
Monotone Cellular Automata in a Random Environment

BÉLA BOLLOBÁS¹, PAUL SMITH² and ANDREW UZZELL³

¹Department of Pure Mathematics and Mathematical Statistics, Wilberforce Road, Cambridge, CB3 0WA, UK
Department of Mathematical Sciences, University of Memphis, Memphis, TN 38152, USA
London Institute for Mathematical Sciences, 35a South Street, London, W1K 2XF, UK
(e-mail: b.bollobas@pmms.cam.ac.uk)

²IMPA, 110 Estrada Dona Castorina, Jardim Botânico, Rio de Janeiro, 22460-320, Brazil
(e-mail: psmith@impa.br)

³Department of Mathematics, Uppsala University, PO Box 480, SE-751 06 Uppsala, Sweden
(e-mail: andrew.uzzell@math.uu.se)

Received 19 September 2013; revised 9 December 2014; first published online 2 February 2015

In this paper we study in complete generality the family of two-state, deterministic, monotone, local, homogeneous cellular automata in \mathbb{Z}^d with random initial configurations. Formally, we are given a set $\mathcal{U} = \{X_1, \dots, X_m\}$ of finite subsets of $\mathbb{Z}^d \setminus \{\mathbf{0}\}$, and an initial set $A_0 \subset \mathbb{Z}^d$ of ‘infected’ sites, which we take to be random according to the product measure with density p . At time $t \in \mathbb{N}$, the set of infected sites A_t is the union of A_{t-1} and the set of all $x \in \mathbb{Z}^d$ such that $x + X \in A_{t-1}$ for some $X \in \mathcal{U}$. Our model may alternatively be thought of as bootstrap percolation on \mathbb{Z}^d with arbitrary update rules, and for this reason we call it \mathcal{U} -bootstrap percolation.

In two dimensions, we give a classification of \mathcal{U} -bootstrap percolation models into three classes – supercritical, critical and subcritical – and we prove results about the phase transitions of all models belonging to the first two of these classes. More precisely, we show that the critical probability for percolation on $(\mathbb{Z}/n\mathbb{Z})^2$ is $(\log n)^{-\Theta(1)}$ for all models in the critical class, and that it is $n^{-\Theta(1)}$ for all models in the supercritical class.

The results in this paper are the first of any kind on bootstrap percolation considered in this level of generality, and in particular they are the first that make no assumptions of symmetry. It is the hope of the authors that this work will initiate a new, unified theory of bootstrap percolation on \mathbb{Z}^d .

2010 *Mathematics subject classification*: Primary 60K35
Secondary 60C05

1. Introduction

Cellular automata are dynamical systems on graphs in which vertices have one of a finite number of states at discrete times t . The states of the vertices change according to an

update rule that is homogeneous (the same rule applies to every vertex) and local (the rule depends on only a finite neighbourhood of the vertex under consideration). These systems were first introduced by von Neumann (see [22], for example) at the suggestion of Ulam [19]. The subject of this paper is a substantial extension of, and has its roots in, *bootstrap percolation*, a particular type of cellular automaton proposed in 1979 by Chalupa, Leath and Reich [11].

In the simplest and most widely studied model of bootstrap percolation, called *r-neighbour bootstrap percolation*, we are given a graph G and a subset $A \subset V(G)$ of the vertices of G . The elements of A are initially *infected*, and at each time t , infected vertices stay infected, and among the uninfected vertices of G , those that have at least r infected neighbours become infected. In keeping with the percolation literature, we use the words ‘vertex’ and ‘site’ interchangeably. The set of all vertices in G that are eventually infected is called the *closure* of A and denoted $[A]$. If A is such that $[A] = V(G)$, then we say that A *percolates* G , or simply that there is *percolation*.

The key problem is to analyse the behaviour of the r -neighbour model acting on an initial set A chosen randomly according to a product measure with density p . In particular, one would like to know how large p should be for there to be percolation with high probability (for a sequence of graphs with size of vertex set tending to infinity) or almost surely (for a single infinite graph). The results in relation to this phase transition are usually framed in terms of the *critical probability*, defined as

$$p_c(G, r) := \inf\{p : \mathbb{P}_p([A] = V(G)) \geq 1/2\}. \quad (1.1)$$

A considerable amount is now known about $p_c(G, r)$ in the case of *lattice graphs*: a graph G is a (d -dimensional) lattice graph if it is (isomorphic to) a (not necessarily planar) translation-invariant locally finite graph with vertex set \mathbb{Z}^d . (Equivalently, G is a lattice graph if (there is an isomorphism between $V(G)$ and \mathbb{Z}^d under which) there exists a finite symmetric set $X \subset \mathbb{Z}^d \setminus \{\mathbf{0}\}$ such that the neighbourhood of any vertex x is the set $x + X$.) The most natural lattice graph, and the one that has attracted the greatest amount of study, is the nearest neighbour graph on \mathbb{Z}^d . Indeed, the first rigorous result in the field of bootstrap percolation, which was proved by van Enter [20], was that $p_c(\mathbb{Z}^2, 2) = 0$. Schonmann [18] later showed that $p_c(\mathbb{Z}^d, r)$ is equal to 0 if $r \leq d$ and 1 otherwise. Aizenman and Lebowitz [1] were the first to recognize that there is much more to say about the model in the context of finite subgraphs of lattice graphs, and they proved that $p_c([n]^d, 2) = \Theta(1)/(\log n)^{d-1}$. Holroyd [15] went considerably further in the case $d = r = 2$, proving that $p_c([n]^2, 2) = (1 + o(1))\pi^2/18 \log n$. Vast generalizations of these results on $[n]^d$ for general d and r were obtained in the weak sense by Cerf and Cirillo [9] and Cerf and Manzo [10], and in the sharp sense by Balogh, Bollobás and Morris [5] and Balogh, Bollobás, Duminił-Copin and Morris [4].

Bootstrap percolation has been studied on a number of other lattice graphs, not only on \mathbb{Z}^d and $[n]^d$; for a small selection of these results, see [3, 6, 7, 8, 13, 14, 17]. In particular, Gravner and Griffeath [13, 14] studied r -neighbour bootstrap percolation on general two-dimensional lattice graphs. (Gravner and Griffeath called their model ‘threshold dynamics’, but it is easily seen to be equivalent to r -neighbour bootstrap

percolation on two-dimensional lattices.) Unfortunately, not all of the results in [13, 14] were rigorously proved, and in some cases they were not even correct (see, for example, [21]).

The first departure from the study of bootstrap percolation on lattice graphs came from Duarte [12], who introduced a two-dimensional directed model (in fact, the introduction by Duarte of this model predates the introduction by Gravner and Griffeath of r -neighbour models on general lattices). He again took the nearest neighbour graph on \mathbb{Z}^2 with threshold $r = 2$, but he took horizontal edges to be directed rightwards, and vertical edges to be directed in both directions (alternatively one could think of the vertical edges as being undirected, or replaced by two edges, one directed each way). Under this model, an uninfected vertex x becomes infected if at least two of its in-neighbours are infected; that is, if we take X to be the set $\{(-1, 0), (0, 1), (0, -1)\}$, then a site x becomes infected if $x + X$ contains at least two infected vertices. It is easy to see that this model causes a drift in the growth of the infected set: new sites only become infected to the right of existing sites. The critical probability for the Duarte model was determined up to a constant factor by Mountford [16]; he showed that $p_c(\vec{\mathbb{Z}}_n^2, 2) = \Theta((\log \log n)^2 / \log n)$, where $\vec{\mathbb{Z}}_n^2$ is the nearest neighbour graph on the torus \mathbb{Z}_n^2 with horizontal edges directed rightwards and vertical edges undirected. The correct constant in place of the ‘ $\Theta(\cdot)$ ’, or even the existence of a constant, is still not known.

However, even general r -neighbour models on directed lattice graphs, without any additional assumptions, are far less general than the models we introduce and analyse in the present paper. Our aims are threefold:

- to define a new family of percolation models, which generalize all previously studied models of bootstrap percolation on \mathbb{Z}^d , and to derive a number of basic properties of these models in two dimensions;
- to divide the two-dimensional models into three classes according to the nature of their large-scale behaviour, and to prove general results about the critical probabilities of models in two of these classes;
- to pose a number of open questions about these new general models.

We hope through achieving these aims to initiate a new, unified understanding of all bootstrap percolation models on the square lattice.

We go considerably further than all previous authors in the following specific sense: rather than taking a single special model or a special class of models, we take all local, homogeneous, monotone models. This means that we fix an s and take all neighbourhoods of x that consist only of sites at distance at most s from x ; for each of these we say whether or not x becomes infected given that all the sites in the neighbourhood are infected, subject to the condition that the neighbourhoods that infect x form an up-set (this is part of what we mean by monotone; we also mean that infected vertices never become uninfected). The same rule is then applied to every site x . The formal definition is as follows.

Definition 1.1. Fix an integer $d \geq 2$. Let $m \in \mathbb{N}$, and for each $i \in [m]$ let $X_i \subset \mathbb{Z}^d \setminus \{\mathbf{0}\}$ be finite and non-empty. Let $\mathcal{U} = \{X_1, \dots, X_m\}$. We call \mathcal{U} an *update family* and each X_i an




Figure 1. An example of an update family \mathcal{U} consisting of six rules. If there is a translation of one of the rules such that all of the sites marked with dots are infected, then the site marked with a cross will become infected too. It is easy to verify that \mathcal{U} is critical. Our methods can be used to show that $p_c(\mathbb{Z}_n^2, \mathcal{U}) = \Theta(1/\log n)$, and therefore that \mathcal{U} has the same large-scale behaviour as two-neighbour bootstrap percolation.

update rule. Let $A \subset \mathbb{Z}^d$. In \mathcal{U} -bootstrap percolation with update family \mathcal{U} and initial set $A = A_0$, for each $t \geq 0$ we set

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

Thus, there is a set A of initially infected sites and at each time step, whenever a translation of a configuration X_i by an element of \mathbb{Z}^d is completely infected, the corresponding translation of the origin also becomes infected. The closure of A under \mathcal{U} is defined to be $[A] := \cup_{t=0}^\infty A_t$. When the closure of A is the whole of \mathbb{Z}^d , we say that A percolates \mathbb{Z}^d under \mathcal{U} , or more usually, when d, \mathcal{U} and A are understood from the context, we say simply that there is percolation.

The r -neighbour model is easily seen to be an example of \mathcal{U} -bootstrap percolation: one takes the update rules X_1, \dots, X_m to be the $m := \binom{|\mathcal{N}|}{r}$ subsets of size r of the graph (in-)neighbourhood \mathcal{N} of the origin. However, these are only a very special subclass of our models, and in general, our models look nothing like r -neighbour models. The family in Figure 1 is an example of the kind of general model we have in mind.

For the remainder of the paper, except for a discussion of open problems in Section 8, we shall restrict our attention to the case $d = 2$. Gravner and Griffeath [13, 14], in their study of two-dimensional r -neighbour bootstrap percolation, took the underlying lattice graphs to be undirected. The result of this was that the bootstrap processes had a strong symmetry property, namely that $X \in \mathcal{U}$ if and only if $-X \in \mathcal{U}$. While this may seem natural, quite the opposite is true: symmetry vastly restricts the range and behaviour of the possible models. We have already seen a glimpse of this with the result of Mountford on the Duarte model, which shows that critical probabilities are not always negative powers of $\log n$. In Section 2 we explain in greater detail why symmetry is such a powerful property for a bootstrap process to have; here we mention briefly that the main reason is that it allows one to constrain growth to one-dimensional strips (and in higher dimensions, to lower-dimensional strips). Gravner and Griffeath imposed an additional restriction on their models, which had the effect of simplifying the analysis even further. The restriction was that the process should be ‘balanced’, meaning in their case roughly

that the four hardest directions in which to grow are equally hard. The reason why this constraint is useful is that, together with symmetry, it implies the existence of bounding parallelograms, all four of whose sides are equally difficult from which to grow, and which can be controlled one-dimensionally in two independent directions.

We emphasize that we impose *none* of these additional constraints. In particular, most importantly, we do not impose any symmetry constraints on our models, and in general our models will not exhibit any symmetry at all. This is one of the key strengths of our results, and one of the key advances over results by other authors.

Let Λ denote either \mathbb{Z}^2 or the discrete torus $\mathbb{Z}_n^2 := (\mathbb{Z}/n\mathbb{Z})^2$. Like most previous authors, we are interested in the random setting, in which sites of the lattice Λ are included in the initial set A independently with probability $p \in (0, 1)$. We describe a set A chosen in this way as p -random, and we write \mathbb{P}_p for the product probability measure. The \mathcal{U} -bootstrap percolation analogue of the critical probability for r -neighbour models defined in (1.1) is as follows:

$$p_c(\Lambda, \mathcal{U}) := \inf\{p : \mathbb{P}_p(A \text{ percolates in } \mathcal{U}\text{-bootstrap percolation on } \Lambda) \geq 1/2\}, \quad (1.2)$$

where A is a p -random subset of Λ .

1.1. The classification of \mathcal{U} -bootstrap percolation processes in two dimensions

For each $u \in S^1$, let $\mathbb{H}_u := \{x \in \mathbb{Z}^2 : \langle x, u \rangle < 0\}$ be the discrete half-plane with boundary perpendicular to u (we use $\langle \cdot, \cdot \rangle$ to denote the usual Euclidean inner product). The following simple definition is one of the most important of the paper; it will form the basis of our tripartition of update families in Definition 1.3.

Definition 1.2. A unit vector $u \in S^1$ is a *stable direction* for \mathcal{U} if $[\mathbb{H}_u] = \mathbb{H}_u$. We denote by $\mathcal{S} = \mathcal{S}(\mathcal{U})$ the set of all stable directions for \mathcal{U} and we call \mathcal{S} the *stable set* for \mathcal{U} . If $u \in \mathcal{S}^c = S^1 \setminus \mathcal{S}$, then we call u an *unstable direction* for \mathcal{U} . A stable direction u is said to be *strongly stable* for \mathcal{U} if it is contained in an open interval of stable directions in S^1 .

One of the key properties of the \mathcal{U} -bootstrap process that we establish in this paper is that the approximate large-scale behaviour of the process is governed primarily by the geometry of the stable set \mathcal{S} alone. This assertion is encapsulated by the following definition, which is our classification of \mathcal{U} -bootstrap percolation models.

Definition 1.3. Let \mathcal{U} be a (two-dimensional) update family with stable set \mathcal{S} . We say that \mathcal{U} is

- *supercritical* if there exists an open semicircle $C \subset S^1$ such that $\mathcal{S} \cap C = \emptyset$;
- *critical* if $\mathcal{S} \cap C \neq \emptyset$ for every open semicircle $C \subset S^1$ and there exists an open semicircle $C \subset S^1$ not containing any strongly stable directions;
- *subcritical* if every open semicircle $C \subset S^1$ contains a strongly stable direction.

The most important of the three classes of \mathcal{U} -bootstrap models is the critical class, which includes all previously studied models of bootstrap percolation on \mathbb{Z}^2 .

A similar classification in the special case of r -neighbour models is given by Gravner and Griffeath [13, 14]. They define a process to be supercritical if there exist finite subsets of \mathbb{Z}^2 with infinite closures, to be subcritical if there exist closed cofinite proper subsets of \mathbb{Z}^2 , and to be critical otherwise. Our definitions coincide precisely with their definitions in the case of the small subclass of our models that they consider. Moreover, the distinction between supercritical and critical is the same in all cases, although not obviously so; in fact, this statement is the content of Theorem 7.1. The distinction between critical and subcritical in general is not the same, however: one can show that the Gravner–Griffeath definition is equivalent to the statement that every direction is stable, while we only require every open semicircle to contain a strongly stable direction. Theorem 1.10 implies that there is an infinite collection of families \mathcal{U} that fall under our definition of subcritical but not under the Gravner–Griffeath definition.

Among the results of this paper, which we state formally next, we prove that the distinction made in Definition 1.3 between supercritical and critical families is the right one, in the specific sense that the critical probabilities of supercritical families are polynomial in t and the critical probabilities of critical families are polylogarithmic in t . We are not able to prove that the distinction between critical and subcritical families given here is also correct, although we conjecture that this is the case, in the specific sense that the critical probabilities of subcritical models should be bounded away from zero¹ (see Conjecture 8.1). We defer further discussion until Section 8.

1.2. Results

As stated earlier in the Introduction, our aims in this paper are threefold: to define and prove basic results about \mathcal{U} -bootstrap percolation; to prove the first results about the critical probabilities of critical and supercritical models; and to pose a number of open questions about these new models.

Our main results concern the second of these aims. We state these results first in terms of the infection time of the origin (Theorems 1.4 and 1.5), and then in their more usual (and essentially equivalent) form in terms of critical probabilities (Theorem 1.7). Given a set $A \subset \mathbb{Z}^2$, define the stopping time

$$\tau = \tau(A, \mathcal{U}) := \min\{t \geq 0 : \mathbf{0} \in A_t\}.$$

Thus, τ is the time at which the origin becomes infected in \mathcal{U} -bootstrap percolation with initial set A . In the statements of all theorems (and throughout the paper), constants, implicit or otherwise, are quantities that depend only on \mathcal{U} .

The following is our main theorem.

Theorem 1.4. *Let \mathcal{U} be a critical two-dimensional bootstrap percolation update family. Then, with high probability as $p \rightarrow 0$,*

$$\tau = \exp(p^{-\Theta(1)}).$$

¹ Since the submission of this paper, Balister, Bollobás, Przykucki and Smith [2] have proved that this conjecture is true.

We also prove the following theorem for supercritical families.

Theorem 1.5. *Let \mathcal{U} be a supercritical two-dimensional bootstrap percolation update family. Then, with high probability as $p \rightarrow 0$,*

$$\tau = p^{-\Theta(1)}.$$

We remark that the lower bound of Theorem 1.5 is a triviality: there must be at least one initially infected site within distance $O(\tau)$ of the origin. The content of the theorem is the upper bound.

As a corollary of Theorems 1.4 and 1.5, we have that for any critical or supercritical update family and any $p > 0$, a p -random subset of \mathbb{Z}^2 almost surely percolates.

Corollary 1.6. *Let \mathcal{U} be a critical or supercritical two-dimensional bootstrap percolation update family, let $p \in (0, 1]$, and let A be a p -random subset of \mathbb{Z}^2 . Then*

$$\mathbb{P}_p([A] = \mathbb{Z}^2) = 1. \quad \square$$

This is immediate from the previous theorems: the probability that any given site is not eventually infected is zero, and there are countably many sites.

As stated above, there are essentially equivalent versions of the first two theorems stated in terms of critical probabilities. We combine these into the following theorem.

Theorem 1.7. *Let \mathcal{U} be a two-dimensional bootstrap percolation update family.*

(i) *If \mathcal{U} is critical, then*

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

(ii) *If \mathcal{U} is supercritical, then*

$$p_c(\mathbb{Z}_n^2, \mathcal{U}) = n^{-\Theta(1)}.$$

We shall not prove Theorem 1.7 explicitly, but we note here that it follows from exactly the same methods as Theorems 1.4 and 1.5.

Remark 1.8. For certain critical update families \mathcal{U} , our methods can be used to be used to determine $p_c(\mathbb{Z}_n^2, \mathcal{U})$ up to a constant factor. This is the case, for example, with the 2-neighbour model, and thus we recover the theorem of Aizenman and Lebowitz [1] as a special case of our results (see Theorem 4.1 and Remark 6.3). However, we do not believe that our methods can be used without further improvements to determine $p_c(\mathbb{Z}_n^2, \mathcal{U})$ up to a constant factor for general critical models.

Remark 1.9. Theorem 1.7 is a statement about critical probabilities of \mathcal{U} -bootstrap percolation on the discrete torus \mathbb{Z}_n^2 . The corresponding statement with \mathbb{Z}_n^2 replaced by

the grid $[n]^2$ is false. This is for trivial reasons: the infection would not in general be able to reach the corners of the grid unless $p = 1 - o(1)$.

Returning to the first of our aims, let us now be clearer about what we mean by proving ‘basic results’ about \mathcal{U} -bootstrap percolation: we mean establishing general properties of stable sets and the action of the processes on half-planes. This is the subject of Section 3. The main result of that section is the following classification of stable sets. Let us say that a unit vector $u \in S^1$ is *rational* if it has either rational or infinite slope relative to the standard basis vectors.

Theorem 1.10. *Let $S \subset S^1$. Then there exists a two-dimensional bootstrap percolation update family \mathcal{U} of which S is the stable set if and only if S can be expressed as a finite union of closed intervals in S^1 with rational end-points.*

1.3. Organization of the paper

In the next section we introduce several key aspects of the proofs and a number of important definitions. Owing to the level of generality of our model, the need for some notational complexity is unfortunately unavoidable. However, we have tried to keep the burden as light as possible, and have gathered together many of the definitions for ease of reference.

In Section 3 we prove some basic properties of stable sets, culminating in the proof of Theorem 1.10. Following that, we move on to the main aim of this paper: the proof of Theorem 1.4 for critical update families. In Section 4, we prove the lower bound of Theorem 1.4. In Section 5 we set out the deterministic framework for the growth of ‘droplets’ that we shall use to prove the upper bounds of both Theorem 1.4 (for critical families) and Theorem 1.5 (for supercritical families). We then complete the proofs of these theorems in Sections 6 and 7 respectively.

Finally, in Section 8, we discuss open problems and conjectures, including the behaviour of subcritical families, sharpening our results, and generalizations to higher dimensions.

2. Some aspects of the proofs

In this section we introduce several important definitions and ideas from the proofs of Theorems 1.4 and 1.5. We begin with two definitions that may be familiar to the reader in the context of certain specific models (such as the 2-neighbour model or ‘balanced’ threshold models); in each case, the definition we give here generalizes previous definitions. The first, that of a set being ‘internally filled’, is our formal way of saying that a set becomes fully infected in \mathcal{U} -bootstrap percolation without external help.

Let us fix an update family \mathcal{U} , so that (for example) the operation of taking the closure of a set is well-defined.

Definition 2.1. A set $X \subset \mathbb{Z}^2$ is *internally filled by A* if $X \subset [X \cap A]$.

In the literature, the expression ‘internally spanned’ is frequently used to describe this property; we feel that the word ‘filled’ better emphasizes that the whole of X must be contained in the closure $[X \cap A]$.

The second definition is that of a ‘ \mathcal{T} -droplet’, which is a finite subset of \mathbb{Z}^2 used as a basis for growth.

Definition 2.2. Let $\mathcal{T} \subset S^1$. A \mathcal{T} -droplet is any finite set of the form

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

where $\{a_u \in \mathbb{Z}^2 : u \in \mathcal{T}\}$ is a collection of sites in \mathbb{Z}^2 . If $u \in \mathcal{T}$ and D is a \mathcal{T} -droplet, then the u -side of D is the set

$$\{x \in D : \langle y, u \rangle \leq \langle x, u \rangle \text{ for all } y \in D\}.$$

Thus, a droplet is an intersection of half-planes. Since the stable directions for the 2-neighbour process are $(\pm 1, 0)$ and $(0, \pm 1)$, if \mathcal{T} is the set of these four unit vectors then a \mathcal{T} -droplet is just a rectangle, precisely as a ‘droplet’ is usually taken to be in the 2-neighbour model.

Let us make two remarks about Definition 2.2. First, we have defined droplets with respect to a set $\mathcal{T} \subset S^1$, not with respect to an update family \mathcal{U} . This is because, for a given update family \mathcal{U} , we shall need to define and work with different sets of directions to the stable set in both the upper and lower bound parts of the proof, and in each case we shall need a correspondingly different notion of a droplet.

Second, as we have just alluded to, droplets will be important in the proofs of both parts of Theorem 1.4. From the point of view of the lower bound, they are important because they restrict growth. If $\mathcal{T} \subset \mathcal{S}$, then the closure of a subset of a \mathcal{T} -droplet is again a subset of the same \mathcal{T} -droplet. From the point of view of the upper bound, droplets form the basis of growth. With the help of a bounded number of new sites along one of its edges, under certain circumstances a \mathcal{T} -droplet will grow into a new, larger \mathcal{T} -droplet. In fact, ensuring that these ‘certain circumstances’ always hold is one of the central components of the proof of the upper bound of Theorem 1.4. We return to this in Section 2.2.

Now that we have Definitions 2.1 and 2.2, let us also make a brief remark about symmetry. One might think that the lack of symmetry should make only trivial, technical differences to the arguments, and that its assumption would cause only a little loss of generality, as was supposed in [13, 14]. However, nothing could be further from the truth: symmetry is the most important instrument at one’s disposal in proofs of lower bounds for critical probabilities in bootstrap percolation. Before explaining why, let us be clear what we mean by ‘symmetry’. We mean principally symmetry of the stable set: the property that $u \in \mathcal{S}$ implies $-u \in \mathcal{S}$. Almost always in the past, though, authors have assumed much more than this, namely that $X \in \mathcal{U}$ if and only if $-X \in \mathcal{U}$. Why should symmetry be so useful? The main reason is that it allows one to restrict growth one-dimensionally. Suppose that u and $-u$ are stable directions, and that the initial set $A \subset \mathbb{Z}^2$ does not intersect an infinite strip with sides orthogonal to u of at least a certain width depending

only on \mathcal{U} . Then $[A]$ cannot intersect that strip either. This allows tight (and easy) control of the growth of a droplet in terms only of sites near to its perimeter. It also gives an easy necessary condition for a droplet to be internally filled, which immediately implies, among other things, a certain extremal lemma (which we discuss in the next subsection). Since we do not have symmetry, we have to work much harder to obtain corresponding properties. Indeed, dealing with the lack of symmetry is the single most important challenge of the general model.

2.1. Lower bounds for critical probabilities

At the most basic level, our proofs of both the upper and lower bounds of Theorem 1.4 follow the corresponding proofs of Aizenman and Lebowitz [1] for the 2-neighbour model. In [1], the starting point for the lower bound is to show that if $\mathbf{0} \in A_t$ then there exists an internally filled droplet at the scale of roughly $\log t$ within distance t of the origin. Our starting point is similar, except that the property of being internally filled is too strong a property to ask for in general, so instead we use a certain notion of being ‘approximately internally filled’; we say a little more precisely what we mean by this shortly.

In order to find and bound the probability of an approximately internally filled droplet of size roughly $\log t$, we need two main ingredients. The first is a statement that says that if a droplet at a certain scale is approximately internally filled then it necessarily contains approximately internally filled droplets at all smaller scales. We call a statement of this form, or a close variant, an *Aizenman–Lebowitz-type lemma*. The second ingredient is a statement that says that if a droplet D of ‘size’ k is approximately internally filled, then $D \cap A$ must have cardinality at least δk , where $\delta > 0$ may depend on \mathcal{U} but not on k . We call a statement of this form an *extremal lemma*.

In the case of the 2-neighbour model, it turns out that one can prove both of these lemmas (with ‘approximately’ replaced by ‘exactly’) using the same idea. (There are many ways of proving the extremal lemma, but this is the only known way of proving the Aizenman–Lebowitz lemma. In fact, if one has a weak symmetry property, such as a pair of stable directions u and $-u$, and one does not care about optimizing δ , then the extremal lemma is immediate.) The idea in both cases is to use an algorithm called the *rectangles process*, which builds up the closure of a finite set of sites by ‘merging’ nearby internally filled droplets one-by-one. To be precise, the algorithm proceeds by looking for two rectangles within ℓ_1 distance 2 of each other, and replacing them by the smallest rectangle containing both. The two crucial properties of the algorithm are, first, that at each step, every rectangle in the current collection is internally filled, and second, that the perimeter length of the largest rectangle in the current collection does not more than double between steps.

The natural generalization of the rectangles process – replacing ‘rectangles’ by ‘ \mathcal{T} -droplets’ for some $\mathcal{T} \subset S^1$ – is of no use. The issue is that the closure of two nearby \mathcal{T} -droplets need not be a larger \mathcal{T} -droplet, so we lose the first of the two crucial properties mentioned above, namely that the \mathcal{T} -droplets at each stage of the process be internally filled. Our solution, as mentioned above, is to weaken the requirement that the droplets be internally filled. The precise way in which we do this is to introduce a new, related algorithm, which we call the *covering algorithm*. The algorithm starts by placing

a \mathcal{T} -droplet (for a certain set \mathcal{T}) around as many disjoint ‘breakthrough blocks’ in a finite initial set $K \subset \mathbb{Z}^2$ as possible, where ‘breakthrough blocks’ are defined as follows.

Definition 2.3. A *breakthrough block* for $X \in \mathcal{U}$ is any set of the form $X \cap (\mathbb{H}_u)^c$, or a translation by an element of \mathbb{Z}^2 of such a set, for some $u \in \mathcal{S}$.

The idea is that a breakthrough block is a minimal set of sites that allows at least one new site to become infected with the help of a \mathcal{T} -droplet, where $\mathcal{T} \subset \mathcal{S}$. The covering algorithm then proceeds by merging nearby \mathcal{T} -droplets in a manner similar to the original rectangles process. One can show that if the definition of the covering algorithm we have just sketched is suitably formalized, and if the final collection of \mathcal{T} -droplets at the end of algorithm is D_1, \dots, D_T , then $[K] \setminus K \subset D_1 \cup \dots \cup D_T$, and therefore that the covering algorithm ‘approximately dominates’ the \mathcal{U} -bootstrap process. Crucially, one can also prove an Aizenman–Lebowitz-type lemma and an extremal lemma for so-called ‘covered droplets’, which are simply droplets D such that the set $\{D\}$ is a possible output of the covering algorithm with input $D \cap A$. Recall that these are the two main lemmas we need to prove the lower bound.

The precise definition of the covering algorithm, statements and proofs of the Aizenman–Lebowitz-type lemma and the extremal lemma, and the proof of the lower bound of Theorem 1.4 are given in Section 4.

2.2. Upper bounds for critical probabilities

For most of this discussion we shall assume that \mathcal{U} is critical. At the end we indicate how our methods for proving the upper bound of Theorem 1.4 for critical families can be applied to prove Theorem 1.5 for supercritical families as well.

The first issue we encounter with the upper bounds is that the stable set \mathcal{S} may be quite large: indeed it may contain a closed semicircle. This means that droplets may not in general grow in all directions. However, in order to show that droplets grow in at least one direction, we only need a weaker property, which follows from the definition of a critical update family and Theorem 1.10. This property is that there exists an open semicircle $C \subset S^1$ such that $|\mathcal{S} \cap C| < \infty$.

We remarked after Definition 2.2 that ensuring that droplets grow into droplets is one of the central parts of the proof of the upper bound of Theorem 1.4. What we mean by ‘droplets growing into droplets’, or more specifically, ‘ \mathcal{T} -droplets growing into \mathcal{T} -droplets’, is that for each \mathcal{T} -droplet D and each $u \in \mathcal{T}$, there exists a set $Z \subset \mathbb{Z}^2 \setminus D$ of bounded size (that is, of size depending on \mathcal{U} and \mathcal{T} but not on D) such that $[D \cup Z]$ contains an entire new ‘row’ of sites on the u -side, extending all the way to the corners of the u -side. Thus, $[D \cup Z]$ is itself a \mathcal{T} -droplet. Once we have a statement of this form, then our proof of the upper bound of Theorem 1.4 follows by methods similar to those used by van Enter and Hulshof [21] to determine the critical probability of a particular so-called ‘unbalanced’ model.

The growth property that we have just described itself hinges on two further properties. The first is that if $u \in S^1$ is such that there exists an open interval in $I \subset S^1$ with $u \in I$ but such that $\mathcal{S} \cap I \setminus \{u\} = \emptyset$, then there exists a finite set $Z \subset \ell_u$ such that $\ell_u \subset [\mathbb{H}_u \cup Z]$,

where

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

(Of course, if $u \notin \mathcal{S}$ then $Z = \emptyset$ will do here.) This property, which is Lemma 5.2, will allow us to grow to within a constant distance of the corners of the \mathcal{T} -droplet. However, if we choose our set \mathcal{T} to be such that $\mathcal{T} \subset \mathcal{S}$ – as in all previous proofs of this type of theorem – then the new row of infected sites *does not* in general grow all the way to the corners. The second property that we need in order to rectify this problem is that there exists a finite set $\mathcal{Q} \subset S^1$, which we call the set of *quasi-stable directions*, such that if \mathcal{T} is the union of \mathcal{Q} and a certain subset of \mathcal{S} , then this corners property holds. This is the content of Lemma 5.3.

Let us briefly observe that both of the properties of the previous paragraph trivially hold in the case of the 2-neighbour model: for the first, we may always take $Z = \{\mathbf{0}\}$, and for the second, we may take $\mathcal{Q} = \emptyset$. This again serves to highlight that many of the issues that arise in connection with the general model do not appear at all in connection with the 2-neighbour model.

In Section 5 we establish the key growth properties described above. We then apply them in Section 6 to prove the upper bound of Theorem 1.4 for critical families. In Section 7, we show that the framework we have set up for critical families may be adapted to prove Theorem 1.5 for supercritical families, by observing that all of the sets Z mentioned above may be taken to be empty if \mathcal{U} is supercritical and the open semicircle $C \subset S^1$ is chosen so that $\mathcal{S} \cap C = \emptyset$.

2.3. Definitions and conventions

Let us collect a few of the most important definitions and conventions that we shall use throughout the paper. We begin with three conventions.

- The set A is always a p -random subset of \mathbb{Z}^2 , unless explicitly stated otherwise.
- The update family \mathcal{U} is fixed throughout. Constants will be allowed to depend on \mathcal{U} (and hence on \mathcal{S} , *etc.*), but not on p .
- The norm $\|\cdot\|$ denotes the Euclidean norm $\|\cdot\|_2$.

Moving on to definitions, we begin with certain subsets of the plane. Recall that we have already defined the half-plane $\mathbb{H}_u = \{x \in \mathbb{Z}^2 : \langle x, u \rangle < 0\}$ and the line $\ell_u = \{x \in \mathbb{Z}^2 : \langle x, u \rangle = 0\}$. We shall also need the shifted half-plane:

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

where $a \in \mathbb{R}^2$. If $a \in \mathbb{Z}^2$ then we have $\mathbb{H}_u(a) = \mathbb{H}_u + a$, but this is not true in general. Next, let $u \in S^1$ be rational (recall that this means that u has either rational or infinite slope with respect to the standard basis vectors of \mathbb{R}^2) and consider the lines orthogonal to u that intersect \mathbb{Z}^2 . These lines are parallel and discrete, and there is a natural integer-valued indexing for them. Let $\ell_u(0)$ be the line intersecting the origin, and for each $i \in \mathbb{Z}$, let $\ell_u(i)$ be the i th line in the direction of u .

Next, let $\mathbb{T} := \mathbb{R}/2\pi\mathbb{Z}$. We shall often need to change between elements of S^1 and elements of \mathbb{T} , so we define the following natural bijection. Let $u : \mathbb{T} \rightarrow S^1$ be the function

$$u(\theta) := (\cos \theta, \sin \theta),$$

and let $\theta : S^1 \rightarrow \mathbb{T}$ be the inverse function, so that if $u = u(\theta)$ then $\theta(u) = \theta$. We extend the domain of the function θ to the whole of $\mathbb{R}^2 \setminus \{\mathbf{0}\}$ in the obvious way: if $x \in \mathbb{R}^2 \setminus \{\mathbf{0}\}$ then we define $\theta(x)$ to be equal to $\theta(x/\|x\|)$. We also use standard function notation, so that for example $u(T) := \{u(\theta) : \theta \in T\}$ for a subset $T \subset \mathbb{T}$, and we abuse notation slightly by abbreviating $u((\theta_1, \theta_2))$ to $u(\theta_1, \theta_2)$.

Now let $\mathcal{T} \subset S^1$ and let $u, v \in \mathcal{T}$. We shall say that u and v are *consecutive* in \mathcal{T} if $\mathcal{T} \cap u(\theta(u), \theta(v)) = \emptyset$. (Note that this is not a symmetric relation: the statements that u and v are consecutive in \mathcal{T} , and that v and u are consecutive in \mathcal{T} , are not equivalent.)

For each update rule X , let

$$T(X) := \{u \in S^1 : X \subset \mathbb{H}_u\},$$

and let

$$\Theta(X) := \{\theta \in \mathbb{T} : u(\theta) \in T(X)\} = \{\theta \in \mathbb{T} : X \subset \mathbb{H}_{u(\theta)}\}.$$

We think of $T(X)$ as the set of directions *destabilized* by the update rule X . If $u \in S^c$ then there must exist a rule $X \in \mathcal{U}$ such that $X \subset \mathbb{H}_u$, which implies that

$$S^c = \bigcup_{X \in \mathcal{U}} T(X). \tag{2.1}$$

Finally, we shall need the following measure of the *diameter* of a finite set $K \subset \mathbb{Z}^2$:

$$\text{diam}(K) := \max\{\|x - y\| : x, y \in K\}.$$

3. The structure of stable sets

In this section we derive a number of basic properties of the stable set and introduce some important definitions relating to the geometry of the stable set. Our main aim is to prove Theorem 1.10, the classification of stable sets.

We begin with a simple lemma that establishes the dichotomy $[\mathbb{H}_u] \in \{\mathbb{H}_u, \mathbb{Z}^2\}$ for every $u \in S^1$.

Lemma 3.1. *Let $u \in S^1$ be an unstable direction for \mathcal{U} . Then $[\mathbb{H}_u] = \mathbb{Z}^2$.*

Proof. Let $u \in S^1$ be unstable and let $a \in \mathbb{Z}^2 \setminus \mathbb{H}_u$ and $X \in \mathcal{U}$ be such that $a + X \subset \mathbb{H}_u$. Then $y + X \subset \mathbb{H}_u$ for all $y \in \ell_u$, since if $x \in X$ and $y \in \ell_u$ then

$$\langle y + x, u \rangle = \langle x, u \rangle \leq \langle a + x, u \rangle < 0.$$

Thus $\ell_u \subset A_1 \setminus A_0$.

If u is rational then we are done: $\ell_u(j) \subset A_{j+1}$ for all integers $j \geq 0$, so $[\mathbb{H}_u] = \mathbb{Z}^2$. If u is irrational, on the other hand, then all we know so far is that $\mathbf{0} \in A_1 \setminus A_0$. We claim

that there exists a site $b \in \mathbb{Z}^2$ such that $\langle b, u \rangle > 0$ and $\mathbb{H}_u(jb) \subset A_j$ for all $j \in \mathbb{N}$. Let

$$\delta := \sup\{\lambda \in \mathbb{R} : X \subset \mathbb{H}_u(-\lambda u)\} = \min\{-\langle x, u \rangle : x \in X\}.$$

If $y \in \mathbb{H}_u(\delta u)$ then $y + X \subset \mathbb{H}_u$, so we have $\mathbb{H}_u(\delta u) \subset A_1$. Also, the set X is finite and contained in \mathbb{H}_u , so $\delta > 0$. Thus, since u is irrational, there exists a site $b \in \mathbb{H}_u(\delta u) \setminus \mathbb{H}_u$, and we have $\mathbb{H}_u(b) \subset \mathbb{H}_u(\delta u) \subset A_1$. Now $\mathbb{H}_u(b)$ has a site on its boundary, namely b , so it is congruent to \mathbb{H}_u , and therefore $\mathbb{H}_u(jb) \subset A_j$ for all $j \in \mathbb{N}$, as claimed. □

The next lemma, together with the observation in (2.1), is one of the two implications in Theorem 1.10.

Lemma 3.2. *Let X be an update rule. Then $T(X)$ is either empty or an open interval in S^1 with rational end-points.*

Proof. For each site $x \in \mathbb{Z}^2 \setminus \{0\}$ define the set $T(x) = \{u : \langle x, u \rangle < 0\}$. It is easy to see that

$$T(x) = u(\theta(x) - \pi/2, \theta(x) + \pi/2). \tag{3.1}$$

Furthermore, since $u(\theta(x)) = x/\|x\|$ is rational, so too are $u(\theta(x) - \pi/2)$ and $u(\theta(x) + \pi/2)$. Thus, $T(x)$ is an open interval with rational end-points.

Now simply notice that

$$T(X) = \bigcap_{x \in X} T(x), \tag{3.2}$$

so $T(X)$ is a finite intersection of open intervals with rational end-points, each interval of length exactly π . The assertion of the lemma follows. □

Remark 3.3. Using identities for S^c in (2.1), for $T(X)$ in (3.2), and for $T(x)$ in (3.1), we have

$$S^c = \bigcup_{X \in \mathcal{U}} \bigcap_{x \in X} u(\theta(x) - \pi/2, \theta(x) + \pi/2). \tag{3.3}$$

From this it follows immediately that there exists a polynomial time (in the sum of the cardinalities of the update rules, say) algorithm for computing the stable set \mathcal{S} of an update family \mathcal{U} .

We are now able to complete the proof of Theorem 1.10. Recall that the theorem is a classification of stable sets: it says that a subset of the circle is the stable set of some update family if and only if it can be expressed as a finite union of closed intervals with rational end-points.

Proof of Theorem 1.10. One implication – that stable sets are finite unions of closed intervals with rational end-points – is a consequence of Lemma 3.2 and (2.1). Therefore our task is just to show the converse. To that end, let $\mathcal{S} \subset S^1$ be such that its complement




Figure 2. The rule X_i consists of the two sites x_i^- and x_i^+ . It destabilizes the interval (θ_i^-, θ_i^+) .

may be written in the form

$$\mathcal{S}^c = \bigcup_{i=1}^m T_i,$$

where

$$T_i = u(\theta_i^-, \theta_i^+)$$

for $i = 1, \dots, m$, and the T_i are disjoint and have rational end-points. Since $u(\theta_i^-)$ and $u(\theta_i^+)$ are rational, for each i , it follows that there exists a non-zero site $x_i^- \in \ell_{u(\theta_i^-)}$ to the right of the origin as one looks in the direction of $u(\theta_i^-)$, and a non-zero site $x_i^+ \in \ell_{u(\theta_i^+)}$ to the left of the origin as one looks in the direction of $u(\theta_i^+)$. Let

$$X_i := \{x_i^-, x_i^+\}$$

be update rules for each $i = 1, \dots, m$, and let $\mathcal{U} := \{X_1, \dots, X_m\}$ be an update family and $\mathcal{S}' = \mathcal{S}'(\mathcal{U})$ its stable set. We claim that $\mathcal{S}' = \mathcal{S}$.

By (2.1),

$$(\mathcal{S}')^c = \bigcup_{i=1}^m T(X_i),$$

so it is enough to show that $T(X_i) = T_i$ for every i . But

$$\begin{aligned} T(X_i) &= \{u : \langle x_i^-, u \rangle < 0\} \cap \{u : \langle x_i^+, u \rangle < 0\} \\ &= u(\theta_i^-, \theta_i^- + \pi) \cap u(\theta_i^+ - \pi, \theta_i^+) \\ &= u(\theta_i^-, \theta_i^+), \end{aligned}$$

which is the definition of T_i . □

4. The lower bound for critical families

The next three sections are the primary focus of the paper. We study critical² \mathcal{U} -bootstrap percolation with the aim of proving Theorem 1.4: the lower bound in this section and the upper bound in Section 6. Recall that an update family is critical if the following two properties of its stable set \mathcal{S} hold. First, $\mathcal{S} \cap C$ is non-empty for every open semicircle $C \subset S^1$, and second, there exists an open semicircle $C \subset S^1$ not containing any strongly stable directions, where a stable direction u is said to be strongly stable if it is contained in an open interval of stable directions.

Recall that a breakthrough block for $X \in \mathcal{U}$ is any set of the form $X \cap (\mathbb{H}_u)^c$, where $u \in \mathcal{S}$. (It is worth observing that there are no empty breakthrough blocks, because the direction u is assumed to be stable.) Let \mathcal{B} be the collection of sets $B \subset \mathbb{Z}^2$ such that B is a breakthrough block for some $X \in \mathcal{U}$, and let

$$\alpha_1 = \alpha_1(\mathcal{U}) := \min\{|B| : B \in \mathcal{B}\}. \quad (4.1)$$

The lower bound of Theorem 1.4 follows from the next theorem. We remind the reader that the stopping time τ is the minimal $t \geq 0$ such that $\mathbf{0} \in A_t$.

Theorem 4.1. *Let \mathcal{U} be a critical update family and let*

$$p \leq \left(\frac{\lambda}{\log t} \right)^{1/\alpha_1},$$

where $\lambda > 0$ is sufficiently small. Then $\tau \geq t$ with high probability as $t \rightarrow \infty$.

It follows from the definition of a critical update family that if \mathcal{U} is critical then there exists a set $\mathcal{S}_L = \mathcal{S}_L(\mathcal{U})$ of 3 or 4 stable directions such that $\mathbf{0}$ lies in the interior of the convex hull of \mathcal{S}_L . (Equivalently, there exist (finite) \mathcal{S}_L -droplets.) Throughout this section we fix the critical update family \mathcal{U} and the set \mathcal{S}_L of 3 or 4 stable directions, and we shall use the convention that all droplets are \mathcal{S}_L -droplets.

In Section 4.1 we define the ‘covering algorithm’, which will be our approximation to the rectangles process of the 2-neighbour model, and we define what it means for a droplet to be ‘covered’, which will be our notion of being ‘approximately internally filled’. We use these definitions to derive an Aizenman–Lebowitz-type lemma (Lemma 4.8) and an extremal lemma (Lemma 4.7) for ‘covered’ droplets. In Section 4.2 we assemble the tools of the previous subsection to complete the proof of Theorem 4.1.

4.1. Covered droplets

We begin by defining the ‘covering algorithm’, which takes as its input a finite set $K \subset \mathbb{Z}^2$ of sites and returns a collection \mathcal{D} of \mathcal{S}_L -droplets that approximately cover $[K]$, in a sense made precise in Lemma 4.5.

Let \hat{D} be a fixed \mathcal{S}_L -droplet such that $X \subset \hat{D}$ for all $X \in \mathcal{U}$. Given a finite set $K \subset \mathbb{Z}^2$, let $D(K)$ be the minimal \mathcal{S}_L -droplet containing K .

² In Section 5 we only assume that \mathcal{U} is critical or supercritical.

Definition 4.2 (covering algorithm). Let \mathcal{U} be critical and $K \subset \mathbb{Z}^2$ be finite. Let B_1, \dots, B_{k_0} be a maximal collection of disjoint breakthrough blocks in K , and let $\mathcal{D}^0 := \{D_1^0, \dots, D_{k_0}^0\}$ be a collection of copies of \hat{D} such that $B_i \subset D_i^0$ for each $i = 1, \dots, k_0$. This is step 0. Now repeat the following procedure for each $t \geq 0$ until STOP. At the start of step $t + 1$ of the algorithm, there is a collection $\mathcal{D}^t = \{D_1^t, \dots, D_{k_t}^t\}$ of droplets, where $k_t = k_0 - t$. If there do not exist indices i and j and a site $x \in \mathbb{Z}^2$ such that

$$D_i^t \cap (\hat{D} + x) \neq \emptyset \quad \text{and} \quad D_j^t \cap (\hat{D} + x) \neq \emptyset, \tag{4.2}$$

then STOP. If there do exist such i, j and x , then choose one such triple. Construct \mathcal{D}^{t+1} from \mathcal{D}^t by deleting from it D_i^t and D_j^t and adding to it $D(D_i^t \cup D_j^t)$:

$$\mathcal{D}^{t+1} := (\mathcal{D}^t \setminus \{D_i^t, D_j^t\}) \cup \{D(D_i^t \cup D_j^t)\}.$$

This is the end of step $t + 1$. The output of the algorithm is the set $\mathcal{D}^T := \{D_1^T, \dots, D_k^T\}$, where $k = k_0 - T$, and T is the last t before STOP.

Remark 4.3. We could instead have defined the covering algorithm by taking our initial collection of droplets \mathcal{D}^0 to consist of a copy of \hat{D} around every *site* of K , rather than around every element of a maximal collection of disjoint breakthrough blocks. The effect of this would have been that we would have obtained $\alpha_1 = 1$ as the bound in Theorem 4.1 for every critical update family. We use the present method because in some cases it gives better bounds, and because it may be useful in potential future refinements of our results.

We note that the covering algorithm is commutative in the sense that the order in which droplets are combined does not affect the output of the algorithm. We omit the proof of this property, since we shall not actually need to use it. On the other hand, the output of the covering algorithm *does* depend on the particular choice of disjoint breakthrough blocks B_1, \dots, B_{k_0} , and on the particular translates of \hat{D} we choose to cover those breakthrough blocks.

Definition 4.4. We say that \mathcal{D} is a *cover* of a finite set $K \subset \mathbb{Z}^2$ if it is a possible output of the covering algorithm with inputs \mathcal{U} and K . An \mathcal{S}_L -droplet D is *covered* if $\mathcal{D} = \{D\}$ for some cover \mathcal{D} of $D \cap A$.

The following lemma allows us to bound the closure of a finite set in terms of an arbitrary output of the covering algorithm applied to that set; it is the sense in which the covering algorithm approximately dominates the bootstrap process.

Lemma 4.5. Let $\mathcal{D} = \{D_1^T, \dots, D_k^T\}$ be an output of the covering algorithm applied to a finite set $K \subset \mathbb{Z}^2$. Then

$$[K] \setminus K \subset D_1^T \cup \dots \cup D_k^T.$$

Proof. Let $L := D_1^T \cup \dots \cup D_k^T$. We claim that $[L \cup K] = L \cup K$. In order to prove this, we must show that there do not exist $X \in \mathcal{U}$ and $x \in \mathbb{Z}^2$ such that $x + X \subset L \cup K$ and $x \notin L \cup K$.

First, it is clear that we cannot have $x + X \subset K \setminus L$. This is simply because $x + X$ contains a breakthrough block, so $x + X \subset K \setminus L$ would contradict the maximality of the initial collection B_1, \dots, B_{k_0} of breakthrough blocks in the covering algorithm.

Second, we cannot have $x + X \subset L \cup K$ and $(x + X) \cap L \neq \emptyset$. To see this, first note that $x + X$ cannot intersect more than one of the D_i^T , because if $t = T$ then (4.2) does not hold for any $1 \leq i < j \leq k$ and \hat{D} contains every update rule. Also, $x + X$ cannot be contained in D_i^T for any i , because each D_i^T is closed. Therefore, if $(x + X) \cap L \neq \emptyset$, then there exists a unique i such that $(x + X) \cap D_i^T \neq \emptyset$ and $(x + X) \setminus D_i^T \subset K \setminus L$. Hence also, there exist $u \in \mathcal{S}_L$ and $y \in \mathbb{Z}^2$ such that

$$\emptyset \neq (x + X) \cap \mathbb{H}_u(y)^c \subset K \setminus L,$$

again contradicting the maximality of the initial collection of breakthrough blocks.

Thus we have $[L \cup K] = L \cup K$ as claimed, and therefore

$$[K] \subset [L \cup K] = L \cup K,$$

which proves the lemma. □

We require one further ingredient for the extremal lemma, which is a subadditivity lemma for the diameters of intersecting droplets. We omit the proof of the lemma because our droplets are always either triangles or parallelograms, and in these cases the lemma is a triviality. However, we note for reference that the lemma holds in much greater generality than is stated here: it holds for droplets with respect to an arbitrary (fixed) finite subset of S^1 . The interested reader may refer to Lemma 23 of version 2 of the present paper on the arXiv for a full proof.³

Lemma 4.6. *Let D_1 and D_2 be \mathcal{S}_L -droplets such that $D_1 \cap D_2 \neq \emptyset$. Then*

$$\text{diam}(D(D_1 \cup D_2)) \leq \text{diam}(D_1) + \text{diam}(D_2). \quad \square$$

We are now ready to prove the extremal lemma for covered droplets.

Lemma 4.7 (extremal lemma). *Let \mathcal{U} be critical and let D be a covered droplet. Then $D \cap A$ contains at least $\Omega(\text{diam}(D))$ disjoint breakthrough blocks.*

Proof. Apply the covering algorithm with input $D \cap A$. The algorithm starts with k_0 droplets containing disjoint breakthrough blocks from $D \cap A$, and it terminates after k_0 steps with a single droplet. At the $(t + 1)$ th step of the algorithm, we replace droplets D_i^t

³ Version 2 of the present paper, which has the alternative name ‘Neighbourhood family percolation’, may be found at <http://arxiv.org/abs/1204.3980v2>.

and D_j^t by $D(D_i^t \cup D_j^t)$, for some i and j , where D_i^t and D_j^t are such that (4.2) holds. It follows by Lemma 4.6 that

$$\text{diam}(D(D_i^t \cup D_j^t)) \leq \text{diam}(D_i^t) + \text{diam}(D_j^t) + \text{diam}(\hat{D}). \tag{4.3}$$

Thus, since $\text{diam}(\hat{D}) = O(1)$, the quantity

$$\sum_{D_i^t \in \mathcal{D}} \text{diam}(D_i^t)$$

increases by at most $O(1)$ at each step of the algorithm, and hence

$$\text{diam}(D) \leq k_0 \text{diam}(\hat{D}) + O(k_0).$$

It follows that $k_0 = \Omega(\text{diam}(D))$, as required. □

Finally in this subsection, we prove the Aizenman–Lebowitz-type lemma for covered droplets.

Lemma 4.8 (Aizenman–Lebowitz-type lemma). *Let D be a covered droplet. Then, for every $1 \leq k \leq \text{diam}(D)$ there exists a covered droplet $D' \subset D$ such that $k \leq \text{diam}(D') \leq 3k$.*

Proof. The proof is similar to (but even simpler than) that of Lemma 4.7. Apply the covering algorithm with input $D \cap A$. By (4.3), the quantity

$$\max\{\text{diam}(D_i^t) : D_i^t \in \mathcal{D}^t\}$$

at most triples at each step of the algorithm, and moreover, it is easy to see that every droplet $D_i^t \in \mathcal{D}^t$ is covered, which proves the lemma. □

4.2. Proof of Theorem 4.1

We require two calculations for the proof of Theorem 4.1. The first says that the probability of there existing a droplet of diameter roughly $\log t$ within distance $O(t)$ of the origin is small. The second says that the probability of there existing a strongly covered droplet containing the origin with diameter at most $\log t$ is also small. These calculations are Lemmas 4.9 and 4.10 respectively. The proof of Theorem 4.1 will split naturally into these two cases.

If $k \in \mathbb{N}$ then let $D(k)$ be the minimal droplet containing the origin all of whose sides have ℓ_2 distance at least λk from the origin, where λ is a large constant. Note that we have previously defined $D(K)$ for K a subset of \mathbb{Z}^2 ; we hope this slight abuse of notation does not cause any confusion.

We begin with the first calculation, which says that it is unlikely that there is a covered droplet of diameter roughly $\log t$ contained in $D(t)$.

Lemma 4.9. *Let \mathcal{U} be critical, let $\epsilon > 0$ be sufficiently small, and let*

$$p \leq \left(\frac{\epsilon}{\log t} \right)^{1/\alpha_1}.$$

Then the probability that there exists a covered droplet $D \subset D(t)$ such that $\log t \leq \text{diam}(D) \leq 3 \log t$ is $o(1)$ as $t \rightarrow \infty$.

Proof. By Lemma 4.7, there exists a constant $\delta > 0$ such that if D is a covered droplet with $\log t \leq \text{diam}(D) \leq 3 \log t$ then $D \cap A$ contains at least $\delta \log t$ breakthrough blocks. Since there are $O(1)$ distinct breakthrough blocks and each breakthrough block has size at least α_1 , the probability that a particular such droplet is covered is at most

$$\binom{O(\log t)^2}{\delta \log t} p^{\alpha_1 \delta \log t} \leq (O(p^{\alpha_1} \log t))^{\delta \log t} \leq \frac{1}{t^3},$$

because we chose ϵ sufficiently small. Now, the number of droplets $D \subset D(t)$ with $\log t \leq \text{diam}(D) \leq 3 \log t$ is $t^2(\log t)^{O(1)}$. Therefore, the probability that there exists a covered droplet $D \subset D(t)$ with $\log t \leq \text{diam}(D) \leq 3 \log t$ is at most

$$t^2(\log t)^{O(1)} \cdot \frac{1}{t^3} = o(1),$$

as required. □

The next calculation says that it is unlikely that there is a covered droplet contained in $D(\log t)$ that itself contains the origin.

Lemma 4.10. *Let \mathcal{U} be critical, let $\epsilon > 0$ be sufficiently small, and let*

$$p \leq \left(\frac{\epsilon}{\log t} \right)^{1/\alpha_1}.$$

Then the probability that the origin is contained in a covered droplet $D \subset D(\log t)$ is $o(1)$ as $t \rightarrow \infty$.

Proof. We begin as in the previous lemma: by Lemma 4.7 there exists $\delta > 0$ such that if D is a covered droplet with diameter at least k then $D \cap A$ contains at least δk disjoint breakthrough blocks. Thus the probability that a droplet with diameter between k and $k + 1$ is covered is at most

$$\binom{O(k^2)}{\delta k} p^{\alpha_1 \delta k} \leq (O(p^{\alpha_1} k))^{\delta k} = \left(\frac{O(\epsilon k)}{\log t} \right)^{\delta k}.$$

The number of droplets of diameter between k and $k + 1$ that contain the origin is at most $k^{O(1)}$. Hence, the probability that there exists a covered droplet $D \subset D(\log t)$ such that $\mathbf{0} \in D$ is at most

$$\sum_{k=1}^{O(\log t)} k^{O(1)} \left(\frac{O(\epsilon k)}{\log t} \right)^{\delta k}.$$

By breaking up the sum at $k = (\log \log t)^2$, this in turn is at most

$$(\log \log t)^{O(1)} \left(\frac{(\log \log t)^2}{\log t} \right)^{\delta} + O(\log t)^{O(1)} O(\epsilon)^{\delta (\log \log t)^2} = o(1),$$

which completes the proof of the lemma. □

We are now ready to prove Theorem 4.1.

Proof of Theorem 4.1. Let p be such that

$$p \leq \left(\frac{\epsilon}{\log t} \right)^{1/\alpha_1},$$

where $\epsilon > 0$ is sufficiently small. We shall show that $\tau \geq t$ with high probability as $t \rightarrow \infty$.

First we observe that if $\mathbf{0} \in A_t$ then we must have $\mathbf{0} \in [D(t) \cap A]$, provided the constant λ in the definition of $D(t)$ is sufficiently large. Therefore, by Lemma 4.5, either $\mathbf{0} \in A$ or $\mathbf{0} \in D$, where $D \subset D(t)$ is a covered droplet. Since the probability of the former event is $p = o(1)$, we need only bound the probability of the latter.

We split into two cases according to the size and position of D . Suppose first that $D \subset D(\log t)$. Then the probability that D is covered is $o(1)$ by Lemma 4.10. On the other hand, if $D \subset D(t)$ but $D \not\subset D(\log t)$, then by Lemma 4.8 there exists a covered droplet $D' \subset D$ such that $\log t \leq \text{diam}(D') \leq 3 \log t$. But the probability of this event is also $o(1)$, by Lemma 4.9, and this completes the proof of the theorem. □

5. Upper bounds: the growth of droplets

The aim of this section is to develop a deterministic framework for growth in \mathcal{U} -bootstrap percolation. In the subsequent two sections, we shall use this framework to bound from below the growth that occurs under the action of a critical (in Section 6) or supercritical (in Section 7) update family. Our aim in constructing this framework can be divided into three sub-aims. The first is to establish a means by which an entire half-plane \mathbb{H}_u can grow by a single row; this we do by introducing the notion of ‘ u -blocks’ in Section 5.1. The second is to transfer from growth of half-planes by a single row to growth of droplets by a single row; this we do by introducing the notion of ‘quasi-stability’ in Section 5.2. The third and final sub-aim is to define sequences of nested droplets and give sufficient conditions for growth to occur between successive droplets; this we do in Section 5.3.

5.1. The infection of new lines

Our first task, then, is to show that there exists a finite set $Z \subset \mathbb{Z}^2$ such that $\ell_u \subset [\mathbb{H}_u \cup Z]$, provided $u \in S^1$ satisfies certain natural conditions. We shall in fact show that one can take $Z \subset \ell_u$.

Given a stable direction u , we define ℓ_u^l to be the set of sites in ℓ_u that lie to the left of the origin as one looks in the direction of u , and ℓ_u^r to be the set of sites in ℓ_u that lie to the right of the origin as one looks in the direction of u .

Definition 5.1. Let $u \in S^1$. A finite set $Z \subset \mathbb{Z}^2$ is a *u -left-block* (respectively, a *u -right-block*) if there exists $x \in \mathbb{Z}^2$ such that $x + Z$ is a subset of ℓ_u and consists of consecutive sites in ℓ_u , and if $\ell_u^l \subset [\mathbb{H}_u \cup (x + Z)]$ (respectively, $\ell_u^r \subset [\mathbb{H}_u \cup (x + Z)]$). It is a *u -block* if it is both a u -left-block and a u -right-block.

These special breakthrough blocks are so-called because we think of them as being the means of growing left or right along the edge of a half-plane or along a side of a droplet. Note that the empty set is a u -block (and also a u -left-block and a u -right-block) if and only if u is unstable.

A stable direction $u(\theta)$ is said to be *left-isolated* if there exists $\delta > 0$ such that for all $\phi \in (\theta, \theta + \delta)$, the direction $u(\phi)$ is unstable. Similarly, $u(\theta)$ is said to be *right-isolated* if there exists $\delta > 0$ such that for all $\phi \in (\theta - \delta, \theta)$, the direction $u(\phi)$ is unstable. We say that u is *isolated* if it is both left-isolated and right-isolated.

The following lemma should be thought of as saying that there exist u -blocks for every $u \in \mathcal{S}$ that will be relevant to us.

Lemma 5.2. *There exists a u -left-block for every left-isolated stable direction u , a u -right-block for every right-isolated stable direction u , and a u -block for every isolated stable direction u .*

Proof. Let $u = u(\theta)$ be a left-isolated stable direction. Then there exists an update rule X and $\delta > 0$ such that $u(\phi)$ is destabilized by X for every ϕ in the open interval $I = (\theta, \theta + \delta)$. Thus

$$X \subset \bigcap_{\phi \in I} \mathbb{H}_{u(\phi)} \subset \mathbb{H}_u \cup \mathcal{L}'_u,$$

so $X \cap (\mathbb{H}_u)^c$ is a u -left-block. If u is right-isolated then a u -right-block exists by symmetry. If u is isolated, then it is left- and right-isolated, so the result follows. \square

5.2. Quasi-stability

Our second task of this section is to transfer the results of the previous subsection from half-planes to droplets: that is, to establish conditions under which a u -block can be used to extend a droplet by a single row along its u -side. The important property that we need is that the infection grows *all the way to the corners*, not just to within a constant distance.

In order to illustrate why this is not straightforward, let us give an example. Suppose X_1 and X_2 are two update rules such that $\Theta(X_i) = (\phi_i, \psi_i)$ for $i = 1, 2$, and $\phi_2 < \phi_1 < \psi_2 < \psi_1$ (see Figure 3). Thus, there are no stable directions between $u(\phi_2)$ and $u(\psi_1)$. Suppose (although it is not necessary) that in fact $u(\phi_2)$ and $u(\psi_1)$ are stable directions, so \mathcal{S} -droplets have consecutive sides with outward normals in the directions of $u(\phi_2)$ and $u(\psi_1)$. Now suppose we are growing leftwards along the $u(\phi_2)$ -side, and we are nearly at the corner of the droplet. How are we going to reach the corner? X_1 does not fit inside the droplet near the corner because $\phi_1 > \phi_2$, so it partly lies outside the $u(\phi_2)$ -side. X_2 does not fit inside the droplet near the corner either, because $\psi_2 < \psi_1$, so it partly lies outside the $u(\psi_1)$ -side (again, see Figure 3). So there is no rule that enables us to reach the corner, and we are stuck. Our solution is to ‘pretend’ that $u(\phi_1)$ and $u(\psi_2)$ are stable directions (we call them ‘quasi-stable directions’). More precisely, we ensure that the set $\mathcal{T} \subset S^1$ with respect to which our droplets are defined includes the directions $u(\phi_1)$ and $u(\psi_2)$. If we do this, and ‘grow’ the $u(\phi_1)$ - and $u(\psi_2)$ -sides of our \mathcal{T} -droplets just as the sides with respect to stable directions, then we find that we are indeed able to grow to the corners.




Figure 3. An illustration of the need for quasi-stable directions. The top part of the figure shows two rules, X_1 and X_2 , that destabilize overlapping intervals in S^1 . Each of the rules consists of just two sites. In the bottom part of the figure, the two long solid lines show one corner of an S -droplet D . Neither X_1 nor X_2 can be used to grow to the corner. The two dashed lines show how the droplet might look if $u(\phi_1)$ and $u(\psi_2)$ were added to the stable set. Now X_1 can be used to grow towards x and y from both directions, and X_2 can be used to grow towards z (and y) from both directions.

Lemma 5.3. *Let \mathcal{U} be an arbitrary update family. Then there exists a finite set $\mathcal{Q} \subset S^1$ such that if u and v are consecutive elements of $S \cup \mathcal{Q}$ then there exists a rule $X \in \mathcal{U}$ such that*

$$X \subset (\mathbb{H}_u \cup \ell_u) \cap (\mathbb{H}_v \cup \ell_v). \tag{5.1}$$

Proof. As suggested above, our construction is to add to S every unit vector u perpendicular to x , for every site $x \in X$ and every update rule $X \in \mathcal{U}$. Thus, we define

$$\mathcal{Q} := \bigcup_{X \in \mathcal{U}} \bigcup_{x \in X} \{u \in S^1 : \langle x, u \rangle = 0\}. \tag{5.2}$$

Now let u and v be consecutive elements of $S \cup \mathcal{Q}$, and suppose that (5.1) does not hold. This means that for every $X \in \mathcal{U}$ we have

$$X \not\subset (\mathbb{H}_u \cup \ell_u) \cap (\mathbb{H}_v \cup \ell_v). \tag{5.3}$$

Choose an arbitrary unit vector $w \in S^1$ such that $\theta(u) < \theta(w) < \theta(v)$. Since w is unstable, there exists an update rule X such that $X \subset \mathbb{H}_w$. Now, X satisfies (5.3), so without loss of generality there exists $x \in X$ such that $x \in \mathbb{H}_u$ and $x \notin (\mathbb{H}_v \cup \ell_v)$. But this implies that there is a unit vector w' such that $\langle x, w' \rangle = 0$ and $\theta(u) < \theta(w') < \theta(v)$, contradicting the construction of \mathcal{Q} . □




Figure 4. Lemma 5.3.

Henceforth we fix a set \mathcal{Q} satisfying the conclusion of Lemma 5.3 (for example, the set defined in (5.2)), and we call this set the *quasi-stable set* and its elements *quasi-stable directions*. The next lemma is a little technical to state, but it should be thought of as saying that, provided we use the quasi-stable directions, for certain directions v , a \mathcal{T} -droplet can grow by a single row in the direction of v (all the way to the corners) with the help of a v -block.

Lemma 5.4. *Let $\mathcal{T} \subset S^1$ be finite and let*

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

be a \mathcal{T} -droplet for some set $\{a_u \in \mathbb{Z}^2 : u \in \mathcal{T}\}$ such that every side of D is sufficiently long. Suppose w_1, v and w_2 are elements of \mathcal{T} and are consecutive in $\mathcal{T} \cup \mathcal{S} \cup \mathcal{Q}$. Let

$$D' := \{x \in \mathbb{Z}^2 : \langle x - a_v, v \rangle \leq 0\} \cap \bigcap_{u \in \mathcal{T} \setminus \{v\}} (\mathbb{H}_u + a_u)$$

be the \mathcal{T} -droplet formed from D by the addition of a single new row along its v -side. Finally, suppose Z is a u -block such that Z is contained in the line segment $D' \setminus D$ and is at least a sufficiently large constant distance from either end of the line segment. Then

$$D' \subset [D \cup Z].$$

Proof. By the definition of a u -block, the set $[D \cup Z]$ contains all of D' except for possibly at most a constant number of sites at either end of the line segment $D' \setminus D$. Suppose the conclusion of the lemma is false, so $D' \not\subset [D \cup Z]$. Without loss of generality, there is a site $z \in (D' \setminus D) \setminus [D \cup Z]$ to the left of Z (as one looks in the direction of v), and we may assume that z is in fact the first such site to the left of Z . Since $z \in D' \setminus D$, we have

$$\langle z - a_v, v \rangle = 0 \quad \text{and} \quad \langle z - a_{w_2}, w_2 \rangle < 0.$$

The directions v and w_2 are consecutive in $\mathcal{T} \cup \mathcal{S} \cup \mathcal{Q}$, so by Lemma 5.3 there exists $X \in \mathcal{U}$ such that

$$X \subset (\mathbb{H}_v \cup \ell_v) \cap (\mathbb{H}_{w_2} \cup \ell_{w_2}).$$

Hence,

$$\langle x, v \rangle \leq 0 \quad \text{and} \quad \langle x, w_2 \rangle \leq 0 \tag{5.4}$$

for all $x \in X$. From these it follows that if $y \in z + X$, then

$$\langle y - a_v, v \rangle \leq 0 \quad \text{and} \quad \langle y - a_{w_2}, w_2 \rangle < 0,$$

and so $y \in D' \setminus D$. Moreover, y does not lie to the left of z on the line $D' \setminus D$, since if it did then we would have

$$\langle z + x, w_2 \rangle > \langle z, w_2 \rangle,$$

where $x = y - z \in X$, which would imply that $\langle x, w_2 \rangle > 0$, contradicting (5.4). But we have proved that $z + X \subset [D \cup Z]$. Therefore, $z \in [D \cup Z]$, which contradicts our assumption that $z \notin [D \cup Z]$. □

5.3. Growth of quasi-droplets

Our final task of this section is to construct a sequence of \mathcal{T} -droplets (for some $\mathcal{T} \subset S^1$) that we can use as a framework for growth in the proof of the upper bound of Theorem 1.4 and in the proof of Theorem 1.5.

In this subsection we shall assume that \mathcal{U} is either supercritical or critical. We fix an open semicircle C as follows:

- if \mathcal{U} is supercritical then we choose C such that $\mathcal{S} \cap C = \emptyset$;
- if \mathcal{U} is critical then instead we choose C such that $\mathcal{S} \cap C$ does not contain any strongly stable directions.

In either case, the existence of such a semicircle follows from the definition of supercritical or critical. Let the left and right endpoints of C be u^l and u^r respectively, let the midpoint of C be u^+ , and define

$$\mathcal{S}_U := ((\mathcal{S} \cup \mathcal{Q}) \cap C) \cup \{u^l, u^r, -u^+\} \quad \text{and} \quad \mathcal{S}'_U := ((\mathcal{S} \cup \mathcal{Q}) \cap C), \tag{5.5}$$

where \mathcal{Q} is a set of quasi-stable directions given by Lemma 5.3. In this subsection, as well as in Sections 6 and 7, all droplets will be \mathcal{S}_U -droplets.

In our construction, the growth of \mathcal{S}_U -droplets will take place predominantly in the u^+ direction, so it will be convenient to rotate the coordinate axes so that $(x, y) := xu^+ + yu^l$ for $x, y \in \mathbb{R}$.

Let $R((a, b), (c, d))$ be the discrete rectangle with opposite corners at (a, b) and (c, d) , so that

$$R((a, b), (c, d)) := \{(x, y) \in \mathbb{Z}^2 : x, y \in \mathbb{R}, a \leq x \leq c, b \leq y \leq d\}.$$

Choose $\lambda > 0$ sufficiently large and choose $\mu > 0$ sufficiently large with respect to λ . (In the applications, λ will always be a constant, but in certain cases μ will be a function of




Figure 5. The rectangle R , the set S , and the \mathcal{S}_U -droplets $(D_m)_{m \geq 0}$.

p , tending to infinity as $p \rightarrow 0$.) Let

$$R := R((0, 0), (\lambda, \mu)) \quad \text{and} \quad S := \bigcup_{x \geq 0} R((0, 0), (x, \mu)). \tag{5.6}$$

Next we define the sequence of \mathcal{S}_U -droplets. For each integer $m \geq 0$, let

$$D_m := S \cap \bigcap_{u \in \mathcal{S}'_U} \mathbb{H}_u(a_u + md_u u), \tag{5.7}$$

where the sets $\{a_u \in \mathbb{Z}^2 : u \in \mathcal{S}'_U\}$ and $\{d_u > 0 : u \in \mathcal{S}'_U\}$ are chosen subject to the following constraints:

- (i) $D_0 \subset R$;
- (ii) for every consecutive pair of directions $u, v \in \mathcal{S}'_U$ there exists a line $L_u^l = L_v^r$ that is parallel to u^+ and intersects S , and is such that, for every $m \in \mathbb{Z}$, the intersection⁴ of $\ell_u + a_u + md_u u$ and $\ell_v + a_v + md_v v$ lies on $L_u^l = L_v^r$;
- (iii) for every $u \in \mathcal{S}'_U$ and $m \geq 0$, the u -side of D_m has cardinality $\Omega(\mu)$.

Note that these constraints do not uniquely specify the D_m , but that they are satisfied by at least one choice of the parameters $\{a_u : u \in \mathcal{S}'_U\}$ and $\{d_u : u \in \mathcal{S}'_U\}$. Any choice of parameters satisfying the constraints will be sufficient for what follows. The droplets are illustrated in Figure 5.

Now that we have constructed the sequence of droplets, we require one final set of definitions. We shall use these to give sufficient conditions for growth to occur from D_{m-1}

⁴ These lines are discrete and therefore may have empty intersection (despite not being parallel). In this case we mean the intersection of the corresponding continuous lines.

to D_m , for each m . For each $u \in S'_U$ and $m \in \mathbb{N}$, define

$$I(u, m) := \{i \in \mathbb{Z} : \ell_u(i) \cap D_{m-1} = \emptyset \text{ and } \ell_u(i) \cap D_m \neq \emptyset\}.$$

Given $u \in S'_U$ and $i \in \mathbb{Z}$, a u -block Z is said to be a (u, i) -block if $Z \subset \ell_u(i)$. A u -block Z is a *suitable* (u, i) -block if it is a (u, i) -block that lies entirely between L_u^r and L_u^l , and at least distance λ from both, where $\lambda > 0$ is sufficiently large.

The following is the main result of this subsection, and it is our main deterministic result about the growth of quasi-droplets. We shall use it in the next section to prove the upper bound of Theorem 1.4 for critical families, and in Section 7 to prove Theorem 1.5 for supercritical families.

Lemma 5.5. *Let $m \in \mathbb{N}$ be such that $D_m \setminus R \neq \emptyset$, and for each $u \in S'_U$ and each $i \in I(u, m)$, let $Z(u, i)$ be a suitable (u, i) -block. Then*

$$D_m \subset \left[R \cup D_{m-1} \cup \bigcup_{u \in S'_U} \bigcup_{i \in I(u, m)} Z(u, i) \right].$$

Proof. For each $u \in S'_U$ and $v, w \in S_U$ such that v, u, w are consecutive in S_U , define

$$\rho(u, i) := \ell_u(i) \cap \mathbb{H}_v(a_v + (m - 1)d_v) \cap \mathbb{H}_w(a_w + (m - 1)d_w),$$

where we set $d_{u^l} = d_{u^r} = 0$. Also, for each $u \in S'_U$, let $i(u) := \min I(u, m)$. Thus $\rho(u, i(u))$ is the set of elements of the first line parallel to ℓ_u that lies outside of D_{m-1} , restricted to the line segment between the v - and w -sides of D_{m-1} when those sides are extended. The situation is illustrated in Figure 6.

We claim that the inclusion

$$\rho(u, i(u)) \subset [R \cup D_{m-1} \cup Z(u, i(u))] \tag{5.8}$$

follows from Lemma 5.4. Indeed, setting the set of directions \mathcal{T} , the \mathcal{T} -droplet D , the direction v , and the set Z in the statement of Lemma 5.4 equal to S_U , D_{m-1} , u , and $Z(u, i(u))$, respectively, it follows that the set $\rho(u, i(u))$ we have just defined is equal to the set $D' \setminus D$ in Lemma 5.4, and that (5.8) holds.

By induction, after iterating the previous step $|I(u, m)|$ times, we have

$$\bigcup_{i \in I(u, m)} \rho(u, i) \subset \left[R \cup D_{m-1} \cup \bigcup_{i \in I(u, m)} Z(u, i) \right].$$

Since $u \in S'_U$ was arbitrary, it follows that

$$\bigcup_{u \in S'_U} \bigcup_{i \in I(u, m)} \rho(u, i) \subset \left[R \cup D_{m-1} \cup \bigcup_{u \in S'_U} \bigcup_{i \in I(u, m)} Z(u, i) \right].$$

This only leaves the remaining parallelograms of sites at the corners of droplet, as shown in Figure 6. Suppose u and v are consecutive directions in S'_U , and let P be the parallelogram at the corner between the u - and v -sides of D_{m-1} and D_m , also as in Figure 6. Let $X \in \mathcal{U}$ be a rule such that the assertion of Lemma 5.3 holds with directions u and v . Choose an




Figure 6. An illustration of the growth mechanism in Lemma 5.5.

arbitrary unit vector $w \in S^1$ such that $\theta(u) < \theta(w) < \theta(v)$, and order the elements x of P in increasing order of $\langle x, w \rangle$ (with ties resolved arbitrarily). Then X can be used to infect the elements of P one-by-one in this order, as required. □

6. The upper bound for critical families

In this section we shall use Lemma 5.5 to prove the upper bound of Theorem 1.4 for critical update families. Throughout we assume that \mathcal{U} is critical, that C is the open semicircle specified at the start of Section 5.3, and that \mathcal{S}_U and \mathcal{S}'_U are as defined in (5.5). Thus, $\mathcal{S} \cap C$ consists only of isolated stable directions, and therefore there exists a u -block for every $u \in \mathcal{S}'_U$. For each $u \in \mathcal{S}'_U$, let $\alpha_2(u)$ be the minimum cardinality of a u -block. Let

$$\alpha_2 := \max\{\alpha_2(u) : u \in \mathcal{S}'_U\}. \tag{6.1}$$

We shall prove the upper bound of Theorem 1.4 in the following form. Recall once again that $\tau = \min\{t : \mathbf{0} \in A_t\}$.

Theorem 6.1. *Let \mathcal{U} be a critical update family, let $\epsilon > 0$, and let*

$$p \geq \left(\frac{1}{\log t}\right)^{1/(\alpha_2 + \epsilon)}.$$

Then $\tau \leq t$ with high probability as $t \rightarrow \infty$.

Remark 6.2. The bound on p in Theorem 6.1 could be replaced by

$$p \geq \frac{\lambda(\log \log t)^2}{(\log t)^{1/\alpha_2}},$$

for a sufficiently large constant λ . However, since the bounds are likely to be far from optimal anyway, we prefer to prove the theorem with the weaker bound for reasons of simplicity.

Remark 6.3. For many natural update families, such as the family for the 2-neighbour model, there exists an arc $C \subset S^1$ of length strictly greater than π (that is, an interval strictly larger than a semicircle) in which there are no strongly stable directions. For such families, provided the definitions of \mathcal{S}'_U and α_2 are changed appropriately to take account of the new set C , it is possible to modify our arguments to show that Theorem 6.1 holds with the bound on p replaced by

$$p \geq \left(\frac{\lambda}{\log t} \right)^{1/\alpha_2},$$

for a sufficiently large constant λ . In certain cases this would result in matching (up to a constant factor) bounds in Theorems 4.1 and 6.1 (recovering, for example, the result of [1]). The necessary modifications to the arguments would be akin to the difference between the (asymptotically one-dimensional) growth in the upper bound of van Enter and Hulshof [21] and the (two-dimensional) growth in the upper bound of Aizenman and Lebowitz [1]. We remark that in an earlier version of this paper (version 2 on the arXiv), these stronger bounds for Theorem 6.1 were proved explicitly.

We now begin the build-up to the proof of Theorem 6.1. First we shall set out the definitions of the various quantities and shapes that we need in order to apply the framework of the previous section to the setting of critical update families. After we have done that, we shall explain how the growth mechanism will work.

Recall that u^l is the left endpoint of C . Observe that u^l is right-isolated, and therefore there exists a u^l -right-block, by Lemma 5.2. Define β to be the minimum cardinality of a u^l -right-block, and define also the following quantities:

$$\begin{aligned} \mu_1(p) &= p^{-\alpha_2 - \epsilon}, & v_1(p) &= p^{-\lambda}, \\ \text{and } \mu_2(p) &= p^{-\lambda + 2\beta}, & v_2(p) &= \exp(p^{-\lambda/3}), \end{aligned}$$

where $\epsilon > 0$ is arbitrary, and $\lambda > 0$ is sufficiently large. Observe that we have the inequalities

$$\mu_1(p) \ll \mu_2(p) \ll v_1(p) \ll v_2(p).$$

Recall that we have rotated the coordinate axes so that $(x, y) = xu^+ + yu^l$. Let $R^{(1)}$ and $R^{(2)}$ be the following rectangles:

$$R^{(1)} := R((0, 0), (\lambda, \mu_1(p))) \quad \text{and} \quad R^{(2)} := R((v_1(p), 0), (v_1(p) + \lambda, \mu_1(p) + \mu_2(p))).$$




Figure 7. The various subsets of \mathbb{Z}^2 involved in the growth mechanism, including the rectangles $R^{(1)}$ and $R^{(2)}$, the S_U -droplets $D^{(1)}$ and $D^{(2)}$, and the triangle T .

We shall use the construction for the droplets $(D_m)_{m \in \mathbb{Z}}$ from the previous section twice: once with $R = R^{(1)}$ and once with $R = R^{(2)}$. Thus, for each $i \in \{1, 2\}$ and each integer $m \geq 0$, let $D_m^{(i)} := D_m$, where $(D_m)_{m \geq 0}$ is the sequence of droplets defined in (5.7) obtained if we set $\mu = \mu_i(p)$, and, in the case $i = 2$, if we translate appropriately so that $R = R^{(2)}$ (see Figure 7). Furthermore, let

$$D^{(1)} := D_{\lambda \cdot v_1(p)}^{(1)} \quad \text{and} \quad D^{(2)} := D_{v_2(p)}^{(2)}.$$

The final subset of \mathbb{Z}^2 that we need to give a name to is the triangle

$$T := \{(x, y) \in \mathbb{Z}^2 : x, y \in \mathbb{R}, 0 \leq x \leq v_1(p) \text{ and } 0 \leq y - \mu_1(p) \leq \lceil p^{2\beta} x \rceil\}.$$

We shall use this triangle to grow in the direction of u^l .

The growth mechanism that we shall use to prove Theorem 6.1 is the following.

- (i) First, we show that there is an internally filled copy of $R^{(1)}$ such that the corresponding copy of $D^{(2)}$ contains the origin.
- (ii) Second, given an internally filled copy of $R^{(1)}$, we use Lemma 5.5 to show that the infection spreads in the direction of u^+ to fill $D^{(1)}$.
- (iii) Next, we show that the infection spreads upwards through T (that is, in the direction of u^l) to fill T and $R^{(2)}$. Since u^l may be at the end-point of an interval in \mathcal{S} , we cannot assume that there exist u^l -blocks, merely that there exist u^l -right-blocks. Thus, it could be that, as the infection spreads through T row-by-row in the direction of u^l , one always has to look for a u^l -right-block to the right of the u^l -right-block found in the previous row.
- (iv) Finally, as in step (ii), we use Lemma 5.5 to show that the infection spreads rightwards from $R^{(2)}$ to fill $D^{(2)}$, and hence infect the origin.

In Lemma 6.4 we show that the events described in (ii) and (iv) each occur with high probability, and in Lemma 6.5 we show that the event described in (iii) occurs with high probability. Thus, the only unlikely event is that we find the internally filled copy of $R^{(1)}$ suitably close to the origin. It is therefore reasonable to think of $R^{(1)}$ as being a ‘critical droplet’.

Lemma 6.4. For each $i \in \{1, 2\}$, the event

$$\{D^{(i)} \subset [R^{(i)} \cup (D^{(i)} \cap A)]\} \tag{6.2}$$

occurs with high probability as $p \rightarrow 0$.

Proof. Fix $i \in \{1, 2\}$. Then for every $m \geq 0$ such that $D_m^{(i)} \neq \emptyset$, we have

$$\mathbb{P}_p(D_m^{(i)} \subset [R^{(i)} \cup D_{m-1}^{(i)} \cup (D_m^{(i)} \cap A)]) \geq (1 - (1 - p^{\alpha_2})^{\Omega(\mu_i(p))})^{O(1)},$$

by Lemma 5.5. (The quantity $\Omega(\mu_i(p))$ in the exponent comes from partitioning the u -side of $D_m^{(i)}$ into segments of length α_2 , for each $u \in S'_U$, and then using condition (iii) of the definition of D_m immediately after (5.7).) Hence,

$$\begin{aligned} \mathbb{P}_p(D^{(i)} \subset [R^{(i)} \cup (D^{(i)} \cap A)]) &\geq (1 - (1 - p^{\alpha_2})^{\Omega(\mu_i(p))})^{O(v_i(p))} \\ &\geq \exp(-O(v_i(p) \exp(-\Omega(p^{\alpha_2} \mu_i(p))))). \end{aligned} \tag{6.3}$$

Now, if $i = 1$ then (6.3) is equal to

$$\exp(-O(p^{-\lambda} \exp(-\Omega(p^{-\epsilon}))) = 1 - o(1),$$

since $\mu_1(p) = p^{-\alpha_2 - \epsilon}$ and $v_1(p) = p^{-\lambda}$. On the other hand, if $i = 2$ then (6.3) is equal to

$$\exp(-O(\exp(p^{-\lambda/3}) \exp(-\Omega(p^{-\lambda/2}))) = 1 - o(1),$$

since $\mu_2(p) = p^{-\lambda + 2\beta}$ and $v_2(p) = \exp(p^{-\lambda/3})$. In either case, (6.2) holds with high probability as $p \rightarrow 0$. □

Lemma 6.5. The event

$$\{R^{(2)} \subset [D^{(1)} \cup (T \cap A)]\}$$

occurs with high probability as $p \rightarrow 0$.

Proof. For each $i \in \mathbb{Z}$, let U_i be the leftmost $p^{-2\beta}$ sites of $\ell_{u^l}(i) \cap T$, and let I be the set of $i \in \mathbb{Z}$ for which U_i is non-empty. For each $i \in I$, let $Z_i \subset U_i$ be a set of β consecutive sites contained in the middle $p^{-2\beta}/2$ sites of U_i .

We claim that

$$R^{(2)} \subset \left[D^{(1)} \cup \bigcup_{i \in I} Z(i) \right]. \tag{6.4}$$

To see this, first let $i_0 := \min I$, and observe that $[D^{(1)} \cup Z(i_0)]$ contains all sites of $\ell_{u^l}(i_0)$ to the right of $Z(i_0)$, at least until within $O(1)$ distance from the right-hand end of the u^l -side of $D^{(1)}$, since $Z(i_0)$ is a u^l -right-block. Losing $O(1)$ sites along each line in this way, we observe that (6.4) easily holds, because the number of horizontal rows (that is, lines parallel to ℓ_{u^l}) intersecting T is at most $O(p^{-\lambda + 2\beta})$, and this is much less than $\Omega(p^{-\lambda})$, which is a lower bound for the number of vertical columns of sites (that is, lines parallel to ℓ_{u^+}) intersecting $D^{(1)}$ to the right of $R^{(2)}$. (For an illustration of the situation, see Figure 7, and in particular the dashed line emerging from the right-hand end of the u^l -side of $D^{(1)}$, which need not be parallel to any of the sides of $D^{(1)}$.)

Hence, to prove the lemma, it is sufficient to show that, with high probability, the middle $p^{-2\beta}/2$ sites of U_i contain a set of β consecutive sites, all contained in A , for every $i \in I$. The probability that this fails is at most

$$p^{-\lambda+2\beta} \cdot (1 - p^\beta)^{\Omega(p^{-2\beta})} \leq p^{-\lambda+2\beta} \cdot \exp(-\Omega(p^{-\beta})) = o(1)$$

as $p \rightarrow 0$, as required. □

We are now ready to complete the proof of Theorem 6.1.

Proof of Theorem 6.1. Let

$$p = \left(\frac{1}{\log t} \right)^{1/(\alpha_2 + \epsilon)}.$$

We shall show that $\tau \leq t$ with high probability as $t \rightarrow \infty$.

It will be convenient to ‘sprinkle’ the probability in two rounds, in order to maintain independence. Formally, let A' and A'' be two independent p -random subsets of \mathbb{Z}^2 , and let $A := A' \cup A''$. Thus, sites in A are infected with probability $2p - p^2$, rather than p ; this is valid (if an abuse of notation) because ϵ was chosen arbitrarily. We shall use the first round to find a suitably positioned internally filled copy of $R^{(1)}$, and the second round to show that this copy of $R^{(1)}$ grows to infect the origin by time t .

There are at least $\Omega(v_2(p))$ sites $x \in \mathbb{Z}^2$ such that the sets $x + R^{(1)}$ are disjoint and $\mathbf{0} \in x + D^{(2)}$. The probability that for no such x do we have $x + R^{(1)} \subset A'$ is at most

$$(1 - p^{O(\mu_1(p))})^{\Omega(v_2(p))} \leq \exp(-\Omega(\exp(-p^{-\alpha_2 - \epsilon}) \exp(p^{-\lambda/3}))) = o(1),$$

since $\mu_1(p) = p^{-\alpha_2 - \epsilon}$ and $v_2(p) = \exp(p^{-\lambda/3})$.

So with high probability there exists an x such that $\mathbf{0} \in x + D^{(2)}$ and $x + R^{(1)} \subset A'$. Fix one such x , and for notational simplicity let us translate our notation by $-x$ so that $\mathbf{0} \in D^{(2)}$ and $R^{(1)} \subset A'$. By Lemmas 6.4 and 6.5, the event

$$\{D^{(2)} \subset [R^{(1)} \cup ((D^{(1)} \cup T \cup D^{(2)}) \cap A'')]\}$$

occurs with high probability.

It only remains to show that $\mathbf{0}$ is infected by time t . But

$$|D^{(1)} \cup T \cup D^{(2)}| \leq p^{-O(1)} \exp(p^{-\lambda/3}) < \sqrt{t},$$

so even if the sites are infected one-by-one, we still have $\mathbf{0} \in A_t$, which completes the proof of the theorem. □

7. Supercritical families

Recall that supercritical families are those for which there exists an open semicircle in S^1 with which the stable set has empty intersection. Our aim in this short section of the paper is to prove Theorem 1.5. Combined with Theorem 1.4, this provides a justification for the distinction made between supercritical and critical update families in Definition 1.3. Using the same methods as those in the proof of Theorem 1.5, we also establish a deterministic

result, Theorem 7.1, which says that an update family is supercritical if and only if there exist finite subsets of \mathbb{Z}^2 with infinite closure.

Almost all of the work that goes into the proof of Theorem 1.5 has already been completed in Section 5. The only further observation required is that if C is an open semicircle such that $\mathcal{S} \cap C = \emptyset$, then the empty-set is a u -block for every $u \in C$. Thus, a sufficient condition for $\mathbf{0} \in A_t$ is that a sufficiently large constant-sized rectangle of initially infected sites is ‘suitably located’ within distance t of the origin. (The only important detail about rectangles being ‘suitably located’ is that there are at least $t^{\Omega(1)}$ disjoint such rectangles – in fact, there are at least $\Omega(t)$.)

Proof of Theorem 1.5. Let \mathcal{U} be supercritical and let p satisfy $p \geq t^{-\epsilon}$, for some sufficiently small $\epsilon > 0$. We shall show that $\tau \leq t$ with high probability as $t \rightarrow \infty$.

Let \mathcal{Q} be a set of quasi-stable directions given by Lemma 5.3, let $C \subset S^1$ be the semicircle specified at the start of Section 5.3, and let \mathcal{S}_U and \mathcal{S}'_U be as specified in (5.5). If $u \in \mathcal{S}'_U$, then we have $u \in \mathcal{Q} \setminus \mathcal{S}$, because $\mathcal{S} \cap C = \emptyset$. Therefore u is a quasi-stable direction, but not a stable direction, and so the empty set is a u -block.

Now let R and S be the sets defined in (5.6), with $\mu > \lambda > 0$ being sufficiently large constants, and let the \mathcal{S}_U -droplets D_1, D_2, \dots be as defined subsequently in (5.7). Then

$$D_m \subset [R] \tag{7.1}$$

for every $m \in \mathbb{N}$, by Lemma 5.5 and the observation we have just made that the empty set is a u -block for every $u \in \mathcal{S}'_U$.

Let $\delta > 0$ be a sufficiently small constant. There are at least $\Omega(t)$ sites $x \in \mathbb{Z}^2$ such that $\|x\| \leq \delta t$, the rectangles $x + R$ are disjoint, and $\mathbf{0} \in x + S$. The probability that we do not have $x + R \subset A$ for any such x is at most

$$(1 - p^{O(1)})^{\Omega(t)} \leq \exp(-p^{O(1)}\Omega(p^{-1/\epsilon})) \rightarrow 0$$

as $t \rightarrow \infty$, since ϵ is sufficiently small. So with high probability there exists $x \in \mathbb{Z}^2$ such that $\|x\| \leq \delta t$, $\mathbf{0} \in x + S$ and $x + R \subset A$. Suppose x has this property. Then $x + D_m \subset [(x + D_m) \cap A]$ for every $m \in \mathbb{N}$, by (7.1). Now take m minimal such that $\mathbf{0} \in x + D_m$. Then $\tau \leq |D_m|$, even if sites in D_m are infected one-by-one, and also $|D_m| < t$, because δ was chosen sufficiently small. Thus $\tau \leq t$, and this completes the proof. □

The next result, which is another easy consequence of Lemma 5.5 (or more specifically, of (7.1)), is an equivalence between supercriticality and the existence of finite sets with infinite closures.

Theorem 7.1. *A two-dimensional bootstrap percolation update family \mathcal{U} is supercritical if and only if there exists a finite set $K \subset \mathbb{Z}^2$ such that $[K]$ is infinite.*

Proof. One implication follows from (7.1). The other is a triviality: if \mathcal{U} is not supercritical, then every open semicircle intersects the stable set. Therefore there exist stable directions u_1, u_2 and u_3 such that the interior of their convex hull contains the origin,

and hence $\{u_1, u_2, u_3\}$ -droplets are finite. Thus, if $K \subset \mathbb{Z}^2$ is an arbitrary finite set, then $[K] \subset D$, where D is the minimal $\{u_1, u_2, u_3\}$ -droplet containing K . \square

Stronger statements than Theorem 7.1 are available for update families with zero or one stable direction(s). The proofs of these statements, which are omitted (but which may be found in version 2 of this paper on the arXiv), are further easy applications of the quasi-stability method.

Theorem 7.2. *Let \mathcal{U} be a two-dimensional bootstrap percolation update family.*

- (i) *There exists a finite set $K \subset \mathbb{Z}^2$ such that $[K] = \mathbb{Z}^2$ if and only if $S = \emptyset$.*
- (ii) *Let $u \in S^1$ be a unit vector. Then there exists a finite set $K \subset \mathbb{Z}^2$ such that $\mathbb{H}_u \subset [K]$ if and only if $S \setminus \{u\} = \emptyset$.* \square

8. Open problems and conjectures

In this paper we have introduced a new, unified model of bootstrap percolation on the square lattice. Owing to the novelty of this general model, we quite naturally end the paper with a number of open problems, of which there are three broad types.

Subcritical families in two dimensions. At present we are not able to prove any results about subcritical families. In particular, we cannot prove that the distinction we have made between critical and subcritical families is the right one, although we conjecture that it is in the following strong sense.⁵

Conjecture 8.1. *Let \mathcal{U} be a two-dimensional subcritical bootstrap percolation update family. Then there exists $p > 0$ such that*

$$\mathbb{P}_p([A] = \mathbb{Z}^2) = 0.$$

This conjecture lies in sharp contrast to Corollary 1.6, which says that if \mathcal{U} is *not* subcritical then for every $p > 0$ we have $\mathbb{P}_p([A] = \mathbb{Z}^2) = 1$. (Note that we always have $\mathbb{P}_p([A] = \mathbb{Z}^2) \in \{0, 1\}$ owing to a standard 0–1 law.)

Sharper results for critical families in two dimensions. As discussed in Remark 6.3, for certain update families our methods could be used with very minor modifications to obtain matching (up to a constant factor) bounds in Theorems 4.1 and 6.1, and thus to determine $p_c(\mathbb{Z}_n^2, \mathcal{U})$ up to a constant factor. However, for many other update families our methods (as they stand) do not give matching upper and lower bounds, and so for those families the order of magnitude of $p_c(\mathbb{Z}_n^2, \mathcal{U})$ remains an open question. Moreover, it is known [16, 21] that the order of magnitude of $p_c(\mathbb{Z}_n^2, \mathcal{U})$ is not always a negative power of $\log n$. It is likely to be a challenging problem to determine $p_c(\mathbb{Z}_n^2, \mathcal{U})$ up to a constant factor in general.

⁵ As noted earlier, since the submission of this paper, Balister, Bollobás, Przykucki and Smith [2] have proved this conjecture.

Higher dimensions. It is natural to ask what can be said about the large-scale behaviour of \mathcal{U} -bootstrap percolation in dimensions $d \geq 3$. The first step towards answering this question is likely to be to give a classification of the models along the lines of Definition 1.3. This will not be a straightforward generalization of the two-dimensional definition, for the following reason. When $d = 3$, the critical probability of the 2-neighbour model is $\Theta(\log n)^{-2}$ [1, 5], while the critical probability of the 3-neighbour model is $\Theta(\log \log n)^{-1}$ [9, 5]. By contrast, in two dimensions we have seen that the critical probabilities of all critical models are of the form $(\log n)^{-\Theta(1)}$, while the critical probabilities of non-critical models have either been proved (in the case of supercritical) or conjectured (in the case of subcritical) to be non-logarithmic. Thus, in three dimensions, the classification of critical \mathcal{U} -bootstrap models ought to divide further: at least into those for which the critical probability is $(\log n)^{-\Theta(1)}$ and those for which it is $(\log \log n)^{-\Theta(1)}$. However, this is of course only the start of the story: once that has been done, the question becomes to determine the critical probability up to a constant factor for all ‘critical’ models in all dimensions. We anticipate that this will be a significant area of future research.

Acknowledgements

The first author is partially supported by NSF grant DMS 1301614 and MULTIPLEX no. 317532, the second author by a CNPq bolsa PDJ, and the third author by the Knut and Alice Wallenberg Foundation. The authors are grateful to Yuval Peres and other members of the Theory Group at Microsoft Research, Redmond, where much of the research in this paper was carried out. The authors are also greatly indebted to the anonymous referee for his or her many detailed comments, which have considerably improved the presentation of this paper and allowed us to greatly streamline many of the proofs.

References

- [1] Aizenman, M. and Lebowitz, J. L. (1988) Metastability effects in bootstrap percolation. *J. Phys. A* **21** 3801–3813.
- [2] Balister, P., Bollobás, B., Przykucki, M. J. and Smith, P. J. Subcritical \mathcal{U} -bootstrap percolation models have non-trivial phase transitions. *Trans. Amer. Math. Soc.* In press. [arXiv:1311.5883](https://arxiv.org/abs/1311.5883)
- [3] Balogh, J. and Bollobás, B. (2006) Bootstrap percolation on the hypercube. *Probab. Theory Rel. Fields* **134** 624–648.
- [4] Balogh, J., Bollobás, B., Duminil-Copin, H. and Morris, R. (2012) The sharp threshold for bootstrap percolation in all dimensions. *Trans. Amer. Math. Soc.* **364** 2667–2701.
- [5] Balogh, J., Bollobás, B. and Morris, R. (2009) Bootstrap percolation in three dimensions. *Ann. Probab.* **37** 1329–1380.
- [6] Balogh, J., Bollobás, B. and Morris, R. (2009) Majority bootstrap percolation on the hypercube. *Combin. Probab. Comput.* **18** 17–51.
- [7] Balogh, J., Peres, Y. and Pete, G. (2006) Bootstrap percolation on infinite trees and non-amenable groups. *Combin. Probab. Comput.* **15** 715–730.
- [8] Biskup, M. and Schonmann, R. H. (2009) Metastable behavior for bootstrap percolation on regular trees. *J. Statist. Phys.* **136** 667–676.
- [9] Cerf, R. and Cirillo, E. N. M. (1999) Finite size scaling in three-dimensional bootstrap percolation. *Ann. Probab.* **27** 1837–1850.

- [10] Cerf, R. and Manzo, F. (2002) The threshold regime of finite volume bootstrap percolation. *Stochastic Process. Appl.* **101** 69–82.
- [11] Chalupa, J., Leath, P. L. and Reich, G. R. (1979) Bootstrap percolation on a Bethe lattice. *J. Phys. C* **12** L31–L35.
- [12] Duarte, A. M. S. (1989) Simulation of a cellular automaton with an oriented bootstrap rule. *Physica A* **157** 1075–1079.
- [13] Gravner, J. and Griffeath, D. (1996) First passage time for threshold growth dynamics on \mathbb{Z}^2 . *Ann. Probab.* **24** 1752–1778.
- [14] Gravner, J. and Griffeath, D. (1999) Scaling laws for a class of critical cellular automaton growth rules. In *Proc. Erdős Center Workshop on Random Walks*, pp. 167–188.
- [15] Holroyd, A. E. (2003) Sharp metastability threshold for two-dimensional bootstrap percolation. *Probab. Theory Rel. Fields* **125** 195–224.
- [16] Mountford, T. S. (1995) Critical length for semi-oriented bootstrap percolation. *Stochastic Process. Appl.* **56** 185–205.
- [17] Schonmann, R. H. (1990) Finite size scaling behavior of a biased majority rule cellular automaton. *Physica A* **167** 619 – 627.
- [18] Schonmann, R. H. (1992) On the behavior of some cellular automata related to bootstrap percolation. *Ann. Probab.* **20** 174–193.
- [19] Ulam, S. (1950) Random processes and transformations. In *Proc. Internat. Congr. Math.*, pp. 264–275.
- [20] van Enter, A. C. D. (1987) Proof of Straley’s argument for bootstrap percolation. *J. Statist. Phys.* **48** 943–945.
- [21] van Enter, A. C. D. and Hulshof, W. J. T. (2007) Finite-size effects for anisotropic bootstrap percolation: logarithmic corrections. *J. Statist. Phys.* **28** 1383–1389.
- [22] von Neumann, J. (1966) *Theory of Self-Reproducing Automata*, University of Illinois Press.