



BOOTSTRAP PERCOLATION IN RANDOM GEOMETRIC GRAPHS

VICTOR FALGAS-RAVRY,* *Umeå University*

AMITES SARKAR,** *Western Washington University*

Abstract

Following Bradonjić and Saniee, we study a model of bootstrap percolation on the Gilbert random geometric graph on the 2-dimensional torus. In this model, the expected number of vertices of the graph is n , and the expected degree of a vertex is $a \log n$ for some fixed $a > 1$. Each vertex is added with probability p to a set A_0 of initially infected vertices. Vertices subsequently become infected if they have at least $\theta a \log n$ infected neighbours. Here $p, \theta \in [0, 1]$ are taken to be fixed constants.

We show that if $\theta < (1 + p)/2$, then a sufficiently large local outbreak leads with high probability to the infection spreading globally, with all but $o(n)$ vertices eventually becoming infected. On the other hand, for $\theta > (1 + p)/2$, even if one adversarially infects every vertex inside a ball of radius $O(\sqrt{\log n})$, with high probability the infection will spread to only $o(n)$ vertices beyond those that were initially infected.

In addition we give some bounds on the (a, p, θ) regions ensuring the emergence of large local outbreaks or the existence of islands of vertices that never become infected. We also give a complete picture of the (surprisingly complex) behaviour of the analogous 1-dimensional bootstrap percolation model on the circle. Finally we raise a number of problems, and in particular make a conjecture on an ‘almost no percolation or almost full percolation’ dichotomy which may be of independent interest.

Keywords: Random geometric graphs; bootstrap percolation; random processes

2020 Mathematics Subject Classification: Primary 60G55

Secondary 60D05; 60K35

1. Introduction

1.1. Background

Bootstrap percolation encompasses a widely studied family of cellular automata on networks. Originally introduced by Chalupa, Leath, and Reich in 1979 [17] in the context of magnetic systems, it has since been used to model a great variety of phenomena—from the spreading of fads or beliefs in social networks [18, 25, 38], to financial contagion and default on obligations in economic networks [3], to the activation of neurons in the brain [2, 31] or the spread of viruses in human populations [18]. This plethora of applications has led to significant work on bootstrap percolation from network scientists, physicists, engineers, and computer scientists as well as mathematicians.

Received 22 October 2021; revision received 25 January 2023.

* Universitetstorget 4, 901 87 Umeå, Sweden. Email address: victor.falgas-ravry@umu.se

** 516 High Street, Bellingham, WA 98225, USA. Email address: amites.sarkar@wwu.edu

© The Author(s), 2023. Published by Cambridge University Press on behalf of Applied Probability Trust.

Formally, an s -threshold bootstrap process on a graph $G = (V, E)$ is defined as follows. At time $t = 0$, an initial set of vertices $A_0 \subseteq V$ is infected (or activated, if one prefers to avoid contagious connotations). Then at each time step $t \geq 0$, the vertices of G having at least s neighbours in A_t become infected and are added to A_t to form A_{t+1} . The infection thus spreads throughout the graph, and results in a set $A_\infty := \bigcup_{t \geq 0} A_t$ of eventually infected vertices.

If G is a finite graph, the key question of interest is then the following: what proportion of vertices of G eventually become infected? This obviously depends on the choice of the set of initially infected vertices A_0 . In most of the work on bootstrap percolation to date, A_0 is chosen according to a Bernoulli process on the vertices of G : each vertex $v \in V$ is included in A_0 with probability p independently of all other vertices. One then asks for which initial infection probability p the infection spreads to ‘most’ of V (almost percolation, $|A_\infty| = |V|(1 - o(1))$) for some sequence of graphs with $|V| \rightarrow \infty$, or to all of V (percolation, $|A_\infty| = |V|$) with high probability (w.h.p., meaning with probability $1 - o(1)$).

Aizenman and Lebowitz [1] were the first to investigate this kind of question when $G = [n]^d$, the d -dimensional $n \times n \times \dots \times n$ grid graph. In a landmark result in 2003, Holroyd [27] showed that for 2-threshold bootstrap percolation on $[n]^2$, $p_c([n]^2, 2) = \frac{\pi^2}{18 \log n}$ was a sharp threshold for percolation, in the sense that if $p = \frac{c}{\log n}$ is the initial infection probability and $c > 0$ is fixed, then if $c < \frac{\pi^2}{18}$, w.h.p. the infection does not spread to the entire square grid, while if $c > \frac{\pi^2}{18}$, then w.h.p. every vertex of $[n]^2$ eventually becomes infected. Surprisingly, this result disproved predictions for the value of the critical threshold that had been made based on numerical simulations for the problem. Holroyd’s results were then extended to other dimensions d and thresholds s by Balogh, Bollobás, Duminil-Copin, and Morris [9].

Motivated by the applications of bootstrap percolation to the modelling of real-life network phenomena, there has been growing interest over the past decade in the study of bootstrap percolation on random graphs. An important work in this vein was the study of bootstrap percolation on Erdős–Rényi random graphs $G_{n,q}$ by Janson, Łuczak, Turova, and Vallier [30] in 2012. Recall that the random graph $G_{n,q}$ is obtained by taking $[n] := \{1, 2, \dots, n\}$ as a vertex set, and including each pair ij as an edge of the random graph with probability q , independently of all the other pairs. The authors of [30] determined inter alia for every threshold $s \geq 2$ and q the critical thresholds $p = p(n, s, q)$ for the infection probability at which the size of A_∞ goes from w.h.p. $o(n)$ to w.h.p. $n - o(n)$ (almost percolation) and from w.h.p. $n - o(n)$ to w.h.p. n (percolation).

Bootstrap percolation has been rigorously studied on several other random graph models: random regular graphs [10, 28], power-law random graphs [4], Bienaymé–Galton–Watson trees [12], random graphs with specified vertex degrees [28], toroidal grids with random long edges added in (a special case of the Kleinberg model) [29], and inhomogeneous random graphs [21], amongst others. A motivation for the latter two models is that they may have degree distributions or spatial characteristics that more closely resemble those of the real-life networks motivating the study of bootstrap percolation.

For similar reasons, there has been interest in bootstrap percolation models on random *geometric* graphs. Indeed, many real-life networks have a distinctly spatial structure that affects their behaviour and properties. Thus it is natural to study bootstrap percolation on random graph models in which geometry plays a role. This was first done in 2014 by Bradonjić and Saniee [15], who introduced a model for bootstrap percolation on random geometric graphs that is the focus of the present paper.

The Gilbert disc model is the most widely studied random geometric graph model, and was first defined by Gilbert [23] in 1961; indeed ‘random geometric graph’ without any further qualifier usually refers to the Gilbert model. Given a measurable metric space Ω , a Gilbert random geometric graph $G_r(\Omega)$ is obtained by taking as the vertex set the point-set \mathcal{P} resulting from a Poisson point process of intensity 1 on Ω . Two vertices v, v' in \mathcal{P} are then joined by an undirected edge if their distance in Ω is less than r . In other words, the neighbours of a vertex v are precisely those points of $\mathcal{P} \setminus \{v\}$ that lie inside the ball of radius r centred at v .

Gilbert studied his model in the case where $\Omega = \mathbb{R}^2$, the 2-dimensional plane equipped with the usual Euclidean distance and Lebesgue measure. The study of the Gilbert model and related random geometric graph models on such unbounded spaces is known as *continuum percolation*, and is the subject of a monograph of Meester and Roy [35]. In a different direction, researchers have been interested in random geometric graph models in bounded, finite-dimensional spaces, in particular when Ω is either $S_n^d := [0, n^{1/d}]^d$, the d -dimensional box of volume n , or $T_n^d := (\mathbb{R}/n^{1/d}\mathbb{Z})^d$, the d -dimensional torus of volume n , where d is fixed and n is large. For both of these choices of Ω , standard results on concentration of the Poisson distribution imply that $G_r(\Omega)$ w.h.p. has $n + o(n)$ vertices.

In the case $\Omega = T_n^d$ (where we can ignore boundary effects and all vertices look the same), $G_r(\Omega)$ can be viewed as a natural geometric analogue of the Erdős–Rényi random graph. Let α_d denote the volume of the d -dimensional unit ball. Then the expected degree of a vertex in $G_r(T_n^d)$ is precisely $\alpha_d r^d$. A classic result of Penrose [36] established that the threshold for connectivity for $G_r(T_n^d)$ occurs at $\alpha_d r^d = \log n$: if $\alpha_d r^d = a \log n$ and $a > 0$ is fixed, then for $a < 1$, w.h.p. $G_r(T_n^d)$ contains isolated vertices and thus fails to be connected, while for $a > 1$, w.h.p. $G_r(T_n^d)$ is connected. Much more is known about Gilbert random geometric graphs (which, together with the closely related k -nearest-neighbour model, have been applied in a variety of contexts, for example to model sensor networks [8] and ad hoc wireless networks [40], and for cluster analysis in spatial statistics [24]). We refer the interested reader to the monograph of Penrose [37] devoted to the topic.

Bradonjić and Saniee considered the Gilbert disc model $G_r(T_n^2)$ where r is given by $\pi r^2 = a \log n$, for some constant $a > 1$ and n large. They studied $\theta a \log n$ -threshold bootstrap percolation on this host graph—i.e. where the threshold for infection is a proportion θ of the expected degree of a vertex. This is somewhat in contrast to previous work, where typically the threshold for infection was fixed rather than growing with the number of vertices n , but may be a more suitable choice of parameter for modelling situations such as the spread of a fad or fashion in a social network.

Bradonjić and Saniee’s paper featured a mixture of rigorous results and simulations. On the theoretical side, they proved two results. First of all, they determined [15, Theorem 1] an explicit function $f_1(a, \theta)$ such that if the initial infection probability p is fixed and satisfies $p < f_1(a, \theta)$, then w.h.p. the infection does not spread at all: $A_\infty = A_0$, and every vertex that is initially uninfected stays uninfected forever. Secondly, Bradonjić and Saniee determined [15, Theorem 2] a second function $f_2(a, \theta)$ such that if p is fixed and satisfies $p > f_2(a, \theta)$, then w.h.p. there is full percolation: every vertex becomes infected. The function $f_1(a, \theta)$ in the first of these results is easily seen to be the best possible (see Proposition 1 below). However, the bound $f_2(a, \theta)$ in the second result seems far from optimal—indeed, the simulations performed by Bradonjić and Saniee suggest as much.

Apart from Bradonjić and Saniee’s 2014 paper, comparatively little mathematical work appears to have been done on bootstrap percolation in random geometric graphs. In a 2016

work, Candellero and Fountoulakis [16] studied s -threshold bootstrap percolation on hyperbolic random geometric graphs for constant s , and determined for their model a critical probability p_c such that if the initial infection probability p satisfies $p \ll p_c$, then w.h.p. the infection does not spread at all, while if $p \gg p_c$, then w.h.p. the infection spreads to a strictly positive proportion of the vertices. More recently, Koch and Lengler [33, 34] studied a localized form of s -threshold bootstrap percolation on geometric inhomogeneous random graphs, with s a fixed constant and where the set of initially infected vertices is located within some bounded source region B (rather than the whole space); they determined a similar critical threshold p_c below which w.h.p. an infection does not spread at all, and above which w.h.p. an infection spreads to a positive proportion of all vertices. Finally, in a very recent Ph.D. thesis, Whittemore [39] studied bootstrap percolation in the Gilbert random geometric graph when the infection threshold is constant, and determined amongst other things the thresholds at which the model’s typical behaviour transitions from almost no percolation to almost percolation. As far as we are aware, this is the (surprisingly limited) extent of rigorous mathematical study of bootstrap percolation on random geometric graph models (though there also exist some experimental and simulation results for bootstrap percolation on geometric scale-free networks; see e.g. [22]).

1.2. The Bradonjić–Saniee model

For the reader’s convenience, we restate here the precise model we shall be studying in this paper.

Let $T_n^d := (\mathbb{R}/n^{1/d}\mathbb{Z})^d$ denote the d -dimensional torus of hypervolume n . Given a parameter r , a Gilbert random geometric graph $G_r(T_n^d)$ on T_n^d is obtained as follows: we let its vertex set \mathcal{P} be the result of a Poisson point process of intensity 1 on T_n^d , and join vertices $u, v \in \mathcal{P}$ by an edge if their distance (in the torus) is less than r . For compactness of notation, we use $G_{n,r}^d$ to denote $G_r(T_n^d)$.

Let $a > 1$ and $d \in \mathbb{N}$ be fixed. Let r be given by the relation $\alpha_d r^d = a \log n$, where α_d is the volume of the d -dimensional unit ball. Note that for this choice of parameters the expected total number of vertices in $G_{n,r}^d$ is n , while the expected degree of a vertex in $G_{n,r}^d$ is $a \log n$. Furthermore, by classical results of Penrose [36], $G_{n,r}^d$ is w.h.p. connected.

Bradonjić and Saniee [15] introduced the following model of bootstrap percolation on $G_{n,r}^d$: let $p, \theta \in [0, 1]$ be fixed. At time $t = 0$, let each vertex of $G_{n,r}^d$ be infected independently at random with probability p . Denote by A_0 this set of initially infected vertices. The infection then spreads through the graph $G_{n,r}^d$ as follows: at each time step $t > 0$, all vertices of $G_{n,r}^d$ which have at least $\theta a \log n$ infected neighbours (i.e. neighbours in the infected set A_{t-1}) become infected themselves and are added to A_{t-1} to form the set A_t . We denote by $A_\infty = \bigcup \{A_t : t \in \mathbb{Z}_{\geq 0}\}$ the set of all vertices of $G_{n,r}^d$ that eventually become infected under this process.

With a, p, θ, d fixed, the main question of interest in this model is the following: what is the typical size of A_∞ for large n ? In this paper we investigate this question in detail in dimensions $d = 1$ and $d = 2$.

1.3. Contributions of this paper

Our main contribution in this paper is identifying in dimension $d = 2$ the threshold at which a sufficiently large *local* outbreak can cascade and lead to a *global infection*. Say that a ball B in T_n^2 is *infected* if all vertices of $G_{n,r}^2$ that lie inside B are infected.

Theorem 1. *Let a, θ, p be fixed. Then the following hold:*

- (i) *If $\theta < \frac{1+p}{2}$, then there exists a constant $C = C(a, \theta, p)$ such that w.h.p., if any ball B in T_n^2 of radius Cr is infected (either artificially or as a result of the bootstrap percolation process), then all but $o(n)$ vertices of $G_{n,r}^2$ eventually become infected. Furthermore, when the infection stops, all connected components of uninfected vertices in $G_{n,r}^2[\mathcal{P} \setminus A_\infty]$ have Euclidean diameter $O(\sqrt{\log n})$ in T_n^2 .*
- (ii) *If $\theta > \frac{1+p}{2}$, then for every constant $C > 0$, w.h.p. even if one adversarially selects a ball B in T_n^2 of radius Cr and infects all the vertices it contains, only $o(n)$ additional vertices of $G_{n,r}^2$ become infected in the bootstrap percolation process starting from the initially infected set $A_0 \cup (B \cap \mathcal{P})$. Furthermore, all connected components of $G_{n,r}^2[A_\infty \setminus (A_0 \cup B)]$ have Euclidean diameter $O(\sqrt{\log n})$ in T_n^2 .*

Note that in our regime $\pi r^2 = a \log n$ and the longest edge of $G_{n,r}^2$ thus has length $O(\sqrt{\log n})$. We therefore view point-sets of diameter $O(\sqrt{\log n})$ as *local configurations*. What Theorem 1 says is thus that for $\theta < \frac{1+p}{2}$, a sufficiently large local infection will, with the help of the initially infected vertices, spread to most of the graph, leaving only isolated local ‘islands’ of uninfected vertices, while for $\theta > \frac{1+p}{2}$, all infectious local outbreaks remain local.

This leads us to conjecture that in the Bradonjić–Saniee model on $G_{n,r}^2$ with $\theta \neq \frac{1+p}{2}$, w.h.p. either an initial infection spreads to at most $o(n)$ new vertices (almost no percolation), or it spreads to all but at most $o(n)$ vertices (almost percolation).

Conjecture 1. *(Almost no percolation/almost full percolation dichotomy.) Let (a, p, θ) be fixed with $\theta \neq \frac{1+p}{2}$ and $a > 1$. Then in the Bradonjić–Saniee model for bootstrap percolation on $G_{n,r}^2$, w.h.p. either $|A_\infty \setminus A_0| = o(n)$ or $|\mathcal{P} \setminus A_\infty| = o(n)$.*

We were unfortunately unable to resolve Conjecture 1 in full, but our results imply it holds if $\theta > \frac{1+p}{2}$ or $\theta < \theta_{\text{local}}$, where $\theta_{\text{local}} = \theta_{\text{local}}(a, p)$ is a quantity that arises as the solution to an explicit continuous optimization problem and whose technical definition (Definition 5) we defer to Section 3.4. Suffice it to say here that θ_{local} is the threshold for the appearance of large local, symmetrically distributed, infectious outbreaks.

Theorem 2. *Let (a, p, θ) be fixed with $a > 1$ and*

$$\theta < \theta_{\text{local}}(a, p). \tag{1}$$

Then w.h.p. almost percolation occurs in the Bradonjić–Saniee model, i.e. $|\mathcal{P} \setminus A_\infty| = o(n)$.

We note here the fact that the ‘symmetric local growth condition’ $\theta < \theta_{\text{local}}(a, p)$ implies the ‘global growth condition’ $\theta < (1+p)/2$ (see Proposition 5): $\theta_{\text{local}} \leq (1+p)/2$ for all fixed $a > 1$ and $p \in [0, 1]$. We believe that θ_{local} gives the threshold for almost percolation in the Bradonjić–Saniee model for bootstrap percolation, and thus that the following strengthening of Conjecture 1 is true.

Conjecture 2. *(Symmetric local growth.) Let (a, p, θ) be fixed with $a > 1$ and*

$$\theta > \theta_{\text{local}}(a, p).$$

Then w.h.p. almost no percolation occurs in the Bradonjić–Saniee model, i.e. $|A_\infty \setminus A_0| = o(n)$.

The content of Conjecture 2 is twofold: the conjecture asserts first of all that completely infecting a large local ball is w.h.p. necessary for the infection to spread globally, and secondly that the likeliest way an infection spreads to a large local ball is if there is an abnormally high concentration of initially infected and of initially non-infected vertices distributed in a symmetric manner around the centre of the said ball.

Theorems 1 and 2 above are stated and proved for the Bradonjić–Saniee bootstrap percolation model in the torus T_n^2 rather than the square $S_n^2 := [0, \sqrt{n}]^2$ to avoid technical complications due to boundary effects. However, as we note in Section 5, our results also hold for their model in the square, modulo a technical modification in the statement of Theorem 1(i). We thus expect Conjectures 1 and 2 to also hold in the square.

Given Theorem 2 and Conjecture 2 on almost percolation, it is natural to ask how much smaller θ needs to be to ensure full percolation: which triples (a, p, θ) guarantee that a global infection w.h.p. infects every vertex of $G_{n,r}^2$? We are unable to answer this question exactly. However, as in Theorem 2 we are able to determine the threshold θ_{islands} for the disappearance of certain symmetric ‘islands’ of uninfected vertices, which provide what we conjecture is the main obstacle to full percolation. Here $\theta_{\text{islands}} = \theta_{\text{islands}}(a, p)$ is an (explicit) solution to a certain optimization problem, whose formal definition we defer to Section 3.5. (Note that θ_{islands} will satisfy the inequality $\theta_{\text{islands}} \leq \frac{1+p}{2}$; see (12).)

Theorem 3. *Let (a, p, θ) be fixed with $a > 1$ and $\theta > \theta_{\text{islands}}(a, p)$. Then w.h.p. some vertices remain uninfected by the end of the bootstrap percolation process in the Bradonjić–Saniee model in the torus T_n^2 , i.e. $|\mathcal{P} \setminus A_\infty| > 0$ and we do not have full percolation.*

Conjecture 3. *Let (a, p, θ) be fixed with $a > 1$ and $\theta < \min(\theta_{\text{islands}}, \theta_{\text{local}})$. Then w.h.p. we have full percolation in the Bradonjić–Saniee model in the torus T_n^2 , i.e. $\mathcal{P} = A_\infty$.*

Theorem 3 carries over immediately to the square setting, but Conjecture 3 does not. In the square setting, one will need to separately compute the threshold for the disappearance of uninfected islands close to the boundary, which will require additional calculations. (This is a rather standard feature in results on random geometric graphs in the square; see e.g. the proof of [6, Theorem 7].) Note that islands close to the boundary will have fewer neighbouring vertices and will be harder to infect.

For the Bradonjić–Saniee bootstrap percolation model in general dimension $d \geq 1$, we also determine the threshold $\theta_{\text{start}}(a, p)$ for the event $A_0 \neq A_\infty$ to hold w.h.p. (Proposition 1), generalizing [15, Theorem 1], and for the event that there exist ‘uninfectable’ initially uninfected vertices with degree strictly less than $\theta a \log n$, giving lower bounds on the threshold for full percolation $f_2(a, \theta)$ given in [15, Theorem 2].

Together with the results and conjectures above, these last results give the following picture for the expected behaviour of the Bradonjić–Saniee model in dimension 2 with (a, p, θ) fixed and $a > 1$ (see Figure 1):

- For $\theta > \theta_{\text{start}}$, w.h.p. we have no percolation: $A_0 = A_\infty$.
- For $\theta_{\text{start}} > \theta > \frac{1+p}{2}$, w.h.p. we have almost no percolation, $|A_\infty \setminus A_0| = o(n)$, even if we adversarially infect a ball of area $O(\log n)$.
- For $\frac{1+p}{2} > \theta > \theta_{\text{local}}$, w.h.p. we have almost no percolation, but adversarially infecting a ball of area $\Omega(\log n)$ w.h.p. leads to almost percolation.

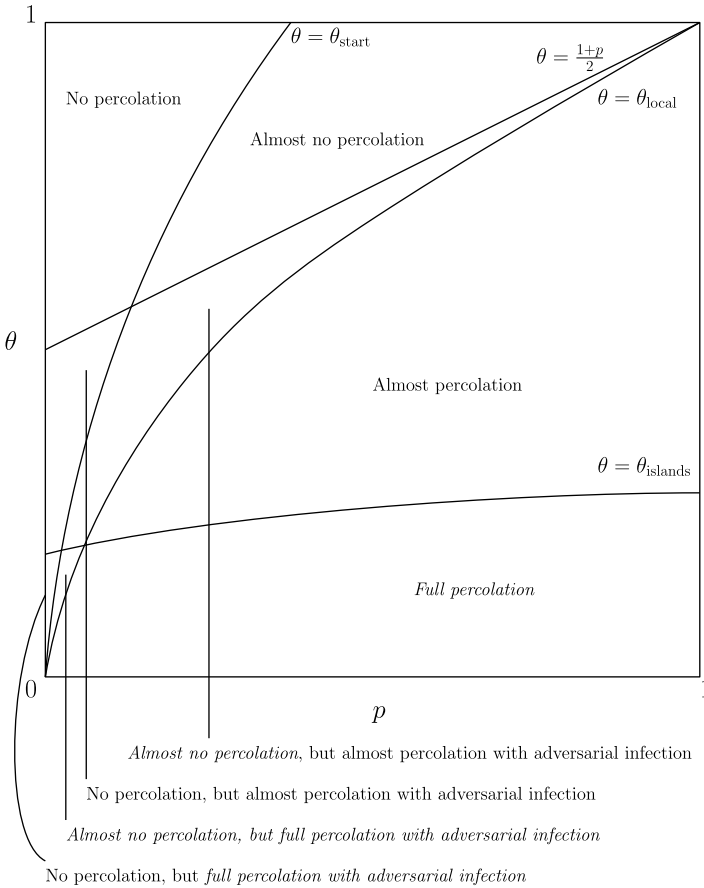


FIGURE 1. Phase diagram in the (p, θ) -plane for $a > 1$ fixed; italics indicate conjectured behaviour.

- For $\theta_{\text{local}} > \theta > \theta_{\text{islands}}$, w.h.p. we have almost percolation but not full percolation, $0 < |\mathcal{P} \setminus A_\infty| \leq o(n)$.
- For $\min(\theta_{\text{local}}, \theta_{\text{islands}}) > \theta$, w.h.p. we have full percolation, $\mathcal{P} = A_\infty$.

Finally, we give a complete picture of the typical behaviour of the Bradonjić–Saniee bootstrap percolation model on the graph $G_{n,r}^1$, i.e. in the 1-dimensional case. This turns out to be surprisingly complex, involving a (p, θ) -phase diagram with six different regions (see the summary in Section 4); however, unlike in the 2-dimensional case, we can compute the various thresholds more or less explicitly. We defer an exact statement of these 1-dimensional results to Section 4.

1.4. Organization of the paper

The remainder of this paper is organized as follows. In Section 2, we prove some basic probabilistic results for Poisson point processes that are required for later results. We also derive thresholds for the events that some initially infected vertices eventually become infected

(Proposition 1), and that some initially uninfected vertices have degree too low to ever become infected (Proposition 2).

In Section 3 we prove our main results for bootstrap percolation on the 2-dimensional torus, while in Section 4 we outline the behaviour of the Bradonjić–Saniee model on the circle. We end the paper in Section 5 with a discussion of some of the many open problems on bootstrap percolation for random geometric graphs.

2. Preliminaries

2.1. Notation

Given a host metric space Ω and a point $\mathbf{x} \in \Omega$, we write $B_r(\mathbf{x})$ for the ball in Ω of radius r centred at \mathbf{x} . We use $|S|$ to denote the size of S if S is a finite set, and the Lebesgue measure of S otherwise. Given $\mathbf{x} \in \mathbb{R}^d$, we denote by $\|\mathbf{x}\|$ the standard Euclidean ℓ_2 -norm of \mathbf{x} .

We use standard graph theoretic terminology and Landau big- O notation throughout the paper. In particular, given a graph H and a subset of its vertex set X , we denote by $H[X]$ the subgraph of H induced by X and by $H - X$ the subgraph of H induced by $V(H) \setminus X$.

2.2. Probabilistic tools

The following lemma, taken from [6, Lemma 1], greatly facilitates calculations of event probabilities, and we shall make extensive use of it throughout the paper.

Lemma 1. (*Estimating probabilities for Poisson point processes.*) *Let $A \subset \mathbb{R}^d$ be measurable, and let $\rho \geq 0$ be a real number such that $\rho|A| \in \mathbb{Z}$. Then the probability that a Poisson process in \mathbb{R}^d with intensity 1 has precisely $\rho|A|$ points in the region A is given by*

$$\exp\{(\rho - 1 - \rho \log \rho)|A| + O(\log_+ \rho|A|)\},$$

with the convention that $0 \log 0 = 0$, and $\log_+ x = \max(\log x, 1)$.

Furthermore, by standard properties of Poisson point processes (see e.g. Kingman [32]), note that the following are equivalent:

- taking as the vertex set of $G_{n,r}^d$ the outcome \mathcal{P} of a Poisson point process of intensity 1 on T_n^d , and then infecting each vertex of \mathcal{P} independently at random with probability $p \in (0, 1)$ to obtain A_0 ;
- letting A_0 be the outcome of a Poisson point process of intensity p on T_n^d , and taking as the vertex set of $G_{n,r}^d$ the union \mathcal{P} of A_0 and of the outcome of a Poisson point process of intensity $1 - p$ on T_n^d .

We will thus be able to jump back and forth between these two equivalent ways of constructing \mathcal{P} , adopting whichever point of view makes our calculations simplest.

At several points in the paper, we will need to count the number of copies of a ‘local’ event inside T_n^d , and the next lemma gives us some simple tools to do this. Let $\Gamma = \mathbb{Z}^d \cap [0, n^{1/d}]^d$; this may be identified with a d -dimensional grid of points inside T_n^d . Given an event E defined for point-sets in T_n^d , for every $\mathbf{x} \in \Gamma$ we let $E(\mathbf{x})$ denote the collection of point-sets in T_n^d whose translates by \mathbf{x} belong to E ; in other words, a configuration of initially uninfected and initially infected points $(\mathcal{P} \setminus A_0, A_0)$ belongs to $E(\mathbf{x})$ if and only if $(\mathcal{P} \setminus A_0 - \mathbf{x}, A_0 - \mathbf{x})$ belongs to E .

We say that an event E is Cr -bounded if it is determined by what happens (i.e. which points are present/initially infected) within a ball of radius Cr of the origin $\mathbf{0}$, for some constant $C > 0$.

For such an event, we define $\mathbf{E} = (\mathbb{1}_{E(\mathbf{x})})_{\mathbf{x} \in \Gamma}$ to be the Γ -dimensional zero–one vector recording for which $\mathbf{x} \in \Gamma$ the event $E(\mathbf{x})$ occurs. Given a Cr -bounded event E , let $N_E := \sum_{\mathbf{x}} \mathbb{1}_{E(\mathbf{x})}$ denote the number of $\mathbf{x} \in \Gamma$ for which $E(\mathbf{x})$ occurs.

Lemma 2. *Suppose E is a Cr -bounded event with $q := \mathbb{P}(E) = e^{-c \log n + o(\log n)}$ for some fixed constant c . The following hold:*

- (i) *If $c > 1$, then with probability $1 - o(1)$ we have $N_E = 0$.*
- (ii) *If $c < 1$, then with probability $1 - o(1)$ we have $N_E = n^{1-c+o(1)}$.*

Proof. Part (i) is immediate from Markov’s inequality. For the lower bound in Part (ii), note that there exists a subset Γ' of Γ satisfying

- (a) $|\Gamma'| = \Omega\left(\frac{n}{\log n}\right)$;
- (b) $\forall \mathbf{x}, \mathbf{y} \in \Gamma'$ with $\mathbf{x} \neq \mathbf{y}$, $\|\mathbf{x} - \mathbf{y}\| > 2Cr$.

Indeed, one can construct such a set Γ' by greedily adding vertices from Γ one by one subject to (b). By (b) and Cr -boundedness, the events $(E(\mathbf{x}))_{\mathbf{x} \in \Gamma'}$ are independent. Thus, N_E stochastically dominates a Binomial($|\Gamma'|, q$) random variable. Since the expectation of this binomial random variable is $|\Gamma'|q = \Omega(nq / \log n) = n^{1-c+o(1)} \gg 1$, a standard Chernoff bound tells us that with probability $1 - o(1)$, $N_E \geq n^{1-c+o(1)}$.

For the upper bound in Part (ii) we use Markov’s inequality: the probability that N_E is greater than $(\mathbb{E}N_E) \log n = O(nq \log n) = n^{1-c+o(1)}$ is $O(1 / \log n) = o(1)$. Thus, with probability $1 - o(1)$, $N_E \leq n^{1-c+o(1)}$, as required. □

2.3. Elementary considerations on the Bradonjić–Saniee model

Consider the Bradonjić–Saniee bootstrap percolation model in dimension $d \geq 1$ with $\alpha_d r^d = a \log n$ and $a > 1$ fixed. When $\theta < p$ and n is large, most vertices will immediately see more than $a\theta \log n < ap \log n$ infected neighbours, so that most uninfected vertices will immediately become infected. However, this is not the whole story. Indeed, even when θ is much smaller than p , there is still a chance that some vertex somewhere might see far fewer than its expected $a \log n$ neighbours (infected or not); if it in fact sees fewer than $\theta a \log n$ neighbours, then it can never become infected. In the other direction, even when $\theta > p$, there could still be a chance that some vertex somewhere will see far more than its expected $pa \log n$ infected neighbours, perhaps as many as $\theta a \log n$, so that the infection could spread first to, and then from, that vertex.

Roughly speaking, an analysis of the Bradonjić–Saniee model must grapple with three separate questions: whether the infection starts to spread at all, whether it continues to spread to most of the graph, and whether it finally infects every vertex. Each of these requires a separate analysis, even in one dimension. As a simple consequence of the probabilistic tools we have introduced, however, we can readily answer here the question of when the infection starts at all, strengthening and generalizing [15, Theorem 1].

Proposition 1. *For $a > 1$ and $0 < p < \theta < 1$ fixed, let*

$$f_{\text{start}}(a, p, \theta) := a(p - \theta + \theta \log(\theta/p)).$$

Then, if $f_{\text{start}}(a, p, \theta) < 1$, w.h.p. at least one initially uninfected vertex is infected in the first round of the bootstrap percolation process. If, however, $f_{\text{start}}(a, p, \theta) > 1$, then w.h.p. no initially uninfected vertex ever becomes infected.

Proof. Write X for the number of vertices which initially see more than $a\theta \log n$ infected neighbours in $G_{n,r}^d$. By Wald’s identity and Lemma 1, we have for $0 < p < \theta < 1$ that

$$\begin{aligned} \mathbb{E}(X) &= n \exp \left\{ pa \log n ((\theta/p) - 1 - (\theta/p) \log (\theta/p)) + O(\log_+ a\theta \log n) \right\} \\ &= \exp \left\{ \log n(1 - f_{\text{start}}(a, p, \theta) + o(1)) \right\}. \end{aligned}$$

Consequently, if $f_{\text{start}}(a, p, \theta) > 1$, then by Markov’s inequality w.h.p. $X = 0$ and no initially uninfected vertex of \mathcal{P} ever becomes infected.

If on the other hand $f_{\text{start}}(a, p, \theta) < 1$, then let E denote the event that the ball of radius $r - \sqrt{d}$ around the origin contains at least $\theta a \log n$ initially infected vertices and that the ball of radius \sqrt{d} around the origin contains at least one initially uninfected vertex of \mathcal{P} . Then by Lemma 1 and standard properties of the Poisson point process, the probability of E is

$$\begin{aligned} \mathbb{P}(E) &= \exp \{-f_{\text{start}}(a, p, \theta) \log n + o(\log n)\} \left(1 - \exp\{-(1-p)\alpha_d(\sqrt{d})^d\} \right) \\ &= \exp \{-f_{\text{start}}(a, p, \theta) \log n + o(\log n)\}. \end{aligned}$$

By Lemma 2(ii) w.h.p. $N_E = n^{1-f_{\text{start}}+o(1)}$. In particular w.h.p. the event $E(\mathbf{x})$ occurs for some $\mathbf{x} \in \Gamma$. This implies that there exists an initially uninfected vertex $v \in B_{\sqrt{d}}(\mathbf{x}) \cap \mathcal{P}$ such that $B_r(v) \cap A_0 \supseteq B_{r-\sqrt{d}}(\mathbf{x}) \cap A_0$ contains at least $\theta a \log n$ initially infected points of \mathcal{P} . Thus v becomes infected in the first round of the bootstrap percolation process, and w.h.p. $A_0 \neq A_1 \subseteq A_\infty$. This concludes the proof of the proposition. \square

Definition 1. For $a > 1$ and $p \in (0, 1)$ fixed, we define $\theta_{\text{start}} = \theta_{\text{start}}(a, p)$ to be the supremum of the $\theta \leq 1$ such that $f_{\text{start}}(a, p, \theta) < 1$.

Since $f_{\text{start}}(a, p, p) = 0$, it follows that $p \leq \theta_{\text{start}} \leq 1$.

Proposition 2. For $a > 1$, $p < 1$, and $0 < \theta < 1$ fixed, set

$$f_{0\text{-stop}}(a, \theta) := a(1 - \theta + \theta \log \theta).$$

(The reason for this choice of notation will become clear later.) Then, if $f_{0\text{-stop}}(a, \theta) < 1$, w.h.p. at least one initially uninfected vertex has degree less than $\theta a \log n$ in $G_{n,r}^d$ and consequently never becomes infected. On the other hand, if $f_{0\text{-stop}}(a, \theta) > 1$, then w.h.p. the minimum degree of $G_{n,r}^d$ is at least $\theta a \log n$.

Proof. Write Y for the number of vertices which have fewer than $a\theta \log n$ neighbours in $G_{n,r}^d$. Again, using Wald’s identity and Lemma 1 we have

$$\begin{aligned} \mathbb{E}(Y) &= n \exp \left\{ a \log n(\theta - 1 - \theta \log (\theta)) + O(\log_+ a\theta \log n) \right\} \\ &= \exp \left\{ \log n(1 - f_{0\text{-stop}}(a, \theta) + o(1)) \right\}. \end{aligned}$$

Thus, if $f_{0\text{-stop}}(a, \theta) > 1$, then $\mathbb{E}(Y) = o(1)$, and by Markov’s inequality w.h.p. the minimum degree of $G_{n,r}^d$ is at least $\theta a \log n$.

On the other hand, if $f_{0\text{-stop}}(a, \theta) < 1$, then let E be the event that the ball of radius $r + \sqrt{d}$ around the origin contains strictly fewer than $\theta a \log n$ vertices of \mathcal{P} and that the ball of radius

\sqrt{d} around the origin contains at least one initially uninfected vertex of \mathcal{P} . Then by Lemma 1 and standard properties of the Poisson point process, the probability of E is

$$\begin{aligned} \mathbb{P}(E) &= \exp \left\{ -f_{0\text{-stop}}(a, p, \theta) \log n + o(\log n) \right\} \left(1 - \exp \left\{ -(1-p)\alpha_d(\sqrt{d})^d \right\} \right) \\ &= \exp \left\{ -f_{0\text{-stop}}(a, p, \theta) \log n + o(\log n) \right\}. \end{aligned}$$

By Lemma 2(ii) w.h.p. $N_E = n^{1-f_{0\text{-stop}}+o(1)}$. In particular, w.h.p. the event $E(\mathbf{x})$ occurs for some $\mathbf{x} \in \Gamma$. This implies that there exists an initially uninfected vertex $v \in B_{\sqrt{d}}(\mathbf{x}) \cap \mathcal{P}$ such that $B_r(v) \subseteq B_{r+\sqrt{d}}(\mathbf{x})$ contains strictly fewer than $\theta a \log n$ points of \mathcal{P} . Thus v never becomes infected and w.h.p. $A_\infty \neq \mathcal{P}$. This concludes the proof of the proposition. \square

A few comments are in order. First, if in Proposition 2 we have $f_{0\text{-stop}}(a, \theta) = 1$, then we cannot apply Lemma 1 directly, since the error term will dominate. However, in this case, $\mathbb{P}(Y > 0)$ will tend to some constant that is neither 0 nor 1. An equivalent remark applies to Proposition 1. More precise results can be established in these special cases using the Stein–Chen method for Poisson approximation; however, we will not pursue such questions here. Second, when, say, $f_{\text{start}}(a, p, \theta) < 1$, not only does the infection start to spread somewhere, it in fact starts to spread in $n^{1-f_{\text{start}}(a,p,\theta)+o(1)} = n^{\alpha+o(1)}$ different places for some $\alpha > 0$, as established in the proof of Proposition 1 (more specifically the lower bound on N_E). Third, in Proposition 2 we have found one obstruction to full infection, but there may (and in fact will) be others.

3. Bootstrap percolation on the torus when $d = 2$

In this section we prove our result for the Bradonjić–Sanjeev model for bootstrap percolation in the Gilbert random geometric graph $G_{n,r}^2$ on the 2-dimensional torus. Throughout the section we let $T_n := T_n^2$ denote the said 2-dimensional torus.

3.1. Almost-percolation from a large local infection: the case $\theta < \frac{1+p}{2}$

Fix $a > 1$, and let $\pi r^2 = a \log n$. In this subsection, our goal is to show that if p, θ are fixed and satisfy the *growing condition*

$$\theta < \frac{1+p}{2}, \tag{2}$$

then w.h.p. any sufficiently large local infection spreads to almost all of T_n . To state our formal result (Theorem 4), we must introduce two tilings of $[0, \sqrt{n}]^2$. Let $K \in \mathbb{N}$ be a large constant to be specified later.

Definition 2. (*Rough tiling, fine tiling.*) *The rough tiling \mathcal{R} partitions $[0, \sqrt{n}]^2$ into interior-disjoint $Kcr \times Kcr$ square tiles, where $c = \frac{\sqrt{n}}{Kr \lfloor \sqrt{n}/Kr \rfloor} = 1 + o(1)$ is chosen to ensure that divisibility conditions are satisfied. The fine tiling \mathcal{F} is a refinement of \mathcal{R} obtained by subdividing each tile of \mathcal{R} into K^4 smaller $\frac{cr}{K} \times \frac{cr}{K}$ square tiles.*

Let p, θ be fixed. Suppose p, θ satisfy (2). Then for any $\eta > 0$ there exist constants $C_\eta > 1$ sufficiently large and $\delta > 0$ sufficiently small such that the following hold: let $\mathbf{0}$ denote the origin in \mathbb{R}^2 , and let $R \geq C_\eta r$ be a real number. Then for any point \mathbf{x} at distance between $R - \delta r$ and $R + \delta r$ of $\mathbf{0}$, the asymmetric lens $B_R(\mathbf{0}) \cap B_r(\mathbf{x})$ has area at least

$$|B_R(\mathbf{0}) \cap B_r(\mathbf{x})| \geq \frac{\pi r^2}{2} (1 - \eta). \tag{3}$$

We can now specify our choice of K . Since a, p, θ are fixed and satisfy (2), there exists a constant $\eta > 0$ such that

$$\pi\theta < \frac{\pi}{2}(1+p)(1-2\eta) - \eta. \tag{4}$$

Fix $\eta > 0$ such that (4) is satisfied. Let C_η and δ be such that (3) is satisfied. Now set $K = \lceil \max(4C_\eta, 1000/\eta, 1000/\delta) \rceil$.

With K fixed (and with it our rough and fine tilings), we can now define tile colourings which we will use as discrete proxies for the spread of an infection in T_n . We assign colours to the tiles of \mathcal{R} and \mathcal{F} as follows: a tile $T \in \mathcal{F}$ is coloured **white** if either it contains fewer than $(1-\eta)p|T|$ initially infected points at the start of the bootstrap percolation process, or it contains fewer than $(1-\eta)|T|$ points in total. Otherwise, we colour T **red** if all its points are infected by the end of the bootstrap percolation process, and **blue** if this is not the case. Furthermore, we colour a tile in \mathcal{R} **white** if one of its subtiles in \mathcal{F} is coloured white, **red** if all its subtiles in \mathcal{F} are coloured red, and **blue** otherwise.

We will be interested in the interface between red and non-red tiles in \mathcal{R} . We thus equip \mathcal{R} with the natural square-grid graph structure by decreeing that two tiles in \mathcal{R} are adjacent if they meet in a side. (Here we identify $[0, \sqrt{n}]^2$ with $(\mathbb{R}/\sqrt{n}\mathbb{Z})^2$ in the natural way to ensure that, as is the case in the torus, the tiles in the rightmost column in \mathcal{R} are adjacent to the corresponding tiles in the leftmost column, and similarly the tiles in the topmost row in \mathcal{R} are adjacent to the corresponding tiles in the bottommost column.) By a *red component* in \mathcal{R} , we mean a connected component of red tiles in the graph H thus defined on \mathcal{R} .

We are now in a position to state the main result of this subsection.

Theorem 4. *Let p, θ be fixed. Suppose p, θ satisfy (2) and let K be as defined above. Then there exist constants $N \in \mathbb{N}$ and $\varepsilon > 0$ such that w.h.p., if there is a connected component of at least N red tiles in the auxiliary graph H on \mathcal{R} , then there exists a giant connected component of $|\mathcal{R}| - o(n^{1-\varepsilon})$ red tiles in H . Furthermore, the non-red tiles consist of a collection of $o(n^{1-\varepsilon})$ vertex-disjoint connected subgraphs of H , each of which has order at most N .*

In other words, once the growing condition is satisfied, any sufficiently large infection spreads to most of T_n , leaving only $o(n^{1-\varepsilon})$ isolated islands of diameter $O(\log n)$ uninfected.

The proof of Theorem 4 relies on two main ingredients. First of all, we shall use (3) and the fine tiling to show that, in the absence of fine white tiles, an infection will spread radially outwards from a sufficiently large infected disc (this is the content of Lemma 3). Next, we shall use the Bollobás–Leader discrete isoperimetric inequality in the toroidal grid to show that any large component of rough red tiles has a large boundary. Combining these results with some probabilistic estimates showing that white tiles are few and far apart will then yield the final result.

In the first part of the proof, we shall use the following technical lemma, which follows from [20, Lemma 8].

Proposition 3. *Let Γ be a continuous, piecewise continuously differentiable curve in T_n . Let $\ell(\Gamma)$ be the length of the curve Γ . Then Γ meets at most $9K\ell(\Gamma)/r$ tiles of \mathcal{F} .*

Lemma 3 (Growing lemma.) *With δ and C_η as defined just before (3), suppose \mathbf{x} is a point in T_n and $R \geq C_\eta r$ is a real number such that the following hold:*

- (i) *all fine tiles that are wholly contained inside $B_R(\mathbf{x}) \setminus B_{R-2r}(\mathbf{x})$ are coloured red;*
- (ii) *no fine tile wholly contained inside $B_{R+2r}(\mathbf{x}) \setminus B_{R-2r}(\mathbf{x})$ is coloured white.*

Then all fine tiles that are wholly contained inside $B_{R+\delta r}(\mathbf{x}) \setminus B_{R-2r}(\mathbf{x})$ are coloured red.

Proof. Any fine tile wholly contained inside $B_{R+\delta r}(\mathbf{x}) \setminus B_{R-2r}(\mathbf{x})$ either is wholly contained inside $B_R(\mathbf{x}) \setminus B_{R-2r}(\mathbf{x})$ (and hence coloured red), or, by our choice of K , is wholly contained inside $B_{R+\delta r}(\mathbf{x}) \setminus B_{R-\delta r}(\mathbf{x})$ (and hence not coloured white).

It is thus enough to show that every vertex $\mathbf{y} \in B_{R+\delta r}(\mathbf{x}) \setminus B_{R-\delta r}(\mathbf{x}) \cap \mathcal{P}$ is eventually infected by our bootstrap percolation process. Now by (3), the lens $L(\mathbf{y}) := B_R(\mathbf{x}) \cap B_r(\mathbf{y})$ has area at least $\frac{\pi r^2}{2}(1 - \eta)$. Since this lens is contained inside the annulus $B_R(\mathbf{x}) \setminus B_{R-2r}(\mathbf{x})$, every fine tile wholly contained inside $L(\mathbf{y})$ is coloured red by Assumption (i). Furthermore, every other tile wholly contained inside the disc $B_r(\mathbf{y})$ is not coloured white by Assumption (ii). We use this information, together with Proposition 3, to show that \mathbf{y} sees strictly more than $\theta\pi r^2$ infected points within distance at most r of itself—which in turn implies that \mathbf{y} must become infected before the end of the process, as required.

First of all, the boundary of $L(\mathbf{y})$ has length at most $2\pi r$, whence by Proposition 3 it intersects at most $18\pi K$ distinct fine tiles. It follows that $L(\mathbf{y})$ must wholly contain a collection of red fine tiles of combined area at least

$$|L(\mathbf{y})| - 18\pi K \frac{c^2 r^2}{K^2} > \frac{\pi r^2}{2}(1 - \eta) - \frac{\eta r^2}{2},$$

where the inequality follows for n large enough from our lower bound on $|L(\mathbf{y})|$, our choice of $K \geq 1000/\eta$, and the fact that $c = 1 + o(1)$.

Similarly, the boundary of $B_r(\mathbf{y})$ has length $2\pi r$, and thus, applying Proposition 3 as above, $B_r(\mathbf{y})$ must wholly contain a collection of non-white fine tiles of combined area at least $\pi r^2 - \frac{\eta r^2}{2}$. Given the definition of our colouring of fine tiles, it follows that $B_r(\mathbf{y})$ contains at least

$$\left(\frac{\pi r^2}{2}(1 - \eta) - \frac{\eta r^2}{2}\right)(1 - \eta) + \left(\frac{\pi r^2}{2}(1 + \eta)\right)p(1 - \eta) > \frac{\pi r^2}{2}(1 + p)(1 - 2\eta) - \frac{\eta r^2}{2}$$

infected points, which by (4) is strictly more than $\theta\pi r^2$. Thus \mathbf{y} is eventually infected by the bootstrap percolation process, and the lemma follows. □

Lemma 3 implies that if a red rough tile is part of the vertex boundary of a connected component of red rough tiles in the graph H on \mathcal{R} , then there must be a white rough tile in the vicinity.

Corollary 1. *Let $T \in \mathcal{R}$ be coloured red. Then either all neighbours of T in the auxiliary graph H on \mathcal{R} are coloured red, or there exists a rough tile T' at graph distance at most 3 of T in H such that T' is coloured white.*

Proof. Suppose that no rough tile T' at graph distance at most 3 of T in H is coloured white. Let \mathbf{x} denote the centre of the tile T . By assumption, every fine tile wholly contained inside the ball of radius $\frac{Kc r}{2} > C_\eta r$ about \mathbf{x} is coloured red. Furthermore, no fine tile wholly contained inside the disc of radius $\sqrt{10}Kc r/2 + (2 + \delta)r$ about \mathbf{x} is coloured white, and the disc $B_{\sqrt{10}Kc r/2}(\mathbf{x})$ wholly contains the four rough tiles adjacent to T in the auxiliary graph H . We may therefore repeatedly apply Lemma 3 to show inductively that the discs of radii $\frac{Kc r}{2} + j\delta r$ for $j = 1, 2, \dots, \lceil (\sqrt{10}/2 - 1/2)Kc/\delta \rceil$ contain only fine red tiles. The last of these discs contains all of $B_{\sqrt{10}Kc r/2}(\mathbf{x})$, and in particular contains all four of the rough tiles adjacent to T in H . Therefore all fine tiles lying inside these four neighbours of T in H are coloured red, so that all the neighbours of T in H are coloured red themselves. □

We shall use Corollary 1 to show that if a connected component of red tiles in H has a large boundary, then we may find a large connected component of white tiles in a sufficiently large

power of H . Recall that the t th power of H , denoted by H^t , is the graph on the vertex set of H in which all pairs of distinct vertices lying at graph distance at most t in H are joined by an edge.

To make this argument formal, we need the standard notions of edge boundary and dual cycles. Given a subset $A \in \mathcal{R}$, the *edge boundary* $\partial_e(A)$ is the collection of pairs (T_1, T_2) such that $T_1 \in A, T_2 \notin A$, and $\{T_1, T_2\}$ is an edge of H . The dual H^* of the graph H is the graph whose vertices are the corners of rough tiles in \mathcal{R} and whose edges are the sides of rough tiles in \mathcal{R} . It can be shown (see e.g. [14, Lemma 1, Chapter 1]) that the edge boundary of a connected component C in H corresponds to a union of cycles in the dual graph H^* .

We can now prove that to each cycle C^* in the edge boundary of a red component C in H , we may associate a (large) connected component of white tiles in the seventh power H^7 of H .

Lemma 4. *Let C be a connected component of red tiles in H . Let C^* be a dual cycle of (graph) length ℓ in the edge boundary of C in H . Let T_0 be an arbitrary red tile in C one of whose sides corresponds to an edge of the dual cycle C^* . Then there exists a connected component of white rough tiles in H^8 of order at least $\min(1, \ell/100)$, one of whose tiles is at graph distance at most 3 of T_0 in H .*

Proof. By Corollary 1, we know that for every pair $e_i = (T_1^i, T_2^i)$ in the edge boundary of C , there exists a white rough tile T_i within graph distance at most 3 of T_1^i . Let $e_1^*, e_2^*, \dots, e_\ell^*$ be the edges of the dual cycle C^* , and let e_1, e_2, \dots, e_ℓ be the corresponding pairs from $\partial_e(C)$. Assume without loss of generality that e_1^* is a side of T_0 .

Traversing the edges $e_1^*, e_2^*, \dots, e_\ell^*$ of C^* in order, we may thus obtain a sequence of rough white tiles T_1, T_2, \dots, T_ℓ , where T_i and T_{i+1} are at graph distance at most 8 of each other in H . Since there are 25 rough tiles within distance at most 3 of a given white rough tile in H , it follows that no white tile may be repeated more than 100 times in our sequence (since each tile has four sides, each of which occurs at most once as an edge e_i^* in C^*). Thus there must be a component of at least $\ell/100$ rough white tiles in H^8 , one of which (namely T_1) is within graph distance at most 3 of T_0 in H . □

Thus, if a red component in H has a large dual cycle in its boundary, there must exist a large white component in H^8 . On the other hand, as we now show, w.h.p. there are no large white components in H^8 . Set $k := \lfloor \sqrt{n}/Kr \rfloor$. Our auxiliary graph H on the set of rough tiles \mathcal{R} is thus a $k \times k$ toroidal graph.

Lemma 5. *There exists a constant $\varepsilon = \varepsilon(a, \eta, K) > 0$ such that the following hold:*

- (i) *the probability that a rough tile is coloured white is at most $n^{-\varepsilon}$;*
- (ii) *w.h.p. there are $O(n^{1-\varepsilon}/\sqrt{\log n})$ white rough tiles in H ;*
- (iii) *w.h.p. connected components of white rough tiles in H^8 have size at most $2/\varepsilon$ in H^8 .*

Proof. Part (i) of the lemma follows from Lemma 1 (applied to the K^4 fine subtiles of a rough tile) and Markov’s inequality. Using Part (i), the expected number of white rough tiles in H is at most $k^2 n^{-\varepsilon} = O(n^{1-\varepsilon}/\log n)$. Applying Markov’s inequality again, we obtain the second part of the lemma.

For the third part, observe that the graph H^8 has maximum degree less than 4^8 , and that each tile is coloured white independently of all other tiles. Now it is well known (see e.g. [11, Problem 45]) that the number of connected subgraphs of order n containing a given vertex v_0 in a graph of maximum degree Δ is at most $(e\Delta)^n$. Consequently, the expected number of connected subgraphs of white rough tiles of size $\ell > 2/\varepsilon$ in H^8 is at most

$$k^2(4^8 e)^\ell n^{-\ell \varepsilon} = o(1),$$

whence Markov’s inequality tells us that w.h.p. no such connected subgraph exists. It follows that w.h.p. all connected components of white rough tiles in H^8 have size at most $2/\varepsilon$ in H^8 , as claimed. \square

Putting Lemmas 4 and 5 together, we have that w.h.p. there is no red component in H with a large dual cycle in its edge boundary. We now bring in the second main ingredient of the proof of Theorem 4, namely a discrete isoperimetric inequality in the toroidal grid due to Bollobás and Leader [13], in order to show that in such circumstances if there is a large red component \mathcal{C} in H , then this component \mathcal{C} is unique and all components in $H - \mathcal{C}$ are small.

The 2-dimensional case of the Bollobás–Leader edge-isoperimetric inequality for toroidal grids [13, Theorem 8] is as follows.

Proposition 4. (*Bollobás–Leader edge-isoperimetric inequality.*) *Let A be a subset of \mathcal{R} with $|A| \leq k^2/2$. Then*

$$|\partial_e(A)| \geq \min \left(4\sqrt{|A|}, 2k \right).$$

Lemma 6. *Let $N_0 \in \mathbb{N}$ be fixed. Then for every n sufficiently large, the following holds: if \mathcal{C} is a connected component of order at least $(N_0)^2$ in H and has the property that every dual cycle in the edge boundary of \mathcal{C} has length at most N_0 , then every connected component in $H[V(H) \setminus \mathcal{C}]$ has order strictly less than $(N_0)^2$.*

Proof. Let N_0 be fixed, and let n be sufficiently large so as to ensure that $2k > N_0$. Let $\{\mathcal{C}_i : i \in I\}$ be the collection of connected components in $H[V(H) \setminus \mathcal{C}]$. Note that each such component \mathcal{C}_i sends an edge to \mathcal{C} in H , and sends no edge to \mathcal{C}_j for $j \neq i$. For every $i_0 \in I$, consider the edge boundary of \mathcal{C}_{i_0} . Since both \mathcal{C}_{i_0} and its complement $\mathcal{C} \cup \bigcup_{i \in I \setminus \{i_0\}} \mathcal{C}_i$ induce connected subgraphs in H , the edge boundary of \mathcal{C}_{i_0} (which is a subset of the edge boundary of \mathcal{C}) consists of a single dual cycle $\mathcal{C}_{i_0}^*$ in H^* .

If $|\mathcal{C}_{i_0}| > k^2/2$, then by Proposition 4 applied to the complement $\mathcal{C} \cup \bigcup_{i \in I \setminus \{i_0\}} \mathcal{C}_i$ of \mathcal{C}_{i_0} , we have that this dual cycle $\mathcal{C}_{i_0}^*$ has length at least

$$|\partial_e(\mathcal{C}_{i_0})| \geq \min \left(4\sqrt{k^2 - |\mathcal{C}_{i_0}|}, 2k \right) \geq \min \left(4\sqrt{|\mathcal{C}|}, 2k \right) > N_0,$$

by our assumptions that $|\mathcal{C}| \geq (N_0)^2$ and $2k > N_0$. This contradicts the fact that every dual cycle in the edge boundary of \mathcal{C} has length at most N_0 . Thus it must be the case that $|\mathcal{C}_{i_0}| \leq k^2/2$.

Then, applying Proposition 4 to \mathcal{C}_{i_0} itself, we have that the dual cycle $\mathcal{C}_{i_0}^*$ has length at least

$$|\partial_e(\mathcal{C}_{i_0})| \geq \min \left(4\sqrt{|\mathcal{C}_{i_0}|}, 2k \right).$$

Since by assumption every dual cycle in the edge boundary of \mathcal{C} has length at most N_0 , and since $2k > N_0$, it follows from the above that $|\mathcal{C}_{i_0}| \leq (N_0)^2/16 < (N_0)^2$, as required. \square

We are now ready to complete the proof of Theorem 4.

Proof of Theorem 4. Let $\varepsilon > 0$ be the constant whose existence is guaranteed by Lemma 5, let $N_0 = \lceil 200/\varepsilon \rceil$, and let $N = (N_0)^2$.

Now suppose that H contains a red component \mathcal{C} of order $|\mathcal{C}| \geq N$. By Lemma 5(iii), we know that w.h.p. all connected components of white rough tiles in H^8 have size at most $N_0/100$. By Lemma 4, this in turn implies that w.h.p. all dual cycles in the edge boundary

of \mathcal{C} have length at most N_0 . This enables us to apply Lemma 6 and deduce that w.h.p. all connected components in $H - \mathcal{C} := H[V(H) \setminus \mathcal{C}]$ have order strictly smaller than $N = (N_0)^2$. In particular, \mathcal{C} is w.h.p. unique, and therefore w.h.p. all other red components must have order strictly less than N .

It remains to bound the number of components in $H - \mathcal{C}$. For this we again apply Lemma 4. For each connected component \mathcal{C}' in $H - \mathcal{C}$, we choose a pair of adjacent tiles (T_0, T') with $T_0 \in \mathcal{C}$ and $T' \in \mathcal{C}'$. Lemma 4 then implies the existence of a rough white tile W at graph distance at most 3 from T_0 in H , and so at graph distance at most 4 from \mathcal{C}' . Each such W can thus be associated with at most 4^4 components \mathcal{C}' . By Lemma 5(iii), w.h.p. there are at most $o(n^{1-\varepsilon})$ rough white tiles. Therefore w.h.p. there can be at most $o(n^{1-\varepsilon})$ components \mathcal{C}' in $H - \mathcal{C}$, each of which has order at most N . It follows that w.h.p.

$$|\mathcal{C}| \geq |\mathcal{R}| - o(n^{1-\varepsilon}),$$

i.e. that w.h.p. \mathcal{C} is a giant red connected component covering all but $o(n^{1-\varepsilon})$ tiles in H , as claimed. This concludes the proof of Theorem 4. □

3.2. Outbreaks stay local: the case $\theta > \frac{1+p}{2}$

Fix $a > 1$, and let $\pi r^2 = a \log n$. In this subsection, our goal is to show that if p, θ are fixed and satisfy the *non-growing condition*

$$\theta > \frac{1+p}{2}, \tag{5}$$

then w.h.p. any outbreak in $T_n := T_n^2$ remains local. To state our formal result (Theorem 5), we must, as in the previous subsection, introduce two tilings of $[0, \sqrt{n}]^2$. Let $K \in \mathbb{N}$ be a large constant to be specified later.

Definition 3. (*Rough tiling, fine tiling.*) *The rough tiling \mathcal{R} partitions $[0, \sqrt{n}]^2$ into disjoint $cKr \times cKr$ square tiles, where $c = \frac{\sqrt{n}}{Kr \lfloor \sqrt{n}/Kr \rfloor} = 1 + o(1)$ is chosen to ensure that divisibility conditions are satisfied. The fine tiling \mathcal{F} is a refinement of \mathcal{R} obtained by subdividing each tile of \mathcal{R} into K^4 smaller $\frac{cr}{K} \times \frac{cr}{K}$ square tiles.*

Remark 1. Note that these tilings are technically distinct from those we used in the previous subsection. (We will pick a different value of K .)

Let p, θ be fixed. Suppose p, θ satisfy (5). Then for any $\eta > 0$ there exists a constant $C_\eta > 1/\sqrt{2}$ sufficiently large so that the area of the lune $B_r((C_\eta r, 0)) \setminus B_{C_\eta r}(\mathbf{0})$ is at most

$$|B_r((C_\eta r, 0)) \setminus B_{C_\eta r}(\mathbf{0})| \leq \frac{\pi r^2}{2}(1 + \eta). \tag{6}$$

We can now specify our choice of K . Since a, p, θ are fixed and satisfy (5), there exists a constant $\eta > 0$ such that

$$\pi\theta > \frac{\pi}{2}(1+p)(1+\eta)^2 + \eta. \tag{7}$$

Fix $\eta > 0$ such that (7) is satisfied. Let C_η be such that (6) is satisfied. Now set $K = \lceil \max(2C_\eta, 10000/\eta) \rceil$.

With K fixed (and with it our rough and fine tilings), we can now define tile colourings which we will use as discrete proxies for the spread of an infection in T_n . We assign colours

to the tiles of \mathcal{R} and \mathcal{F} as follows: a tile $T \in \mathcal{F}$ is coloured **black** if either it contains strictly more than $(1 + \eta)p|T|$ points of A_0 , or it contains strictly more than $(1 + \eta)|T|$ points of \mathcal{P} . Otherwise, we colour T **red** if some of its initially uninfected points become infected at some stage in the bootstrap percolation process, and **blue** if this is not the case. Furthermore, we colour a tile in \mathcal{R} **black** if one of its subtiles in \mathcal{F} is coloured black, **red** if one of its subtiles in \mathcal{F} is coloured red, and **blue** otherwise.

We equip \mathcal{R} with the natural square-grid graph structure by decreeing that two tiles in \mathcal{R} are adjacent if they meet in a side, and as in the previous subsection, we identify $[0, \sqrt{n}]^2$ with $(\mathbb{R}/\sqrt{n}\mathbb{Z})^2$ in the natural way. We thus obtain an auxiliary toroidal grid-graph H on \mathcal{R} . We can now state the main result of this subsection.

Theorem 5. *Let p, θ be fixed. Suppose p, θ satisfy (5), and let K and our rough tiling be as defined above. Then there exists $\varepsilon > 0$ such that for any constant N , even if one fully infects all points inside N adversarially chosen tiles of \mathcal{R} , w.h.p. the following will hold:*

- (i) *all but $o(n^{1-\varepsilon})$ tiles of \mathcal{R} are coloured blue;*
- (ii) *the non-blue tiles form a collection of $o(n^{1-\varepsilon})$ vertex-disjoint connected subgraphs of H^2 , each of which has order at most $100N^2/\varepsilon^2$.*

The key to Theorem 5 is the following lemma, showing that a non-black rough tile surrounded by non-black tiles will be coloured blue.

Lemma 7. *Suppose T is a tile in \mathcal{R} . Suppose none of the tiles in the 3×3 square grid of tiles of \mathcal{R} centred at T is coloured black. Then T is coloured blue.*

Proof. Let \mathbf{x} denote the centre of the tile T . We shall prove the stronger claim that no initially uninfected vertex in the ball $B_{cKr/\sqrt{2}}(\mathbf{x}) \supseteq T$ ever becomes infected.

Indeed, suppose $t \geq 0$ and no point in $B_{cKr/\sqrt{2}}(\mathbf{x}) \setminus A_0$ has yet become infected. Consider any such point v . Which infected points does v see within distance r of itself? In a worst-case scenario, every point in the lune $\text{Lune}(v) = B_r(v) \setminus B_{cKr/\sqrt{2}}(\mathbf{x})$ has become infected. We show that even if this was the case, v does not become infected in the next round of the bootstrap percolation process.

Observe first of all that the length of the boundary $\partial\text{Lune}(v)$ of the lune $\text{Lune}(v)$ is at most $2\pi r$, and thus $\partial\text{Lune}(v)$ meets at most $18\pi K$ tiles of \mathcal{F} by Proposition 3. Similarly, the boundary of the asymmetric lens $\text{Lens}(v) = B_r(v) \cap B_{cKr/\sqrt{2}}(\mathbf{x})$ has length at most $2\pi r$ and thus meets at most $18\pi K$ tiles of \mathcal{F} .

Since both T and the eight tiles around it are not coloured black, each fine tile wholly contained in $\text{Lens}(v)$ contains at most $(1 + \eta)p\frac{c^2r^2}{K^2}$ points of A_0 , and contains no other point of A_t by our assumption. Furthermore, every other fine tile having non-empty intersection with $B_r(v)$ contains at most $(1 + \eta)\frac{c^2r^2}{K^2}$ points of \mathcal{P} in total, and thus at most that many points of A_t .

By (6) and our choices of η and K , we have that the area of $\text{Lune}(v)$ is at most $\frac{\pi r^2}{2}(1 + \eta)$ (since the area of the lune is maximized if v lies on the circle of radius $cKr/\sqrt{2}$ about \mathbf{x} , and $cK/\sqrt{2} > C_\eta$). It follows that for n large enough the number of infected points of A_t within distance r of v is at most

$$\begin{aligned} |B_r(v) \cap A_t| &\leq (1 + \eta)|\text{Lune}(v)| + (1 + \eta)p(\pi r^2 - |\text{Lune}(v)|) + 36\pi K(1 + \eta)\frac{c^2r^2}{K^2} \\ &< \pi r^2 \left(\frac{1+p}{2}\right) (1 + \eta)^2 + \eta r^2 < \theta \pi r^2. \end{aligned}$$

Here, the first strict inequality follows from our bound on the area of $\text{Lune}(v)$, the facts that $\eta \leq \pi$ and $c < 2$ for n large enough, and from our choice of K ensuring that $144(1 + \pi)\pi/K < \eta$. The second strict inequality follows from (7).

In particular, v sees strictly fewer than $\theta a \log n$ infected points from A_t , and does not become infected in the next round of the bootstrap percolation process. Since $v \in B_{cKr/\sqrt{2}}(\mathbf{x}) \setminus A_0$ was arbitrary, it follows by induction on t that $B_{cKr/\sqrt{2}}(\mathbf{x}) \setminus A_0 = B_{cKr/\sqrt{2}}(\mathbf{x}) \setminus A_\infty$. Since the rough tile T is a subset of $B_{cKr/\sqrt{2}}(\mathbf{x})$ and is not coloured black, it follows that T is coloured blue as claimed. \square

Similarly to Lemma 5, we now prove the following.

Lemma 8. *Let $p \in (0, 1)$ be fixed. Then there exists a constant $\varepsilon = \varepsilon(a, \eta, K) > 0$ such that w.h.p. the following hold:*

- (i) *there are $O(n^{1-\varepsilon}/\sqrt{\log n})$ black rough tiles in H ;*
- (ii) *connected components of black rough tiles in H^7 have diameter at most $1/\varepsilon$.*

Proof. By Lemma 1 (applied to the K^4 fine subtiles of a rough tile) and Markov’s inequality, there exists a constant $\varepsilon = \varepsilon(a, \eta, K) > 0$ such that the probability that a rough tile is coloured black is at most $n^{-\varepsilon}$ for all n sufficiently large. Thus the expected number of black rough tiles in H is at most $\left(\frac{\sqrt{n}}{Kr}\right)^2 n^{-\varepsilon} = O(n^{1-\varepsilon}/\log n)$. Applying Markov’s inequality again, we obtain the first part of the lemma.

For the second part, observe that the graph H^7 has maximum degree less than 4^7 , and that each rough tile is coloured black independently of all other tiles. In particular, the expected number of paths of black rough tiles of length $\ell > 1/\varepsilon$ in H^7 is at most

$$\left(\frac{\sqrt{n}}{Kr}\right)^2 4^{7\ell} n^{-(\ell+1)\varepsilon} = O(n^{-\varepsilon}) = o(1),$$

whence Markov’s inequality tells us that w.h.p. no such path exists. It follows that w.h.p. all connected components of black rough tiles in H^7 have diameter at most $1/\varepsilon$, as claimed. \square

We can now prove Theorem 5.

Proof of Theorem 5. Let $\varepsilon = \varepsilon(\eta, K) > 0$ be as in Lemma 8. Consider the non-blue tiles in \mathcal{R} . By Lemma 7, every non-blue tile must be within distance at most 2 in H of either a black tile or of one of the at most N adversarially infected tiles.

By Lemma 8(i), this immediately implies that w.h.p. all but $o(n^{1-\varepsilon})$ tiles of \mathcal{R} are coloured blue. This establishes Part (i) of Theorem 5.

Furthermore, by Lemma 8(ii), w.h.p. every connected component of black rough tiles has diameter at most $1/\varepsilon$ in H^7 . It follows that every component of black or adversarially infected rough tiles has diameter at most $(N + 1)/\varepsilon$ in H^7 , and hence order at most $4N^2/\varepsilon^2$.

Now, by Lemma 7, to every connected component of non-blue tiles in H^2 , one may associate a connected component of black or adversarially infected rough tiles in H^7 (since every non-blue tile must be within distance 2 of a black or adversarially infected rough tile). Since there are fewer than 25 rough tiles within distance at most 2 in H of a given rough tile, it follows from our earlier bound on the order of connected components of black or adversarially infected rough tiles in H^7 that w.h.p. every connected component of non-blue tiles in H^2 must have order at most $100N^2/\varepsilon^2$. This establishes Part (ii) and concludes the proof of the theorem. \square

3.3. Proof of Theorem 1

With our tiling results Theorems 4 and 5 in hand, we can prove the main results of this paper.

Proof of Theorem 1(i). Suppose (a, p, θ) is fixed with $a > 1$ and (p, θ) satisfying the growing condition (2). Then there exists $\eta = \eta(p, \theta) > 0$ such that (4) is satisfied. We can then define $K = K(\eta)$ and the rough tiling \mathcal{R} of T_n as in Theorem 4. Let $N \in \mathbb{N}$ and $\varepsilon > 0$ be constants such that the conclusions of Theorem 4 hold (note that our choice of these constants depends only on a, p, θ). By Lemma 5 we may also ensure with our choice of ε that the probability that a rough tile in \mathcal{R} is coloured white is at most $n^{-\varepsilon}$, provided n is taken sufficiently large.

Now we select a constant $C = C(a, p, \theta)$ sufficiently large so that any ball of radius Cr in T_n wholly contains a connected component in H of at least N non-white rough tiles of \mathcal{R} . This requires a little calculation. Set $M = \lceil \frac{2N}{\varepsilon} \rceil$. Observe that there are at most n ways of choosing an $M \times M$ square grid of rough tiles in \mathcal{R} . Taking a simple union bound, the probability that there are at least $\frac{2}{\varepsilon}N$ white rough tiles in such an $M \times M$ grid is at most

$$2^{M^2} (n^{-\varepsilon})^{\frac{2}{\varepsilon}N} = O(n^{-2N}).$$

By Markov’s inequality, it follows that w.h.p. every $M \times M$ grid of rough tiles in \mathcal{R} contains fewer than $\frac{2}{\varepsilon}N$ white rough tiles. In particular, every such grid contains a column of $\frac{2}{\varepsilon}N > N$ non-white rough tiles.

We now take C sufficiently large so as to ensure that every ball of radius Cr wholly contains an $M \times M$ square grid of rough tiles. This is easily done: $C = 4MKc$ will certainly do, for instance. It follows that if we infect any ball B of radius Cr in T_n , then, changing the colour of the non-white rough tiles wholly contained inside B to red, this yields w.h.p. a connected component of red rough tiles of order at least N in H . Applying Theorem 4, we obtain that all but $o(n^{1-\varepsilon})$ of the rough tiles from \mathcal{R} are coloured red following the bootstrap percolation process started from the initially infected set $A_0 \cup (B \cap \mathcal{P})$.

By Lemma 1, w.h.p. every rough tile in \mathcal{R} contains at most $O(\log n)$ vertices, so we deduce from the above that all but $O(n^{1-\varepsilon} \log n) = o(n)$ vertices of \mathcal{P} eventually become infected.

This leaves only the ‘Furthermore’ part of Theorem 1(i) to establish. Consider a component of non-red rough tiles in H^2 . Such a component is a union of disjoint components of non-red rough tiles in H , connected by paths of length 2 in H where the middle tile T_0 in the path is red and the two other tiles are non-red and belong to distinct components. By Corollary 1, we know that there must be a white tile at distance at most 3 from T_0 . Much as in the proof of Theorem 4, it then follows that to the edge boundary in H of an H^2 -connected component of non-red tiles we may associate an H^8 -connected component of white tiles W . By Lemma 5, we know that w.h.p. all such components W have diameter at most $1/\varepsilon$.

There are at most $|W|4^3$ red tiles T_0 within graph distance 3 in H of a white tile in W , each red tile T_0 is the middle point of a path of length 2 in H between at most 4 distinct non-red components in H , and by Theorem 4 w.h.p. all non-red components in H have order at most N . It follows that all connected components of non-red tiles in H^2 have order at most $(4^8)^{1/\varepsilon}4^34N$.

Now since $K > 1$, to each connected component of forever-uninfected vertices in $G_{n,r}[\mathcal{P} \setminus A_0]$ is associated a collection of non-red rough tiles forming a (subset of a) connected component in H^2 . It follows that the Euclidean diameter in T_n of such a component is bounded above by $(4^8)^{1/\varepsilon}4^34N\sqrt{2Kcr} = O(\sqrt{\log n})$, as claimed. This concludes the proof of Theorem 1(i). \square

Proof of Theorem 1(ii). Suppose (a, p, θ) is fixed with $a > 1$ and (p, θ) satisfying the non-growing condition (5). Then there exists $\eta = \eta(p, \theta) > 0$ such that (7) is satisfied. We can then




FIGURE 2. A lens and a lune.

define $K = K(\eta)$ and the rough tiling \mathcal{R} of T_n as in Theorem 5. Let $\varepsilon > 0$ be such that the conclusion of Theorem 5 holds (note that our choice of ε depends only on a, p, θ).

Let $C > 1$ be fixed. Let B be an arbitrarily chosen ball of radius Cr in T_n . By Proposition 3, B meets at most

$$N = \left\lceil \frac{\pi(Cr)^2}{(Kcr)^2} + 9K(2\pi Cr)/r \right\rceil < 100KC^2$$

tiles of \mathcal{R} . Fully infect all points inside these at most N tiles and apply Theorem 5 to deduce that even if all points in $\mathcal{P} \cap B$ are infected, then w.h.p. all but $o(n^{1-\varepsilon})$ tiles of \mathcal{R} are coloured blue.

By Lemma 1 and a union bound, w.h.p. every tile of \mathcal{R} contains at most $O(\log n)$ points of \mathcal{P} . Since the point-set $A_\infty \setminus (A_0 \cup B)$ is contained inside the union of the non-blue tiles of \mathcal{R} , it follows that $A_\infty \setminus (A_0 \cup B)$ contains at most $o(n^{1-\varepsilon} \log n) = o(n)$ points of \mathcal{P} .

This leaves only the ‘Furthermore’ part of Theorem 1(ii) to establish. We have shown in Theorem 5(ii) that w.h.p. all connected components of non-blue tiles in H^2 have order at most $100N^2/\varepsilon^2 = O(1)$. Since $K > 2$, the tiles containing points of a connected component in $G_{n,r}^2[A_\infty \setminus (A_0 \cup B)]$ must form a connected component of non-blue tiles in H^2 . It immediately follows that each such connected component has Euclidean diameter $O(\sqrt{\log n})$ in T_n . This concludes the proof of Theorem 1(ii). \square

3.4. Starting a sufficiently large local outbreak: proof of Theorem 2

With Theorem 1 in hand, we now turn our attention to the problem of determining when (for which triples (a, p, θ)) a large local outbreak will occur w.h.p. in the Bradonjić–Saniee model for bootstrap percolation on random geometric graphs. We are unfortunately unable to answer this question rigorously, but we can relate it to the solution of an optimization problem which we conjecture gives the correct threshold for ‘large’ local outbreaks.

Definition 4. For $t \in \mathbb{R}_{\geq 0}$, let $R_{\text{lens}}(t)$ denote the asymmetric lens $B_t(\mathbf{0}) \cap B_{1/\sqrt{\pi}}((t, 0))$. Also, let $R_{\text{lune}}(t)$ denote the lune $B_{1/\sqrt{\pi}}((t, 0)) \setminus B_t(\mathbf{0})$ (see Figure 2 for an illustration). Given functions $f, g: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 1}$, we define the quantities

$$\begin{aligned}
 I(f, g)(t) &:= \int_{\mathbf{x} \in R_{\text{lens}}(t)} (pf(\|\mathbf{x}\|) + (1-p)g(\|\mathbf{x}\|)) \, d\mathbf{x} + \int_{\mathbf{x} \in R_{\text{lune}}(t)} pf(\|\mathbf{x}\|) \, d\mathbf{x}, \\
 Q(f, g)(t) &:= \int_{\mathbf{x} \in \mathbb{R}^2: \|\mathbf{x}\| \leq t} p(f(\|\mathbf{x}\|) - 1 - f(\|\mathbf{x}\|) \log[f(\|\mathbf{x}\|)]) \\
 &\quad + (1-p)(g(\|\mathbf{x}\|) - 1 - g(\|\mathbf{x}\|) \log[g(\|\mathbf{x}\|)]) \, d\mathbf{x}, \tag{8}
 \end{aligned}$$

and $q(f, g) := \lim_{t \rightarrow \infty} Q(f, g)(t)$, where $\|\cdot\|$ denotes the standard Euclidean ℓ_2 -norm.

Note that for $f, g \geq 1$, we have $f - 1 - f \log f \leq 0$ and $g - 1 - g \log g \leq 0$. Thus $Q(f, g)$ is a non-increasing function of t and either converges to a limit or goes to $-\infty$.

The motivation for these quantities is as follows. We imagine an infection spreading from the origin $\mathbf{0}$. Around $\mathbf{0}$, at time $t=0$, we see increased densities of both infected and uninfected points. It is natural to assume that these increased densities are radially symmetric. If so, we may specify them by the functions $f \geq 1$ and $g \geq 1$, so that the intensities of (initially) infected and uninfected points at \mathbf{x} are given by $pf(\|\mathbf{x}\|)$ and $(1-p)g(\|\mathbf{x}\|)$ respectively. These functions must both decay to 1 as $\|\mathbf{x}\| \rightarrow \infty$, or else the probability of the associated configuration \mathbf{C} will be zero.

Now we consider the situation as the infection spreads. First we normalize so that the area of the neighbourhood of \mathbf{x} has area 1, instead of $a \log n$. Next suppose that the infection has spread to radius t , i.e., to all of $B_t(\mathbf{0})$. Consider a point \mathbf{x} on the boundary of this disc; without loss of generality $\mathbf{x} = (t, 0)$. Every point inside $R_{\text{lens}}(t)$ will now be infected as a result of the bootstrap percolation process, but in the lune $R_{\text{lune}}(t)$, only those points that were initially infected will be infected. Summing these contributions and integrating, we see that $I(f, g)(t)$ represents the (normalized) number of infections seen by \mathbf{x} ; if this exceeds θ , \mathbf{x} itself will be infected, and the infection will thus propagate radially outwards from $\mathbf{0}$.

Given f and g , and a configuration \mathbf{C} specified by f and g as above, what is the probability $p_{\mathbf{C}}$ that \mathbf{C} actually occurs? For this, we turn to Lemma 1. Although this lemma only applies to finite unions of disjoint regions, the function $q(f, t)$, derived from Lemma 1, represents the (normalized) logarithm of the sought probability $p_{\mathbf{C}}$. Maximizing $q(f, t)$ over all admissible functions f and g , subject to the constraint that $I(f, g)(t) > \theta$ for all $t > 0$, yields an optimal pair (f, g) , and if p and θ are such that the (un-normalized) associated probability $p_{\mathbf{C}}$ exceeds n^{-c} , for some $c < 1$, then such a configuration is bound to occur somewhere in T_n . This is the content of Definition 5 below.

The use of the integral in (8) needs to be justified. For this, we use a tiling argument. Tiling the neighbourhood of (a truncated version of) an optimal configuration \mathbf{C} , we apply Lemma 1 to the disjoint union of finitely many tiles, to define a ‘discretized’ event E , based on \mathbf{C} , whose rigorously calculated probability p_E is well approximated by $p_{\mathbf{C}}$, and whose existence rigorously guarantees that an infection will spread. This is achieved in Claims 1 and 2 below.

Definition 5. (*Symmetric local growth threshold.*) Let $Q_{\max} = Q_{\max}(p, \theta)$ denote the supremum of $q(f, g)$ over all continuous functions $f, g: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 1}$ satisfying

$$I(f, g)(t) > \theta \quad \forall t \in \mathbb{R}_{\geq 0}.$$

Furthermore, for $a > 1$ and $p \in (0, 1)$ fixed, let $\theta_{\text{local}} = \theta_{\text{local}}(a, p)$ be the supremum of all $\theta \leq 1$ such that

$$Q_{\max}(p, \theta) > -1/a. \tag{9}$$

Observe that θ_{local} depends only on the values of a and p , and that it is well-defined and greater than or equal to p , since $Q_{\max}(p, p) = 0$, as can be seen by taking f and g to be constant and equal to 1.

Proof of Theorem 2. Since θ satisfies (1), it follows from the definitions of θ_{local} and Q_{\max} that for any fixed $T \in \mathbb{R}_{\geq 0}$ there exist functions $f, g: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 1}$ and a real $\varepsilon > 0$ such that the following hold:

- (i) for every $t \in [0, T], I(f, g)(t) > \theta + 2\varepsilon$;
- (ii) $Q(f, g)(T + 2) > -1/a + 2\varepsilon$.

We use this information to construct for any constant $C > 0$ a $(C + 2)r$ -bounded event occurring w.h.p. and guaranteeing the existence of a ball of radius Cr that becomes wholly infected over the course of our bootstrap percolation process. As we note in Proposition 5 below, this implies that $\theta_{\text{local}} \leq (1 + p)/2$, and thus that for $\theta < \theta_{\text{local}}$ the growing condition (2) is satisfied. Combining this information with Theorem 1(i) then yields the desired almost percolation.

We now give the details. Pick $T = (C + 2)/\sqrt{\pi}$. Fix $\varepsilon > 0$ and let $K > 0$ be a sufficiently large positive real number to be specified later. Partition \mathbb{R}^2 into a fine grid of interior-disjoint $\frac{\sqrt{a \log n}}{K} \times \frac{\sqrt{a \log n}}{K}$ square tiles with axis-parallel sides. Let \mathcal{F} be the finite family of tiles that are wholly contained inside the ball of radius $(C + 2)r = T\sqrt{a \log n}$ about the origin. Note that $|\mathcal{F}| = O(1)$. Let f, g be functions $f, g: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 1}$ such that Assumptions (i) and (ii) above are satisfied.

Let \mathcal{P}_p and \mathcal{P}_{1-p} be independent Poisson point processes on $B_{T\sqrt{a \log n}}(\mathbf{0})$ with intensities p and $1 - p$ respectively. Let E be the event that for every tile F in \mathcal{F} , F contains

- at least $N_f(F) := \int_{\mathbf{x} \in F} pf(\|\mathbf{x}\|/\sqrt{a \log n})d\mathbf{x}$ points of \mathcal{P}_p , and
- at least $N_g(F) := \int_{\mathbf{x} \in F} (1 - p)g(\|\mathbf{x}\|/\sqrt{a \log n})d\mathbf{x}$ points of \mathcal{P}_{1-p} .

Claim 1. For any $\varepsilon > 0$ fixed, there exists $K_1 > 0$ such that picking $K \geq K_1$ ensures that

$$\mathbb{P}(E) \geq e^{-c \log n + o(\log n)} \quad \text{for some fixed constant } c < 1.$$

Proof. Set M_1 to be the area of a 2-dimensional ball of radius T . Since f, g are continuous functions from the compact set $[0, T]$ to $\mathbb{R}_{\geq 1}$, we have that f and g are bounded above by some $M_2 > 0$ on $[0, T]$ and further that $\log f$ and $\log g$ are uniformly continuous over $[0, T]$. In particular, for every $\varepsilon > 0$ there exists $K_1 > 0$ sufficiently large so that for all $K \geq K_1$, and all $x, y \in [0, T], |x - y| \leq 2/K$ implies that both $|\log [f(x)] - \log [f(y)]|$ and $|\log [g(x)] - \log [g(y)]|$ are less than $\varepsilon/M_1 M_2$.

For each tile $F \in \mathcal{F}$, set $\rho_f(F) := N_f(F)/|F|$ and $\rho_g(F) := N_g(F)/|F|$. By the consequence of uniform continuity observed above, picking $K \geq K_1$ ensures that

$$f\left(\frac{\|\mathbf{x}\|}{\sqrt{a \log n}}\right) \left(\log \left[f\left(\frac{\|\mathbf{x}\|}{\sqrt{a \log n}}\right) \right] - \log [\rho_f(F)] \right) \geq -\varepsilon/M_1$$

for all $\mathbf{x} \in F$. Combining this with Lemma 1 (rescaled by a factor of p), the probability that F contains at least $N_f(F)$ points of \mathcal{P}_p is

$$\begin{aligned} & \exp \{ (\rho_f(F) - 1 - \rho_f(F) \log (\rho_f(F))) p|F| + O(\log \log n) \} \\ & \geq \exp \left\{ \int_{\mathbf{x} \in F} p \left(f\left(\frac{\|\mathbf{x}\|}{\sqrt{a \log n}}\right) - 1 - f\left(\frac{\|\mathbf{x}\|}{\sqrt{a \log n}}\right) \log f\left(\frac{\|\mathbf{x}\|}{\sqrt{a \log n}}\right) \right) d\mathbf{x} \right\} \\ & \cdot \exp \left\{ -p \frac{\varepsilon}{M_1} |F| + O(\log \log n) \right\}. \end{aligned}$$

One may obtain a similar expression for the probability that F contains at least $N_g(F)$ points of \mathcal{P}_{1-p} , substituting $1 - p$ for p and g for f .

Since $\{\mathcal{P}_p \cap F, F \in \mathcal{F}\}$ and $\{\mathcal{P}_{1-p} \cap F, F \in \mathcal{F}\}$ together form a collection of independent random variables, it follows that the probability that every tile $F \in \mathcal{F}$ contains at least $N_f(F)$

points of \mathcal{P}_p and at least $N_g(F)$ points of \mathcal{P}_{1-p} —in other words, the probability of E —is at least

$$\mathbb{P}(E) \geq \exp \{ (Q(f, g)(T) - \varepsilon) a \log n + O(\log \log n) \},$$

which by Assumption (ii) is at least $e^{-(1-a\varepsilon) \log n + o(\log n)}$, proving our claim with the constant $c = 1 - a\varepsilon$. □

Given a non-negative integer i , let $A_i = A_i(\mathcal{F})$ denote the collection of tiles of \mathcal{F} that meet the annulus $B_{(i+1)\sqrt{a \log(n)}/K}(\mathbf{0}) \setminus B_{i\sqrt{a \log(n)}/K}(\mathbf{0})$. Furthermore, let $D_i = D_i(\mathcal{F})$ denote the union of the tiles of \mathcal{F} that are wholly contained inside the disc $B_{i\sqrt{a \log(n)}/K}(\mathbf{0})$. Observe that both these tile families depend on \mathcal{F} (and thus on our choice of K).

Claim 2. *For every $\varepsilon > 0$ fixed, there exists $K_2 > 0$ such that picking $K \geq K_2$ ensures that at least one of the following holds:*

- the event E fails to occur;
- for every $i \leq KT - 1$, every tile $F \in A_i$, and every vertex $\mathbf{x} \in F$, we have that

$$N_{\mathbf{x}} := |B_r(\mathbf{x}) \cap D_i \cap \mathcal{P}_{1-p}| + |B_r(\mathbf{x}) \cap \mathcal{P}_p| > \theta a \log n.$$

Proof. As in the previous claim, we note that the continuous functions f, g are bounded above on the compact set $[0, T]$ by some $M_2 > 0$.

Assume the event E occurs. Fix $i \leq KT - 1$, and $F \in A_i$. Let $\mathbf{x} \in F$ and set $t := \|\mathbf{x}\|/\sqrt{a \log n}$. Denote by $S^1(\mathbf{x})$ the collection of tiles of F wholly contained inside $D_i \cap B_r(\mathbf{x})$, and by $S^2(\mathbf{x})$ the collection of tiles of F wholly contained inside $B_r(\mathbf{x})$. Since E occurs, we have

$$N_{\mathbf{x}} \geq \sum_{F \in S^2(\mathbf{x})} N_f(F) + \sum_{F \in S^1(\mathbf{x})} N_g(F).$$

Now since f, g are bounded above by M_2 , the expression on the right-hand side is at least

$$\begin{aligned} & I(f, g)(t) a \log n - M_2 p \left(a \log n - \sum_{F \in S^2(\mathbf{x})} |F| \right) \\ & - M_2(1 - p) \left(|R_{\text{lens}}| a \log n - \sum_{F \in S^1(\mathbf{x})} |F| \right). \end{aligned}$$

It readily follows from an application of Proposition 3 that for any $\varepsilon > 0$, there exists $K_2 > 0$ such that picking $K \geq K_2$ ensures that the above is at least $(I(f, g) - \varepsilon) a \log n$. Assumption (i) then tells us that

$$N_{\mathbf{x}} \geq (\theta + \varepsilon) a \log n,$$

proving our claim. □

With Claims 1 and 2 in hand, the proof of Theorem 2 is now straightforward.

Consider (a, p, θ) satisfying $a > 1$ and (1). As recorded in Proposition 5 below, this implies $\theta < (1 + p)/2$. Let $C = C(a, \theta, p) > 0$ be a constant such that by Theorem 1(i), the infection of

a ball of radius Cr in T_n leads w.h.p. to almost percolation in our bootstrap percolation process on T_n . This then specifies our choice of $T = (C + 2)/\sqrt{\pi} > 0$.

Let K_1, K_2 be the constants from Claims 1–2. Pick $K \geq \max(K_1, K_2)$. As remarked in Section 2, by standard properties of Poisson point processes we may view the vertex set in the Bradonjić–Saniee model in the torus as being the union of a Poisson process \mathcal{P}_p of initially infected points of intensity p with a Poisson process \mathcal{P}_{1-p} of initially uninfected points of intensity $1 - p$. Partitioning the torus T_n into a fine grid of interior-disjoint square tiles with side length $\sqrt{a \log n}/c'K$, where $c' \geq 1$ is chosen to ensure that divisibility conditions are met, we can define a natural analogue E' of our event E inside T_n .

Clearly if E' occurs, then by applying Claim 2 to A_0, \dots, A_{KT-1} , we see that all vertices inside a ball of radius Cr about the origin in T_n become infected, which then leads to almost percolation. Since E' is $(C + 2)r$ -bounded and has probability $q = e^{-c \log n + o(\log n)}$ for some fixed constant $c < 1$ (by Claim 1), it follows from Lemma 2(ii) that w.h.p. some translate of E' occurs in T_n . Thus, for $a > 1$ fixed, w.h.p. we have almost percolation in the (p, θ) regime satisfying (1), as claimed. □

We conclude this subsection by recording some trivial inequalities between θ_{local} and other quantities of interest, justifying the phase diagram we give in Figure 1.

Proposition 5. *For all $a > 1$ fixed and $p \in [0, 1]$, the following inequalities hold:*

$$p \leq \theta_{\text{local}} \leq \min\left(\frac{1+p}{2}, \theta_{\text{start}}\right).$$

Proof. For the lower bound, observe that taking f, g to be identically 1 we have $I(f, g) \geq p$ for all $t \geq 0$ and $q(f, g) = 0$. It then follows from Definition 5 that $\theta_{\text{local}} \geq p$ as claimed.

For the upper bound, observe first of all that if p, θ are fixed and $\theta > \theta_{\text{start}}$, then w.h.p. no percolation occurs, which by Theorem 2 implies $\theta_{\text{start}} \geq \theta_{\text{local}}$. Next, note that the proof of Theorem 2 implies that for a, p, θ fixed with $a > 1$ and $\theta < \theta_{\text{local}}$, for any fixed $C > 0$, w.h.p. there will be a ball of radius $C\sqrt{\log n}$ in T_n such that every vertex inside it eventually becomes infected. On the other hand, Theorem 1(ii) implies that if (a, p, θ) are fixed with $a > 1$ and $\theta > \frac{1+p}{2}$, then w.h.p. there exists a constant $C > 0$ such that no ball of radius $C\sqrt{\log n}$ in T_n becomes infected in this way. This immediately yields that $\theta_{\text{local}} \leq \frac{1+p}{2}$. □

Remark 2. The results in Appendix C show that $\frac{d}{dp}(\theta_{\text{local}}(p)) \rightarrow \infty$ as $p \rightarrow 0$ and $\frac{d}{dp}(\theta_{\text{local}}(p)) \rightarrow 1/2$ as $p \rightarrow 1$, so that θ_{local} is tangent to θ_{start} at $p = 0$ and to $\theta = \frac{1+p}{2}$ at $p = 1$. In particular, the tangencies at $p = 0$ and $p = 1$ of the curve $\theta = \theta_{\text{local}}(p)$ given in Figure 1 are correct.

3.5. Islands

Beyond the almost percolation guaranteed by Theorem 2, it is natural to ask when full percolation occurs in the Bradonjić–Saniee model. One obstacle for full percolation is the presence of initially uninfected vertices of degree strictly less than $\theta a \log n$. Proposition 2 gives us the (implicit) θ threshold for the w.h.p. disappearance of such vertices. However there may be other, likelier, obstacles to full percolation, namely ‘islands’ of initially uninfected vertices with few initially infected vertices, surrounded by sparsely populated ‘lakes’ containing unusually few vertices.

Much as in the previous subsection, we are unable to rigorously determine the threshold for the w.h.p. disappearance of islands of uninfected vertices. We are, however, able to determine

the threshold for the disappearance of symmetric islands, which is given by the solution of an explicit continuous optimization problem and can be determined explicitly using the method of Lagrange multipliers.

As the arguments involved are very similar to those in the previous section, we give only a minimal level of detail. The basic idea is quite simple: we look for a radially symmetric distribution of infected/non-infected point densities which guarantees that some island cannot be infected from the outside. To such a distribution we associate a local $O(r)$ -bounded event, from which we can in turn derive a threshold for the w.h.p. disappearance of such islands—the proof essentially follows that of Theorem 2, mutatis mutandis.

To make this more precise, let us give analogues of Definition 4 tailored to the island (rather than local outbreak) setting.

Definition 6. Let $T > 0$ be fixed. For $t \in [0, T]$, write $R_{\text{inner}}(t)$ for the asymmetric lens $B_T(\mathbf{0}) \cap B_{1/\sqrt{\pi}}((t, 0))$ and $R_{\text{outer}}(t)$ for the (possibly empty) lune $B_{1/\sqrt{\pi}}((t, 0)) \setminus B_T(\mathbf{0})$. A symmetric T -island distribution is a pair (f, g) of continuous functions $f, g: \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ such that for every $t \in [0, T]$, the following holds:

$$\int_{\mathbf{x} \in R_{\text{inner}}(t)} pf(\|\mathbf{x}\|)d\mathbf{x} + \int_{\mathbf{x} \in R_{\text{outer}}(t)} (pf(\|\mathbf{x}\|) + (1 - p)g(\|\mathbf{x}\|)) d\mathbf{x} < \theta. \tag{10}$$

Given a symmetric T -island distribution (f, g) , we define its weight $q(f, g)$ as in (8) (just the equation, not the limit).

Definition 7. (Symmetric islands threshold.) For every $T \geq 0$, the optimal T -island weight $q_{\text{max}}(T) = q_{\text{max}}(T)(p, \theta)$ is the supremum of the weight $q(f, g)$ over all symmetric T -island distributions (f, g) . The symmetric islands threshold $\theta_{\text{islands}} = \theta_{\text{islands}}(a, p)$ is then defined to be the supremum of the $\theta \leq 1$ such that

$$\sup_{T \geq 0} q_{\text{max}}(T) > -1/a. \tag{11}$$

We note that the quantity θ_{islands} can in principle be computed using Euler–Lagrange equations—see Appendix A. However, the solution will be implicit rather than explicit. Let us record here, however, the simple fact that

$$\theta_{\text{islands}} \leq \frac{1 + p}{2}. \tag{12}$$

Indeed, let (a, p) and $\varepsilon > 0$ be fixed. Pick a constant $T > 0$ sufficiently large so that the area of $R_{\text{inner}}(T)$ is at least $\frac{1}{2} - \varepsilon$. Then it is easily checked that taking the functions f, g to be identically 1 gives a symmetric T -island distribution (f, g) with weight $q(f, g) = 0 > -1/a$ for any $\theta > \frac{1+p}{2} + \varepsilon(1 - p)$. The inequality (12) follows immediately.

Proof of Theorem 3. This is a simple modification of the proof of Theorem 2: if $\theta > \theta_{\text{islands}}$, then there exists a symmetric T -island distribution (f, g) such that $q(f, g) > -1/a + 2\varepsilon$. As in the proof of Theorem 2, passing to a fine tiling of a ball of radius $(T + 2)\sqrt{a \log n}$, one may use (f, g) to define a tiling event E in T_n^2 such that (i) E is an $O(r)$ -bounded event with probability $\mathbb{P}(E) = e^{-(c+o(1)) \log n}$, where $c < 1$ is a fixed constant, and (ii) if E occurs then none of the points in a ball of radius $T\sqrt{a \log n}$ around the origin that are initially uninfected ever become infected in our bootstrap percolation process. Applying Lemma 2(ii), we have that w.h.p. some translate of E occurs, whence w.h.p. we do not have full percolation in this regime. We leave the details to the reader. □

The content of Conjecture 3 is that the last obstruction to full percolation that will vanish as we decrease the infection threshold θ will correspond to the local distribution of initially infected/uninfected points that are well approximated by a symmetric T -island, for some $T = T(a, p)$. Motivation for the conjectured circular symmetry comes from the isoperimetric inequality in the plane as well as probabilistic considerations: for an uninfected island configuration to remain both uninfected and likely to occur in T_n^2 , one expects it is best to minimize the length of its boundary, and to spread out unlikely low densities of points outside the island and of initially uninfected points inside the island as uniformly as possible.

Theorem 3 gives an upper bound on the values of θ for which full percolation may occur w.h.p. in the Bradonjić–Saniee model. While we are unable to prove a matching lower bound and thereby prove Conjecture 3, we note here that one can nevertheless prove some rigorous lower bounds on the threshold $\theta_{\text{percolation}}(a, p)$ below which percolation occurs w.h.p., which we believe refine the earlier simple bounds due to Bradonjić and Saniee [15, Theorem 2] (in both cases, the exact value of the threshold is given implicitly rather than explicitly, which makes a comparison difficult).

Theorems 1(ii) and 2 imply that once θ falls below θ_{local} , w.h.p. the only obstructions to full percolation are components of uninfected vertices of Euclidean diameter $O(\sqrt{\log n})$ in T_n^2 . Let us consider what point configurations make the existence of such components possible. Given a component of never-infected points of diameter $C\sqrt{\log n}$, for some $C > 0$, consider a pair of uninfected points \mathbf{u}, \mathbf{v} from that component with $\|\mathbf{u} - \mathbf{v}\| = C\sqrt{\log n}$. Then the whole component of never-infected points lies inside the lune L formed by the two discs of radius $C\sqrt{\log n}$ centred at \mathbf{u} and \mathbf{v} respectively. Since \mathbf{u}, \mathbf{v} do not become infected, it must be the case that the number of points in $B_r(\mathbf{u}) \setminus L$ plus the number of initially uninfected points in $B_r(\mathbf{u}) \cap L$ is less than $\theta a \log n$, and a similar statement holds for \mathbf{v} . This implies that either $(B_r(\mathbf{u}) \cup B_r(\mathbf{v})) \setminus L$ contains an abnormally low number of points from our Poisson point process, or that $(B_r(\mathbf{u}) \cup B_r(\mathbf{v})) \cap L$ contains an abnormally low number of initially infected points. By performing a case analysis and some Lagrangian optimization, one can upper-bound the probability of such an event. Once this upper bound becomes $o(1/n)$, one can then apply Markov’s inequality to show that such unlikely point configurations w.h.p. do not occur, and thus that we have w.h.p. entered the full percolation regime. However, we do not believe the bounds coming from this argument are optimal (since they will not match those from Conjecture 3) or particularly helpful, so we relegate a sketch of the aforementioned case analysis and optimization to Appendix B.

4. Bootstrap percolation on the circle

In this section, we discuss the behaviour of the Bradonjić–Saniee model in dimension $d = 1$, i.e. bootstrap percolation on a Gilbert random geometric graph on the circle. In the regime we are considering, this means we have a Poisson process of intensity 1 on a circle of circumference n providing us with the vertex set of a Gilbert geometric graph with parameter r given by $2r := a \log n$, where $a > 1$ is fixed. A related model of bootstrap percolation on the Cartesian product of a cycle with a complete graph has previously been studied by Gravner and Sivakoff in [26].

As our main interest in this (already long) paper is the behaviour of the Bradonjić–Saniee model in the torus, we give a very informal, high-level discussion of the behaviour of the 1-dimensional case; our arguments can be made fully rigorous using fine tiling arguments similar to those deployed in Section 3.

The 1-dimensional case reveals a broadly similar picture to what we conjecture holds in the 2-dimensional case, with two exceptions: (i) the fact that starting an infection is enough to ensure a large local outbreak, doing away with the necessity of computing an analogue of θ_{local} , and (ii) the existence of *blocking sets* (see below) that can stop a large local outbreak from becoming global, since in one dimension you cannot ‘go round an obstacle’. A pleasant feature of the 1-dimensional case, however, is that the optimization problems involved are far simpler and can be resolved explicitly using the method of Lagrange multipliers.

We break up our discussion of the 1-dimensional model below in subsections on starting and growing an infection, blocking sets, and islands, before summarizing the behaviour in the various regimes thus identified in Section 4.5 with a picture. Throughout this section, (a, p, θ) is fixed with $a > 1$ and $0 < p < \theta < 1$.

4.1. Technical remarks: uniform distribution and approximate densities

While the discussion that follows will be quite informal, let us set down here a few remarks on how it can be formalized. Let $C_1 > C_2 > 0$ be fixed constants. Fix $c = c(C_1, C_2) > 0$. Let I be a fixed interval of length $|I| \in [C_1r, C_2r]$ on the circle. We say that a finite point-set \mathcal{P} with $|\mathcal{P} \cap I| = x|I|$ is *uniformly distributed on I* if every subinterval J of I of length at least $c\sqrt{r}$ contains $|J|(x + o(1))$ points from \mathcal{P} .

If \mathcal{P} is the outcome of a Poisson process of intensity bounded away from 0 on the circle, then it is an easy consequence of Lemma 1 that *conditional on $|\mathcal{P} \cap I| = x|I|$* , w.h.p. \mathcal{P} is uniformly distributed on I (for a suitable choice of $c = c(C_1, C_2)$). In particular, the probability of seeing that many points in I is, up to a $(1 + o(1))$ factor, the probability of seeing that many uniformly distributed points.

This remark, together with a partition of the circle into interval tiles of length C_1r , for some suitably small constant C_1 , allows us to reason in terms of continuous point-density functions rather than discrete point-sets, as we do below.

To fully formalize some of our arguments, e.g. to rule out a family of obstructions by showing that the likeliest event in this family is unlikely to occur, one should also approximate the initially infected/uninfected point density inside each interval tile as a fraction j/C_3 , where C_3 is a sufficiently large constant and

$$j = \min(\lceil C_3 \#\{\text{points in the interval}\} / C_1r \rceil, 100C_3/C_1).$$

Doing so then allows us, when considering events supported on a bounded number of tiles, to take union bounds over a finite number of tile density configurations and simply apply Markov’s inequality; see e.g. the proof of [7, Lemma 5] for an example of this kind of argument.

4.2. Growing an infection

Our first observation is that if $f_{\text{start}}(a, p, \theta) < 1$, then, for some ε , some vertex v (at position 0, say) will see at least $a(\theta + \varepsilon) \log n$ infected neighbours, rather than just $a\theta \log n$. This follows from the continuity of the function f_{start} . Using the notation $I_x = (-x, x)$, we may also assume that both the infected and the uninfected vertices in $I_{100r} = (-100r, 100r)$ are uniformly distributed on both $I_{100r} \setminus I_r$ and I_r , so that any interval J of length $c\sqrt{\log n}$ inside I_r contains $|J|(1 - p)(1 + o(1))$ initially uninfected vertices and $|J|(\theta + \varepsilon)(1 + o(1))$ initially infected ones, while any such interval inside $I_{100r} \setminus I_r$ contains $|J|(1 - p)(1 + o(1))$ initially uninfected vertices and $|J|p(1 + o(1))$ initially infected ones. (The densities $1 - p$, $\theta + \varepsilon$, and p come from the colouring theorem for Poisson processes [32], which implies in this case that

the initially infected and uninfected points can be regarded as *independent* Poisson processes of intensities p and $1 - p$ respectively. The uniform distribution assumption, for its part, holds w.h.p. as discussed in Section 4.1.)

One consequence of this is that not only does v become infected (if it was not already infected), but all vertices in the interval $(-\delta \log n, \delta \log n)$ become infected, for some $\delta = \delta(\varepsilon)$. After the first round of new infections, the interval $I_{\delta \log n}$ contains $|I_{\delta \log n}|(1 - p + \theta + \varepsilon)(1 + o(1))$ vertices, all of which are infected; $I_r \setminus I_{\delta \log n}$ contains $|I_r \setminus I_{\delta \log n}|(1 - p)(1 + o(1))$ uninfected vertices and $|I_r \setminus I_{\delta \log n}|(\theta + \varepsilon)(1 + o(1))$ infected ones; and w.h.p. $I_{100r} \setminus I_r$ contains infected and uninfected vertices with approximate densities p and $1 - p$ respectively.

Now we show that w.h.p. the ‘infected interval’ $I_{\delta \log n}$ grows until it infects every vertex in I_{100r} . For simplicity, we shall say that the interval I_x is infected if every vertex within it is infected. Suppose that, after some rounds of the bootstrap process, I_x is in fact infected. (Initially, we take $x = \delta \log n$.) We examine the neighbours of the first uninfected vertex u to the right of I_x , with a view to showing that u gets infected next, by virtue of having many infected neighbours in I_x . An identical argument will apply on the left.

Assume first that $x \leq r/2$. Then w.h.p. u will have $N_1(x)(1 + o(1))$ infected neighbours, where

$$\begin{aligned} N_1(x) &= (\theta + \varepsilon)(2r - 3x) + xp + 2x(1 - p + \theta + \varepsilon) \\ &= 2r\theta + \varepsilon(2r - x) + x(2 - \theta - p) \\ &> 2r\theta = a\theta \log n, \end{aligned}$$

so that u does indeed become infected w.h.p.

Next assume that $r/2 < x \leq r$, and write $y = x - r/2$. This time w.h.p. u will have $N_2(y)(1 + o(1))$ infected neighbours, where

$$\begin{aligned} N_2(y) &= (1 - p + \theta + \varepsilon)r + \left(\frac{r}{2} - y\right)(\theta + \varepsilon) + \left(\frac{r}{2} + y\right)p \\ &= \frac{(2 - p + 3\theta + 3\varepsilon)r}{2} + y(p - \theta - \varepsilon) \\ &> 2r\theta = a\theta \log n. \end{aligned}$$

Now since $N_2(y)$ is a linear function satisfying

$$N_2(0) = \frac{(2 - p + 3\theta + 3\varepsilon)r}{2} > \frac{(1 + 3\theta)r}{2} > 2r\theta$$

and

$$N_2\left(\frac{r}{2}\right) = \frac{(2 + 2\theta + 2\varepsilon)r}{2} > 2r\theta,$$

it follows that $N_2(y) > 2r\theta$ and hence that u becomes infected w.h.p.

Finally we show that the infection spreads beyond I_r , to at least the entire interval I_{100r} . When $r < x \leq 100r$, u will have at least $(r + rp)(1 + o(1))$ infected neighbours (this will be an underestimate if $x < 2r$), and this exceeds $2r\theta$ when $2\theta < 1 + p$. Accordingly, we name the inequality $2\theta < 1 + p$ (which is already familiar to us from Theorem 1) the *global growth condition* and display it for convenience:

Global growth condition: $\theta < \frac{1 + p}{2}$.

We also have the condition from Proposition 1 for the infection to start:

Starting condition: $a(p - \theta + \theta \log(\theta/p)) < 1$ or $\theta < p$.

Next, we describe the spread of the infection beyond I_{100r} . As a consequence of Lemma 2(ii), if the starting condition is met, then the infection will actually start in $n^{\alpha+o(1)}$ sites, where $\alpha = 1 - f_{\text{start}}(a, p, \theta)$. If the global growth condition is also met, these local infections will grow on both sides from each such site, until one of them is met by a (left or right) *blocking set*, which we describe below.

4.3. Blocking sets

In this subsection, we assume that both the starting and the global growth conditions are satisfied. What can then stop an infection from spreading around the circle? We shall distinguish two possible obstructions: blocking sets (where initially uninfected vertices in an interval of length at least $2r$ stay uninfected) and islands (where initially uninfected vertices in an interval of length at most $2r$ stay uninfected). In this subsection, we deal with the former.

A *clockwise*, or *right*, *blocking set* consists of two contiguous intervals, both of length r . The left interval contains zr vertices, and the right interval contains xpr initially infected vertices, both sets of vertices being uniformly distributed within their respective intervals (note that as remarked in Section 4.1, uniform distribution occurs w.h.p. conditional on the sizes of the two vertex sets). Even if all the vertices in the left interval become infected, the infection will not spread to the right interval (or, more precisely, clockwise) as long as

$$(z + xp)r < a\theta \log n = 2r\theta, \tag{13}$$

i.e. as long as $xp + z < 2\theta$. In other words, if this strict inequality is satisfied then the right interval in a right blocking set can only be infected ‘from the right’—the vertices in the left interval do not on their own suffice to spread the infection further to the right.

We wish to maximize the probability that such a right blocking set configuration occurs somewhere in the circle. Using Lemma 1, we see that the probability $q = q(x, z)$ of a blocking set with parameters x and z is given by

$$q = \exp\{r(z - 1 - z \log z) + pr(x - 1 - x \log x) + o(r)\} \\ = \exp \left\{ \frac{a \log n}{2} ((z - 1 - z \log z) + p(x - 1 - x \log x) + o(1)) \right\}.$$

Consequently, to maximize the probability that a blocking set occurs, we must maximize

$$f(x, z) = z - 1 - z \log z + p(x - 1 - x \log x)$$

subject to the constraint

$$g(x, z) = xp + z \leq 2\theta,$$

with p and θ fixed. Using the method of Lagrange multipliers, we see that $f(x, z)$ is maximized when

$$(x, z) = \left(\frac{2\theta}{1+p}, \frac{2\theta}{1+p} \right),$$

and, since the growing condition is satisfied, we note that $x = z < 1$ at the maximum. Substituting, the maximum f_{\max} of $f(x, z)$ subject to $g(x, z) \leq 2\theta$ is given by

$$f_{\max} = 2\theta - (1 + p) - 2\theta \log \left(\frac{2\theta}{1 + p} \right),$$

so that the maximum of $q(x, z)$ is

$$q_{\max} = \exp \left\{ -a \log n \left(\frac{1 + p}{2} - \theta + \theta \log \left(\frac{2\theta}{1 + p} \right) + o(1) \right) \right\}.$$

We conclude that the expected number of blocking sets is $n^{\beta+o(1)}$, where

$$\beta = 1 - a \left(\frac{1 + p}{2} - \theta + \theta \log \left(\frac{2\theta}{1 + p} \right) \right).$$

Recall also that, as long as $\theta > p$, the infection starts in $n^{\alpha+o(1)}$ places, where

$$\alpha = 1 - f_{\text{start}}(a, p, \theta) = 1 - a(p - \theta + \theta \log(\theta/p)).$$

Now assume that the starting and growing conditions are satisfied, and consider the case $0 < \alpha < \beta < 1$. Infections will start to spread in $n^{\alpha+o(1)}$ places, and these will be blocked by $n^{\beta+o(1)} \gg n^{\alpha+o(1)}$ blocking sets (a right blocking set stops infections spreading clockwise, and a left blocking set stops infections spreading anticlockwise). Accordingly, each of the $n^{\alpha+o(1)}$ growing infections will spread for distance $n^{1-\beta+o(1)}$ before being blocked in both the clockwise and the anticlockwise direction, so that the infections will cover $n^{1+\alpha-\beta+o(1)} = o(n)$ of the circle. We call this the regime of ‘polynomial growth’; in this regime, the spread of infection is largely contained.

Next, assume again that the starting and growing conditions are satisfied, but suppose that $0 < \alpha/2 < \beta < \alpha < 1$. The $n^{\alpha+o(1)}$ spreading infections will encounter $n^{\beta+o(1)} \ll n^{\alpha+o(1)}$ blocking sets, and since these blocking sets will usually be isolated, with spreading infections on each side, most of the circle will become infected. The only likely obstructions (save for the ‘islands’ to be described in the next subsection) are pairs of blocking sets (reading clockwise, a right blocking set followed by a left blocking set) with no growing infection between them.

To be slightly more precise, suppose there is some interval I of length at least r such that (a) every initially uninfected vertex in I stays uninfected, and (b) every vertex in the intervals of length r to the left and to the right of I become infected. Let v denote the leftmost uninfected vertex in I . Let zr denote the number of vertices in the interval of length r to the left of v , and xpr the number of initially infected vertices in the interval of length r to the right of v . Then x, z satisfy (13). Now, conditional on having zr vertices in the interval to the left of r and xpr initially infected vertices in the interval to the right of r , we know that w.h.p. these points are uniformly distributed on these intervals. Thus, up to a $(1 + o(1))$ multiplicative factor, the probability of seeing such a point configuration is the probability of seeing a right blocking set with parameter (x, z) , and hence is at most $n^{\beta+o(1)}$. (Formally, one must also use approximate densities as sketched in Section 4.1 to partition the collection of all (x, z) into a finite collection of rational pairs, as well as the continuity of the functions f and q in order to show that the contribution from configurations with (x, z) close to $\left(\frac{2\theta}{1+p}, \frac{2\theta}{1+p}\right)$ dominates.)

There will be $n^\beta \cdot n^{\beta-\alpha+o(1)} = n^{2\beta-\alpha+o(1)}$ such pairs, typically separated by distance $n^{1+\alpha-2\beta+o(1)}$, so that the entire circle except for regions of length totalling $n^{1-2\beta+\alpha+o(1)} = o(n)$

will become infected. We call this the regime of ‘polynomial obstructions’: when $2\beta > \alpha$, blocking sets by themselves cannot prevent the spreading infections from covering most of the circle.

There are thus two thresholds: the first, separating the regime of polynomial growth from that of polynomial obstructions, occurring at $\alpha = \beta$, and the second, separating the regime of polynomial obstructions from that of ‘logarithmic obstructions’ (caused by *islands*—see the subsection below), occurring at $\alpha = 2\beta$.

When finding these thresholds, we recall that they lie entirely inside the region $\theta \geq p$, since $\alpha = 1$ when $\theta \leq p$. For the first threshold, we have $\alpha > \beta$ exactly when

$$\frac{1+p}{2} - \theta + \theta \log\left(\frac{2\theta}{1+p}\right) > p - \theta + \theta \log(\theta/p),$$

which yields the following expression for the threshold:

First threshold condition: $\frac{1-p}{2} > \theta \log\left(\frac{1+p}{2p}\right).$

Note that this region contains the region $\theta \leq p$ (as can be seen from the Taylor expansion of the right-hand side), and is independent of the choice of $a > 1$. For the second threshold, we have $\alpha > 2\beta$ exactly when

$$1 - a(p - \theta + \theta \log(\theta/p)) > 2 \left(1 - a \left(\frac{1+p}{2} - \theta + \theta \log\left(\frac{2\theta}{1+p}\right) \right) \right),$$

which yields the following expression for the threshold:

Second threshold condition: $a \left(1 - \theta + \theta \log\left(\frac{4\theta p}{(1+p)^2}\right) \right) > 1$ or $\theta < p.$

This time, we need to include the additional condition $\theta < p$.

4.4. Islands

Next assume that the starting, global growth, and first and second threshold conditions are all satisfied. What can then prevent full percolation? In this subsection, we sketch how the final obstructions for the infection to overcome are *islands*, i.e. intervals of length at most r in which no new infection is recorded.

For $c \in [0, 1]$, a cr -*island*, illustrated in Figure 3, consists of five contiguous intervals whose lengths, from left to right, are $cr, (1 - c)r, cr, (1 - c)r, cr$. Inside these intervals there are,




FIGURE 3. An island.

from left to right, zcr vertices, $y(1 - c)r$ vertices, $xpcr$ initially infected vertices, $y(1 - c)r$ vertices, and zcr vertices, all uniformly distributed within their respective intervals. We note that even if all the vertices in the four outermost intervals become infected, no vertices inside the middle interval will become infected, as long as

$$cz + 2y(1 - c) + xpc < 2\theta.$$

We wish to maximize the probability that such a set occurs.

Remark 3. Formally, we should consider islands where we allow for a non-uniform distribution of vertices within the five intervals. However, the number of points around such islands must still satisfy the displayed inequality above (just consider the leftmost uninfected vertex in the middle interval), whence, as we remarked in Section 4.1, the probability of seeing such a point configuration is, up to a $(1 + o(1))$ multiplicative factor, the probability of seeing an island as defined above.

Remark 4. Similarly to the above, formally we should also consider the case where the proportions of vertices and infected vertices vary across the five intervals, i.e. where these contain z_1cr vertices, $y_1(1 - c)r$ vertices, $xpcr$ initially infected vertices, $y_2(1 - c)r$ vertices, and z_2cr vertices respectively, and where we have two constraints, namely $cz_1 + (y_1 + y_2)(1 - c) + xpc < 2\theta$ and $cz_2 + (y_1 + y_2)(1 - c) + xpc < 2\theta$. However, it is easily checked that the probability of such configurations is maximized by symmetric configurations with $y_1 = y_2 = y$ and $z_1 = z_2 = z$, so that the cr -islands defined above that we consider are the only ones we need to worry about in an asymptotic analysis. Here again the argument can be formalized using approximate densities.

Using Lemma 1, we see that the probability $q = q(x, y, z, c)$ of an island with parameters x, y, z and c is given by

$$q = \exp\{2cr(z - 1 - z \log z) + 2(1 - c)r(y - 1 - y \log y) + pcr(x - 1 - x \log x) + o(r)\} \\ = \exp \left\{ \frac{a \log n}{2} (2c(z - 1 - z \log z) + 2(1 - c)(y - 1 - y \log y) + pc(x - 1 - x \log x) + o(1)) \right\}.$$

Consequently, to maximize the probability that an island occurs, we must maximize

$$f(x, y, z, c) = 2c(z - 1 - z \log z) + 2(1 - c)(y - 1 - y \log y) + pc(x - 1 - x \log x)$$

subject to the constraints

$$g(x, y, z, c) = cz + 2y(1 - c) + xpc \leq 2\theta$$

and $c \in [0, 1]$ (for the island to exist and not to be a blocking set), with p and θ fixed.

Using the method of Lagrange multipliers, and noting that the growing condition is satisfied (which excludes the solution $x = y = z = 1$), we see that $f(x, y, z, c)$ is maximized when

$$(x, y, z, c) = \left(\left(\frac{p}{2 - p} \right)^2, \left(\frac{p}{2 - p} \right)^2, \frac{p}{2 - p}, \frac{2(\theta(2 - p)^2 - p^2)}{p(1 - p)(2 - p)} \right),$$

where we also require that $c \in [0, 1]$ (for the island to exist). We separate the analysis into three cases, depending on the value of $c_\star = \frac{2(\theta(2-p)^2 - p^2)}{p(1-p)(2-p)}$.

Case 1: $c_\star \in (0, 1)$. The maximum f_{\max} of $f(x, y, z, c)$ subject to our constraints is then given by

$$f_{\max} = \frac{8(p-1)}{(2-p)^2} - 4\theta \log\left(\frac{p}{2-p}\right),$$

so that the maximum of $q(x, y, z, c)$ is

$$q_{\max} = \exp\left\{-a \log n \left(\frac{4(1-p)}{(2-p)^2} + 2\theta \log\left(\frac{p}{2-p}\right) + o(1)\right)\right\}.$$

Set

$$f_{c_\star\text{-stop}}(a, p, \theta) := a \left(\frac{4(1-p)}{(2-p)^2} + 2\theta \log\left(\frac{p}{2-p}\right)\right).$$

By Lemma 2(ii), we have that if $f_{c_\star\text{-stop}}(a, p, \theta) < 1$, then w.h.p. a c_\star -island occurs somewhere on the circle, and we do not have full percolation. On the other hand, if $f_{c_\star\text{-stop}}(a, p, \theta) > 1$, then by Lemma 2(i) w.h.p. there are no c -islands for any $c \in [0, 1]$. (Formally, we take an equipartition of $[0, 1]$ into a large but finite number of rational subintervals; we use these to partition the collection of c -islands, for $c \in [0, 1]$, into a finite number of families, to which we can apply Markov’s inequality.) Thus w.h.p. we have full percolation (since the starting, growing, and first and second threshold conditions are satisfied, and since any uninfected vertex must correspond to an island).

Case 2: $c_\star \leq 0$. When $c_\star \leq 0$, the optimum legitimate island is the 0-island (i.e. an uninfected vertex with degree too low to ever be infected), and we recover the necessary condition for full percolation from Proposition 2, namely that $f_{0\text{-stop}}(a, p, \theta) > 1$. (This explains the earlier choice of notation.) Note that the analysis in this section reveals that this necessary condition for full percolation $f_{0\text{-stop}}(a, p, \theta) > 1$ is sufficient when $c_\star \leq 0$, but not when $c_\star \in (0, 1)$. This is because, when $c_\star \in (0, 1)$, the condition $f_{0\text{-stop}}(a, p, \theta) > 1$ is strictly weaker than $f_{c_\star\text{-stop}}(a, p, \theta) > 1$, since a non-degenerate c_\star -island is more likely to occur than the degenerate 0-island from Proposition 2.

Case 3: $c_\star \geq 1$. When $c_\star \geq 1$, the optimum island has $c = 1$, and a separate calculation with Lagrange multipliers shows that the optimum choices of x and z (y no longer features in the island) are

$$(x, z) = \left(\frac{1 + 4\theta p - 2\sqrt{1 + 8\theta p}}{2p^2}, \frac{\sqrt{1 + 8\theta p} - 1}{2p}\right) = \left(\left(\frac{\sqrt{1 + 8\theta p} - 1}{2p}\right)^2, \frac{\sqrt{1 + 8\theta p} - 1}{2p}\right).$$

These choices lead to a maximum of $q(x, y, z, c)$ of

$$q_{\max} = \exp\left\{-a \log n \left(1 - \theta + \frac{p}{2} - \frac{\sqrt{1 + 8\theta p} - 1}{4p} + 2\theta \log\left(\frac{\sqrt{1 + 8\theta p} - 1}{2p}\right) + o(1)\right)\right\},$$

so that, writing

$$f_{1\text{-stop}}(a, p, \theta) = a \left(1 - \theta + \frac{p}{2} - \frac{\sqrt{1 + 8\theta p} - 1}{4p} + 2\theta \log\left(\frac{\sqrt{1 + 8\theta p} - 1}{2p}\right)\right),$$

we see that when $c_\star \geq 1$, if $f_{1-\text{stop}}(a, p, \theta) < 1$, then by Lemma 2(ii), w.h.p. an r -island occurs somewhere on the circle, and we do not have full percolation. On the other hand, if $f_{1-\text{stop}}(a, p, \theta) > 1$, then by Lemma 2(i) w.h.p. there are no cr -islands for any $c \in [0, 1]$ (as in Case 1), and w.h.p. we have full percolation (since the starting, growing, and first and second threshold conditions are satisfied).

To summarize, suppose we are given p and θ , and we wish to check whether full percolation occurs, assuming that the starting, growing, and second threshold conditions are satisfied. First, we calculate

$$c_\star = c_\star(p, \theta) = \frac{2(\theta(2-p)^2 - p^2)}{p(1-p)(2-p)}.$$

Then, depending on the value of c_\star , the condition for full percolation is as displayed below:

Full percolation: $1 < \begin{cases} a(1 - \theta + \theta \log \theta), & c_\star \leq 0, \\ a \left(\frac{4(1-p)}{(2-p)^2} + 2\theta \log \left(\frac{p}{2-p} \right) \right), & 0 < c_\star < 1, \\ a \left(1 - \theta + \frac{p}{2} - \frac{\sqrt{1+8\theta p-1}}{4p} + 2\theta \log \left(\frac{\sqrt{1+8\theta p-1}}{2p} \right) \right), & c_\star \geq 1. \end{cases}$

The lack of islands is necessary for full percolation, and, given that the other four conditions are met, it is w.h.p. also sufficient. To see this, suppose that the second threshold condition is satisfied. Then, once the infection has stopped spreading, the Gilbert graph on the uninfected vertices splits into several components. Take one such component C , with extreme vertices u and v . If u and v lie at distance at least $2r$, then the probability that this occurs is of the same order as the probability that there exists a pair of blocking sets with no infection between them, which is $o(1)$ by the second threshold condition. Thus u and v lie at distance less than $2r$. If they lie at distance less than r , then the probability that this occurs is of the same order as the probability that they are part of an island. If they lie at distance $(2-c)r$ with $r \leq (2-c)r \leq 2r$, then we must have five intervals of lengths $r, (1-c)r, cr, (1-c)r$, and r (from left to right, with u the rightmost endpoint of the first interval and v the leftmost endpoint of the last interval), containing (respectively) z_1r points, $y_1p(1-c)r$ initially infected points, $xpcr$ initially infected points, $y_2p(1-c)r$ initially infected points, and z_2r points, where

$$z_1 + (y_1 + y_2)p(1-c) + xpc < 2\theta, \quad z_2 + (y_1 + y_2)p(1-c) + xpc < 2\theta.$$

It is easily checked that the probability of such a configuration is maximized by taking $y_1 = y_2 = y$ and $z_1 = z_2 = z$. A short calculation with Lagrange multipliers then shows that the most likely way this can happen is when $c = 0$ (corresponding to a blocking set) or $c = 1$ (corresponding to an r -island). In summary, if the second threshold condition is satisfied, then the most likely way for a small component of uninfected vertices to arise is in the shape of an island, and the threshold for the disappearance of islands is also the threshold for full percolation.

4.5. Summary

Figure 4 is a phase diagram for the case $a = 2$, showing the different regimes discussed above. If the starting condition is not met, there is no growth—no new infections take place. If the starting condition is met, but the global growth condition is not met, then the growth




FIGURE 4. Phase diagram for $a = 2$.

will be confined to regions of width $\Theta(\log n)$ —we term this ‘logarithmic growth’. If these two conditions, but not the first threshold condition, are met, we are in the region of ‘polynomial growth’; if the first (but not the second) threshold condition is met, then we have ‘polynomial obstructions’. Finally, if all these conditions are met, the threshold for full percolation separates the region of ‘logarithmic obstructions’ (islands) from that of full percolation. Note that the term ‘logarithmic obstruction’ (respectively, ‘polynomial obstruction’) refers to the size of the obstruction, rather than to the total proportion of the circle occupied by such obstructions. It is possible that logarithmic obstructions will dominate polynomial ones for certain values of the parameters, but, to keep things simple, we do not pursue this question further here.

It is convenient to visualize the spread of infection as a continuous process, starting in certain random places, and then spreading in the form of expanding arcs, before possibly being blocked by blocking sets or islands. Indeed, consider a typical ‘interface’ I at the place where an infection is blocked. To the left of I , we have full infection, while to the right of I , only initially infected points are infected. The interface I itself consists of uninfected points alternating with (initially infected points and) newly infected points. Now the probability that the length d_I of I exceeds x decays exponentially in x , so long interfaces are comparatively rare. Thus

they may be conveniently visualized as single points marking the boundaries of infections, and the frozen state of the model may be visualized as a collection of random arcs on a circle.

5. Concluding remarks

We end this paper with some remarks.

- We stated and proved our main results for the Bradonjić–Sanicee model on the 2-dimensional torus $T_n := T_n^2$ rather than the square $S_n := [0, \sqrt{n}]^2$ to avoid having to consider boundary effects. We note here that the boundary effects do not play any significant role in the Bradonjić–Sanicee model on S_n , apart from complicating the analysis relative to T_n .

Indeed, since the presence of a boundary does not help the infection, the analogue of Theorem 1(ii) immediately follows for S_n .

For the analogue of Theorem 1(i) in the square, one must modify the proof. Instead of considering the whole of S_n , one should instead focus on the subsquare S_n' consisting of all rough tiles at graph distance at least 7 from a ‘boundary tile’ in the auxiliary graph H , and one must employ the Bollobás–Leader edge-isoperimetric inequality for the square grid rather than the toroidal grid. Since $o(|\mathcal{R}|)$ rough tiles lie close to the boundary of S_n , this results in some changed constants as well as the presence of a connected component in $G_{n,r}[\mathcal{P} \setminus A_\infty]$ of diameter $\Omega(\sqrt{n})$ and order $O(\sqrt{n \log n})$ all around the boundary of the square. Thus the bound on the Euclidean diameter of components of forever-uninfected points will only hold for components all of whose points are at a sufficiently large multiple of $\sqrt{\log n}$ away from the boundary.

- We expect that Theorem 1 also holds in dimension $d \geq 3$. However, some work will be required to adapt our arguments to the higher-dimensional setting. For instance, we note that the proof of Proposition 3, and both the statement and the proof of Lemma 4 (involving dual cycles), use 2-dimensional ideas and would need to be handled differently in dimension $d \geq 3$.
- In contrast with the constant infection threshold work of Candellero and Fountoulakis [16] on hyperbolic random geometric graphs and of Koch and Lengler [34] on inhomogeneous random graphs, the behaviour for the Bradonjić–Sanicee model appears to be rather different: for $p < \theta$ it is possible for some vertices to be infected in the first round of the process but for the process to only spread the infection to $|A_\infty \setminus A_0| = o(n)$ vertices. Indeed, if $\frac{1+p}{2} < \theta$ and $f_{\text{start}} < 1$ both hold, this occurs w.h.p. as a consequence of Theorem 1(ii) and Proposition 1. Both of these conditions can be satisfied simultaneously provided we pick $a > 1$ sufficiently small and $p < 1$ sufficiently large.

Our work also leaves a number of questions open.

- Foremost among these is the problem of proving Conjecture 1. Our intuition is that one may need to consider far more subtle tile colourings to achieve this. Explicitly, rather than colour a fine tile T red if all its points eventually become infected, one should plausibly instead fix some large integer constant $Q = Q(K)$ and label T with $(i_T/Q, j_T/Q)$ if it contains between $\frac{i_T}{Q}|T|$ and $\frac{i_T+1}{Q}|T|$ points of A_∞ , and between $\frac{j_T}{Q}|T|$ and $\frac{j_T+1}{Q}|T|$ points of $\mathcal{P} \setminus A_\infty$, where i_T, j_T are non-negative integers. A similar ‘approximate density’ idea was used in [7]. One would then need to show that for most fine tiles, $(i_T/Q, j_T/Q)$ is close to either $(1,0)$ or $(p, 1 - p)$, and would need to analyse the component structure

of fine tiles T with $(i_T/Q, j_T/Q)$ close to $(1,0)$, deploying more refined versions of the arguments from the proof Theorem 4 to show such tiles form either an overwhelming majority or an overwhelming minority of tiles.

- Another intriguing problem left open by Theorem 1 is, what happens if $\theta = \frac{1+p}{2} \pm o(1)$ and we adversarially infect a ball of radius Cr ? Here it is not clear to us what kind of behaviour one should expect. Could it be, for example, that both $|A_\infty \setminus A_0|$ and $|\mathcal{P} \setminus A_\infty|$ have order $\Omega(n)$? The analysis required to understand the typical behaviour in this regime is likely to be delicate.
- Given that bootstrap percolation has been studied on \mathbb{Z}^d , it would be natural to consider bootstrap percolation on a supercritical Gilbert disc graph in the plane, i.e. on the host graph $G_r(\mathbb{R}^2)$, where πr^2 is a sufficiently large constant (to ensure that $G_r(\mathbb{R}^2)$ almost surely contains an infinite connected component), and the infection threshold is some constant T . For what $p \geq 0$ does an initial infection probability guarantee the almost sure emergence of an infinite connected component of eventually infected vertices?
- It would also be natural to study an analogue of the Bradonjić–Saniee model on the k -nearest-neighbour random geometric graph model, or on models of random geometric graphs in the torus allowing for the presence of some long-distance edges by superimposing e.g. a sparse Erdős–Rényi random graph or a configuration model on top of the Gilbert random geometric graph.
- Conjectures 2 and 3 provide an obvious area where there is considerable room for improvement on the results of the present paper. Progress on these fronts may require a better understanding of the behaviour of the solutions to the optimization problems used to define the thresholds θ_{local} and θ_{islands} (our conjectured thresholds for almost percolation and full percolation, respectively).

One question of particular interest to us is whether, given a fixed triple (a, p, θ) , one can identify the ‘critical radius’ for local infections or islands. To be more precise, we expect that there may be a constant $C > 0$ such that spreading a local infection to radius Cr is ‘harder’ (less likely) than both spreading it to a radius $(C - \varepsilon)r$ and spreading an infection from a ball of radius Cr to a ball of radius $(C + \varepsilon)r$. Our heuristic is that while the local outbreak is small, it requires unlikely point configurations to spread radially outwards, but that once it gets sufficiently large, ‘global’ behaviour kicks in and the infection is carried forward by its momentum without requiring a high density of infected points near its boundary.

A motivation for determining such a critical radius would be the possibility of explicitly determining and computing θ_{local} . Similarly, we expect that there is an optimal radius for islands resisting full percolation, and determining that optimal radius would help give a more explicit form for θ_{islands} .

- Finally, given the motivations for studying bootstrap percolation, it would be natural to consider variants of the Bradonjić–Saniee model where e.g. some vertices are vaccinated or have a higher threshold for infection. See for example [19] for some recent work in this vein.

Appendix A. Analysis of θ_{islands} via the Euler–Lagrange equations

Recall that $\theta_{\text{islands}} = \theta_{\text{islands}}(a, p)$ is the threshold for the appearance of symmetric islands. In this appendix, we bound this threshold using the Euler–Lagrange equations, applied to (10).




FIGURE 5. A simple obstruction to full percolation. Here f and g denote the density of initially infected and initially uninfected points in the various regions. The centre of the largest circle is $\mathbf{0}$, and \mathbf{u} sits on the boundary of the shaded red disc of radius τ about $\mathbf{0}$; circles of radius t about $\mathbf{0}$ and radius 1 about \mathbf{u} are represented in the picture.

As a warm-up, consider the following island, illustrated in Figure 5. We consider two concentric circles of radii $r/2$ and $3r/2$, forming a disc and an annulus of areas A_x and A_z respectively. Suppose that the inner disc contains $x p A_x$ initially infected points, and that the annulus contains $z A_z$ points. Then, even if all the points in the annulus become infected, no points in the inner disc become infected, as long as

$$G(x, z) = 3z + px < 4\theta.$$

The probability of the configuration is

$$q = \exp \left\{ \frac{a \log n}{4} F(x, z) \right\},$$

where

$$F(x, z) = 8(z - 1 - z \log z) + p(x - 1 - x \log x),$$

and so we must maximize $F(x, z)$ subject to $G(x, z) = 4\theta$. A short calculation with Lagrange multipliers shows that $z = x^{3/8}$, while x is determined by the equation

$$3x^{3/8} + px = 4\theta.$$

Once x has been determined, the threshold is given by

$$4 + a \left\{ 8(x^{3/8} - 1 - \frac{3}{8}x^{3/8} \log x) + p(x - 1 - x \log x) \right\} = 0.$$

The lack of such islands is a necessary condition for full percolation. To refine this condition, we will consider a radially symmetric island, with continuously varying densities of infected and uninfected points, as in Section 3.5. Such an island generalizes the one just considered. To simplify the exposition and calculations, we will use a slightly different normalization (scaling by r rather than $\sqrt{a \log n}$) and notation from those in Section 3.5.

Specifically, consider the following island, centred at $\mathbf{0}$. At distance tr from $\mathbf{0}$, the densities of infected and uninfected points are $pf(t)$ and $(1 - p)g(t)$ respectively, where f and g satisfy $f(t) \rightarrow 1$ and $g(t) \rightarrow 1$ as $t \rightarrow \infty$. Rescaling by r , for $\mathbf{u} \in \partial B_t(\mathbf{0})$, write

$$A(t) = B_1(\mathbf{u}) \cap B_t(\mathbf{0}) \quad \text{and} \quad B(t) = B_1(\mathbf{u}) \setminus B_t(\mathbf{0}),$$

exactly as before. Then the condition for an infection to stop spreading in towards $\mathbf{0}$ around the circle $\partial B_{\tau r}(\mathbf{0})$ is that

$$p \int_{A(\tau) \cup B(\tau)} f \, dA + (1 - p) \int_{B(\tau)} g \, dA < \pi\theta.$$

At this point we need to make some simplifying assumptions, which, while restricting the dimensions of the island (and thus possibly rendering it suboptimal), greatly facilitate calculations. First, we will assume that $\tau \leq 1$. Second, we will assume that the optimal function f is increasing for $t \geq 0$, and that the optimal g is increasing for $t \geq \tau$ (an assumption which will be consistent with the solution we obtain). The reason for these assumptions is that, without them, points in the interior of $B_{\tau r}(\mathbf{0})$ might see more infected neighbours than points on the boundary $\partial B_{\tau r}(\mathbf{0})$. It is then conceivable that the infection could spread from the inside of the island outwards, causing the entire island to succumb to infection.

Given these assumptions, we must maximize

$$q(f, g) = p \int (f - 1 - f \log f) \, dA + (1 - p) \int (g - 1 - g \log g) \, dA$$

subject to the above constraint, for fixed $\tau \in [0, 1]$. The sets $A(\tau)$ and $B(\tau)$, as well as the optimal solutions f and g , are illustrated in Figure 6 (for $\tau \leq 1/2$) and Figure 7 (for $1/2 \leq \tau \leq 1$).

Using the method of Lagrange multipliers, we maximize

$$L(f, g) = p \int (f - 1 - f \log f) \, dA + (1 - p) \int (g - 1 - g \log g) \, dA - \lambda \left(p \int_{A(\tau) \cup B(\tau)} f \, dA + (1 - p) \int_{B(\tau)} g \, dA \right).$$




FIGURE 6. An obstruction with $\tau \leq 1/2$.




FIGURE 7. An obstruction with $1/2 \leq \tau \leq 1$.

To evaluate the last two integrals, we introduce the function $h(t)$, defined by the equation

$$h(t) = \begin{cases} 1 & \text{if } t \leq 1 - \tau, \\ \frac{1}{\pi} \cos^{-1} \left(\frac{t^2 - \tau^2 - 1}{2\tau} \right) & \text{if } 1 - \tau < t \leq 1 + \tau, \\ 0 & \text{if } 1 + \tau < t. \end{cases}$$

The function $h(t)$ is the proportion of the (1-dimensional) circle of radius t , centred at $\mathbf{0}$, that lies in the (2-dimensional) ball $B_1(\mathbf{u})$, where $\mathbf{u} \in \partial B_\tau(0)$. Using this notation, we have

$$\begin{aligned} (2\pi)^{-1}L(f, g) &= p \int_0^\infty f(t) - 1 - f(t) \log [f(t)] dt + (1 - p) \int_0^\infty g(t) - 1 - g(t) \log [g(t)] dt \\ &\quad - \lambda \left(p \int_0^\infty f(t)h(t) dt + (1 - p) \int_\tau^\infty g(t)h(t) dt \right) \\ &= p \int_0^\infty f(t) - 1 - f \log [f(t)] dt + (1 - p) \int_0^\infty g(t) - 1 - g(t) \log [g(t)] dt \\ &\quad - \lambda \left(p \int_0^{1-\tau} f(t) dt + p \int_{1-\tau}^{1+\tau} f(t)h(t) dt + (1 - p) \int_\tau^{1+\tau} g(t)h(t) dt \right) \\ &= \int_0^\infty F(f, g, t) dt. \end{aligned}$$

The Euler–Lagrange equations reduce in this case to $\partial F/\partial f = 0$ and $\partial F/\partial g = 0$. These equations have the solution

$$f(t) = \begin{cases} \exp(-\lambda) & t \leq 1 - \tau, \\ \exp(-\lambda h(t)) & 1 - \tau \leq t \leq 1 + \tau, \\ 1 & t \geq 1 + \tau, \end{cases}$$

and

$$g(t) = \begin{cases} 1 & t \leq \tau, \\ \exp(-\lambda) & \tau \leq t \leq \max(\tau, 1 - \tau), \\ \exp(-\lambda h(t)) & \max(\tau, 1 - \tau) \leq t \leq 1 + \tau, \\ 1 & t \geq 1 + \tau, \end{cases}$$

where λ is determined from the constraint

$$p \int_0^\infty 2\pi f(t)h(t) dt + (1 - p) \int_\tau^\infty 2\pi g(t)h(t) dt = \pi\theta.$$

For a given τ , this allows us to compute $q_{\max}(\tau)$. Finally, we optimize over τ , and set

$$\sup_{\tau} aq_{\max}(\tau) = -\pi$$

to determine the θ -value of the threshold for the disappearance of islands. Note that the warm-up example (in which $\tau = 1/2$) is in some sense a step-function approximation to this solution. As $p \rightarrow 1$, we expect that the optimal value of τ tends to zero.

Note that the analysis here provides an alternative derivation of the islands in the case $d = 1$. We need only replace $h(t)$ by its 1-dimensional version $h_1(t)$, given by

$$h_1(t) = \begin{cases} 1 & \text{if } \tau \leq 1 - t, \\ 1/2 & \text{if } 1 - \tau < t \leq 1 + \tau, \\ 0 & \text{if } 1 + \tau < t. \end{cases}$$

Appendix B. Lower bounds for the threshold for full percolation

In this appendix, we sketch out some more details of the case analysis and Lagrangian optimization that can be used to give rigorous (but almost certainly non-optimal) lower bound on the full percolation threshold, continuing the discussion at the end of Section 3.5 (with the same notation).

Draw discs of radius r around \mathbf{u} and \mathbf{v} , resulting in one of the configurations in Figure 8, according to whether $\delta := C\sqrt{\log n}/r$ lies in the range $(0,1],(1,2]$, or $(2, \infty)$. Here x and zp respectively denote the density of points and of initially infected points of the process in the corresponding region, and similarly y and py denote the density of points and of initially infected points in the corresponding region. We know the points \mathbf{u} and \mathbf{v} do not become infected as part of the bootstrap percolation process, even though all points outside the lune L are infected. Our aim is to find the likeliest set of point densities in the appropriate regions making this event possible. We will then calculate the probabilities of these likeliest configurations, and deduce that if the fixed triple (a, p, θ) is such that all such configurations have probability $o(1/n)$, no island can exist w.h.p.

Assuming that the densities of infected points in the various regions are as indicated, the most likely obstructions can be identified by optimizing x, y, z and δ for each case. Write $L(\delta)$ for the area of the lune formed by two unit discs whose centres lie at distance δ , so that

$$L(\delta) = \pi - \gamma - \sin \gamma,$$




FIGURE 8. Stopping an infection in the cases $0 < \delta < 1$, $1 < \delta < 2$, $\delta > 2$.

where γ is given by

$$\frac{\delta}{2} = \sin\left(\frac{\gamma}{2}\right).$$

Also, write $M(\delta)$ for the area of the lune formed by a disc B_δ of radius δ and a unit disc D whose centre lies on the perimeter of B_δ , so that

$$M(\delta) = \delta^2(\pi - \beta - \sin \beta) + \frac{\beta}{2},$$

where β is given by

$$2\delta = \sec\left(\frac{\beta}{2}\right).$$

Case 1: $\delta > 2$. In this case, the island occurs with probability at most

$$q_1 = \exp(2r^2 f_1(x, y, \delta) + o(r^2)) = \exp\left\{\frac{2a \log n}{\pi} f_1(x, y, \delta) + o(\log n)\right\},$$

where

$$f_1(x, y, \delta) = (\pi - M(\delta))(x - 1 - x \log x) + pM(\delta)(y - 1 - y \log y),$$

and where we also need

$$g_1(x, y, \delta) = (\pi - M(\delta))x + pM(\delta)y < \pi\theta$$

to prevent \mathbf{u} and \mathbf{v} from getting infected. It is easy to see that this configuration is likeliest when δ is as large as possible and when $x = y$. A quick calculation with Lagrange multipliers yields the threshold

$$a \left\{ 1 + p - 2\theta + 2\theta \log\left(\frac{2\theta}{1+p}\right) \right\} = 1.$$

Case 2: $1 < \delta < 2$. In this case, the island occurs with probability at most

$$q_2 = \exp(r^2 f_2(x, y, z, \delta) + o(r^2)) = \exp\left\{\frac{a \log n}{\pi} f_2(x, y, z, \delta) + o(\log n)\right\},$$

where

$$f_2(x, y, z, \delta) = 2(\pi - M(\delta))(x - 1 - x \log x) + 2p(M(\delta) - L(\delta))(y - 1 - y \log y) + pL(\delta)(z - 1 - z \log z),$$

and where we also need

$$g_2(x, y, z, \delta) = (\pi - M(\delta))x + p(M(\delta) - L(\delta))y + pL(\delta)z < \pi\theta$$

to prevent **u** and **v** from getting infected. A quick calculation with Lagrange multipliers yields that at the optimum $(x, y, z) = (x, x, x^2)$, and that x and δ are obtained by solving the equations

$$pL'(\delta)(x - 1) + 2M'(\delta)(p - 1) = 0, \\ x(\pi + (p - 1)M(\delta) - pL(\delta)) + x^2 pL(\delta) = \pi\theta,$$

after which we set $q_2 = 1/n$ to get the bound on the threshold.

Case 3: $0 < \delta < 1$. In this case, the island occurs with probability at most

$$q_3 = \exp(r^2 f_3(x, y, z, \delta) + o(r^2)) = \exp\left\{\frac{a \log n}{\pi} f_3(x, y, z, \delta) + o(\log n)\right\},$$

where

$$f_3(x, y, z, \delta) = 2(\pi - L(\delta))(x - 1 - x \log x) + (L(\delta) - \delta^2 L(1))(y - 1 - y \log y) + p\delta^2 L(1)(z - 1 - z \log z),$$

and where we also need

$$g_3(x, y, z, \delta) = (\pi - L(\delta))x + (L(\delta) - \delta^2 L(1))y + p\delta^2 L(1)z < \pi\theta$$

to prevent **u** and **v** from getting infected. A quick calculation with Lagrange multipliers yields that at the optimum $(x, y, z) = (x, x^2, x^2)$, and that x and δ are obtained by solving the equations

$$L'(\delta)(x - 1) + 2(p - 1)(x + 1)\delta L(1) = 0, \\ x(\pi - L(\delta)) + x^2(L(\delta) + \delta^2 L(1)(p - 1)) = \pi\theta,$$

after which we set $q_3 = 1/n$ to get the bound on the threshold.

Behaviour as $p \rightarrow 0$ and $p \rightarrow 1$

When $p \rightarrow 0$, a routine analysis shows that the optimum δ tends to infinity (so that the threshold in Case 1 serves as the lower bound). The threshold is thus tangent to the line

$$a(1 - 2\theta + 2\theta \log(2\theta)) = 1,$$

which in turn shows that $\theta_{\text{islands}}(a, 0) > 0$ for $a > 1$, as illustrated in Figure 1.

When $p \rightarrow 1$, a routine analysis shows that the optimum δ tends to zero (so that the threshold in Case 3 serves as the lower bound). The threshold is thus tangent to the line

$$a(1 - \theta + \theta \log \theta) = 1,$$

which matches the upper bound from Appendix A.

Appendix C. Tangency results

Tangency at $p = 0$

In this subsection, we show that the starting threshold and the local growth threshold θ_{local} are tangent at $p = 0$. The idea is simple. When p and $\theta = \theta_{\text{start}}(p)$ are small, we consider the effect of lowering the infection threshold from θ to $\tilde{\theta} = \theta - \delta$, with $\delta = o(\theta)$. Around a point \mathbf{u} that was newly (i.e., not initially) infected under the higher threshold θ , we now see a small disc $D = B_{\varepsilon r}(\mathbf{u})$ of newly infected points for some small constant $\varepsilon > 0$ (depending on δ). Next, consider a point $\mathbf{v} \in \partial D$. Once again, \mathbf{v} now sees increased infection in D , but also an infection rate of p , instead of θ , in the large lune $L = B_r(\mathbf{v}) \setminus B_r(x)$. However, since θ and p are both very small, the vast increase in the infection rate in D more than compensates for the greater area of the lune L , in which the infection rate is only a little lower than that in $B_r(\mathbf{u})$. Consequently, the infection grows outward from \mathbf{u} .

Going into more detail, note that $\theta_{\text{start}} = (-a \log p)^{-1}(1 + o(1))$ as $p \rightarrow 0$. A more careful analysis now yields the following result.

Proposition 6. *Let $C > \pi^{-2}$, and write $\theta = \theta_{\text{start}}(p)$. Then, for sufficiently small p , depending on a and C , we have*

$$\theta - C\theta^2 \leq \theta_{\text{local}}(p) \leq \theta = (-a \log p)^{-1}(1 + o(1)).$$

In particular, the starting threshold and local growth threshold are tangent at $p = 0$.

Proof. Write $\tilde{\theta} = \theta - C\theta^2 = \theta - \delta$. Suppose that an infection, started at \mathbf{u} , has spread to the disc $D = B_{\varepsilon r}(\mathbf{u})$. For a point $\mathbf{v} \in \partial D$, the area of the lune $L = B_r(\mathbf{v}) \setminus B_r(\mathbf{u})$ is $(2\varepsilon + O(\varepsilon^3))r^2$. Therefore, ignoring second-order terms, the condition for \mathbf{v} to be infected (and the infection to spread) under the bootstrap percolation model with threshold $\tilde{\theta}$ is

$$\pi \varepsilon^2 + (\pi(1 - \varepsilon^2) - 2\varepsilon)\theta + 2\varepsilon p \geq \tilde{\theta}\pi = (\theta - \delta)\pi.$$

Noting that $\theta \gg p$ (since $p \log p \rightarrow 0$ as $p \rightarrow 0$), we may replace this by

$$F(\theta, \varepsilon, \delta) = \pi(1 - \theta)\varepsilon^2 - 2\varepsilon\theta + \delta\pi \geq 0.$$

This holds for all $\varepsilon \geq 0$ as long as

$$(1 - \theta)C\theta^2 = (1 - \theta)\delta \geq \pi^{-2}\theta^2.$$

For sufficiently small p , this last inequality is guaranteed by the hypothesis $C > \pi^{-2}$, proving the first inequality in the theorem. The remaining inequalities follow from the definitions. □

An analysis of the argument reveals that, with $\delta = \pi^{-2}\theta^2$, a small disc $B_1 = B_{\varepsilon_1 r}(\mathbf{u})$ is immediately infected, where $\varepsilon_1 = \theta/2\pi$. After that, growth is progressively more difficult, in that the function F decreases, until the critical radius $\varepsilon_2 = 2\varepsilon_1 = \theta/\pi$, at which point $F = 0$. After that, F increases, and growth proceeds more and more easily.

Tangency at $p = 1$

Next we show that the local growth threshold θ_{local} is tangent to the limiting growth threshold $\theta = \frac{1+p}{2}$ at $p = 1$. Again, the idea is simple. Let us take $p = 1 - \delta$, so that, along the limiting growth threshold, $\theta = 1 - \delta/2$. Since we are well away from the starting threshold, the initial

infection rate of p will result in large circular regions $B_{Kr}(\mathbf{u})$ (where $K = \Theta(1/\delta)$) in which the initial infection density is $\theta = 1 - \delta/2$, not $p = 1 - \delta$, so that every point in $B_{(K-1)r}(\mathbf{u})$ will immediately become infected.

Let us now reduce the infection threshold θ from $1 - \delta/2$ to $1 - \delta/2 - C\delta^2$, and assume that the infection has spread to a disc $D = B_{Lr}(\mathbf{u})$, where $L \geq K - 1$. Consider a point $\mathbf{v} \in \partial D$. The point \mathbf{v} will see an infection density of 1 in $A = B_r(\mathbf{v}) \cap D$, and a density of at least p in $B = B_r(\mathbf{v}) \setminus D$. Because of the curvature of D , the area of A is slightly less than that of B , so that the average infection density in $B_r(\mathbf{v})$ will be $1 - \delta/2 - C\delta^2$ instead of $\frac{1+p}{2} = 1 - \delta/2$. However, we have lowered the threshold to $1 - \delta/2 - C\delta^2$ for precisely this reason, so that \mathbf{v} becomes infected, and the infection continues to spread.

Making these estimates rigorous is just a matter of bounding the Poisson distribution, as in the following proof.

Proposition 7. *Let $C > (6\sqrt{2}\pi)^{-1}$. Then, if $\delta = 1 - p$ is sufficiently small (depending on a and C), we have*

$$\frac{1+p}{2} - C\sqrt{a}\delta^2 = \frac{1+p}{2} - C(1-p)^2 \leq \theta_{\text{local}}(p) \leq \frac{1+p}{2}.$$

In particular, the local growth threshold and limiting growth threshold are tangent at $p = 1$.

Proof. Let C and δ be as in the statement of the proposition, and let $\theta = \frac{1+p}{2} - C\sqrt{a}\delta^2$. Consider a disc $D = B_{Kr}(\mathbf{u})$ of radius Kr , where $K = K(\delta)$ is large but to be determined. With an initial infection parameter of p , we expect to see

$$p\pi K^2 r^2 = paK^2 \log n$$

infected points in D . If we see instead

$$\theta\pi K^2 r^2 = \theta aK^2 \log n = paK^2 \log n(1 + \delta/2)(1 + o(1))$$

infections, uniformly distributed across D , then every point in $D' = B_{(K-1)r}(x)$ will immediately become infected. Setting $\rho = 1 + \delta/2$, the probability q of this occurring is given by

$$q = (1 + o(1))e^{paK^2 \log n(\rho - 1 - \rho \log \rho)} = e^{-(1+o(1))pa\delta^2 K^2 \log n/8} = n^{-(1+o(1))pa\delta^2 K^2/8},$$

by Lemma 1. Thus we should expect to see some fully infected discs of radius Kr , where

$$K = \frac{1}{\delta} \sqrt{\frac{8}{a}}(1 + o(1)). \tag{14}$$

Next we show that if K is sufficiently large, the infection will continue to spread indefinitely. For $K \geq 1/2$, write $M(K)$ for the area of the lune formed by a disc B_K of radius K and a unit disc whose centre lies on the perimeter of B_K . Exact formulas are given in Appendix B, but asymptotically

$$M(K) = \frac{\pi}{2} - \frac{1}{3K} + O\left(\frac{1}{K^2}\right).$$

Now, if the infection has already spread to all of $B_{Lr}(\mathbf{u})$, where $L \geq K - 1$, then the condition for it to grow further (and indefinitely) is that

$$M(L) + (\pi - M(L))p \geq \pi\theta.$$

Recall that $p = 1 - \delta$ and $\theta = 1 - \delta/2 - C\sqrt{a\delta^2}$. Using the above approximation for $M(K)$, and ignoring second-order terms (so that we may replace $L \geq K - 1$ by $L \geq K$, for instance), we can write the condition for the infection to spread as

$$\left(\frac{1}{2} - \frac{1}{3\pi K}\right) + \left(\frac{1}{2} + \frac{1}{3\pi K}\right)(1 - \delta) \geq 1 - \frac{\delta}{2} - C\sqrt{a\delta^2},$$

or

$$K \geq \frac{1}{3\pi C\sqrt{a\delta}}. \tag{15}$$

Combining (14) and (15) with the hypothesis $C > (6\sqrt{2}\pi)^{-1}$ yields the result. \square

Acknowledgements

We are grateful to two anonymous referees for their careful work on the paper, which led to improvements of both its presentation and its correctness.

Funding information

The first author’s research was supported by Swedish Research Council grants VR 2016-03488 and VR 2021-03687. Research on this project was done while the second author visited the first author in Umeå in spring 2018 with financial support from STINT Initiation grant IB 2017-7360, which the authors gratefully acknowledge.

Competing interests

There were no competing interests to declare which arose during the preparation or publication process of this article.

References

- [1] AIZENMAN, M. AND LEBOWITZ, J. (1988). Metastability effects in bootstrap percolation. *J. Phys. A* **21**, 3801–3813.
- [2] AMINI, H. (2010). Bootstrap percolation in living neural networks. *J. Statist. Phys.* **141**, 459–475.
- [3] AMINI, H., CONT, R. AND MINCA, A. (2016). Resilience to contagion in financial networks. *Math. Finance* **26**, 329–365.
- [4] AMINI, H. AND FOUNTOLAKIS, N. (2014). Bootstrap percolation in power-law random graphs. *J. Statist. Phys.* **155**, 72–92.
- [5] ARRATIA, R., GOLDSTEIN, L. AND GORDON, L. (1989). Two moments suffice for Poisson approximations: the Chen–Stein method. *Ann. Probab.* **17**, 9–25.
- [6] BALISTER, P., BOLLOBÁS, B., SARKAR, A. AND WALTERS, M. (2005). Connectivity of random k -nearest-neighbour graphs. *Adv. Appl. Probab.* **37**, 1–24.
- [7] BALISTER, P., BOLLOBÁS, B., SARKAR, A. AND WALTERS, M. (2009). A critical constant for the k nearest-neighbour model. *Adv. Appl. Probab.* **41**, 1–12.
- [8] BALISTER, P., BOLLOBÁS, B., SARKAR, A. AND KUMAR, S. (2007). Reliable density estimates for coverage and connectivity in thin strips of finite length. In *Proc. 13th Annual ACM International Conference on Mobile Computing and Networking*, Association for Computing Machinery, New York, pp. 75–86.
- [9] 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.
- [10] BALOGH, J. AND PITTEL, B. (2007). Bootstrap percolation on the random regular graph. *Random Structures Algorithms* **30**, 257–286.
- [11] BOLLOBÁS, B. (2006). *The Art of Mathematics: Coffee Time in Memphis*. Cambridge University Press.
- [12] BOLLOBÁS, B. *et al.* (2014). Bootstrap percolation on Galton–Watson trees. *Electron J. Probab.* **19**, article no. 13.
- [13] BOLLOBÁS, B. AND LEADER, I. (1991). Edge-isoperimetric inequalities in the grid. *Combinatorica* **11**, 299–314.

- [14] BOLLOBÁS, B. AND RIORDAN, O. (2006). *Percolation*. Cambridge University Press.
- [15] BRADONJIĆ, M. AND SANIEE, I. (2014). Bootstrap percolation on random geometric graphs. *Prob. Eng. Inf. Sci.* **28**, 169–181.
- [16] CANDELLERO, E. AND FOUNTOLAKIS, N. (2016). Bootstrap percolation and the geometry of complex networks. *Stoch. Process. Appl.* **126**, 234–264.
- [17] CHALUPA, J., LEATH, P. AND REICH, G. (1979). Bootstrap percolation on a Bethe lattice. *J. Phys. C* **12**, L31–L35.
- [18] DREYER, P. AND ROBERTS, F. (2009). Irreversible k -threshold processes: graph-theoretical threshold models of the spread of disease and of opinion. *Discrete Appl. Math.* **157**, 1615–1627.
- [19] EINARSSON, H. *et al.* (2019). Bootstrap percolation with inhibition. *Random Structures Algorithms* **55**, 881–925.
- [20] FALGAS-RAVRY, V. AND WALTERS, M. (2012). Sharpness in the k -nearest-neighbours random geometric graph model. *Adv. Appl. Prob.* **44**, 617–634.
- [21] FOUNTOLAKIS, N., KANG, M., KOCH, C. AND MAKAI, T. (2018). A phase transition regarding the evolution of bootstrap processes in inhomogeneous random graphs. *Ann. Appl. Prob.* **28**, 990–1051.
- [22] GAO, J., ZHOU, T. AND HU, Y. (2015). Bootstrap percolation on spatial networks. *Sci. Reports* **5**, 1–10.
- [23] GILBERT, E. (1961). Random plane networks. *J. Soc. Indust. Appl. Math.* **9**, 533–543.
- [24] GONZÁLEZ-BARRIOS, J. AND QUIROZ, A. (2003). A clustering procedure based on the comparison between the k nearest neighbors graph and the minimal spanning tree. *Statist. Prob. Lett.* **62**, 23–34.
- [25] GRANOVETTER, M. (1978). Threshold models of collective behavior. *Amer. J. Sociol.* **83**, 1420–1443.
- [26] GRAVNER, J. AND SIVAKOFF, D. (2017). Bootstrap percolation on products of cycles and complete graphs. *Electron. J. Prob.* **22**, article no. 29.
- [27] HOLROYD, A. (2003). Sharp metastability threshold for two-dimensional bootstrap percolation. *Prob. Theory Relat. Fields* **125**, 195–224.
- [28] JANSON, S. (2009). On percolation in random graphs with given vertex degrees. *Electron. J. Prob.* **14**, 87–118.
- [29] JANSON, S., KOZMA, R., RUSZINKÓ, M. AND SOKOLOV, Y. (2019). A modified bootstrap percolation on a random graph coupled with a lattice. *Discrete Appl. Math.* **258**, 152–165.
- [30] JANSON, S., ŁUCZAK, T., TUROVA, T. AND VALLIER, T. (2012). Bootstrap percolation on the random graph $G_{n,p}$. *Ann. Appl. Prob.* **22**, 1989–2047.
- [31] KOZMA, R. (2007). Neuropercolation. *Scholarpedia* **2**, 1360.
- [32] KINGMAN, J. (1993). *Poisson Processes*. Oxford University Press.
- [33] KOCH, C. AND LENGLER, J. (2016). Bootstrap percolation on geometric inhomogeneous random graphs. In *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, pp. 147:1–147:15.
- [34] KOCH, C. AND LENGLER, J. (2021). Bootstrap percolation on geometric inhomogeneous random graphs. *Internet Mathematics*. Available at <https://arxiv.org/abs/1603.02057>.
- [35] MEESTER, R. AND ROY, R. (1996). *Continuum Percolation*. Cambridge University Press.
- [36] PENROSE, M. (1997). The longest edge of the random minimal spanning tree. *Ann. Appl. Prob.* **7**, 340–361.
- [37] PENROSE, M. (2003). *Random Geometric Graphs*. Oxford University Press.
- [38] RAMOS, M. *et al.* (2015). How does public opinion become extreme?. *Sci. Reports* **5**, article no. 10032.
- [39] WHITTEMORE, A. (2021). Bootstrap percolation on random geometric graphs. Doctoral Thesis. University of Nebraska.
- [40] XUE, F. AND KUMAR, P. (2004). The number of neighbors needed for connectivity of wireless networks. *Wireless Networks* **10**, 169–181.