Normalization Phenomena in Asynchronous Networks

Amin Karbasi[1], Johannes Lengler[2], and Angelika Steger[2]

[1] School of Engineering & Applied Science, Yale University, New Haven, USA
amin.karbasi@yale.edu
[2] Department of Computer Science, ETH Zurich, Zurich, Switzerland
johannes.lengler|angelika.steger@inf.ethz.ch

Abstract. In this work we study a diffusion process in a network that consists of two types of vertices: *inhibitory* vertices (those obstructing the diffusion) and *excitatory* vertices (those facilitating the diffusion). We consider a continuous time model in which every edge of the network draws its transmission time randomly. For such an asynchronous diffusion process it has been recently proven that in Erdős-Rényi random graphs a *normalization phenomenon* arises: whenever the diffusion starts from a large enough (but still tiny) set of active vertices, it only percolates to a certain level that depends only on the activation threshold and the ratio of inhibitory to excitatory vertices. In this paper we extend this result to all networks in which the percolation process exhibits an explosive behaviour. This includes in particular inhomogeneous random networks, as given by Chung-Lu graphs with degree parameter $\beta \in (2, 3)$.

1 Introduction

One of the main goals in studying complex networks (e.g., social, neural networks) is to better understand the interconnection between the elements of such systems, and as a result being able to reason about their accumulated behavior. Of particular interest is to make a connection between the network's *structure* and *function*: once we have quantified the configuration of a network, how can we turn the results into predictions on the overall system behaviour?

A natural setting in which network structure plays a central role is the *diffusion of innovation* where a new product is introduced to the market (e.g., a new search engine), a new drug is suggested to doctors, or a new political movement has gained power in an unstable society. Once an innovation appears, people may have different reactions: some embrace it and try to promote the innovation, and others may refute it and try to obstruct the circulation of innovation. As a result, depending on the nature of the innovation (how much intrinsically people like or dislike the new product, idea, etc) and the structure of the network, the diffusion may die out quickly or spread explosively through the population. In order to understand to what extent an innovation is accepted it is important to understand the dynamics of the diffusion within the underlying network.

We make the following assumptions: after trying an innovation, an individual's reaction (i.e., like or dislike of the innovation) is parametrized by a binary random variable that takes -1 with probability τ (indicating that she dislikes the innovation) and $+1$ with probability $1-\tau$ (indicating that she likes the innovation). We assume that each individual's reaction is *independent* of that of its neighbors. However, an individual's tendency to become *active*, meaning that she tries the new innovation, depends on input from neighbors who have already tried the innovation and either promote or bash it. More precisely, we adopt the so-called *linear threshold model* [16]: an individual switches her state from inactive to active if the difference between its active supporting and its active inhibiting neighbors is above a threshold k. Given an initial set of early active individuals (those who were first exposed to the innovation and tried it), the diffusion process unfolds according to the aforementioned model. What we are interested in is to understand under what conditions an initially active set will spread through a non-trivial portion of the population.

Related Work. Classically, such diffusion processes have been studied without inhibition under the name of *bootstrap percolation* [9], [11]. With hindsight we note that in such processes the exact timing of the transmission of information from an active vertex to its neighbors plays no role and we can thus assume that all transmission times are exactly one. The activation then takes place in rounds. Such process were first studied by Chalupa et al. [9] on a 2-dimensional lattice in the context of magnetic disordered systems. Since then it has been the subject of intense research and found numerous applications for modeling complex systems' behaviors such as social influence [16], infectious disease propagation [18], jamming transitions [20], and neuronal activity [8], to name a few. The main focus of the previous work has been to understand the final number of active vertices. In the case of the 2-dimensional lattice, Holroyd [15] determined a sharp asymptotic threshold: as the size of the initial active set goes above a threshold, the process percolates to almost all vertices. This result was then generalized to 3-dimensional [4] and recently to any d-dimensional grid [3]. Bootstrap percolation has also been studied on a variety of graphs such as trees [5], random regular graphs [6], Erdős-Rényi graphs [17], and power-law graphs [2]. All this work exhibits an all-or-nothing phenomenon: either the initial set is small and the diffusion stops quickly or it is large enough and percolates to almost all vertices.

In [12] the authors introduced the bootstrap percolation process with inhibitory vertices as a model for the spread of activity through a population of neurons. That paper analyzed the bootstrap percolation process for Erdős-Rényi random graphs. The authors showed that on the one hand, if activation takes place in rounds inhibition makes the network highly susceptible to small changes in the starting set, but that on the other hand it does not so in the *asynchronous case* (where the spread of activity requires an exponentially distributed, random transmission time). However, the proofs relied substantially on the symmetry of Erdős-Rényi random graphs. The aim of the paper at hand is to extend the results from [12] (in the asynchronous setting) to a much larger class of graphs that do not exhibit a uniform degree distribution as the Erdős-Rényi random graphs do. This graph class contains in particular power-law distributed random graphs as provided by the so-called Chung-Lu model.

It has been empirically observed that many real networks follow a power-law degree distribution, including internet topology [13], World Wide Web [1], and Facebook social network [21]. That is, for most real-world graphs, the fraction of vertices with degree d scale like $d^{-\beta}$ where β is usually between 2 and 3. As a result, power-law graphs represent real-world networks more realistically than Erdős-Rényi graphs [7]. There are many generative models that construct such power-law graphs. In the present work we use the model of Chung and Lu [10]. It is known that for a graph of size n and in the absence of any inhibition, $\theta(n) := n^{(\beta-2)/(\beta-1)}$ is a coarse threshold for bootstrap percolation in such scale-free graphs [2].

Main Results. We say that a process looks *explosive* to an individual if in a short time a large number of her neighbors become active (for a formal definition, see Section 3). In [12] it was shown (without using this terminology) that asynchronous bootstrap percolation is explosive for Erdős-Rényi random graphs. In Theorem 2, we show that explosive percolation is more prevalent and happens also for *power-law graphs*, if we restrict ourselves to vertices of large degree.

In Theorem 1 we prove that an explosive process is automatically *normalizing*: of all individuals to which the process looks explosive, a certain fraction will turn active that can be accurately estimated as a function of τ and k that is *independent* of the structure of the graph. In contrast to all-or-nothing phenomena where the size of the final active set can either be very small or very large, our result provides a middle ground for asynchronous percolation processes. Note that to some vertices the process does not look explosive for a rather trivial reason: they may not have many neighbors at all. For these low degree vertices we can still give a heuristic prediction on how many of them will become active (cf. Section 5.1).

To further support our theoretical results, we perform experiments on real-world networks such as the Epinions social network. We observe that already for a small number of vertices our theoretical estimates for the final set of active vertices matches the numbers provided by the experimental data.

2 Formal Definitions and Notation

Let $G = (V, E)$ be a finite graph, and let $A \subset V$. Let $k \geq 2$ be an integer, and let $\tau \in [0, 1]$. Then the (k, τ)-*bootstrap percolation process* on G with starting set A is defined as follows. We first split the set of vertices randomly into two subsets V^+ and V^-. Each vertex is independently assigned to V^- with probability τ, and to V^+ otherwise. We call the vertices in V^+ *excitatory*, and the vertices in V^- *inhibitory*. We start the percolation process at time $t = 0$. At this time, all vertices in A turn *active*, and all other vertices are *inactive*. Whenever a vertex becomes active, it sends out signals to all its neighbors. Each signal takes a random *transmission time* to travel to its target; all transmission times are independently drawn from the exponential distribution Exp(1) with expectation 1. For every vertex v and time t, let $S^+(v,t)$ and $S^-(v,t)$ be the number of signals that have arrived at v until time t and that originated from excitatory and inhibitory vertices, respectively. Let $S(v,t) := S^+(v,t) - S^-(v,t)$. Whenever there is an inactive vertex v and a time $t > 0$ such that $S(v,t) \geq k$ then v immediately turns active, and sends out signals to all its neighbors. Active vertices do not turn inactive again. In particular, during the process each active vertex sends out exactly one signal to each of its neighbors. For $t \geq 0$, we denote by $a(t)$ the number of active neighbors at time t, and by a^* the number of active neighbors at termination, i.e., $a^* = \lim_{t\to\infty} a(t)$. As further notation, let $\Gamma_s^+(v)$ and $\Gamma_s^-(v)$ be the number of active excitatory and inhibitory neighbors of v among the first s active neighbors, respectively. (We only consider these random variables when there are at least s active neighbors.) Note that $\Gamma_s^+(v,t) + \Gamma_s^-(v,t) = s$. Finally, let $X_i(v) \in \{\pm 1\}$ be the random variable that describes whether the i-th signal arriving at v is excitatory or inhibitory.

Chung-Lu model. To generate power law graphs of size n, we use the model of Chung and Lu [10]. In this model, each vertex i is assigned a positive weight w_i. The probability p_{ij} of having an edge between vertices i and j is $\min\{1, w_i w_j / z\}$ where $z = \sum_{i=1}^n w_i$. Note that $G(n,p)$ can be viewed as a special case where for all vertices i we set $w_i = pn$. To generate a power-law graph with exponent $2 < \beta < 3$ and (constant) average degree $\bar{d}$, we set $w_i = \bar{d}\frac{\beta-2}{\beta-1}(\frac{n}{i})^{1/(\beta-1)}$.

Basic notation. For a sequence of events $\mathcal{E} = \mathcal{E}(n)$ we say that $\mathcal{E}$ holds *with high probability* (w.h.p.) if $\Pr[\mathcal{E}(n)] \to 1$ for $n \to \infty$. For any $x, y, z \in \mathbb{R}$, we use $x \in y \pm z$ to abbreviate the inequalities $y - z \leq x \leq y + z$.

3 Results

The main result of our paper is that the normalization phenomenon occurs for a large class of graphs, i.e., there is a universal constant $\alpha = \alpha(\tau, k)$ that does not depend on the structure of the graph such that bootstrap percolation activates an α-fraction of all vertices to which the process appears *fast*. We start with a precise definition of what we mean by "fast".

Definition 1. *Let $G = (V,E)$ be a graph*[3] *with n vertices, and let $A, S \subset V$. Furthermore, let $C \in \mathbb{N}$ and $\eta, \delta > 0$. We say that percolation on G with starting set A is* locally (C,η,δ)-explosive *for the set S if the following holds with probability at least $1-\delta$. For all but at most $\eta|S|$ vertices $v \in S$ there are times $t = t(v)$ such that at time t the vertex v has no active neighbor and at time $t+\eta$ it has at least C active neighbors.*

In the above definition, the probability is taken with respect to the random choices of inhibitory/excitatory signs and random delays. Also note that the

[3] All results and proofs carry over immediately to directed graphs as well. In this case, explosiveness needs to be defined with respect to the in-neighbors of G.

time t may well depend on the run of the diffusion and how it unfolds. If t can be chosen independently of v (but possibly still depending on the run), then we call the process *globally* (C, η, δ)*-explosive for* S. In this case, we call the time t the *start of the explosion*.[4]

Intuitively, a locally (C, η, δ)-explosive process looks fast from the point of view of individual vertices $v \in S$: the time that is needed to go from 0 active neighbors to C active neighbors is at most η. A *globally* (C, η, δ)-explosive process also looks fast from a global point of view, since the process needs only time η to go from a situation where almost no vertex in S has any active neighbors to a point where almost all of them have many neighbors.

We have defined (C, η, δ)-explosive processes for all values of C and η, but we will use them only in the case that C is large and $\eta > 0$ is very small. It turns out that for many standard graph models, percolation processes are locally (C, η, δ)-explosive for a large enough set S. For example, for the Erdős-Rényi random graph model $G(n, p)$ (with $p \gg 1/n$) it was observed in [12] that this is the case for $S = V$ and for all constants $C, \eta > 0$ (see also Theorem 3 in Section 4.1 of the present paper). In Theorem 2 below we prove the same for power law graphs, and in Section 5 we perform experiments indicating that percolation is also explosive on real-world networks.

The following theorem tells us that every process that is locally (C, η, δ)-explosive for some set S is also *normalizing* for that set, i.e., the number of active vertices in S at the end of the process is roughly $\alpha|S|$, where the constant α given by

$$\alpha = \min\left\{\left(\frac{1-\tau}{\tau}\right)^k, 1\right\}. \tag{1}$$

Note crucially that α does not depend on the structure of the graph or on the size of S.

Theorem 1. *For every $\varepsilon > 0$ there exist positive constants C_S, C, η, $\delta > 0$ such that the following holds. For a graph $G = (V, E)$ with n vertices, and sets $A, S \subset V$ such that $|S| \geq C_S$ and $|A \cap S| \leq \eta|S|$, if G with starting set A is locally (C, η, δ)-explosive for S, then with probability at least $1-\varepsilon$ the percolation process will terminate with $(\alpha \pm \varepsilon)|S|$ active vertices in S. In particular, for $S = V$ the final active set has size $a^* = (\alpha \pm \varepsilon)n$.*

Theorem 1 has vast implications since many bootstrap percolation processes are (C, η, δ)-explosive – mostly even globally explosive. In the latter case, our proofs imply that we truly have an explosive behaviour in the set S: for every $\varepsilon, \varepsilon' > 0$ there is a constant C and a time t (that may depend on the process at hand) such that at time t there are at most $\varepsilon|S|$ active vertices in S, while at time $t + \varepsilon'$ there are at least $(\alpha - \varepsilon)|S|$ active vertices in S. Since we know that the final number of active vertices is at most $(\alpha + \varepsilon)|S|$, we may informally restate this fact as follows. For sufficiently large C and sufficiently small η, in a (C, η, δ)-explosive process all but an arbitrarily small fraction of the activations in S will happen in an arbitrarily short time interval.

[4] By slight abuse of notation, as t is not unique.

Recall that the C'-core of a graph is defined as the largest subgraph in which all vertices have degree at least C'. In the following, we show that for every $C, \eta > 0$, *scale free networks* are locally (C, η, δ)-explosive if we choose S to be the C'-core for some constant $C' = C'(C, \eta, \delta)$.

Theorem 2. *Let $G = G_n = (V_n, E_n)$ be a Chung-Lu power law graph with exponent $\beta \in (2, 3)$. Moreover, let $a = a_n$ be such that $a_n \in \omega(n^{(\beta-2)/(\beta-1)})$ and $a_n \in o(n)$, and let $A = A_n \subset V_n$ be a random set of size a. Then for all constants $C, \eta, \delta > 0$ there exists $C' > 0$ such that w.h.p. G with starting set A is globally (C, η, δ)-explosive for the C'-core of G. In particular, by Theorem 1 for every $\varepsilon > 0$ w.h.p. the fraction of active vertices in the C'-core is $\alpha \pm \varepsilon$ for sufficiently large C', where α is given by Equation* (1).

It should not be surprising to see that for normalization in power-law graphs we need to restrict ourselves to large degree vertices. Most generative models for power-law graphs (including the Chung-Lu model) contain linearly many vertices with degrees strictly less than k. Typically, there are even linearly many isolated vertices. It is clear that none of these vertices can be activated by a bootstrap percolation process unless they are in the initial active set. In Section 5, we develop a heuristic estimate on the fraction of low degree vertices that finally become active. The fraction of active vertices can be as small as 0 (for vertices with degree less than k), and approaches α as the degree grows. Note, however, that the exact fraction for low degree vertices that become active depends on the degree distribution of the graph, whereas for high degree vertices this fraction is a universal constant that is independent of the graph structure.

4 Proofs

This section contains the proofs of Theorem 1 and Theorem 2. Due to space restrictions, we only give rough sketches. Full proofs can be found in the appendix. The proof of Theorem 1 takes up ideas from [12], where it was proven that bootstrap percolation with inhibition is normalizing on Erdős-Rényi graphs.

Proof of Theorem 1

We start with some basic facts about the percolation process. Fix some vertex $v \in V$, and recall that $\Gamma_s^+(v)$ and $\Gamma_s^-(v)$ are the number of excitatory and inhibitory vertices among the first s active neighbors of v, respectively. When the process starts, we do not need to decide right away for the signs of all the vertices in V. Rather, we can postpone the decision until a vertex becomes active. So whenever a neighbor of v turns active, it flips a coin to decide on its sign, and this coin flip is independent of any other coin flips. Hence, the number of inhibitory vertices among its first s active neighbors is binomially distributed $\mathrm{Bin}(s, \tau)$. By the Chernoff bounds, if s is sufficiently large then $\Gamma_s^+(v)$ and $\Gamma_s^-(v)$ are concentrated around their expectations $(1-\tau)s$ and τs, respectively.

We will link the probability that a vertex $v \in S$ becomes active with a random walk on $\mathbb{Z}$, using the following fact about random walks (see [14], Problem 5.3.1.).

Lemma 1. *Let $X_1, X_2, \dots, X_n$ be a sequence of independent random variables, each of which is equal to* 1 *with probability* $p \in [0,1]$ *and* -1 *otherwise. Consider the* biased random walk $Z_i = X_1 + X_2 + \cdots + X_i$. *Then there exists for every* $\varepsilon > 0$ *and* $k \in \mathbf{N}$ *a constant* $C_0 = C_0(\varepsilon, k)$ *such that the following is true:*

$$\Pr[\exists i \le C_0 \text{ s.t. } Z_i = k] \in (1 \pm \varepsilon) \cdot \min\left\{1, p^k/(1-p)^k\right\}.$$

Recall that $X_i(v)$ is 1 if the i-th signal arriving in v is excitatory, and -1 otherwise, and let $Z_i(v) := X_1(v) + X_2(v) + \cdots + X_i(v)$. We know that the vertex v becomes active with the arrival of the first signal that causes $Z_i(v)$ to become k, if such a signal exists. We will show that $Z_i(v)$ follows (essentially) a one-dimensional random walk with bias τ.

There are two problems which complicate the analysis: the first being that the processes $(Z_i(v))_{i\in\mathbb{N}}$ and $(Z_i(u))_{i\in\mathbb{N}}$ are not independent for different vertices u and v, and the second being that for a fixed vertex v, the variables $X_i(v)$ and $X_j(v)$ are not independent for $i \neq j$, meaning that $(Z_i(v))_{i\in\mathbb{N}}$ is not a true random walk.

We overcome these problems as follows. Fix some large constant $\tilde{C} > 0$. Since the process is explosive, for the typical vertex v there exists a time at which v has at least $\tilde{C}$ active neighbors, but has not yet received any signals (if all $\tilde{C}$ neighbors become active in a sufficiently small time interval, then the signals are very unlikely to arrive within this interval). If $\tilde{C}$ is sufficiently large, then the fraction of positive signals among all signals on their way will be roughly $1-\tau$. Since the transmission delays are distributed with an exponential distribution, which is *memoryless*, all the signals on their way are equally likely to arrive first. In particular, $X_1(v)$ is positive with probability roughly $1-\tau$, and this holds *independent* of the sign of the first incoming signal of other vertices. After the first signal has arrived, the fraction of positive signals on their way will still be roughly $1-\tau$, since removing a single signal has only a very small impact on this fraction. Thus $X_2(v)$ is also positive with probability $\approx 1-\tau$, and the same holds for the first few incoming signals. Therefore, for small i the random variable $Z_i(v)$ resembles a one-dimensional random walk as in Lemma 1. As i grows larger, $Z_i(v)$ fails to follow a random walk, but then $Z_i(v)$ is typically already very negative. Thus we can show directly that most likely $Z_i(v)$ never becomes positive again. The details are rather involved, and we omit them due to space restrictions.

We remark that if we would assign positive or negative labels to the *edges* instead of the vertices, then similar, but substantially simpler arguments apply. E.g., the one-dimensional random walks for different vertices are independent of each other, so it is not necessary to condition on the history of the process.

4.1 Proof of Theorem 2

In this section we prove Theorem 2, which states that bootstrap percolation on a power law random graph G is locally (C, η, δ)-explosive for the C'-core of G. We remark that the condition $\mathfrak{a} \gg n^{(\beta-2)/(\beta-1)}$ is necessary since the threshold for bootstrap percolation without inhibition in Chung-Lu graphs is $n^{(\beta-2)/(\beta-1)}$, see [2, Theorem 2.3]. We want to apply Theorem 1. The main idea to see that bootstrap percolation is explosive is to observe that the C-core contains an Erdős-Rényi random graph $G(n', p)$, where n' and p depend on C. In order to prove Theorem 2, we will use the following statement about bootstrap percolation in Erdős-Rényi random graphs $G(n, p)$ with n vertices, where each edge is present independently of each other with probability p. For convenience, let

$$\Lambda := \left(\frac{(k-1)!}{(1-\tau)^k np^k}\right)^{1/(k-1)}. \tag{2}$$

In [12, Theorem 2] it was shown that the threshold for bootstrap percolation is $(1-1/k)\Lambda$, i.e., w.h.p. a random set of size $a_0 = (1+\varepsilon)(1-1/k)\Lambda$ will activate almost all of the graph, while for a random set of size $a_0 = (1-\varepsilon)(1-1/k)\Lambda$ the bootstrap percolation process dies out with $a^* \leq 2a_0$. Moreover, it was shown that the bootstrap percolation process on $G(n, p)$ is globally explosive and normalizing. More precisely, we have the following.

Theorem 3. *For every $\varepsilon > 0$ there is a constant $D > 0$ such that the following holds. Assume $D/n \leq p \ll n^{-1/k}$. Let G contain an Erdős-Rényi random graph $G(n, p)$. Let $x \geq D$, and let A be a random set of size $|A| = x\Lambda$. Then with probability at least $1 - O(1/\log x)$ the bootstrap percolation process on G with starting set A activates at least $(\alpha - \varepsilon)n$ vertices in time $O(x^{-1/(2k)} + (pn)^{-1})$.*

Apart from some technical differences, Theorem 3 differs from the statement proven in [12] in an important aspect: there the statement was only proven for the case $G = G(n, p)$, while we only require the graph to *contain* a $G(n, p)$. Note that this a non-trivial extension since additional edges may obstruct percolation due to inhibition. For this reason the proof becomes more subtle at some points even though the main line of argument remains similar to the proof in [12].

We are now ready to prove Theorem 2. The key observation is that the vertices of weight at least w induce a graph that contains an Erdős-Rényi graph as a subgraph, and thus by Theorem 3 percolation is explosive. However, we cannot immediately use $w = C'$ since then the size of the initial set would be below threshold, so we need to iterate the argument for several values of w. For convenience, let the *w-weight core* $G_{\geq w}$ of a Chung-Lu power law graph be the subgraph induced by all vertices of weight at least w. Since the weight corresponds to the expected degree, this notion is closely related to the w-core.

Proof (of Theorem 2). Let $C, \eta > 0$. Further let $\varepsilon > 0$, and let $D > 0$ be the constant given by Theorem 3. For sake of exposition, we assume that the i-th weight is given by $(n/i)^{1/(\beta-1)}$, i.e., that the average degree is $\bar{d} = (\beta-1)/(\beta-2)$. Other values of $\bar{d}$ will change the calculations below only by constant factors. Since $2 < \beta < 3$, we may choose $\gamma > 0$ such that $(\beta-1)/2 < \gamma < 1$. Let $\gamma' := (\beta-1)(k-1)/(2k+1-\beta)$. It is easy to check that $0 < \gamma' < (\beta-1)/2 < \gamma < 1$. For all $i \in \mathbb{N}$, let $w_i := n^{\gamma^i/(\beta-1)}$. Then the w_i-weight core G_i has size n_i, where n_i is given by the equation $(n/n_i)^{1/(\beta-1)} = w_i$. We easily deduce $n_i = n^{1-\gamma^i}$.

Any two vertices in G_i have weight at least w_i, so they are connected with probability at least $p_i := \min\{1, w_i^2/n\} = \min\{1, n^{-1+2\gamma^i/(\beta-1)}\}$.

Let θ_i be the threshold for percolation in an Erdős-Rényi random graph $G(n_i, p_i)$. By Theorem 3 we know that there is an absolute constant $\tilde{C}$ (depending only on k and τ) such that $\theta_i = \tilde{C}^{-1} \cdot (n_i p_i^k)^{-1/(k-1)} = \tilde{C}^{-1} \cdot (n^{1-\gamma^i/\gamma'})$. Therefore, $n_{i-1}/\theta_i = \tilde{C} \cdot ((n/n_{i-1})^{(-1+\gamma/\gamma')})$ for all i. In particular, for $\tilde{D} := ((\alpha-\varepsilon)\tilde{C}/D)^{\gamma'/(\gamma-\gamma')}$ note that $(\alpha-\varepsilon)n_{i-1} > D\theta_i$ whenever $n_{i-1} \leq \tilde{D}n$. Let i_0 be the some index such that $n_{i_0-1} \leq \tilde{D}n$. Note that equivalently i_0 satisfies $n^{-\gamma^{i_0-1}} < \tilde{D}$.

We show inductively that we have explosive percolation on G_i for all $1 \leq i \leq i_0$. For $i = 1$, the number of vertices in G_1 that are initially active is $n^{-1+(\beta-2)(\beta-1)}n_1 = n^{\gamma-(\beta-1)/2} = n^{\Omega(1)} = n_1^{\Omega(1)}$. Any two vertices in G_1 are connected with probability $p_1 = 1$. We apply Theorem 1 (with some probability $p_1' = n_1^{-1/k-\varepsilon} < p_1$ to ensure the condition $p_1' \ll n_1^{-1/k}$), and obtain that an $(\alpha-\varepsilon)$-fraction of G_1 is activated after time $o(1)$. For convenience, set the time $t = 0$ to be the time *after* this phase, since we may safely ignore an additional time of $o(1)$.

For the inductive step, assume that an $(\alpha-\varepsilon)$-fraction of G_{i-1} is activated at some time $t_{i-1} < \eta$. Obviously G_{i-1} is a subgraph of G_i, so there are at least $(\alpha-\varepsilon)n_{i-1}$ active vertices in G_i. Since $(\alpha-\varepsilon)n_{i-1} > D\theta_i$ for $i-1 < i_0$, we may apply Theorem 3 to deduce that bootstrap percolation activates with probability at least $1-q_i$ at least an $(\alpha-\varepsilon)$-fraction of G_{i-1} until time $t_{i-1}+\Delta_i$, where $q_i = O(1/\log(n_{i-1}/\theta_i))$ and $\Delta_i = O((n_{i-1}/\theta_i)^{-1/(2k)} + (p_i n_i)^{-1})$. Since $n_{i-1}/\theta_i = \Theta(n^{\gamma^i(\gamma-\gamma')/\gamma'})$ and $p_i n_i = \Omega(n^{\gamma^i(3-\beta)/(\beta-1)})$, let $c := \min\{(\gamma-\gamma')/(2k\gamma'), (3-\beta)/(\beta-1)\} > 0$. Then

$$t_{i_0} = \sum_{i=1}^{i_0} \Delta_i = O\left(\sum_{i=1}^{i_0} n^{-c\gamma^i}\right) = O\left(n^{-c\gamma^{i_0}}\right) = O((n_{i_0}/n)^c).$$

In particular, by choosing i_0 such that n_{i_0}/n is a sufficiently small constant, we can achieve $t_{i_0} \leq \eta$. Similarly, since $q_i = O((\gamma^i \log n)^{-1})$, the accumulated error probability is $\sum_{i=1}^{i_0} q_i = O(q_{i_0}) = O(1/\log(n/n_{i_0}))$, and again by choosing i_0 such that n_{i_0}/n is a sufficiently small constant, we can achieve $\sum_{i=1}^{i_0} q_i \leq \delta$.

This proves that with probability at least $1-\delta$ bootstrap percolation activates an $(\alpha-\varepsilon)$-fraction of the w_{i_0}-weight core of G. Finally observe that w_{i_0} is a constant if n_{i_0}/n is constant. Now choose $C' := c\max\{2C/(\alpha-\varepsilon), w_{i_0}\}$ for some $c > 1$. Observe that for suitable c at least a $(1-\eta)$-fraction of the C'-core is also in the w_{i_0}-weight core, and each of those vertices has at time η an expected number of at least $(1-\eta)C'(\alpha-\varepsilon) \geq cC$ active neighbors in the w_{i_0}-weight core. Thus the bootstrap percolation process is (C, η, δ)-explosive for the C'-core.

5 Experiments

Note that Theorem 2 explains what happens to the set of high degree vertices. However, many interesting graphs have a quite large number of low-degree vertices. To this end, we first develop a heuristic estimate for what happens in the remainder of the graph.

5.1 Heuristic Estimation

The proof of Theorem 1 shows that if a process is explosive for a set S then the vertices in S follow closely a random walk with bias $1-\tau$ (cf. Lemma 1). For our estimation we will assume that this is in fact true for all vertices in G. So consider a random walk with bias $1-\tau$. Let $p_{i,k,j}$ denote the probability of reaching k within i steps if we start from j. Clearly we have $p_{0,k,k}=1$ and $p_{0,k,j}=0$ for all $j<k$. For all $i \geq 1$ we can recursively compute $p_{i,k,j}$ as follows:

$$p_{i,k,j} = \begin{cases} 1; & j=k \\ (1-\tau)\cdot p_{i-1,k,j+1} + \tau \cdot p_{i-1,k,j-1} & \text{otherwise} \end{cases}$$

Assume we know that (after percolation has terminated) a $\hat{p}$-fraction of all edges starts at an active vertex. Then we can assume that the probability that a vertex in $V\setminus A$ with degree i has been activated is given by $\sum_{j=k}^{i} \Pr[\text{Bin}(i,\hat{p})=j]\cdot p_{j,k,0}$. Using this idea we can now set up an approximation for a given graph $G=(V,E)$ with n vertices and m edges. Let n_i denote the number of vertices with degree i, for $0 \leq i \leq n$, and let $p_{\text{boot}} = a_0/n$ be the probability that a vertex belongs to the initial active set A. Then $\hat{p}$ satisfies the following equation:

$$\hat{p}\cdot m = \sum_{i=0}^{n}\left((p_{\text{boot}} i n_i) + (1-p_{\text{boot}})\cdot \sum_{j=k}^{i} \Pr[\text{Bin}(i,\hat{p})=j]\cdot p_{j,k,0}\cdot i \cdot n_i\right). \quad (3)$$

By Equation (3), we can numerically compute $\hat{p}$. Afterwards, we can estimate the number a^* of vertices that will turn active:

$$a^* \approx \sum_{i=0}^{n}\left(p_{\text{boot}}\cdot n_i + (1-p_{\text{boot}})\cdot \sum_{j=k}^{i} \Pr[\text{Bin}(i,\hat{p})=j]\cdot p_{j,k,0}\cdot n_i\right). \quad (4)$$

Note that the i-th summand in (4) predicts the number of vertices of degree i, so we also obtain an estimate for the number of active vertices of a given degree.

5.2 Simulations

To test Equation (4), we compare it with simulations. We use the Epinion social network [19] that describes the trust relationship between its members. The size of the network is 75879, but we only consider its largest connected component with size $n = 75877$. We use $k = 4$ as activation parameter, and we start with a random active set A of size $a_0 = 2000$. In all the experiments, we simulate the asynchronous percolation process for different values of τ ranging from 0.1 to 0.9 with the step size of 0.1. For each value of τ we run the diffusion process 20 times and report the average size of the final active vertices along with standard deviations. We plot the fraction of active vertices within $V \setminus A$ (or within $\{v \in V \setminus A \mid \deg(v) \geq 16\}$ for Figures 1a and 1b).

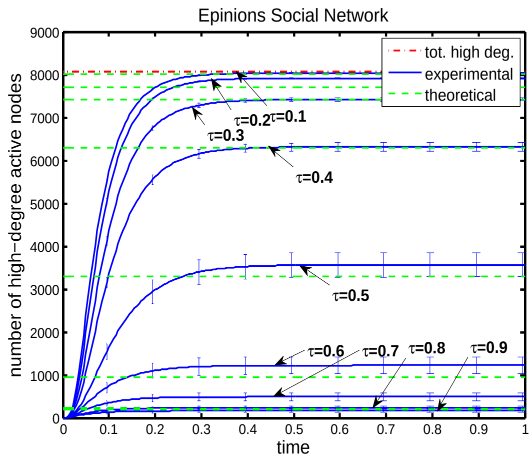

(a) Percolation process (high-degree vertices)

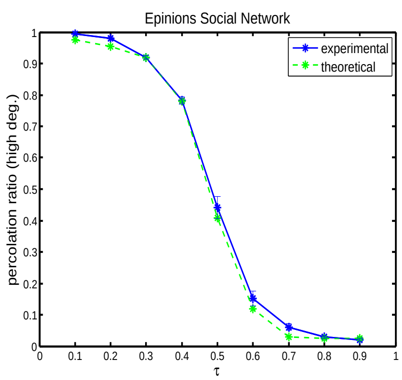

(b) Percolation ratio (high-degree)

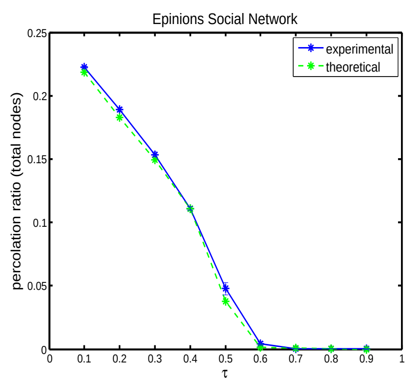

(c) Percolation ratio (all).

Fig. 1: Asynchronous bootstrap percolation in the Epinions social network. *(a)* shows the evolution of the process for different values of τ among high degree vertices, i.e., those with degrees at least 16. *Red dashed line:* number of high-degree vertices in $V \setminus A$; *blue lines:* growth of number of active high-degree vertices in $V \setminus A$ over time; *green dashed lines:* prediction from Theorem 1 (Equation (1)) for the C'-core if C' is sufficiently large. *(b)* shows the active fraction (blue lines) among all high-degree vertices in $V \setminus A$ after termination, and the prediction from Equation (1) (green lines), cf. above. *(c)* shows the corresponding fraction among all vertices in $V \setminus A$, and the prediction from equation (4).

6 Acknowledgements

We thank Qi Zhang for her help in preparing Figure 1.

References

1. Adamic, L. A., and Huberman, B. A. Power-law distribution of the world wide web. *Science 287*, 5461 (2000), 2115–2115.
2. Amini, H., and Fountoulakis, N. Bootstrap Percolation in Power-Law Random Graphs. *Journal of Statistical Physics 155*, 1 (2014), 72–92.
3. Balogh, J., Bollobás, B., Duminil-Copin, H., and Morris, R. The sharp threshold for bootstrap percolation in all dimensions. *Transactions of the American Mathematical Society 364*, 5 (2012), 2667–2701.
4. Balogh, J., Bollobás, B., and Morris, R. Bootstrap percolation in three dimensions. *The Annals of Probability* (2009), 1329–1380.
5. Balogh, J., Peres, Y., and Pete, G. Bootstrap percolation on infinite trees and non-amenable groups. *Combinatorics, Probability & Computing 15*, 5 (2006), 715–730.
6. Balogh, J., and Pittel, B. G. Bootstrap percolation on the random regular graph. *Random Structures & Algorithms 30*, 1-2 (2007), 257–286.
7. Barabási, A.-L., and Albert, R. Emergence of scaling in random networks. *Science 286*, 5439 (1999), 509–512.
8. Breskin, I., Soriano, J., Moses, E., and Tlusty, T. Percolation in living neural networks. *Physical review letters* (2006).
9. Chalupa, J., Leath, P., and Reich, G. Bootstrap percolation on a bethe lattice. *Journal of Physics C: Solid State Physics 12*, 1 (1979), L31.
10. Chung, F., and Lu, L. The average distances in random graphs with given expected degrees. *PNAS 99*, 25 (2002), 15879–15882.
11. Dorogovtsev, S. N., Goltsev, A. V., and Mendes, J. F. Critical phenomena in complex networks. *Reviews of Modern Physics 80*, 4 (2008), 1275.
12. Einarsson, H., Lengler, J., Mousset, F., Panagiotou, K., and Steger, A. Bootstrap percolation with inhibition. arXiv:1410.3291.
13. Faloutsos, M., Faloutsos, P., and Faloutsos, C. On power-law relationships of the internet topology. *ACM SIGCOMM Computer Communication Review 29*, 4 (1999), 251–262.
14. Grimmett, G., and Stirzaker, D. *One Thousand Exercises in Probability.* OUP Oxford, 2001.
15. Holroyd, A. E. Sharp metastability threshold for two-dimensional bootstrap percolation. *Probability Theory and Related Fields 125* (2003), 195–224.
16. Kempe, D., Kleinberg, J., and Tardos, É. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2003), ACM, pp. 137–146.
17. Lelarge, M. Diffusion and cascading behavior in random networks. *Games and Economic Behavior 75*, 2 (2012), 752–775.
18. Moore, C., and Newman, M. E. Epidemics and percolation in small-world networks. *Physical Review E 61*, 5 (2000), 5678.
19. Richardson, M., Agrawal, R., and Domingos, P. Trust management for the semantic web. In *The Semantic Web-ISWC 2003*. 2003, pp. 351–368.
20. Toninelli, C., Biroli, G., and Fisher, D. S. Jamming percolation and glass transitions in lattice models. *Physical Review Letters 96*, 3 (2006), 035702.
21. Ugander, J., Karrer, B., Backstrom, L., and Marlow, C. The anatomy of the facebook social graph. *arXiv:1111.4503* (2011).