 3.1). Thus, to improve the bounds above, we must use some properties of the Poisson distribution other than its mean.

Morris [47] proposed the following approach to this problem in high dimensions. Let us label a vertex v as ‘infected’ if $A(v) \geq 2d$, as ‘vulnerable’ if $d < A(v) < 2d$, and as ‘removed’ if $A(v) \leq d$. Note that infected vertices topple immediately, and that vulnerable vertices with at least $d - 1$ neighbours that topple at some point will also topple. The vertices that eventually topple therefore contain the closure of the infected set of vertices in $(d - 1)$ -neighbour bootstrap percolation on the subgraph of \mathbb{Z}^d induced by the non-removed vertices. Note that if d is large and $\lambda = (1 + \varepsilon)d$, then only a few vertices will be removed, and even fewer vertices will be infected.

The coupling above motivates the following problem, which was first studied by Gravner and McDonald [38]. Let us write $\mathbb{Z}^d(q)$ for the random induced subgraph of \mathbb{Z}^d obtained by removing each vertex independently with probability q , and say that a set $A \subseteq \mathbb{Z}^d(q)$ *percolates in the r -neighbour polluted bootstrap process* if the closure of A in the r -neighbour bootstrap process on $\mathbb{Z}^d(q)$ contains an infinite component. Define

$$p_c^\infty(\mathbb{Z}^d(q), r) := \inf \left\{ p : \mathbb{P}_p(A \text{ percolates in the } r\text{-neighbour polluted bootstrap process}) \geq 1/2 \right\}$$

for each $d \geq r \geq 2$. It was proved by Gravner and McDonald [38] that

$$p_c^\infty(\mathbb{Z}^2(q), 2) = \Theta(\sqrt{q})$$

almost surely as $q \rightarrow 0$, and a Cerf–Cirillo-type argument (see Section 3) can be used to show that moreover $p_c^\infty(\mathbb{Z}^d(q), d) > 0$ for every $d \in \mathbb{N}$ and $q > 0$. On the other hand, the following conjecture was stated in [47].

Conjecture 10.6 *For each $d > r \geq 1$, there exists $q_0(d, r) > 0$ such that*

$$p_c^\infty(\mathbb{Z}^d(q), r) = 0$$

almost surely for every $0 < q < q_0(d, r)$.

Note that when $r = 1$ we can take $q_0(d, 1) = 1 - p_c^{\text{site}}(\mathbb{Z}^d)$, but the conjecture is open (and seems to be very difficult) for every $d > r \geq 2$. Conjecture 10.6 motivates the following conjecture, also stated in [47].

Conjecture 10.7

$$\frac{\lambda_c^S(\mathbb{Z}^d)}{d} \rightarrow 1$$

as $d \rightarrow \infty$.

The extension of these conjectures to more general update rules is yet another fascinating, and likely very hard, open problem.

Acknowledgements

The author would like to thank Béla Bollobás for introducing him to bootstrap percolation, and for his encouragement and support over many years. He would also like to thank his frequent collaborators, Józsi Balogh, Béla Bollobás, Hugo Duminil-Copin and Paul Smith, for many interesting conversations over the years, and for their many important contributions to the area, only a few of which we have had space to describe; Paul Smith for Figures 1 and 2, and for several helpful discussions during the preparation of this survey; and Teeradej Kittipassorn and the anonymous referee for a number of helpful comments on the manuscript.

References

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

- [2] R. Arratia, Site recurrence for annihilating random walks on \mathbb{Z}^d , *Ann. Prob.*, **11** (1983), 706–713.
- [3] P. Bak, C. Tang and K. Wiesenfeld, Self-organized criticality: an explanation of $1/f$ noise, *Phys. Rev. Letters*, **59** (1987), 381–384.
- [4] P. Balister, B. Bollobás, M.J. Przykucki and P.J. Smith, Subcritical \mathcal{U} -bootstrap percolation models have non-trivial phase transitions, *Trans. Amer. Math. Soc.*, **368** (2016), 7385–7411.
- [5] J. Balogh and B. Bollobás, Bootstrap percolation on the hypercube, *Prob. Theory Rel. Fields*, **134** (2006), 624–648.
- [6] J. Balogh, B. Bollobás, H. Duminil-Copin and R. Morris, The sharp threshold for bootstrap percolation in all dimensions, *Trans. Amer. Math. Soc.*, **364** (2012), 2667–2701.
- [7] J. Balogh, B. Bollobás and R. Morris, Majority bootstrap percolation on the hypercube, *Combin. Probab. Computing*, **18** (2009), 17–51.
- [8] J. Balogh, B. Bollobás and R. Morris, Bootstrap percolation in three dimensions, *Ann. Prob.*, **37** (2009), 1329–1380.
- [9] J. Balogh, B. Bollobás and R. Morris, Bootstrap percolation in high dimensions, *Combin. Probab. Computing*, **19** (2010), 643–692.
- [10] J. Balogh, Y. Peres and G. Pete, Bootstrap percolation on infinite trees and non-amenable groups, *Combin. Prob. Computing*, **15** (2006), 715–730.
- [11] J. Balogh and B. Pittel, Bootstrap percolation on the random regular graph, *Random Structures Algorithms*, **30** (2007), 257–286.
- [12] J. van den Berg and H. Kesten, Inequalities with applications to percolation and reliability, *J. Appl. Prob.*, **22** (1985), 556–569.
- [13] B. Bollobás, H. Duminil-Copin, R. Morris and P. Smith, Universality of two-dimensional critical cellular automata, *Proc. London Math. Soc.*, to appear, arXiv:1406.6680.
- [14] B. Bollobás, H. Duminil-Copin, R. Morris and P. Smith, The sharp threshold for the Duarte model, *Ann. Prob.*, to appear.
- [15] B. Bollobás, K. Gunderson, C. Holmgren, S. Janson and M. Przykucki, Bootstrap percolation on Galton-Watson trees, *Electron. J. Prob.*, **19** (2014), 1–27.

- [16] B. Bollobás, S. Griffiths, R. Morris, L. Rolla and P. Smith, Nucleation and growth in two dimensions, submitted, arXiv:1508.06267.
- [17] B. Bollobás, P.J. Smith and A.J. Uzzell, Monotone cellular automata in a random environment, *Combin. Probab. Computing*, **24** (2015), 687–722.
- [18] N. Cancrini, F. Martinelli, C. Roberto and C. Toninelli, Kinetically constrained spin models, *Prob. Theory Rel. Fields*, **140** (2008), 459–504.
- [19] R. Cerf and E.N.M. Cirillo, Finite size scaling in three-dimensional bootstrap percolation, *Ann. Prob.*, **27** (1999), 1837–1850.
- [20] R. Cerf and F. Manzo, The threshold regime of finite volume bootstrap percolation, *Stochastic Proc. Appl.*, **101** (2002), 69–82.
- [21] R. Cerf and F. Manzo, A d -dimensional nucleation and growth model, *Prob. Theory Rel. Fields*, **155** (2013), 427–449.
- [22] R. Cerf and F. Manzo, Nucleation and growth for the Ising model in d dimensions at very low temperatures, *Ann. Prob.*, **41** (2013), 3697–3785.
- [23] J. Chalupa, P.L. Leath and G.R. Reich, Bootstrap percolation on a Bethe lattice, *J. Phys. C.*, **12** (1979), L31–L35.
- [24] P. Dehghanpour and R.H. Schonmann, A nucleation-and-growth model, *Prob. Theory Rel. Fields*, **107** (1997), 123–135.
- [25] P. Dehghanpour and R.H. Schonmann, Metropolis dynamics relaxation via nucleation and growth, *Comm. Math. Phys.*, **188** (1997), 89–119.
- [26] A.M.S. Duarte, Simulation of a cellular automaton with an oriented bootstrap rule, *Phys. A*, **157** (1989), 1075–1079.
- [27] H. Duminil-Copin and A. Holroyd, Sharp metastability for threshold growth models, manuscript, available at <http://www.unige.ch/~duminil>.
- [28] H. Duminil-Copin and A.C.D. van Enter, Finite volume Bootstrap Percolation with balanced threshold rules on \mathbb{Z}^2 , *Ann. Prob.*, **41** (2013), 1218–1242.
- [29] A.C.D. van Enter, Proof of Straley’s argument for bootstrap percolation, *J. Stat. Phys.*, **48** (1987), 943–945.

- [30] A.C.D. van Enter and A. Fey, Metastability threshold for anisotropic bootstrap percolation in three dimensions, *J. Stat. Phys.*, **147** (2012), 97–112.
- [31] A.C.D. van Enter and W.J.T. Hulshof, Finite-size effects for anisotropic bootstrap percolation: logarithmic corrections, *J. Stat. Phys.*, **128** (2007), 1383–1389.
- [32] P. Erdős and P. Ney, Some Problems on Random Intervals and Annihilating Particles, *Ann. Prob.*, **2** (1974), 828–839.
- [33] A. Fey, R. Meester and F. Redig, Stabilizability and percolation in the infinite volume sandpile model, *Ann. Prob.*, **37** (2009), 654–675.
- [34] L.R. Fontes, R.H. Schonmann and V. Sidoravicius, Stretched Exponential Fixation in Stochastic Ising Models at Zero Temperature, *Commun. Math. Phys.*, **228** (2002), 495–518.
- [35] J. Gravner and A.E. Holroyd, Slow convergence in bootstrap percolation, *Ann. Appl. Prob.*, **18** (2008), 909–928.
- [36] J. Gravner, A.E. Holroyd and R. Morris, A sharper threshold for bootstrap percolation in two dimensions, *Prob. Theory Rel. Fields*, **153** (2012), 1–23.
- [37] J. Gravner and D. Griffeath, Scaling laws for a class of critical cellular automaton growth rules, In: Random Walks (Budapest, 1998), *Bolyai Soc. Math. Stud.*, **9** (1999), 167–186.
- [38] J. Gravner and E. McDonald, Bootstrap percolation in a polluted environment, *J. Stat. Phys.* **87** (1997), 915–927.
- [39] A. Holroyd, Sharp Metastability Threshold for Two-Dimensional Bootstrap Percolation, *Prob. Theory Rel. Fields*, **125** (2003), 195–224.
- [40] A. Holroyd, The metastability threshold for modified bootstrap percolation in d dimensions, *Electron. J. Prob.*, **11** (2006), 418–433.
- [41] A.E. Holroyd, T.M. Liggett and D. Romik, Integrals, partitions, and cellular automata, *Trans. Amer. Math. Soc.*, **356** (2004), 3349–3368.
- [42] S. Janson, T. Łuczak, T. Turova and T. Vallier, Bootstrap percolation on the random graph $G(n, p)$, *Ann. Appl. Prob.*, **22** (2012), 1989–2047.

- [43] F. Martinelli, Lectures on Glauber dynamics for discrete spin models, Lectures on Probability Theory and Statistics, Springer Lecture Notes in Mathematics, **1717** (1998), 93–191.
- [44] F. Martinelli and C. Toninelli, Towards a universality picture for the relaxation to equilibrium of kinetically constrained models, arXiv:1701.00107.
- [45] R. Morris, Zero-temperature Glauber dynamics on \mathbb{Z}^d , *Prob. Theory Rel. Fields*, **149** (2011), 417–434.
- [46] R. Morris, The second term for bootstrap percolation in two dimensions, manuscript, available at <http://w3.impa.br/~rob/>.
- [47] R. Morris, Bootstrap percolation and other automata, *European J. Combin.*, to appear.
- [48] N. Morrison and J.A. Noel, Extremal Bounds for Bootstrap Percolation in the Hypercube, arXiv:1506.04686.
- [49] T.S. Mountford, Critical length for semi-oriented bootstrap percolation, *Stochastic Process. Appl.*, **56** (1995), 185–205.
- [50] S. Nanda, C.M. Newman and D. Stein, Dynamics of Ising spin systems at zero temperature, In *On Dobrushin's way (From Probability Theory to Statistical Mechanics)*, eds. R. Minlos, S. Shlosman and Y. Suhov, *Am. Math. Soc. Transl.*, **198** (2000), 183–194.
- [51] J. von Neumann, Theory of Self-Reproducing Automata. Univ. Illinois Press, Urbana, 1966.
- [52] C.M. Newman and D. Stein, Zero-temperature dynamics of Ising spin systems following a deep quench: results and open problems, *Physica A*, **279** (2000), 159–168.
- [53] D. Reimer, Proof of the van den Berg–Kesten Conjecture, *Combin. Prob. Computing*, **9** (2000), 27–32.
- [54] R.H. Schonmann, On the behaviour of some cellular automata related to bootstrap percolation, *Ann. Prob.*, **20** (1992), 174–193.
- [55] A. Vespagnani, R. Dickman, M. Muñoz and S. Zapperi, Absorbing-state phase transitions in fixed-energy sandpiles, *Phys. Rev. E*, **62** (2000), 45–64.

IMPA
Estrada Dona Castorina, 110
Jardim Botânico
Rio de Janeiro, Brazil
rob@impa.br

