Combinatorics, Probability and Computing (2006) 15, 715–730. © 2006 Cambridge University Press
doi:10.1017/S0963548306007619 Printed in the United Kingdom

Bootstrap Percolation on Infinite Trees and Non-Amenable Groups

JÓZSEF BALOGH[1†], YUVAL PERES[2‡] and GÁBOR PETE[3§]

[1]Department of Mathematics, The Ohio State University, 231 W 18th Ave, Columbus, OH 43235, USA
(e-mail: jobal@sol.cc.u-szeged.hu www.math.ohio-state.edu/~jobal)

[2]Departments of Statistics and Mathematics, 367 Evans Hall,
University of California, Berkeley, CA 94720, USA
(e-mail: peres@stat.berkeley.edu www.stat.berkeley.edu/~peres)

[3]Department of Statistics, 367 Evans Hall, University of California, Berkeley, CA 94720, USA
(e-mail: gabor@stat.berkeley.edu www.stat.berkeley.edu/~gabor)

Received 27 January 2004; revised 25 September 2004

Bootstrap percolation on an arbitrary graph has a random initial configuration, where each vertex is occupied with probability p, independently of each other, and a deterministic spreading rule with a fixed parameter k: if a vacant site has at least k occupied neighbours at a certain time step, then it becomes occupied in the next step. This process is well studied on $\mathbb{Z}^d$; here we investigate it on regular and general infinite trees and on non-amenable Cayley graphs. The critical probability is the infimum of those values of p for which the process achieves complete occupation with positive probability. On trees we find the following discontinuity: if the branching number of a tree is strictly smaller than k, then the critical probability is 1, while it is $1-1/k$ on the k-ary tree. A related result is that in any rooted tree T there is a way of erasing k children of the root, together with all their descendants, and repeating this for all remaining children, and so on, such that the remaining tree T' has branching number $\mathrm{br}(T') \leqslant \max\{\mathrm{br}(T)-k, 0\}$. We also prove that on any $2k$-regular non-amenable graph, the critical probability for the k-rule is strictly positive.

1. Introduction and results

Consider a countable, connected, locally finite graph $G = G(V, E)$, with two possible states for each site in the vertex set V: vacant (0) or occupied (1). Start with a configuration picked according to the product Bernoulli measure $\mathbb{P}_p$, *i.e.*, each site is occupied randomly

† Partially supported by NSF grant DMS-0302804 and OTKA (Hungarian National Foundation for Scientific Research) grant T34475.
‡ Partially supported by NSF grants DMS-0104073 and DMS-0244479.
§ Partially supported by NSF grants DMS-0104073 and DMS-0244479, and OTKA grant T30074.

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

and independently with probability p. Then fix a parameter k, and consider the following deterministic spreading rule: if a vacant site has at least k occupied neighbours at a certain time step, then it becomes occupied in the next step. This process is called *bootstrap percolation*. *Complete occupation* is the event that every vertex becomes occupied during the process. The main problem is to determine the *critical probability* $p(G,k)$ for complete occupation: for infinite graphs G this is the infimum of the initial probabilities p that make $\mathbb{P}_p(\text{complete occupation}) > 0$. This model has a rich history in statistical physics, mostly on $G = \mathbb{Z}^d$ and finite boxes; we will give some references later.

For infinite trees the most important characteristic of growth is the *branching number* $\mathrm{br}(T)$ of the tree: see [18] or [20]. It is defined as the supremum of real numbers $\lambda \geqslant 1$ such that T admits a positive flow from the root to infinity, where on every edge $e \in E(T)$, the flow is bounded by $\lambda^{-|e|}$, and $|e|$ denotes the number of edges (including e) on the path from e to the root. This supremum does not depend on the root, and remains unchanged if we modify a finite portion of the tree. Two basic examples are $\mathrm{br}(T_k) = k$ for the $(k+1)$-regular tree, and $\mathrm{br}(T_\xi) = \mathbb{E}\xi$ a.s. given non-extinction for the Galton–Watson tree T_ξ with offspring distribution ξ. For finite trees, the branching number is 0.

On T_k, k-neighbour bootstrap percolation has $p(T_k,k) = 1 - 1/k$: see (1.4) in Proposition 1.2 below. In contrast, we have the following.

Theorem 1.1. *Let T be an infinite tree. If $\mathrm{br}(T) < k$, then $p(T,k) = 1$.*

The above results show a somewhat surprising discontinuity of the function

$$f_k(b) := \inf\{p(T,k) \,:\, \mathrm{br}(T) \leqslant b,\ T \text{ has bounded degree}\} \tag{1.1}$$

at the value $b = k$. If we omit the condition of bounded degree, the discontinuity is even sharper: it is easy to construct a tree with $\mathrm{br}(T) = k$ and $p(T,k) = 0$. A possible explanation of this discontinuity is given by Theorem 1.3 below.

For regular trees we give an equation for the critical probability, from which the actual value is more-or-less computable.

Proposition 1.2. *Let $2 \leqslant k \leqslant d$. The critical probability $p(T_d,k)$ is the supremum of all p for which the equation*

$$\mathbb{P}(\mathrm{Binom}(d,(1-x)(1-p)) \leqslant d-k) = x \tag{1.2}$$

has a real root $x \in (0,1)$. In particular, for any constant $\gamma \in [0,1]$ and a sequence of integers k_d with $\lim_{d\to\infty} k_d/d = \gamma$,

$$\lim_{d\to\infty} p(T_d,k_d) = \gamma. \tag{1.3}$$

Furthermore, for the extreme values of the parameter k,

$$p(T_d,d) = 1 - \frac{1}{d} \qquad \text{and} \qquad p(T_d,2) = 1 - \frac{(d-1)^{2d-3}}{d^{d-1}(d-2)^{d-2}} \sim \frac{1}{2d^2}. \tag{1.4}$$

There is a generalization of a weaker form of Theorem 1.1. For this we first have to introduce the following simple notion, which will also be central to our proofs.

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

Definition. A finite or infinite connected subset $F \subseteq V$ of vertices is called a *k-fort* if each $v \in F$ has outdegree $\deg_{V \setminus F}(v) \leqslant k$. Here $\deg_H(v) = |\{w \in H : (v,w) \in E\}|$, for any $H \subseteq V$.

A key observation is that the failure of complete occupation by the k-neighbour rule is equivalent to the existence of a vacant $(k-1)$-fort in the initial configuration.

Theorem 1.3. *Let T be an infinite tree. Then every vertex $x \in T$ is contained in a k-fort F with* $\mathrm{br}(F) \leqslant \max\{\mathrm{br}(T) - k, 0\}$.

This means that after fixing any vertex as the root, we can erase k children of it, together with all their descendants, and can repeat this for all the remaining children, and so on, so that this pruning process results in a required subtree F. It is interesting to note that the natural idea of pruning off the k subtrees with the largest branching numbers at each generation does not work in general.

For $\mathrm{br}(T) < k$ we get a $(k-1)$-fort with $\mathrm{br}(F) < 1$, which can happen only if F is finite, so $\mathrm{br}(F) = 0$. In fact, in Theorem 1.1 we prove that there are infinitely many finite $(k-1)$-forts of bounded size, which implies $p(T,k) = 1$. The impossibility of $0 < \mathrm{br}(F) < 1$ might be viewed as the reason for the discontinuity of $f_k(b)$ at $b = k$, though we do not actually know continuity at other points. See Section 5 for more discussion and open problems.

An infinite graph G has the *anchored expansion property* if, for some fixed vertex $o \in V(G)$, the *anchored Cheeger constant* is positive:

$$0 < \imath^*(G) := \liminf \left\{ \frac{|\partial_e S|}{|S|} : o \in S \subset V,\ S \text{ is finite and connected} \right\}, \tag{1.5}$$

where $\partial_e S$ is the set of edges in $E(G)$ with exactly one endpoint in S. It is easy to see that the value of $\imath^*(G)$ does not depend on the vertex o. This notion is implicit in [24], and was defined explicitly in [5]. For transitive graphs (such as Cayley graphs of finitely generated infinite groups) it coincides with the more familiar but less robust concept of non-amenability, where the infimum is taken over all finite connected subsets S. For background on non-amenability see [20] or [19], and on anchored expansion [13] and [26].

Theorem 1.4. *Let G_d be a d-regular graph. If $\imath^*(G_d) + 2k > d$, then $p(G_d,k) > 0$. In particular, if G_d has the anchored expansion property, then $p(G_d, \lceil d/2 \rceil) > 0$.*

This result is sharp in the sense that there exists a 6-regular non-amenable Cayley graph G_6 with $p(G_6,2) = 0$: see Section 4. We will pose a possible characterization of amenability in Section 5.

The issue of positivity of the critical probability is simpler for the case of trees. For this, let us denote by $q(G,k)$ the infimum of initial probabilities for which, following the k-neighbour rule on G, there will be an infinite connected component of occupied vertices in the final configuration with positive probability. Clearly, $q(G,k) \leqslant p(G,k)$.

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

Proposition 1.5. *For any integer d, and $k \geqslant 2$, if T is an infinite tree with maximum degree $d+1$, then $p(T,k) \geqslant q(T_d,k) > 0$.*

The first inequality of this proposition follows immediately from viewing T as a subgraph of T_d. The positivity of the critical probability $q(T_d,k)$ will be proved using our proof of Proposition 1.2 and an idea from [16].

Bootstrap percolation was first defined in the statistical physics literature in [6], where the formulae of (1.4) were given. A variant of the model appeared in [7]. The problem of complete occupation on $\mathbb{Z}^2$ was solved by [10]. Schonmann proved [23] that the critical probability $p(\mathbb{Z}^d,k)$ for bootstrap percolation is 0 for $k \leqslant d$ and is 1 for $k > d$. The process can also be considered on finite graphs: see, *e.g.*, [2], [3] and [15]. A short recent physics survey is [1]. Bootstrap percolation also has connections to the dynamics of the Ising model at zero temperature: see [11] for $\mathbb{Z}^d$, and [16] for T_2.

We conclude this Introduction by some basic observations.

If a graph G satisfies $\mathbb{P}_p(\text{complete occupation with the } k\text{-rule}) \in \{0,1\}$ for all $p \in [0,1]$, and so $\mathbb{P}_p(\text{complete occupation of } G) = 1$ for any $p > p(G,k)$, then we will say that the 0–1 *law holds* for G with the k-rule.

For example, if the orbit of each vertex under the automorphism group of G is infinite, then the product probability measure of the initial configuration is ergodic [20, Proposition 6.3], while complete occupation is an invariant property; hence it has probability 0 or 1. Furthermore, if there is a finite $(k-1)$-fort in such a G, we immediately have infinitely many copies of this, so $p(G,k) = 1$. On the other hand, we find the following.

Lemma 1.6. *If there are no finite $(k-1)$-forts in a graph G, then $p(G,k) \leqslant 1 - p_c(G)$, where $p_c(G)$ denotes the critical probabilityfor standard site percolation on G.*

Proof. In the case of no complete occupation, the vacant $(k-1)$-fort has to be infinite, thus we have an infinite connected vacant component in the initial configuration. To have this event with positive probability, the density of initial vacant sites has to be at least the critical probability $p_c(G)$. □

Therefore, if $p_c(G) > 0$ holds for a graph without finite $(k-1)$-forts, which is usually the case (*e.g.*, if the degrees of vertices are bounded – see [20, Proposition 6.9]), then $p(G,k) < 1$. For instance, on any tree T we have $p_c(T) = 1/\mathrm{br}(T)$, as was shown in [18].

We will say that a graph G is *uniformly bigger* than a graph H if every vertex of G is contained in a subgraph of G that is isomorphic to H.

Lemma 1.7 (Monotonicity). *If a graph G is uniformly bigger than H, and H satisfies the 0–1 law for some k-rule, then we have $p(G,k) \leqslant p(H,k)$.*

Proof. For any $p > p(H,k)$, any fixed vertex v of G becomes occupied almost surely, because of the copy of H containing v. There are countably many vertices of G, so we have $\mathbb{P}_p(\text{complete occupation of } G) = 1$ with this p. □

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

In particular, if T is a tree with maximal degree $d+1$ and it satisfies the 0–1 law, then we get $p(T,k) \geqslant p(T_d,k)$. Proposition 1.5 is a generalization of this fact. We thank Ádám Timár for pointing out the importance of considering $q(T_d,k)$ for the generalization.

2. Regular trees

Proof of Proposition 1.2. Consider the $(d+1)$-regular tree T_d, and fix $2 \leqslant k \leqslant d$. This tree has no finite $(k-1)$-forts, and it is easy to see that any infinite fort of it contains a complete $(d+2-k)$-regular subtree. Hence, unsuccessful complete occupation for the k-rule is equivalent to the existence of a $(d+2-k)$-regular vacant subtree in the initial configuration.

Note that complete occupation on T_d obeys the 0–1 law. So incomplete occupation has probability 1 if and only if a fixed origin is contained in a $(d+2-k)$-regular vacant subtree with positive probability. Now a simple use of Harris's inequality (see [20, Section 6.2]) gives that this is equivalent to having the following event with positive probability: a d-ary tree, rooted at the fixed origin that is declared to be vacant, has a vacant $(d+1-k)$-ary subtree starting from the same root. Therefore, we need to determine when the connected component of vacant sites of the root, which is a random Galton–Watson tree with offspring distribution $\mathrm{Binom}(d,1-p)$, contains a $(d+1-k)$-ary subtree with positive probability. If the probability of *not* having such a subtree is denoted by $y = y(p)$, then each of the d children of the root has probability $1-p$ of being vacant, and given this event, has probability $1-y$ of being the root of a vacant $(d+1-k)$-ary subtree. Therefore, y clearly satisfies the equation (1.2), *i.e.*, it is a fixed point of the function

$$\begin{aligned} x \mapsto B_{d,k,p}(x) &:= \mathbb{P}(\mathrm{Binom}(d,(1-x)(1-p)) \leqslant d-k) \\ &= \sum_{j=0}^{d-k} \binom{d}{j} (1-x-p+xp)^j (p+x-px)^{d-j}. \end{aligned}$$

One fixed point in $[0,1]$ is $x=1$; we are going to show that y is actually the smallest one in $[0,1]$. It is easy to see that

$$\frac{\partial}{\partial x} B_{d,k,p}(x) = d(1-p)\mathbb{P}(\mathrm{Binom}(d-1,(1-x)(1-p)) = d-k),$$

which is positive for $x \in [0,1)$, with at most one extremal point (a maximum) in $(0,1)$. Thus $B_{d,k,p}(x)$ is a monotone increasing function with $B_{d,k,p}(0) > 0$ and with at most one inflection point in $(0,1)$. If y_n denotes the probability that the required vacant subtree does not even reach the nth level below the root, then $y_0 = 0$, $y_{n+1} = B_{d,k,p}(y_n)$, and $y_n \to y$. On the other hand, the sequence y_n clearly approaches the smallest fixed point of $B_{d,k,p}(x)$, which so coincides with y. Thus, the infimum of the probabilities p for which equation (1.2) has no positive real root $x < 1$ is indeed the criticalprobability $p(T_d,k)$.

If $\lim_{d\to\infty} k_d/d = \gamma$, then for any fixed p and x, by the Weak Law of Large Numbers:

$$B_{d,k_d,p}(x) = \mathbb{P}\left(\frac{\mathrm{Binom}(d,(1-x)(1-p))}{d} \leqslant \frac{d-k_d}{d}\right) \to \begin{cases} 1, & \text{if } (1-x)(1-p) < 1-\gamma, \\ 0, & \text{if } (1-x)(1-p) > 1-\gamma, \end{cases}$$

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

as $d \to \infty$. Solving the equation $(1-x)(1-p) = 1-\gamma$ for x gives a critical value $x_c = (\gamma - p)/(1-p)$. Thus for $p > \gamma$ we have $\lim_{d\to\infty} B_{d,k_d,p}(x) \to 1$ for all $x \in [0,1]$, while for large enough d, $B_{d,k,p}(x)$ is convex in $[0,1]$, so there is no positive root $x < 1$ of $B_{d,k_d,p}(x) = x$. On the other hand, for $p < \gamma$ there must be a root $x = x(d)$ for large enough d, clearly satisfying $\lim_{d\to\infty} x(d) = x_c$. These prove (1.3).

The first equality of (1.4) follows immediately from (1.2). The second equality can be deduced by a standard calculus argument from our above formula for the first derivative of $B_{d,k,p}(x)$. □

Remark 1. We will use later that the extinction probability $y(p)$ introduced in the above proof satisfies $y(p) \to 0$ as $p \to 0$. This follows from the facts that $B_{d,k,0}(x) < x$ for $x > 0$ small enough, that the functions $B_{d,k,p}(x)$ converge uniformly to $B_{d,k,0}(x)$ as $p \to 0$, and that $B_{d,k,p}(0) > 0$.

Remark 2. A Galton–Watson tree with offspring distribution $\mathrm{Binom}(d, 1-p)$ can contain a $(d+1-k)$-ary subtree only if its mean is $(1-p)d \geqslant d+1-k$. Thus $p(T_d,k) \leqslant \frac{k-1}{d}$ follows immediately.

Remark 3. The problem of finding regular subtrees in certain Galton–Watson trees was first considered in [8], where the formula of (1.4) for $p(T_9,2)$ was used. For general Galton–Watson processes, see [21]. From (1.3) it follows that the critical mean value for a binomial offspring distribution to produce an N-ary subtree in the Galton–Watson tree is asymptotically N. In [21] it was shown that this critical mean value is $\sim eN$ for a geometric offspring distribution, and $\sim N$ for a Poisson offspring. An interesting feature of these phase transitions is that, unlike the case of usual percolation $N = 1$, for $N \geqslant 2$ the probability of having the N-ary subtree is already positive at criticality. For bootstrap percolation this means that the probability of complete occupation is still 0 at $p = p(T_d,k)$ if $k < d$.

Remark 4. The critical probability $p(T,k)$ can be computed also for quasi-transitive (periodic) trees and Galton–Watson trees, as we will see, for example, in Section 5.

Proof of Proposition 1.5. To prove the second inequality, $q(T_d,k) > 0$, we will first show that, for any non-backtracking path $v_0, v_1, \ldots, v_n$ in T_d,

$$\mathbb{P}_p\big(\{v_0,\ldots,v_n\} \text{ does not intersect any vacant } (k-1)\text{-fort}\big) \leqslant [1 - z(p)^2]^{\lfloor n/2 \rfloor}, \qquad (2.1)$$

where $z(p)$ is the probability that an infinite rooted tree Z with $d-1$ children at the root, and d children everywhere else, has a $(d+1-k)$-ary vacant subtree containing the root. Before proving (2.1), note that $z(p) \to 1$ as $p \to 0$. This follows easily from the fact $y(p) \to 0$ of Remark 1 above.

To prove (2.1), for each v_i consider a copy Z_i of Z inside T_d, rooted at v_i, disjoint from the path $v_0,\ldots,v_n$. Then the subtrees Z_i are also disjoint from each other. The $(d+1-k)$-ary vacant subtrees rooted at v_{2i} and v_{2i+1} inside Z_{2i} and Z_{2i+1} join together

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

to give a vacant $(d+2-k)$-regular tree, *i.e.*, a $(k-1)$-fort, inside T_d. The probability that this does not happen for any of the pairs v_{2i}, v_{2i+1} is exactly the right-hand side of (2.1).

The number of different paths $v_0, v_1, \ldots, v_n$ from a fixed vertex $v_0 = o$ is $(d+1)d^{n-1}$. Therefore, if p is so small that $\sqrt{1-z(p)^2} < 1/d$, then the probability that there is at least one such path that does not intersect any vacant $(k-1)$-forts in the initial configuration is exponentially small in n. By the Borel–Cantelli lemma, any infinite non-backtracking path from o eventually intersects a vacant $(k-1)$-fort almost surely, hence the bootstrap percolation process will not be able to form an infinite occupied cluster containing o. There are countably many possible o vertices in T_d, so we have the same for all vertices with probability 1. Thus $q(T_d, k) \geqslant p$ with the above small $p > 0$. □

3. General infinite trees

To start our discussion of the connection between branching number and bootstrap percolation, let us prove a simple combinatorial lemma, which implies Theorem 1.3 for the special case of $\mathrm{br}(T) < k$, but is not yet enough to prove Theorem 1.1.

Lemma 3.1 (Red Lemma). *If some vertex x of a tree T is not contained in any finite $(k-1)$-fort, then there is a k-ary subtree containing x.*

Proof. Consider the tree T as rooted at the vertex x. First colour red all vertices with at most $k-1$ children. In the second step, colour red each vertex with at most $k-1$ non-red children, and repeat this over and over again: see Figure 1. In the limiting final colouring, if the root x is red, then it obtained its colour in a finite number of steps, so there is a finite set F of vertices such that x becomes red even if we fix all the vertices outside F to be uncoloured forever. If we take this F to be minimal, then it is a finite connected subtree of T, with all leaves painted red in the first step, and all vertices becoming eventually red. But now, this red F is clearly a finite $(k-1)$-fort in T, contradicting the choice of x. Therefore, x is not red in the final limiting colouring. This means it has at least k non-red children, and each of these children also has at least k non-red children, and so on. Hence, the non-red component of x in T contains the k-ary subtree we wanted. □

Theorem 1.1 is not obvious from this lemma because if we do not forbid all finite $(k-1)$-forts, but only the appearance of too many small ones, we can get $p(T,k) < 1$, while we can have vertices with at most $k-1$ children lying close to each other (see the tree on the right in Figure 1), so $\mathrm{br}(T) < k$ could possibly occur, as well. Thus we need a quantitative version of Lemma 3.1. For this, fix a root o for T, and denote by $L_r(x)$ the set of the vertices y from which the shortest path to o contains x, and $\mathrm{dist}(o,y) = \mathrm{dist}(o,x) + r$.

Lemma 3.2 (Blue Lemma). *Let x be a vertex with $|L_R(x)| < (k-1)k^{R-1}$ for some positive integer R. Then $\cup_{0 \leqslant r \leqslant R} L_r(x)$ contains a $(k-1)$-fort of T.*

Proof. Let x be a vertex of the tree satisfying the conditions of the lemma, and label its level by 0. Colour a vertex on level R blue, if it has at least k children. In general, colour

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

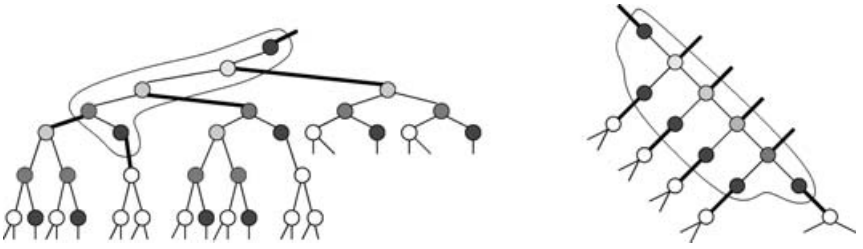


Figure 1. Some finite 1-forts.

a vertex on level r blue, if it has at least k blue children ($r < R$). The vertex x is definitely not blue, otherwise $|L_R(x)| \geqslant k^R$ would hold. Moreover, x has at most $k-2$ blue children, since $|L_R(x)| < (k-1)k^{R-1}$. If x has less than $k-1$ children, then it is a fort by itself. Otherwise, the non-blue component containing x contains at least 2 vertices. We claim that the non-blue connected component containing x is a $(k-1)$-fort. First of all, x has outdegree at most $k-1$, counting its mother and its possible $k-2$ blue children. Any other vertex from this set has a non-blue mother, and being non-blue means that it has at most $k-1$ blue neighbours. A non-blue vertex in this component in the level R has at most $k-1$ children, and its mother is not blue. □

Note that almost the same argument for $x = o$ gives that $|L_R(o)| < k^R$ implies a $(k-1)$-fort inside $\cup_{0 \leqslant r \leqslant R} L_r(o)$. (The reason for the strengthening is simple: o does not have a mother.)

Proof of Theorem 1.1. We will prove that if there is no $(k-1)$-fort with at most N vertices, then $\mathrm{br}(T) \geqslant k - \frac{2k \log k}{\log N}$. This suffices because destroying a finite number of forts of size at most N does not affect $\mathrm{br}(T)$, so $\mathrm{br}(T) < k$ will imply the existence of infinitely many $(k-1)$-forts of bounded size, which shows $p(T,k) = 1$.

Any leaf of T would be a $(k-1)$-fort with one vertex, so there are no leaves, and we have $|L_r(x)| \leqslant |L_{r+1}(x)|$ for any r and x. Hence, the fort that the Blue Lemma finds for us has less than $Rk^R < k^{2R}$ vertices. Thus having no $(k-1)$-forts of size at most N implies that $|L_r(x)| \geqslant (k-1)k^{r-1}$ and $|L_r(o)| \geqslant k^r$ for every $r \leqslant R = \frac{\log N}{2 \log k}$. We will prove that if $\lambda > 0$ is such that

$$\lambda^R \leqslant (k-1)k^{R-1}, \tag{3.1}$$

then $\mathrm{br}(T) \geqslant \lambda$. For example, if $\lambda = k - \frac{c_k}{\log N}$, then $\lambda^R \leqslant N^{1/2} \exp(-\frac{c_k}{2k \log k})$, while $(k-1)k^{R-1} = \frac{k-1}{k} N^{1/2}$, so $c_k = 2k \log k$ is good for any $k \geqslant 2$.

Now we have to show that if the capacity of an edge e is $\lambda^{-|e|}$, then the network admits a positive flow from the root to infinity. Start the flow with an amount k^{-R} at the root o. On level R there are at least k^R vertices; divide the initial amount equally among them, and build the flow from o to $L_R(o)$ according to these amounts. Then, through each edge before level $L_R(o)$ the amount that flows is at most the initial k^{-R}, while the capacity of

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

such an edge is at least λ^{-R}, which is bigger because of (3.1). So this is an admissible flow from o to $L_R(o)$.

The value of the flow at each vertex in $L_R(o)$ is at most k^{-2R}. For each such vertex x, we have $|L_R(x)| \geqslant (k-1)k^{R-1}$. Divide the amount at x equally among these vertices, do the same for all $x \in L_R(o)$, and continue the flow from $L_R(o)$ to $L_{2R}(o)$ according to this. Through each edge between $L_R(o)$ and $L_{2R}(o)$ the amount that should flow is at most k^{-2R}, while the capacity of an edge is at least λ^{-2R}. Thus we have an admissible flow constructed from o to $L_{2R}(o)$.

Continuing in this manner: between the levels $L_{(n-1)R}(o)$ and $L_{nR}(o)$ the amount that should flow through an edge is at most $(k-1)^{-(n-2)}k^{-(nR-n+2)}$, and the capacity of such an edge is at least λ^{-nR}. This second is bigger because of (3.1), thus we have constructed an admissible flow from o to infinity with the positive value k^{-R}. □

The converse is clearly false, as shown for example by a $(d+1)$-regular tree with an additional vertex placed on each of the edges. This tree T'_d has branching number $\sqrt{d}$, while any vertex of the original T_d together with its $d+1$ neighbours form a 1-fort of size $d+2$, so $p(T'_d, 2) = 1$. Furthermore, one might ask how sharp the above result $\mathrm{br}(T) \geqslant k - \frac{2k\log k}{\log N}$ is: see Section 5.

Proof of Theorem 1.3. Fix a vertex x as the root of T. It is enough to prove the theorem for $k = 1$, and thus find a small 1-fort F_1 in T containing x, because then we can inductively find a 1-fort F_2 inside F_1 with $\mathrm{br}(F_2) \leqslant \mathrm{br}(F_1) - 1 \leqslant \mathrm{br}(T) - 2$, which will also be a 2-fort in T, and so on.

By the Max-Flow-Min-Cut theorem, the branching number is characterized by

$$\mathrm{br}(T) = \sup\left\{\lambda : \inf_{\Pi} \sum_{e\in\Pi} \lambda^{-|e|} > 0\right\}, \tag{3.2}$$

where the inf is over cutsets Π of edges separating x from ∞. The expression $\mu_\lambda(\Pi) := \sum_{e\in\Pi} \lambda^{-|e|}$ will be called the *λ-content* of the cutset (or of an arbitrary set of edges).

Fix some $\beta > 1$, and take an arbitrary finite tree $\mathscr{T}$ with root r. By its boundary $\partial\mathscr{T}$ we mean the set of edges with a leaf as an endpoint. If $\mathscr{T} = \{r\}$, then let $\mu_\beta(\partial\mathscr{T}) = 1$. Otherwise, denote the children of r by $r_1, \ldots, r_\ell$. Deleting the edge (r, r_i) from $\mathscr{T}$ results in two connected components; the subtree that contains r_i will be denoted by $\mathscr{T}_i$, and $\mathscr{T}_i$ together with r by $\hat{\mathscr{T}}_i$. We have the disjoint union $\cup_{i=1}^{\ell} \partial\hat{\mathscr{T}}_i = \partial\mathscr{T}$, hence $m_1 + \cdots + m_\ell = \mu_\beta(\partial\mathscr{T})$, where $m_i := \mu_\beta(\partial\hat{\mathscr{T}}_i)$. We may assume $m_1 \leqslant m_2 \leqslant \cdots \leqslant m_\ell$. Now let us delete from $\mathscr{T}$ the entire 'β-largest' subtree $\mathscr{T}_\ell$. Then look at the subtrees $\mathscr{T}_1, \ldots, \mathscr{T}_{\ell-1}$, and repeat the whole procedure with each $\mathscr{T}_i$ instead of $\mathscr{T}$, with root r_i, deleting the 'β-largest' subtree from each $\mathscr{T}_i$. Repeat this procedure until reaching the boundary of $\mathscr{T}$ in all subtrees. The remaining subtree $\mathscr{F}$ is clearly a 1-fort inside $\mathscr{T}$. We claim that

$$\mu_{\beta-1}(\partial\mathscr{F} \cap \partial\mathscr{T}) \leqslant \mu_\beta(\partial\mathscr{T})^\alpha, \tag{3.3}$$

where $\alpha = \beta/(\beta-1)$. Equality holds only for finite β-ary trees $\mathscr{T}$, for integer β.

Before proving this claim, we show how it implies the existence of an infinite 1-fort F inside T, rooted at x, with $\mathrm{br}(F) \leqslant \mathrm{br}(T) - 1$.

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

Take a strictly decreasing sequence of positive numbers $\{\beta_n\}$ converging to $\mathrm{br}(T) \geqslant 1$. Let $\alpha_n = \beta_n/(\beta_n - 1) > 1$. We can suppose that T has no leaves. We have $\beta_1 > \mathrm{br}(T)$, so by the characterization (3.2), for any $\epsilon_1 > 0$ there exists a cutset Π^1 separating x from ∞ with $\mu_{\beta_1}(\Pi^1) < \epsilon_1$. If the finite subtree between x and Π^1 is called $\mathscr{T}^1$, then our above procedure finds a 1-fort $\mathscr{F}^1$ of $\mathscr{T}^1$, with $\mu_{\beta_1-1}(\partial\mathscr{F}^1 \cap \Pi^1) < \epsilon_1^{\alpha_1}$. This upper bound is less than $1/2$ if we choose $\epsilon_1 = 1/2$. Now denote the lower endvertices of the edges in $\partial\mathscr{F}^1 \cap \Pi^1$ by $x_1, \ldots, x_\ell$. The infinite subtree of T starting at x_i, called T_i, has branching number less than β_2. Hence, for each i and any $\epsilon_2 > 0$, we can take a cutset Π_i^2 separating x_i from ∞ with $\mu^*_{\beta_2}(\Pi_i^2) < \epsilon_2$, where $\mu^*_{\beta_2}$ denotes β_2-content with distances $|e|$ measured from the new root x_i. That is, $\mu_{\beta_2}(\Pi_i^2) < \epsilon_2 \beta_2^{-|x_i|}$. If the finite subtree between x_i and Π_i^2 is called $\mathscr{T}_i^2$, then our pruning procedure yields a 1-fort $\mathscr{F}_i^2$ inside each $\mathscr{T}_i^2$, satisfying $\mu^*_{\beta_2-1}(\partial\mathscr{F}_i^2 \cap \Pi_i^2) < \epsilon_2^{\alpha_2}$. If we take the union $\Phi^2 := (\partial\mathscr{F}_1^2 \cap \Pi_1^2) \cup \cdots \cup (\partial\mathscr{F}_\ell^2 \cap \Pi_\ell^2)$, then $\mu_{\beta_2-1}(\Phi^2) = \sum_{i=1}^{\ell} \mu_{\beta_2-1}(\partial\mathscr{F}_i^2 \cap \Pi_i^2) < \sum_{i=1}^{\ell} \epsilon_2^{\alpha_2}(\beta_2 - 1)^{-|x_i|} = \epsilon_2^{\alpha_2}\mu_{\beta_2-1}(\partial\mathscr{F}^1 \cap \Pi^1)$. Since $\mu_{\beta_2-1}(\partial\mathscr{F}^1 \cap \Pi^1)$ is a finite number independent of ϵ_2, we can choose ϵ_2 so small that the last upper bound is less than $1/4$. Now we repeat everything with the infinite subtrees of T starting at the lower endvertices $\{y_i, i = 1, \ldots, m\}$ of Φ^2, using β_3 and some $\epsilon_3 > 0$. This gives a collection of cutsets $\{\Pi_i^3\}$ and finite 1-forts $\{\mathscr{F}_i^3\}$. Take the union $\Phi^3 := (\partial\mathscr{F}_1^3 \cap \Pi_1^3) \cup \cdots \cup (\partial\mathscr{F}_m^3 \cap \Pi_m^3)$, and choose ϵ_3 sufficiently small so that $\mu_{\beta_3-1}(\Phi^3) < 1/8$. Repeat this ad infinitum, choosing ϵ_n such that $\mu_{\beta_n-1}(\Phi^n) < 2^{-n}$.

The union of all the finite 1-fort-pieces, $F := \mathscr{F}^1 \cup \mathscr{F}_1^2 \cup \cdots \cup \mathscr{F}_\ell^2 \cup \cdots$, is an infinite 1-fort of T, and each Φ^n is a cutset of F separating x from ∞. For any fixed $\beta > \mathrm{br}(T)$, if n is large enough to have $\beta_n < \beta$, then $\mu_{\beta-1}(\Phi^n) < \mu_{\beta_n-1}(\Phi^n) < 2^{-n}$. Thus, by definition (3.2), $\mathrm{br}(F) \leqslant \beta - 1$. Since this holds for all $\beta > \mathrm{br}(T)$, we have proved $\mathrm{br}(F) \leqslant \mathrm{br}(T) - 1$.

We prove (3.3) by induction on the depth of $\mathscr{T}$. If this depth is 1, *i.e.*, each child r_i of r is a leaf, then $\mathscr{F}$ is just obtained by deleting r_ℓ, so we need to prove $(\ell - 1)/(\beta - 1) \leqslant (\ell/\beta)^\alpha$. By taking derivatives with respect to ℓ it is easy to check that the only value of α for which this inequality holds for all real $\ell \geqslant 1$ is the chosen $\alpha = \beta/(\beta - 1)$. Equality holds only for $\ell = \beta$.

Suppose inductively that inside each subtree $\mathscr{T}_i$, $i = 1, \ldots, \ell$, we have our 1-fort $\mathscr{F}_i$ with $\mu^*_{\beta-1}(\partial\mathscr{F}_i \cap \partial\mathscr{T}_i) \leqslant m_i^\alpha$, where μ^*_λ denotes λ-content measured inside $\mathscr{T}_i$ with root r_i, and $m_i = \mu^*_\beta(\partial\mathscr{T}_i)$. We get $\mathscr{F}$ by joining the subtrees $\mathscr{F}_1, \ldots, \mathscr{F}_{\ell-1}$ at r, where $m_1 \leqslant \cdots \leqslant m_{\ell-1} \leqslant m_\ell$. Note that for $\ell = 1$ the claim is obvious. Now $\mu_{\beta-1}(\partial\mathscr{F} \cap \partial\mathscr{T}) = \big(\sum_{i=1}^{\ell-1} \mu^*_{\beta-1}(\partial\mathscr{F}_i \cap \partial\mathscr{T}_i)\big)/(\beta - 1) \leqslant (m_1^\alpha + \cdots + m_{\ell-1}^\alpha)/(\beta - 1)$, while $\mu_\beta(\partial\mathscr{T}) = (m_1 + \cdots + m_\ell)/\beta$. Therefore, we would like to prove that

$$\frac{m_1^\alpha + \cdots + m_{\ell-1}^\alpha}{\beta - 1} \leqslant \left(\frac{m_1 + \cdots + m_\ell}{\beta}\right)^\alpha \tag{3.4}$$

for all possible values of the m_is. Let $y = (m_1 + \cdots + m_\ell)/m_\ell$. Because of $\alpha > 1$ we have $m_i^\alpha \leqslant m_i m_\ell^{\alpha-1}$. Adding these inequalities up, we get $(m_1^\alpha + \cdots + m_{\ell-1}^\alpha)/(\beta - 1) \leqslant (y - 1)m_\ell^\alpha/(\beta - 1)$. Now recall that we proved $(y - 1)/(\beta - 1) \leqslant (y/\beta)^\alpha$ in the previous paragraph, hence our last upper bound is at most $(ym_\ell/\beta)^\alpha$. But this is just the right-hand side of (3.4), thus the proof of Theorem 1.3 is complete. □

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

4. Regular graphs with anchored expansion

A simple generalization of the result [23] for the Cayley graph $\mathbb{Z}^2$ with standard generators is proved by [12, Proposition 2.6]: for any symmetric generating set of $\mathbb{Z}^2$, the $2k$-regular Cayley graph Γ_{2k} has $p(\Gamma_{2k}, k) = 0$ and $p(\Gamma_{2k}, k+1) = 1$. As we have seen, the critical probabilities for regular trees all lie strictly between 0 and 1. Theorem 1.4 suggests that this contrast between $\mathbb{Z}^d$ and the free groups might have a geometric reason (see also the end of Section 5). Indeed, the proof of the theorem will be based on the 'perimeter method': see [4].

Proof of Theorem 1.4. Given an initial configuration of occupied vertices, a set $S \subseteq V(G)$ is called *internally spanned* if it becomes completely occupied even in the process restricted to S, *i.e.*, if we set all vertices in $V(G) \setminus S$ to be vacant forever. First of all, we claim that if complete occupation of G occurs, then for any fixed vertex $o \in V(G)$ there exists a strictly increasing sequence of finite connected internally spanned sets $o \in V_1 \subset V_2 \subset \cdots \subset V(G)$.

If the vertex o becomes occupied, then it does so in finite time, so there exists a finite vertex set V_1 such that o becomes occupied even in the finite process restricted to V_1. If we choose V_1 to be minimal, then it is clearly a connected internally spanned set containing o. Then let $V_1^* = V_1 \cup \{v\}$ for some vertex v neighbouring V_1. Each vertex of V_1^* becomes occupied in finite time, so there is a minimal finite set V_2 such that all of V_1^* becomes occupied even if the process is restricted to V_2. This finite set V_2 is internally spanned, connected, and strictly larger than V_1. Repeating this construction, we get the desired sequence of random sets $o \in V_1 \subset V_2 \subset \cdots$. Let us note that, with a bit more care, one can achieve $\cup_{n \geqslant 1} V_n = V(G)$, as well, but we will need only that there are arbitrarily large finite connected internally spanned sets containing o.

Denote $v_n = |V_n|$ and $w_n = |\partial_e V_n|$, and take some $0 < h < i^*(G)$ such that $h + 2k - d > 0$ still holds. The anchored expansion property ensures that $w_n / v_n \geqslant h$ for all sufficiently large n.

Look at the k-neighbour process restricted to an internally spanned V_n. If there are x_n initially occupied vertices in V_n, then the number of edges between these occupied vertices and all the vacant vertices of G (*i.e.*, the boundary of the occupied part) is at most dx_n initially. When a vacant vertex becomes occupied, the boundary will have at most $d - k$ new edges, while at least k old edges disappear, so the boundary increases by at most $d - 2k$. By the end of the complete occupation of V_n, we have occupied $v_n - x_n$ initially vacant vertices, and have ended up with a boundary w_n. Therefore,

$$\begin{aligned} dx_n + (v_n - x_n)(d - 2k) \geqslant w_n > h v_n, \text{ so} \\ x_n > v_n \frac{h - d + 2k}{2k} =: c v_n. \end{aligned} \tag{4.1}$$

Now take an i.i.d. Bernoulli(p) initial configuration on the whole infinite graph, with $0 < p < c$. Then, for any finite set $S \subset V(G)$,

$$\mathbb{P}_p(S \text{ contains at least } c|S| \text{ initially occupied vertices}) < e^{-I_p(c)|S|}, \tag{4.2}$$

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

by the Large Deviation Principle (see [9, Theorem 2.1.14]), where

$$I_p(c) = c \log \frac{c}{p} + (1-c) \log \frac{1-c}{1-p} \sim c \log \frac{1}{p}$$

when c is fixed and $p \to 0$.

By a beautiful, now well-known, percolation argument from [17], in a d-regular graph there are at most $((d-1)e)^m$ possible connected sets $S \ni o$ (usually called 'lattice animals') of size $|S| = m$. Therefore, putting everything together, for all large enough $M > 0$,

$$\begin{aligned}\mathbb{P}_p(\text{complete occupation}) &\leqslant \mathbb{P}_p(\exists \text{ internally spanned } S \ni o \text{ with } |S| > M) \\ &\leqslant \sum_{m=M}^{\infty} e^{-I_p(c)m + (\log(d-1)+1)m} \\ &\to 0, \text{ as } M \to \infty, \text{ if } I_p(c) > \log(d-1)+1.\end{aligned}$$

Thus $\mathbb{P}_p(\text{complete occupation}) = 0$ for $I_p(c) > \log(d-1)+1$, which holds for all small enough $p > 0$, in particular, for $p < K(c)/(de-e)^{1/c}$, where $K(c) = c(1-c)^{(1-c)/c}$. Therefore, $p(G_d, k) \geqslant K(c)/(de-e)^{1/c} > 0$. □

For the $(d+1)$-regular T_d, the above upper bound on the number of lattice animals roughly coincides with the true asymptotics $Cm^{-3/2}\left[d^d/(d-1)^{d-1}\right]^m$: see, *e.g.*, [22]. However, for $d+1 = 2k$, $i^*(T_d) = d-1$, $c = \frac{d-1}{d+1}$, the resulting estimate $p(T_d, \lceil (d+1)/2 \rceil) > \frac{1-o(1)}{(d-1)e}$ is very weak compared to the true value $\sim 1/2$ coming from (1.3).

The sharpness of our theorem is shown by the free product $\mathbb{Z}^2 * \mathbb{Z}$ with its natural 6-regular non-amenable Cayley-graph: from $p(\mathbb{Z}^2, 2) = 0$ it follows immediately that $p(\mathbb{Z}^2 * \mathbb{Z}, 2) = 0$. This also shows that the positivity result Proposition 1.5 cannot be generalized to graphs with fast growth.

Theorem 1.4 can easily be used to give examples of nontrivial critical probabilities for regular graphs that are not trees. For instance, the natural 4-regular Cayley graph of $\mathbb{Z}_3 * \mathbb{Z}_3$ has no finite 1-forts, so by Lemma 1.6 we get $0 < p(\mathbb{Z}_3 * \mathbb{Z}_3, 2) < 1$. For a more general result, see the end of Section 5.

5. Concluding remarks and open problems

The Red Lemma (Lemma 3.1) and the Monotonicity Lemma (Lemma 1.7) give that having no finite $(k-1)$-forts implies $p(T,k) \leqslant p(T_k,k) = 1 - 1/k$. There are examples showing that, in general, having no $(k-1)$-forts with $\mathrm{br}(F) < b-k+1$, where $b > k$ integer, does not imply $p(T,k) \leqslant p(T_b,k)$. On the other hand, the 1-fort F found by Theorem 1.3 is the largest possible (in terms of the $e^{\mathrm{br}(T)-1}$-dimensional Hausdorff measure of the boundary space ∂F, where the distance between two infinite rays $\xi, \eta \in \partial F$ is $e^{-|\xi \wedge \eta|}$ – see [20]) when T is a $\mathrm{br}(T)$-regular tree. This suggests that regular trees might play the role of extreme cases in the sense that $f_k(b) = p(T_b,k)$ for all $b \in \mathbb{N}$ for the function in (1.1), *i.e.*, they might be the trees with a fixed branching number which are the easiest to occupy. This would also nicely coincide with similar results for random walks: see [25] and [27]. However, as we will show below, for $b > k$ this is not the case, even for Galton–Watson

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

trees, for which Theorem 1.3 holds even with random pruning. So we are left with the following open problem.

The easiest trees to occupy. *Determine the function $f_k(b)$. Is it strictly positive for all real $b \geqslant 1$? Is it continuous apart from $b = k$?*

It is possible that $f_k(k) = 1 - \frac{1}{k}$. Also note that requiring a fixed bound on the degrees instead of the branching number also implies strict positivity, by Proposition 1.5.

Galton–Watson trees. One can study the same problems on a Galton–Watson tree T_ξ with offspring distribution ξ. For any p, the event $\{\mathbb{P}_p(\text{complete occupation of } T_\xi) > 0\}$ is an inherited event, so it has probability 0 or 1 (see [20, Proposition 4.6]), which shows that $p(T_\xi, k)$ is a constant almost surely, given non-extinction. If $\mathbb{P}(\xi < k) > 0$, then infinitely many finite $(k-1)$-forts of bounded size occur, so $p(T_\xi, k) = 1$. Otherwise, T_ξ can be built up from copies of T_k, and we get $p(T_\xi, k) \leqslant p(T_k, k) = 1 - 1/k$. We also have $\mathbb{P}_p(\text{complete occupation of } T_\xi) = 1$ a.s., given non-extinction, for $p > p(T_\xi, k)$. Just as above, this shows the following monotonicity property. If two offspring distributions ξ and η satisfy $\mathbb{P}(\xi < m) \geqslant \mathbb{P}(\eta < m)$ for all $m = 1, 2, \ldots$, *i.e.*, η stochastically dominates ξ, then there is natural coupling between the trees T_ξ and T_η such that T_η is uniformly bigger than T_ξ a.s., and so we get $p(T_\xi, k) \geqslant p(T_\eta, k)$.

A GW tree beating a regular tree. Consider the GW tree T_ξ with root r and offspring distribution $\mathbb{P}(\xi = 2) = \mathbb{P}(\xi = 4) = 1/2$. Then $\mathrm{br}(T_\xi) = \mathbb{E}\xi = 3$ a.s. [18], there are no finite 1-forts in T_ξ, and $0 < p(T_\xi, 2) < 1$ is an almost sure constant. We claim that $p(T_\xi, 2) < p(T_3, 2) = 1/9$.

Let $\mathscr{R}(x, T_\xi)$ be the event $\{$the vertex x of T_ξ is in an infinite vacant 1-fort$\}$, and set $q(T_\xi) = \mathbb{P}_p(\mathscr{R}(r, T_\xi))$. This is not an almost sure constant, so let us take expectation over all GW trees: $q = \mathbb{E}(q(T_\xi))$. Now

$$q = \frac{1}{2}\mathbb{E}(q(T_\xi) \mid \xi_r = 2) + \frac{1}{2}\mathbb{E}(q(T_\xi) \mid \xi_r = 4).$$

Regarding the first term, $\mathbb{P}_p(\mathscr{R}(r, T_\xi) \mid \xi_r = 2) = \mathbb{P}_p\big(r \text{ is initially vacant, and at least one of } \mathscr{R}(r_1, T'_\xi) \text{ and } \mathscr{R}(r_2, T''_\xi) \text{ does not fail}\big)$, where r_1, r_2 are the two children of r, and T'_ξ, T''_ξ are the corresponding subtrees. By the independence of initial configurations in T'_ξ and T''_ξ, this is equal to $(1-p)\big(\mathbb{P}_p(\mathscr{R}(r_1, T'_\xi)) + \mathbb{P}_p(\mathscr{R}(r_2, T''_\xi)) - \mathbb{P}_p(\mathscr{R}(r_1, T'_\xi))\mathbb{P}_p(\mathscr{R}(r_2, T''_\xi))\big)$. Now, by the recursive structure of T_ξ, and the independence of the subtrees T'_ξ and T''_ξ, taking the conditional expectation gives

$$\mathbb{E}(q(T_\xi) \mid \xi_r = 2) = (1-p)(2q - q^2).$$

A similar argument for the second term gives

$$\mathbb{E}(q(T_\xi) \mid \xi_r = 4) = (1-p)(4q^3 - 3q^4).$$

Altogether, we have the equation $q = \frac{1}{2}(1-p)(2q - q^2 + 4q^3 - 3q^4)$, and need to determine the infimum of ps for which there is no solution $q \in (0, 1]$ – that infimum will be $p(T_\xi, 2)$. Setting $f(q) = 2 - q + 4q^2 - 3q^3$, an examination of $f'(q)$ gives that

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

$\max\{f(q) : q \in [0,1]\} = f((4+\sqrt{7})/9) = 2.2347\cdots$. So there is no solution $q > 0$ if and only if $2/(1-p) > 2.2347\cdots$, which gives $p(T_\xi, 2) = 0.10504\cdots < 1/9$. □

Small k-forts. If we define $\Gamma_k(N)$ as the set of trees without k-forts of size at most N, and $\gamma_k(N) = \inf\{\mathrm{br}(T) : T \in \Gamma_k(N)\}$, then we know only

$$2 - 2^{-(1+o(1))N/2} \geqslant \gamma_1(N) \geqslant 2 - \frac{c}{\log N}. \tag{5.1}$$

The upper bound is achieved by trees analogous to the tree on the left in Figure 1. (The vertices of degree two are distributed following a greedy strategy: let the root be the first one of them, and then, at the generic step, put them on the highest level possible, in the highest-possible number at that level, subject to not forming a 1-fort of size at most N.) Actually, this gives asymptotically the smallest branching number that a tree in $\Gamma_1(N)$ with maximal degree 3 can have. The lower bound comes from the proof of Theorem 1.1.

Amenable and non-amenable groups. As we discussed in Section 4, for the free Abelian groups $\mathbb{Z}^d$ the critical probabilities are almost completely determined by [23] and [12]. The simplest non-Abelian group, the *Heisenberg group*, can be considered with natural generator sets of 2 or 3 elements [14], and it seems reasonable to conjecture that the corresponding 4- or 6-regular Cayley graphs H_4 and H_6 have $p(H_{2k}, k) = 0$. One can easily find finite k-forts to prove $p(H_{2k}, k+1) = 1$.

The most famous amenable groups with exponential growth are the *lamplighter groups* $\mathbb{Z}_r \wr \mathbb{Z}^d$. With a natural generating set, the Cayley graph of $\mathbb{Z}_r \wr \mathbb{Z}$ is the *Diestel–Leader graph* $\mathrm{DL}(r,r)$, where $\mathrm{DL}(r,s)$ is the 'horocyclic product' of two regular trees T_r and T_s: see [28]. These transitive graphs with degree $r+s$ are amenable if and only if $r = s$, and it is conjectured that for $r \neq s$ they are not quasi-isometric to any Cayley graph. It is not difficult to see that k-neighbour bootstrap percolation on $\mathrm{DL}(r,s)$, where $r \leqslant s$, has critical probability 1 for $k > s$, while strictly between 0 and 1 for $r+1 \leqslant k \leqslant s$, if such k exists. However, it is unclear if $p(\mathrm{DL}(r,s), r) = 0$ holds or not. A positive answer, together with our proof of Theorem 1.4, would have the interesting consequence that, as p gets closer and closer to 0, complete occupation of $\mathrm{DL}(r,r)$ by the r-neighbour rule will happen more and more through Følner sets, rather than through the exponentially growing balls.

If a finitely generated non-amenable group G contains a free subgroup on two elements, then, as David Revelle pointed out, there exists a generating set of k elements with $0 < p(G_k, \lceil k/2 \rceil) < 1$. The reason is that if G is originally defined by a symmetric generating set of t elements, then taking $2t$ free symmetric generators inside the free subgroup, we arrive at a $k = t + 2t$-regular graph, in which each vertex is contained in a $2t$-regular subtree. So our results give $0 < p(G_k, \lceil 3t/2 \rceil) \leqslant p(T_{2t-1}, \lceil 3t/2 \rceil) < 1$.

An open question inspired by the above results: *Is a group amenable if and only if, for any finite generating set, the resulting k-regular Cayley graph has $p(G_k, \ell) \in \{0,1\}$ for any ℓ-neighbour rule?*

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

Acknowledgements

We are grateful to Dayue Chen, Manjunath Krishnapur, Fabio Martinelli, Robin Pemantle, David Revelle, Ádám Timár, Bálint Virág and the referee for helpful discussions and comments.

References

[1] Adler, J. and Lev, U. (2003) Bootstrap percolation: Visualizations and applications. *Brazil. J. Phys.* **33** 641–644.
[2] Aizenman, M. and Lebowitz, J. L. (1988) Metastability effects in bootstrap percolation. *J. Phys. A* **21** 3801–3813.
[3] Balogh, J. and Bollobás, B. (2006) Bootstrap percolation on the hypercube. *Probab. Theory Relat. Fields* **134** 624–648.
[4] Balogh, J. and Pete, G. (1998) Random disease on the square grid. *Random Struct. Alg.* **13** 409–422.
[5] Benjamini, I., Lyons, R. and Schramm, O. (1999) Percolation perturbations in potential theory and random walks. In *Random Walks and Discrete Potential Theory* (Cortona, 1997), *Sympos. Math. XXXIX* (M. Picardello and W. Woess, eds), Cambridge University Press, Cambridge, pp. 56–84.
[6] Chalupa, J., Reich, G. R. and Leath, P. L. (1979) Bootstrap percolation on a Bethe lattice. *J. Phys. C* **12** L31–L35.
[7] Chalupa, J., Reich, G. R. and Leath, P. L. (1982) Inverse high-density percolation on a Bethe lattice. *J. Statist. Phys.* **29** 463–473.
[8] Chayes, J. T., Chayes, L. and Durrett, R. (1988) Connectivity properties of Mandelbrot's percolation process. *Probab. Theory Rel. Fields* **77** 307–324.
[9] Dembo, A. and Zeitouni, O. (1998) *Large Deviations Techniques and Applications*, 2nd edn, Springer, New York.
[10] van Enter, A. C. D. (1987) Proof of Straley's argument for bootstrap percolation. *J. Statist. Phys.* **48** 943–945.
[11] Fontes, L. R., Schonmann, R. H. and Sidoravicius, V. (2002) Stretched exponential fixation in stochastic Ising models at zero temperature. *Comm. Math. Phys.* **228** 495–518.
[12] Gravner, J. and Griffeath, D. (1996) First passage times for threshold growth dynamics on $\mathbb{Z}^2$. *Ann. Probab.* **24** 1752–1778.
[13] Häggström, O., Schonmann, R. H. and Steif, J. (2000) The Ising model on diluted graphs and strong amenability. *Ann. Probab.* **28** 1111–1137.
[14] de la Harpe, P. (2000) *Topics in Geometric Group Theory*, Chicago Lectures in Mathematics, University of Chicago Press.
[15] Holroyd, A. E. (2003) Sharp metastability threshold for two-dimensional bootstrap percolation. *Probab. Theory Rel. Fields* **125** 195–224.
[16] Howard, C. D. (2000) Zero-temperature Ising spin dynamics on the homogeneous tree of degree three. *J. Appl. Probab.* **37** 736–747.
[17] Kesten, H. (1982) *Percolation Theory for Mathematicians*, Birkhäuser, Boston.
[18] Lyons, R. (1990) Random walks and percolation on trees. *Ann. Probab.* **18** 931–958.
[19] Lyons, R. (2000) Phase transitions on nonamenable graphs. *J. Math. Phys.* **41** 1099–1126.
[20] Lyons, R., with Peres, Y. (2004) *Probability on Trees and Networks*, book in preparation for Cambridge University Press, available at `http://mypage.iu.edu/~rdlyons`.
[21] Pakes, A. G. and Dekking, F. M. (1991) On family trees and subtrees of simple branching processes. *J. Theoret. Probab.* **4** 353–369.

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press

[22] Pitman, J. (1998) Enumerations of trees and forests related to branching processes and random walks. In *Microsurveys in Discrete Probability* (Princeton, 1997), Vol. 41 of *DIMACS Ser. in Discrete Math. and Theoret. Comput. Sci.*, AMS, Providence, pp. 163–180.
[23] Schonmann, R. H. (1992) On the behavior of some cellular automata related to bootstrap percolation. *Ann. Probab.* **20** 174–193.
[24] Thomassen, C. (1992) Isoperimetric inequalities and transient random walks on graphs. *Ann. Probab.* **20** 1592–1600.
[25] Virág, B. (2000) On the speed of random walks on graphs. *Ann. Probab.* **28** 379–394.
[26] Virág, B. (2000) Anchored expansion and random walk. *Geom. Funct. Anal.* **10** 1588–1605.
[27] Virág, B. (2002) Fast graphs for the random walker. *Probab. Theory Rel. Fields* **124** 50–74.
[28] Woess, W. (2005) Lamplighters, Diestel–Leader graphs, random walks, and harmonic functions. *Combin. Probab. Comput.* **14** 415–433

https://doi.org/10.1017/S0963548306007619 Published online by Cambridge University Press