

A Note on Bootstrap Percolation Thresholds in Plane Tilings using Regular Polygons

Neal Bushaw*

Daniel W. Cranston[†]

March 29, 2019

Abstract

In *k*-bootstrap percolation, we fix $p \in (0, 1)$, an integer k , and a plane graph G . Initially, we infect each face of G independently with probability p . Infected faces remain infected forever, and if a healthy (uninfected) face has at least k infected neighbors, then it becomes infected. For fixed G and p , the *percolation threshold* is the largest k such that eventually all faces become infected, with probability at least $1/2$. For many infinite graphs, we show that this threshold is independent of p .

We consider bootstrap percolation in tilings of the plane by regular polygons. A *vertex type* in such a tiling is the cyclic order of the faces that meet a common vertex. First, we determine the percolation threshold for each of the Archimedean lattices. More generally, let \mathcal{T} denote the set of plane tilings T by regular polygons such that if T contains one instance of a vertex type, then T contains infinitely many instances of that type. We show that no tiling in \mathcal{T} has threshold 4 or more. Further, the only tilings in \mathcal{T} with threshold 3 are four of the Archimedean lattices. Finally, we describe a large subclass of \mathcal{T} with threshold 2.

vertex type

1 Introduction

In *k*-bootstrap percolation, we fix $p \in (0, 1)$, an integer k , and a plane graph G . Initially, we infect each face of G independently with probability p ; call the set of initially infected faces \mathcal{I} . Infected faces remain infected forever, and if a healthy (uninfected) face has at least k infected neighbors, then it becomes infected. We say that \mathcal{I} *percolates* if eventually all faces become infected. For short, we call this the *k*-bootstrap model. For fixed G and p , the *percolation threshold*¹, or simply *threshold*, is the largest k such that in the *k*-bootstrap model \mathcal{I} percolates with probability at least $1/2$. For a large class of infinite graphs, we show that the threshold is independent of p .

\mathcal{I} percolates
k-bootstrap
model
threshold

The *k*-bootstrap model has a long, rich history. Introduced by Chalupa, Leath, and Reich [11] in 1979 as a way to model magnetic materials, it is an example of a monotone cellular automata (introduced by von Neumann [18] in 1966). Most of the work in this field has focused on finding thresholds for growing families of graphs. For example, if we infect each face of the $n \times n$ square grid independently with some probability p , how large must p be so the infection percolates almost surely, as $n \rightarrow \infty$? The answer to this question, and the first sharp result in

*Department of Mathematics and Applied Mathematics, Virginia Commonwealth University, Richmond, VA; nobushaw@vcu.edu

[†]Department of Mathematics and Applied Mathematics, Virginia Commonwealth University, Richmond, VA; dcranston@vcu.edu; This research is partially supported by NSA Grant H98230-16-0351.

¹Note that this is different than the *probability* thresholds often considered for sequences of finite graphs.

the area, was proved by Holroyd [14]. While Holroyd’s result is striking on its own, it has been extended greatly: studying the problem in higher dimensions, finding more terms of the critical probability function, and much more (see, e.g., [1, 2, 3, 9, 10]). These bootstrap models have been generalized significantly in recent years, with the advent of graph bootstrap percolation [8].

Outside the realm of grids, bootstrap percolation has been studied on many different families of graphs. This includes work determining critical probabilities for random regular graphs [6], the Erdős-Renyi random graph $G_{n,p}$ [13, 15], the hypercube [4], infinite trees [5], and others. Largely ignored, however, has been percolation on infinite lattices (aside from the square lattice [16, 17], discussed below). We explore this direction here.

The *length* of a face of a plane graph is its number of sides. A *configuration* is a finite plane graph. A configuration H *appears in* G if there is a map from faces of H to faces of G that preserves both face length and the number of edges shared by every pair of faces. When H appears in G , we also say that G *contains a copy of* H . The following observation is straightforward, but it is our main tool for proving upper bounds on percolation thresholds.

length
configuration
appears in G
contains a
copy of H

Observation 1. *Let C be a configuration such that each face of C has at most k neighboring faces outside of C . If G contains infinitely many copies of C , then G has percolation threshold at most k .*

Proof. Suppose we are in the $(k + 1)$ -bootstrap model. Note that if some copy of C has no initially infected face, then \mathcal{I} does not percolate, since no face in that copy of C ever becomes infected. Since G has infinitely many copies of C (and each face of G is infected independently), with probability 1 at least one copy of C in G has no face initially infected. So, in the $(k + 1)$ -bootstrap model, \mathcal{I} percolates with probability 0. \square

An immediate consequence of Observation 1 is that the (infinite) square lattice has percolation threshold at most 2, since we can take as our configuration C four square faces that meet at a common vertex. van Enter [17] famously proved a matching lower bound. That is, the percolation threshold of the square lattice is 2. In this note, we extend this result, using the same approach, to determine the percolation thresholds for many tilings of the plane by regular polygons. Beyond this, we prove that, somewhat surprisingly, for a large class of graphs (those whose vertex types repeat infinitely often) the percolation threshold is never more than four, and the only tilings achieving this are Archimedean lattices. We also determine a large class of tilings for which the threshold is exactly two.

2 Archimedean Lattices

A function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ on a tiling is a *tiling translation* if it has the form $f : (x, y) \mapsto (x+a, y+b)$ for some $a, b \in \mathbb{R}$ and it maps the center of every d -gon to the center of a d -gon. These are simply translations of the plane which map our polygons to congruent polygons. As an example, for any $a, b \in \mathbb{Z}$, $f : (x, y) \mapsto (x + a, y + b)$ is a tiling translation for the (unit) square lattice, but is not a tiling translation for the hex lattice when a and b are both nonzero, since the height of a regular hexagon, with one side axis-aligned, is not a rational multiple of its width. An event E is called *translation-invariant* if for every initially infected set \mathcal{I} and every tiling translation f we have $f(\mathcal{I}) \in E$ if and only if $\mathcal{I} \in E$. For example, the event $E_1 = \{\mathcal{I} : \mathcal{I} \text{ percolates to the entire plane}\}$ is translation invariant; if a set \mathcal{I} percolates, it will certainly also percolate when that set is translated to another location in the plane, since our percolation process is independent of a face’s location in the plane. At the other extreme, as an example of a *non-translation-invariant* event consider $E_2 = \{\mathcal{I} : \mathcal{I} \text{ infects the origin eventually}\}$. Let $\mathcal{I}_2 = \{\text{only the face containing the origin is infected}\}$. Now $\mathcal{I}_2 \in E$, but for any nontrivial tiling translation f we have $f(\mathcal{I}_2) \notin E_2$. An event is *weakly translation-invariant* if there exist infinitely many distinct tiling translations f such that for each initially infected set \mathcal{I} we have

tiling
translation

translation-
invariant

weakly
translation-
invariant

$f(\mathcal{I}) \in E$ if and only if $\mathcal{I} \in E$. Our main tool for proving lower bounds on percolation thresholds is the following lemma of Kolmogorov about translation invariant events. This result is quite general, so we state it in a simple form which is enough for our purposes.

Kolmogorov’s 0–1 Law. *Let T be an infinite graph that is locally finite. If E is a weakly translation-invariant event, then $\Pr(E) \in \{0, 1\}$.*

The proof is not hard, but requires enough machinery that we do not reproduce it here. Nevertheless, this lemma is crucial to our work, so we give a brief description for the probabilistically-minded reader. Our probability space is constructed as a countable product space (whose fundamental events are whether or not individual faces are infected). As such, any event – in particular, our ‘initial set percolates’ event – can be approximated arbitrarily well by a cylinder set (all hexagons within some fixed distance of a specified hexagon). Since we have weak translation-invariance, if we translate sufficiently far we can find another approximating cylinder set which is disjoint from the first; thus, events within one cylinder set are independent of those within the other. By repeating this process, we find infinitely many disjoint copies of our approximating cylinder set. Each of these translations of our cylinder is initially entirely infected with positive probability. Hence, with probability one, at least one of these approximating events will occur².

To show that the square lattice has threshold 2, Chalupa, Leath, and Reich [11] defined an event A with the following three properties: (1) if A occurs, then in the 2-bootstrap model the initially infected set \mathcal{I} percolates on the square lattice, (2) A occurs with positive probability, and (3) A is translation invariant. Properties (1) and (2) clearly imply that in the 2-bootstrap model on the square lattice, \mathcal{I} percolates with positive probability. Now Kolmogorov’s 0–1 Law shows that this probability is 1.

An *Archimedean Lattice* is a vertex transitive (infinite) plane graph in which each face is a regular polygon. It is well-known that there are 11 such lattices³, including the three regular tilings (by the triangle, square, and hexagon). To describe an Archimedean Lattice, we write $(f_1 \dots f_s)$, where f_1, \dots, f_s are the face lengths, in cyclic order, that meet at each vertex. For instance, the regular tilings by triangle, square, and hexagon are denoted (3.3.3.3.3.3), (4.4.4.4), and (6.6.6). In Figure 1 we show the other eight Archimedean Lattices, along with configurations that bound their percolation thresholds, via Observation 1. Clearly, every lattice has percolation threshold at least 1. For lattices (3.3.3.3.3.3), (3.3.3.3.3.6), (3.3.3.4.4), (3.3.4.3.4), and (3.4.6.4) we obtain a matching upper bound, using Observation 1 and the configurations in Figure 1. For (3.6.3.6) our upper bound is 2, and for each of (3.12.12), (4.6.12), (4.8.8), and (6.6.6) it is 3. So, to determine the bootstrap threshold for each of these five lattices, the interesting work is proving a matching lower bound. We first present a proof for (4.8.8). Since the proofs of all five lower bounds are similar, we will just outline the differences for the remaining four lattices.

Archimedean
Lattice

Theorem 1. *For every $p \in (0, 1)$, the percolation threshold for the lattice (4.8.8) is 3.*

Proof. By applying Observation 1, using the configuration in Figure 1, we get an upper bound of 3. So we need only to prove a matching lower bound.

We draw (4.8.8) so each 8-gon has its center at a lattice point (and each lattice point is the center of an 8-gon). We fix $p \in (0, 1)$, and initially infect each face independently with probability p . Call this set of initially infected faces \mathcal{I} . We show that in the 3-bootstrap model \mathcal{I} percolates with positive probability. Let \mathcal{D}_t denote the set of faces with centers at (x, y) such that $|x| \leq t$ and $|y| \leq t$. So \mathcal{D}_t contains $(2t + 1)^2$ 8-gons and $(2t)^2$ squares.

\mathcal{I}
 \mathcal{D}_t

Suppose that all faces in \mathcal{D}_t are infected. We want to prove a lower bound on the probability that eventually all faces in \mathcal{D}_{t+1} become infected. Suppose that some infected face f is adjacent to a face in the top row of 8-gons of \mathcal{D}_t . Now the infection at f will spread to every face in the

²The details are available, for example, in [7, p. 118ff].

³At the start of Section 3, we outline a proof of this fact.

same row as f that is adjacent to some face in \mathcal{D}_t . This spread happens as follows. First, we infect the two 4-gons that are adjacent to f and also each have two neighbors in \mathcal{D}_t . Now we




Figure 1: The 8 non-regular Archimedean lattices, along with the configurations used to prove upper bounds on their percolation thresholds.

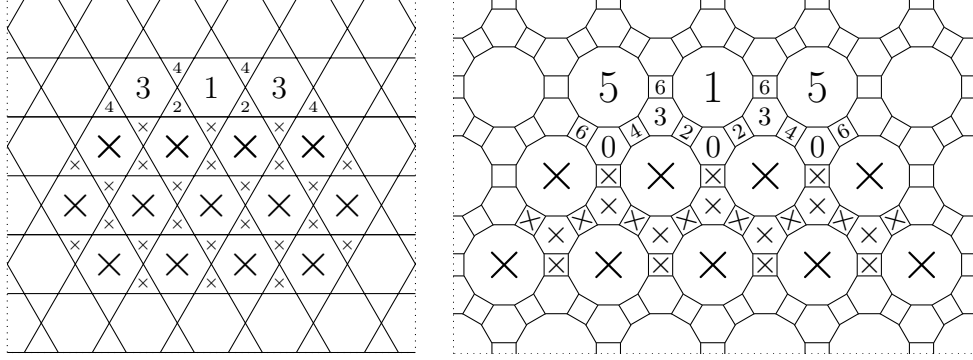


Figure 2: The order in which faces in the next row become infected, when proving that the lattice (3.6.3.6) has threshold 2 and that the lattice (4.6.12) has threshold 3.

infect each 8-gon f' that is adjacent to f and also adjacent to a face \mathcal{D}_t (since f' is also adjacent to an infected 4-gon). By repeating this argument with f' in place of f , we see that the infection spreads along the row above the top row of \mathcal{D}_t (to the full width of \mathcal{D}_t , which is $2t + 1$). Since each face in the row just above \mathcal{D}_t is initially infected with probability p , the probability that none is infected is $(1 - p)^{2t+1}$. The same argument applies to the row beneath the bottom of \mathcal{D}_t and to the columns to the left and right of \mathcal{D}_t . So the probability that at least one of these two rows and two columns has no infected face is at most $4(1 - p)^{2t+1}$. If we infect all of both columns and both rows, then we also infect each square with exactly one neighbor in \mathcal{D}_t (these are at the corners). Finally, we infect each corner 8-gon, since it now has two adjacent infected 8-gons and one adjacent infected square. Thus, all of \mathcal{D}_{t+1} becomes infected.

For each $s \geq 0$, call $\mathcal{D}_{s+1} \setminus \mathcal{D}_s$ a *ring* around \mathcal{D}_0 . We partition the faces of each ring into top and bottom rows, left and right columns, and four corners. To infect the whole plane, it is enough to have all faces in \mathcal{D}_t infected (for some t) and for each $s \geq t$ to have at least one infected face in each of its top and bottom rows and right and left columns. The probability of having at least one ring without the necessary infected faces is at most $\sum_{s \geq t} 4(1-p)^{2s+1} = 4(1-p)^{2t+1}/(1-(1-p)^2)$. For t sufficiently large, this probability is less than 1. The probability that every 8-gon in \mathcal{D}_t is initially infected is $p^{(2t+1)^2}$; if the 8-gons are all infected, then the squares immediately become infected. Since each face is infected independently, the probability of infecting the whole plane is at least $p^{(2t+1)^2}(1 - 4(1-p)^{2t+1})/(1 - (1-p)^2)$, which is positive for t sufficiently large. So, in the 3-bootstrap model, with positive probability, the whole plane becomes infected. Now we use Kolmogorov's 0-1 Law to show that, in fact, the whole plane becomes infected with probability 1. To apply the 0-1 Law, we only need to note that the event that the initial set \mathcal{I} percolates is weakly translation-invariant. \square

The proofs of the lower bounds for (3.6.3.6), (4.6.12), and (6.6.6) are similar. The only noticeable difference is the shapes of the sets analogous to \mathcal{D}_t and the details of how \mathcal{D}_t grows to \mathcal{D}_{t+1} when we have at least one infected face on each side of the ring $\mathcal{D}_{t+1} \setminus \mathcal{D}_t$. In each case, the shape of the set \mathcal{D}_t is closer to a hexagon than a square, so the ring $\mathcal{D}_{t+1} \setminus \mathcal{D}_t$ has six sides, rather than four; this is most obvious for (6.6.6). In Figure 2 we show examples of how one side of \mathcal{D}_t grows to \mathcal{D}_{t+1} for (3.6.3.6) and (4.6.12). The faces marked with \times are already infected, and the integers denote the order that new faces become infected. For (4.6.12), the faces labeled 0 become infected immediately, since each has three infected neighbors.

The fact that (3.12.12) has bootstrap threshold at least 3 follows directly from the fact that (6.6.6) does. We inflate each 12-gon in (3.12.12) to include one third of each incident triangle. This produces (6.6.6), as shown in Figure 3. When we inflate a 12-gon, it does not become incident to any new face. Since (6.6.6) has threshold 3, we conclude that in the 3-bootstrap




Figure 3: Inflating 12-gons to hexagons shows that (3.12.12) has threshold at least 3, since (6.6.6) has threshold 3.

model, with probability 1 every 12-gon in (3.12.12) becomes infected. And once the three 12-gons incident to a triangle become infected, so does the triangle. Thus, (3.12.12) has threshold at least 3. Finally, recall that the lattice (4.4.4.4) has threshold 2. (This was proved by van Enter [17]; it is this proof which inspired the present paper.)

3 More General Tilings

In a plane tiling by regular polygons, the *vertex type* for a vertex v is the cyclicly ordered list of the lengths of faces that meet at v . Since the interior angle of a regular t -gon is known (its measure in degrees is $180(t - 2)/t$), determining the set of all possible vertex types is a simple exercise in diophantine equations. Up to reflection, we have 21 types. These are 3.3.3.3.3.3, 3.3.3.3.6, 3.3.3.4.4, 3.3.4.3.4, 3.3.6.6, 3.6.3.6, 3.3.4.12, 3.4.3.12, 3.4.4.6, 3.4.6.4, 4.4.4.4, 3.7.42, 3.8.24, 3.9.18, 3.10.15, 3.12.12, 4.5.20, 4.6.12, 4.8.8, 5.5.10, 6.6.6. (Analyzing these 21 possibilities gives a straightforward, albeit tedious, proof that there are only 11 Archimedean lattices.) Grünbaum and Shephard [12] give nice pictures of the 21 types, as well as many plane tilings by regular polygons.

Let \mathcal{T} denote the set of plane tilings such that if $T \in \mathcal{T}$ and some vertex type appears in T , then that type appears in T infinitely often. It is easy to see that \mathcal{T} contains more tilings than just the Archimedean Lattices. A portion of such a tiling is shown in Figure 5. We prove the following.

Main Theorem. *No tiling in \mathcal{T} has threshold 4 or more, and the only tilings in \mathcal{T} with threshold 3 are the lattices (3.12.12), (4.6.12), (4.8.8), and (6.6.6).*

Proof. Fix $T \in \mathcal{T}$. As a warmup, we show that T has threshold at most 4. Suppose T has a vertex v of type other than 5.5.10 and 6.6.6. Note, by examining the 21 types above, that v has an incident 3-face or 4-face. So, by definition, T has infinitely many 3-faces or 4-faces. Now Observation 1 shows that T has threshold at most 4. As we show in the next paragraph, type 5.5.10 cannot appear in any plane tiling. Finally, if T has only vertex type 6.6.6, then T is the lattice (6.6.6), which has threshold 3.

The rest of the proof simply refines the idea in the previous paragraph. We first show that six types cannot appear in T at all. Suppose that T contains a vertex of type 3.7.42. Since no other type contains 7-gons or 42-gons, the lengths of faces incident to this 3-gon must alternate between 7 and 42. But this is impossible, since 3 is odd. So T contains no vertex of type 3.7.42. Similar arguments show that T contains no vertex of any of types 3.8.24, 3.9.18, 3.10.15, 4.5.20, and 5.5.10.

For the remaining types t other than 3.6.3.6, 3.12.12, 4.6.12, 4.8.8, and 6.6.6, we show that if T contains type t , then T contains a configuration H where each face of H has at most 2 adjacent faces outside H . Since H appears infinitely often, by Observation 1 the threshold of T is at most 2, as desired. The details follow.

If T contains two adjacent triangles, then we take these as H . This handles six types, leaving only 3.4.4.6, 3.4.6.4, 3.4.3.12, and 4.4.4.4. If v has type 4.4.4.4, then H is its four incident squares. If v has type 3.4.3.12, then H is the two incident triangles and the incident square. If v has type 3.4.4.6, then a short analysis shows that T contains one of the configurations on the left in Figure 4. Finally, if v has type 3.4.6.4, then a (slightly longer) proof shows that T contains the configuration on the right in Figure 4 (or else contains two triangles linked by one or two squares, similar to the cases on the left of Figure 4).




Figure 4: Left: The three possibilities for C when T contains a vertex of type 3.4.4.6. Right: A configuration, C , of 31 faces in which each face has at most two neighbors outside of C .

So the vertex types in T are some subset of 3.6.3.6, 3.12.12, 4.6.12, 4.8.8, and 6.6.6. To see that T must be an Archimedean lattice, note that none of these types agree in two or more successive face lengths. So it is impossible for T to “switch” from one type to another. \square

It is worth noting that we cannot relax the hypothesis in the Main Theorem to require only that *some* vertex type appears infinitely often. For example, suppose we start with the hex lattice and replace finitely many hexagons each with 6 triangles. If any of the resulting vertices of type 3.3.3.3.3.3 has no incident faces initially infected, then the percolation threshold drops from 3 to 1. Hence, the percolation threshold depends heavily on p , the probability that each face is initially infected.

To conclude, we briefly discuss a family of tilings we call $\mathcal{T}_{\text{strips}}$. These tilings are formed by “stacking” infinite horizontal strips of polygons above and below each other to fill the entire plane. The two types of strips that we use are *hex strips*, consisting of hexagons and triangles, and *square strips*, consisting just of squares. Figure 5 shows an example. Since the hex strips can be shifted left or right, this family contains uncountably many tilings.

Despite the variety in the tilings of $\mathcal{T}_{\text{strips}}$, they all have the same threshold. The proof is similar to our proof for the lattice (4.8.8), with a little difficulty added by the irregularly shaped rings we use now (what were previously $\mathcal{D}_{t+1} \setminus \mathcal{D}_t$). Two hex strips are *offset* if the centers of their hexagons are not directly above one another.

Theorem 2. *Every tiling in $\mathcal{T}_{\text{strips}}$ has percolation threshold 2.*

Proof. Let T be a tiling in $\mathcal{T}_{\text{strips}}$. Again the upper bound follows from Observation 1. The main step is to show that T contains a configuration C such that each face of C has at most two adjacent faces outside C . A short analysis yields that T contains infinitely many copies of one of the following: (a) adjacent triangles, (b) a hexagon with six adjacent triangles, (c) four

hex strips
square strips

offset




Figure 5: A tiling in $\mathcal{T}_{\text{strips}}$, along with a marked face and \mathcal{A}_4 .

squares incident to a common vertex, (d) two triangles adjacent to a common square, or (e) two triangles linked by two squares (as in Figure 4).

Now we show that for every p with $0 < p < 1$, if $k = 2$, then our random set \mathcal{I} percolates with positive probability. By combining this with the 0–1 Law, we conclude that the bootstrap threshold for T is 2.

First we must find an analogue of \mathcal{D}_t from our proof for the lattice (4.8.8). Consider a face f of T which is not a triangle. We let \mathcal{A}_t denote a collection of faces that is centered on f and that is shaped somewhere between a square and a hexagon (depending on the number of offset rows involved). In the strip containing f , \mathcal{A}_t contains $2t$ consecutive faces to the left of f (including triangles), and $2t$ consecutive faces to the right of f . For the strip above this, \mathcal{A}_t contains the faces directly above, if the two strips are not offset, and the faces above and slightly towards the center, if the strips are offset. Similarly for the strip below, \mathcal{A}_t contains the faces directly below if the two strips are not offset, and the faces below and slightly toward the center when the faces are offset. We continue this for the t rows above f and the t rows below f . This means that \mathcal{A}_t always consists of $2t + 1$ rows of faces, but the number of faces in the rows decreases slightly as we move away from the center row (whenever successive strips are offset).

Now \mathcal{A}_t looks like a square when it has no offset strips, and looks closer to a hexagon when it has many. Even when \mathcal{A}_t looks like a rectangle, we think of $\mathcal{A}_{t+1} \setminus \mathcal{A}_t$ as having six sides. The top and bottom sides are easy to see; they consist of faces directly above/below the faces in the top/bottom row of \mathcal{A}_t . The top-left side consists of faces directly left of an end-face of \mathcal{A}_t and which are in a strip above f . The bottom-left, top-right, and bottom-right sides are defined similarly.

The key insight is that, just like for the lattice (4.8.8), if \mathcal{A}_t is infected and $\mathcal{A}_{t+1} \setminus \mathcal{A}_t$ has even a single infected (non-triangular) face in one of its sides, then that entire side becomes infected. By repeatedly applying this idea, we see that the infection spreads along the entire

top-left side. Once two adjacent sides are infected (e.g. top-left and top, or bottom-right and top-right), the corner face lying between them also has two infected neighbors, so it becomes infected. The important consequence of all this is the following. If \mathcal{A}_t is completely infected, and at least one face on each of the six sides of $\mathcal{A}_{t+1} \setminus \mathcal{A}_t$ is infected, then \mathcal{A}_{t+1} also becomes completely infected.

Now we bound the probability that this happens. Each side has at least t non-triangular faces⁴, so the probability that none of the faces on a side are infected is at most $(1-p)^t$. Thus, the probability that at least one side of $\mathcal{A}_{t+1} \setminus \mathcal{A}_t$ has no infected face is no more than $6(1-p)^t$.

Now our argument exactly follows that for (4.8.8). The probability that at least one ring around \mathcal{A}_t does not become infected is at most $\sum_{j=0}^{\infty} 6(1-p)^{t+j} = \frac{6(1-p)^t}{p}$, and for large enough t we have $\frac{6(1-p)^t}{p} < 1$. Now the probability that \mathcal{A}_t is initially entirely infected and that every ring around \mathcal{A}_t contains an infected face on each of the six sides is $p^{|\mathcal{A}_t|} (1 - \frac{6(1-p)^t}{p}) > 0$. Since this event is translation-invariant in the horizontal direction, it is weakly translation-invariant. So the 0–1 Law tells us that \mathcal{I} percolates with probability 1. \square

Acknowledgment

We are indebted to Branko Grünbaum and Geoffrey Shephard for their article *Tilings by Regular Polygons* [12], which inspired much of this work. Thanks also to a referee, whose feedback helped improve our presentation.

References

- [1] M. Aizenman and J. L. Lebowitz. Metastability effects in bootstrap percolation. *J. Phys. A*, 21(19):3801–3813, 1988.
- [2] J. Balogh, B. Bollobás, H. Duminil-Copin, and R. Morris. The sharp threshold for bootstrap percolation in all dimensions. *Trans. Amer. Math. Soc.*, 364(5):2667–2701, 2012.
- [3] J. Balogh, B. Bollobás, and R. Morris. Bootstrap percolation in three dimensions. *Ann. Probab.*, 37(4):1329–1380, 2009.
- [4] J. Balogh, B. Bollobás, and R. Morris. Majority bootstrap percolation on the hypercube. *Combin. Probab. Comput.*, 18(1-2):17–51, 2009.
- [5] J. Balogh, Y. Peres, and G. Pete. Bootstrap percolation on infinite trees and non-amenable groups. *Combin. Probab. Comput.*, 15(5):715–730, 2006.
- [6] J. Balogh and B. G. Pittel. Bootstrap percolation on the random regular graph. *Random Structures Algorithms*, 30(1-2):257–286, 2007.
- [7] B. Bollobás and O. Riordan. *Percolation*. Cambridge University Press, New York, 2006.
- [8] B. Bollobás, P. Smith, and A. Uzzell. Monotone cellular automata in a random environment. *Combin. Probab. Comput.*, 24(4):687–722, 2015.
- [9] R. Cerf and E. Cirillo. Finite size scaling in three-dimensional bootstrap percolation. *Ann. Probab.*, 27(4):1837–1850, 1999.
- [10] R. Cerf and F. Manzo. The threshold regime of finite volume bootstrap percolation. *Stochastic Process. Appl.*, 101(1):69–82, 2002.

⁴ When \mathcal{A}_t contains two successive hex strips that are offset, the row further from f contributes to \mathcal{A}_t two fewer faces than the row nearer f (including one fewer hex faces). Thus, the top and bottom sides can each have as many as $2t + 1$ adjacent non-triangular faces. But, this only helps us, since a side with more faces is *more* likely to have an infected face.

- [11] J. Chalupa, P. Leath, and G. R. Reich. Bootstrap percolation on a bethe lattice. *J. Phys. C: Solid State Phys.*, (12):L31–L35, 1979.
- [12] B. Grünbaum and G. C. Shephard. Tilings by regular polygons. Patterns in the plane from Kepler to the present, including recent results and unsolved problems. *Math. Mag.*, 50(5):227–247, 1977.
- [13] C. Holmgren, T. Juškevičius, and N. Kettle. Majority bootstrap percolation on $G(n, p)$. *Electron. J. Combin.*, 24(1):Paper 1.1, 32, 2017.
- [14] A. E. Holroyd. Sharp metastability threshold for two-dimensional bootstrap percolation. *Probab. Theory Related Fields*, 125(2):195–224, 2003.
- [15] S. Janson, T. Łuczak, T. Turova, and T. Vallier. Bootstrap percolation on the random graph $G_{n,p}$. *Ann. Appl. Probab.*, 22(5):1989–2047, 2012.
- [16] R. H. Schonmann. On the behavior of some cellular automata related to bootstrap percolation. *Ann. Probab.*, 20(1):174–193, 1992.
- [17] A. C. D. van Enter. Proof of Straley’s argument for bootstrap percolation. *J. Statist. Phys.*, 48(3-4):943–945, 1987.
- [18] J. von Neumann. *Theory of self-reproducing automata*. Univ. Illinois Press, Champaign, 1966.