

How to distribute antidote to control epidemics

Christian Borgs* Jennifer Chayes* Ayalvadi Ganesh† Amin Saberi‡

1 Introduction

Motivated by the control of epidemics on technological and social networks, we give a rigorous analysis of a variant of the contact process on a finite graph in which the cure rate is allowed to vary from one vertex to the next, and even to depend on the current state of the system. We think of the cure rate at each vertex as representing the amount of antidote available to that vertex. In order to model a system with a fixed amount of antidote, we limit the sum of the cure rates over all vertices, so that the sum scales at most linearly with the number of vertices in the network.

In the standard contact process, the rate at which an infected vertex becomes healthy is the same for all vertices. In other words, it is implicitly assumed that all the vertices of the graph are monitored uniformly or receive the same dose of antidote.

It is known that, on the high-degree graphs typical of real-world networks, the standard contact process is very susceptible to epidemics. In the case of a star graph, it has been proved [7, 3] that if the cure rate is constant from one vertex to the next, then no matter how small the rate of infection transmission between neighbors, the infection has a positive probability of surviving for a time super-polynomial in the number of vertices. This result was used to show that the epidemic threshold of the standard contact process goes to zero with the number of vertices, both on preferential attachment graphs [3] and on the related power-law configuration model graphs [7], thereby rigorously establishing a theoretical physics result of [17].

In this paper, we allow the cure rate to be non-uniform, i.e., the amount of antidote to vary from one vertex to the next. The questions we would like to answer in this context are: What is the efficacy of standard epidemiological protocols for non-uniform distribution of antidote? Are there better ways of distributing antidote non-uniformly? In particular, is it possible to distribute antidote in such a way that the epidemic threshold does not go to zero with the number of vertices?

To answer these questions, we consider two non-trivial methods for distributing antidote non-uniformly among the vertices. First, we study the widely-used contact tracing method

*Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA. Email: {borgs, jchayes}@microsoft.com.

†Department of Mathematics, University of Bristol, Bristol BS8 1TW, UK. Email: a.ganesh@bristol.ac.uk

‡Department of Management Science and Engineering, Institute for Computational and Mathematical Engineering, Stanford University. Email: saberi@stanford.edu.

from epidemiology: we augment the cure rates of all neighbors of an infected node. We prove that *contact tracing is insufficient to contain epidemics* in the sense that, on a star graph, the epidemic threshold still goes to zero no matter how much we enhance the cure rate of neighbors of infected vertices, provided that the overall cure rate scales like the number of vertices. In other words, using contact tracing to control epidemics would require a total amount of antidote which is super-linear in the number of vertices. We also give quantitative estimates on how quickly the threshold goes to zero, or alternatively on the scaling of the amount of antidote with the number of vertices. This still leaves open the possibility that contact tracing could be effective if applied at a sufficiently early stage of an infection. We also answer this question, in the sense that we estimate the maximum size of an infection at which it becomes impossible to control via contact tracing.

Next, we propose an *a priori* simpler method for distributing antidote: we augment the cure rate proportional to the degree of the vertex — independent of the current state of the model — again keeping the total antidote proportional to the number of vertices. We show that, *on a general graph with bounded average degree, a curing mechanism proportional to degree controls epidemics*, in the sense that the epidemic threshold remains bounded below by a strictly positive number, independent of the number of vertices. In other words, on a general graph, curing proportional to degree requires an amount of antidote that scales only linearly with the number of vertices, provided the average degree is bounded. For graphs in which the average degree is much smaller than the maximum degree, this could result in a dramatic shift in behavior.

Finally, we show that for *expander graphs*, the total amount of antidote used by the above mechanism (i.e., curing proportional to degree) cannot be reduced by more than a constant factor. It has been shown [15, 8] that the standard power-law graphs have constant conductance and eigenvalue gap, so this lower bound applies to these graphs.

The organization of this paper is as follows. In the next section, we give detailed motivation for studying this problem in the context of technological and epidemiological networks, and we review previous results on the contact process on various graphs. In Section 3, we state our main results precisely, in particular giving the scaling of our results with the number of vertices of the graph. In Section 4, we provide the proofs of our results.

2 Motivation, Models, and Previous Results

There has recently been much interest in the spread of viruses and worms on general networks — especially so-called preferential attachment networks used to model both technological networks, such as the Internet, the World Wide Web and Instant Messaging Networks, and social networks. Assuming that the virus does not mutate, an appropriate model is the so-called *susceptible-infected-recovered (SIR) model*, in which each vertex is in one of three states — depending on whether it is healthy (but susceptible), infected and actively infecting neighboring vertices, or cured (and no longer susceptible).

If the virus mutates, it is much more dangerous from an epidemiological perspective. In this case, even if the vertex is cured, and therefore not susceptible to the original virus, it is still susceptible to a mutated virus. In the case of human disease, the occurrence of

mutating viruses is well-documented; rapidly mutating viruses are very difficult to cure permanently, and tend to lead to epidemics.

In the case of technological networks, such effects have been observed with rapidly mutating worms. In contrast to a computer virus which spreads when the user executes a certain program, a so-called worm can infect a computer without the user taking any action. Instead, the worm exploits a vulnerability unintentionally left in the operating system; it enters the computer the moment the machine is connected to the network. Worms spread either by infecting vulnerable randomly generated IP addresses, or by infecting vulnerable addresses on the contact lists of infected computers. When a computer is patched, it is no longer vulnerable to the given worm.

A mutating worm has a list, often a long list, of vulnerabilities. It uses the first vulnerability to infect a set of machines. Then at a random time, or at a specific pre-arranged time, or in response to a message from a central authority, it begins to use the second vulnerability on its list to infect other computers on the network. Computers which have been patched against the first vulnerability are still susceptible to the second. Especially if the worm has a long list of vulnerabilities, and if the time scale of changing the vulnerability is shorter than the typical patching time, these worms can lead to very serious epidemics.

If the virus or worm mutates very rapidly through many states, an approximate model for the system is the so-called *susceptible-infected-susceptible (SIS) model*, also known as the *contact process*. Here each vertex is in one of two states — healthy (but susceptible) or infected. State transitions at vertices occur independently as follows: An infected vertex becomes healthy with rate ρ , independently of the status of its neighbors. A healthy vertex becomes infected at a rate equal to the propagation rate of the disease, β , times the number of its infected neighbors. The relevant parameter is $\lambda = \beta/\rho$.

The contact process has been studied extensively in the probability community [12], where it is usually studied on bounded-degree graphs. The most important general result in that context is the existence of epidemic thresholds. For infinite graphs it has been shown that there exist two epidemic thresholds $\lambda_1 \leq \lambda_2$. If the infection ratio $\lambda = \beta/\rho$ is larger than λ_2 , then with positive probability the epidemic can spread and survive at any point of the graph. If $\lambda_1 < \lambda < \lambda_2$, the epidemic survives with positive probability, but every vertex almost surely eventually heals without being reinfected. If $\lambda < \lambda_1$, the epidemic dies out almost surely. As it turns out, $\lambda_1 = \lambda_2$ for \mathbb{Z}^d , while $\lambda_1 < \lambda_2$ for regular trees (see [12] and [20, 21]).

It is easy to see that, in a finite graph, the infection will eventually die out with probability 1. However, there is still a natural definition of epidemics in the finite case, as can be seen by considering finite subsets of well-studied infinite graphs, such as \mathbb{Z}^d . It turns out that, for the cube $[-\ell, \ell]^d$, there is a λ_c such that if $\lambda > \lambda_c$ then with probability bounded away from zero the infection survival time is exponential in ℓ^d , while if $\lambda < \lambda_c$ the infection dies out before time $\log(\ell)$ with probability $1 - o(1)$. Moreover, this λ_c is equal to the epidemic threshold for \mathbb{Z}^d . (See [12] for proofs of these statements.) Therefore, it is natural to say that the infection becomes an epidemic if the time that it takes for the infection to die out is super-polynomial in the number of vertices of the graph.

In spite of the fact that most of the rigorous work on the contact process has concerned regular lattices or mean-field models, many epidemiologists have observed that the struc-

ture of contacts in populations has important implications for the propagation of infection. Empirical studies have shown the relevance of network structure for epidemiological networks in the cases of human contact networks [13, 14, 17], animal contact networks [5] and computer networks [16, 22, 10]. For other recent works in mathematical epidemiology on models for controlling infections see [6, 11, 9].

For many real-world networks, the underlying graph G is a power-law graph, i.e., a graph whose degree distribution follows a power law [1, 19, 2, 4], such as the preferential attachment graph. The contact process on power-law graphs was first analyzed by Pastor-Satorras and Vespignani [17, 18]. Using simulation and mean-field approximations, they argued that the epidemic threshold λ_c in power-law networks should tend to zero as the number of vertices tends to infinity. This behavior is not hard to understand: the maximal degree of a power-law graph grows like a power of the number of vertices n (for the original Barabasi-Albert power-law graph, this power is $n^{1/2}$), so these graphs contain high-degree stars as subgraphs. Once the infection reaches such a star, it will survive for a long time, unless β/ρ is smaller than the critical value for this star. For preferential attachment models, this behavior was rigorously established in [3], and for the configuration model it was first established in [7]. In both cases, it has been shown that λ_c decays like a power in n (this is explicit in [7], and implicit in [3]). Formulated differently, these results say that the amount of antidote needed to cure an infection on power-law graphs grows faster than the number of vertices, making it very hard to control the infection on large graphs.

A slightly different aspect on the spread of viruses was studied by Wagner and Anantharam [24]. They raised the question whether an inhomogeneous spread rate β_{xy} for the different edges $\{x, y\}$ of the underlying graph might be more efficient for the virus. Specifically, they analyzed the contact process on the real line, comparing a constant infection rate β to a varying infection rate β_{xy} , normalized in such a way that the average infection rate is the same. They showed that no matter how low the original rate β is, it is possible to redistribute it in such a way that the infection becomes supercritical.

3 Our Model and Statement of Results

Let $G = (V, E)$ be an undirected graph with vertex set V and edge set E . The contact process is a continuous-time Markov process $X(t)$ on $\{0, 1\}^V$, where $X_i(t)$ is the indicator that node i is infected at time t . State transitions at nodes occur independently as follows: each infected node becomes healthy at some constant rate ρ , while each healthy node becomes infected at rate β times the number of its infected neighbors. Note that a node becomes susceptible again as soon as it is cured.

In our model, we will assume that, as in the contact process, the virus is spreading at a constant infection rate β . However, we assume that we can distribute different amounts of antidote between different nodes. The rate at which an infected node becomes healthy is proportional to the amount of antidote it receives.

We will assume that there is a fixed amount of antidote $R = \sum_x \rho_x$ available which can be distributed non-uniformly between the nodes, even depending on their current state of the infection. These are some of the questions that we would like to answer in this context:

what is the best policy for distributing the antidote? Is there a way to reduce the amount of antidote needed to control an infection on a power law graph or star to a constant times the number of nodes?

An obvious choice, employed in disease control, is to inoculate neighbors (or contacts) of infected individuals. In the contact process, this is modeled by augmenting the cure rate ρ_x for all neighbors of an infected vertex. This amounts to setting

$$\rho_x = \rho + \rho' d_x^*(t), \quad (1)$$

for specified constants ρ and ρ' , where $d_x^*(t)$ is the number of infected neighbors of x at time t .

To evaluate the effectiveness of this method, let us consider the star with n leaves, which has a critical infection rate $\beta_c = \rho n^{-1/2+o(1)}$ under uniform inoculation rate ρ . The next theorem shows that inoculating according to (1) is not enough to result in a positive β_c . Instead, we get that $\beta_c = O(n^{-1/3+o(1)})$. Formulated differently, the amount of antidote needed to cure an infection with constant infection rate β still grows super-linearly in the number of nodes, at least like $n^{4/3-o(1)}$.

Theorem 1. *Given $\rho > 0$, $\rho' > 0$ and $\varepsilon > 0$ there are constants $c > 0$, $C < \infty$ and $n_0 < \infty$ such that the following holds for the star network with one central hub and $n \geq n_0$ leaves.*

(i) *For $n^{-1/3+\varepsilon}\rho \leq \beta \leq \rho$ and arbitrary initial conditions, with probability $1 - O(n^{-c})$ the survival time of the epidemic is either smaller than $\tau_1 = C\rho^{-1} \log n$ or larger than $\tau_2 = \rho^{-1} e^{c\beta^3 n / \rho^3}$.*

(ii) *Let β be as above. If the epidemic starts with an infected center, or with a healthy center and at least $(\rho/\beta) \log n$ infected leaves, then with probability $1 - O(n^{-c})$, the survival time is at least τ_2 , and the number of nodes touched by the infection is at least n^c .*

(iii) *Let β be as above, and let $(1 + \log k)/\rho < \tau < \tau_2$. If the epidemic starts with a healthy center and k infected leaves, then the probability that the survival time is smaller than τ lies between*

$$\left(1 - \frac{\beta}{\beta + \rho}\right)^k \left(1 - \frac{1 + \log k}{\rho\tau}\right) \quad \text{and} \quad \left(1 - \frac{\beta}{\beta + \rho}\right)^k + O(n^{-c}).$$

(iv) *If $\beta = \rho n^{-1/3-\epsilon}$, then the epidemic dies out in expected time $O(\log n)$, starting from any initial condition.*

For the star, contact tracing is therefore only effective as long as the center is not infected and the number of infected leaves is $O(\rho/\beta)$. After that point, with high probability, it can't prevent an epidemic with very long survival time, even if the amount of antidote grows like $\beta n^{4/3-o(1)}$.

Is there a better way to distribute the antidote? The answer is yes. This is the result of our main theorem. To formulate it, we define τ to be the expected survival time, starting from all vertices infected. We also denote the degree of a vertex x in the underlying graph G by d_x .

Theorem 2. *Let G be an arbitrary graph on n nodes, and let $\rho_x = d_x$. If $\beta < 1$, then $\tau = O(\log n)$.*

For graph with bounded average degree \bar{d} (like the power law graphs used to describe the WWW), this implies that the total amount of antidote needed is just $\beta\bar{d}n$, corresponding to a finite λ_c . (For varying ρ , we define λ as the ratio of β and the average amount of antidote, $\frac{1}{n}\sum_x \rho_x$.)

This raises the question of whether we could do significantly better, so that $\lambda_c \rightarrow \infty$ as $n \rightarrow \infty$. The answer is no, even if we allow adapting the curing rate ρ_x to the current state of the infection. This is the content of our next theorem.

Recall that a graph $G = (V, E)$ is an (α, η) -expander, if for each set of nodes $W \subset V$ with $|W| \leq \alpha|V|$, the number of edges joining W to its complement $V \setminus W$ is at least $\eta|W|$. We use X_t to denote the set of infected vertices at time t , and allow for an arbitrary allocation $\rho_x = \rho_x(X_t, t)$ depending on X_t and maybe on the time t as well. For comparison purposes, we impose the normalization condition

$$\sum_x \rho_x(W, t) \leq \sum_x d_x \quad \text{for all } W \subset V \text{ and } t. \quad (2)$$

Theorem 3. *Let $\epsilon > 0$ and let G_n be a sequence of (α, η) -expanders on n nodes. Let $\rho_x(X_t, t)$ be an arbitrary distribution obeying the constraint (2). If $\beta \geq \frac{(1+\epsilon)\bar{d}}{\alpha\eta}$, then $\tau \geq e^{\Theta(n \log n)}$.*

The theorem implies that for expanders with bounded average degree, inoculating according to degrees is a constant factor competitive inoculation scheme. We do not expect Theorem 3 to hold for general graphs. Indeed, we expect that there are graphs for which other inoculation schemes may out-perform the proportional-to-degree inoculation scheme by more than a constant factor. For example, if a graph contains densely connected clusters that are loosely coupled to one another, it may be better to “quarantine” the clusters from one another.

Perhaps the right qualitative interpretation of our work is that contact tracing is effective only when the infection has not reached high degree nodes. It seems that a more effective strategy is to use “intelligent tracing” [23] which uses a global knowledge of the topology. We leave an the question of deriving or approximating the optimal adaptive strategy for controlling an infection spreading in a network as an important open problem.

4 Proofs

4.1 Proof of Theorem 2

Without loss of generality, we may assume that G is connected (otherwise, we apply the arguments below separately to every component), so that $d_x \geq 1$ for all x . Each infected vertex x will infect each of its uninfected neighbors with rate β , but recovers with rate $\rho_x = d_x$. Let $N_t = |X_t|$. Then

$$\frac{\mathbb{E}[N_{t+dt} | X_t] - N_t}{dt} \leq \sum_{x \in X_t} d_x(\beta - 1) \leq N_t(\beta - 1),$$

implying that $d\mathbb{E}[N_t]/dt \leq -(1 - \beta)\mathbb{E}[N_t]$ and hence

$$\mathbb{E}[N_t] \leq N_0 e^{-(1-\beta)t} \leq n e^{-(1-\beta)t}.$$

Thus, with high probability, the epidemic dies out in time $O(\log n)$, as claimed.

4.2 Proof of Theorem 3

Choose $\tilde{\varepsilon}$ in such a way that $(1 - \tilde{\varepsilon}) = (1 - \varepsilon)(1 + \tilde{\varepsilon})$, let ∂X denote the edge boundary of X , and as before, let $N_t = |X_t|$. For N_t in the interval $[(1 - \tilde{\varepsilon})\alpha n, \alpha n]$ we then have

$$\frac{\mathbb{E}[N_{t+dt} | X_t] - N_t}{dt} = \beta |\partial X_t| - \sum_{x \in X_t} \rho_x \geq (\beta \eta (1 - \tilde{\varepsilon}) \alpha - \bar{d}) n \geq \tilde{\varepsilon} n. \quad (3)$$

Here the first inequality comes from the fact that for $(1 - \tilde{\varepsilon})\alpha n \leq |X_t| \leq \alpha n$ we have $|\partial X_t| \geq \eta |X_t| \geq \eta (1 - \tilde{\varepsilon}) \alpha n$, while the second follows from our definition of $\tilde{\varepsilon}$ and the assumption of the theorem on β .

Given the strong drift established in (3), we see that until the first time that $N_t < (1 - \tilde{\varepsilon})\alpha n$, N_t stochastically dominates a continuous-time biased random walk on the interval $(-\infty, \alpha n)$. Let τ be the first time that this biased random walk, when started in stationarity, is smaller than $(1 - \tilde{\varepsilon})\alpha n$. Now $\Pr[\tau = 0] = \exp(-\Theta(n \log n))$, and the probability density of seeing a transition crossing $(1 - \tilde{\varepsilon})\alpha n$ is $\exp(-\Theta(n \log n))$, so $\mathbb{E}[\tau] = \exp(-\Theta(n \log n))$, so we conclude that if the infection starts out large enough, it takes exponentially long for it to die out.

4.3 Proof of Theorem 1

4.3.1 Proof Idea

Consider the star graph on n leaves, and let k denote the number of them that are infected. We shall assume that ρ and ρ' are fixed while $n \rightarrow \infty$, and that β tends to 0 as $n \rightarrow \infty$. We start by describing the course of a typical infection, which should convey to the reader the reason for the $1/3$ exponent in the theorem, and then proceed with the proof. The course of the epidemic may be divided into center-healthy phases and center-infected phases.

Suppose that the center is healthy. No new leaves become infected, leaves become healthy at rate ρk , and the center becomes infected at rate βk . The number of leaves that become healthy is the minimum of k and a geometric random variable with parameter $\beta/(\beta + \rho)$ (and expectation ρ/β). This is well-approximated by

$$\frac{\rho}{\beta} \text{Exp}(1),$$

where Exp denotes an exponential random variable.

Suppose that the center is infected. Let us assume that $1 \ll k \ll n$. Then the center becomes healthy with rate $\rho' k + \rho \approx \rho' k$, new leaves become infected at rate $(n - k)\beta \approx n\beta$,

and leaves become cured at rate $k(\rho + \rho')$. By the time the center becomes healthy again, we expect $O(1)$ leaves to become healthy, but about

$$\frac{n\beta}{\rho'k} \text{Exp}(1)$$

new leaves to become infected.

If the number of leaves that become well during a center-healthy phase about balances the number of leaves that become infected during a center-infected phase, then we have $\rho/\beta \approx n\beta/(\rho'k)$. So let us define

$$k_0 := \frac{n\beta^2}{\rho\rho'}.$$

If the infection persists for a long time, then we expect that typically there will be about k_0 infected leaves. We now distinguish two cases:

Survival of the Infection: In order for all these k_0 infected leaves not to become cured during a single center-healthy phase, it must be that $k_0 \gg \rho/\beta$, i.e. $\frac{n\beta^3}{\rho^2\rho'} \gg 1$. So let us assume that this inequality holds, i.e., let us assume that

$$\beta \gg n^{-1/3}.$$

Next we compare the number of new infections during a center-infected phase to k when $k \approx k_0$. This number is approximately $\frac{n\beta}{\rho'k_0} = \frac{\rho}{\beta}$, of the same order as the number of cured leaves during a center-healthy phase. Both are $\gg 1$ and $\ll k_0$, implying that for a long time, the number of infected leaves stays near to k_0 .

Fast Curing: If

$$\beta \ll n^{-1/3},$$

we have a good chance that all these k_0 infected leaves become cured during a single center-healthy phase, so that the infection dies out. We thus expect fast curing.

4.3.2 Proof of long survival if $\beta = \rho n^{-1/3+\varepsilon}$.

We denote the state of the chain by $(0, k)$ when the center is healthy, and there are k infected leaves, and by $(1, k)$, when the center and k of the leaves are infected. The evolution of the chain is then a process with rates

$$\begin{aligned} (1, k) &\rightarrow (1, k+1) && \text{at rate } \beta(n-k) \\ (1, k) &\rightarrow (1, k-1) && \text{at rate } (\rho + \rho')k \\ (1, k) &\rightarrow (0, k) && \text{at rate } \rho + \rho'k \\ (0, k) &\rightarrow (0, k-1) && \text{at rate } \rho k \\ (0, k) &\rightarrow (1, k) && \text{at rate } \beta k, \end{aligned} \tag{4}$$

so that, in particular, the number of infected leaves reduces by a geometric random variable with parameter $\beta/(\beta + \rho)$ whenever the center is cured.

Our next lemma is the main technical lemma on which the proof of Theorem 1 is based.

Lemma 1. Let $\beta, \rho, \rho' > 0$, let

$$k_1 \leq \frac{\beta^2 n}{4(\rho + \rho')(\rho + \beta) + \beta^2},$$

and let $\theta = \frac{1}{2} \frac{\beta}{\beta + \rho}$.

i) Assume that the chain (4) starts in a state $(1, k)$ with $1 \leq k < k_1$. Then the probability $P_{k, k_1}^{(1)}$ that the chain reaches a state with no infected leaves before it reaches $(1, k_1)$ is bounded by $e^{-k\theta}$.

ii) Assume that the chain (4) starts in a state $(0, k)$ with $1 \leq k < k_1$. Then the probability $P_{k, k_1}^{(0)}$ that the chain reaches a state with no infected leaves before it reaches $(1, k_1)$ is bounded by $3e^{-k\theta}$.

Proof. (i) In order to determine whether the chain (4) reaches the state $(1, k_1)$ before the number of leaves drops to zero when starting from $(1, k)$, it is clearly enough to consider the one dimensional chain

$$\begin{aligned} k &\rightarrow k + 1 && \text{at rate } \beta(n - k) \\ k &\rightarrow k - 1 && \text{at rate } (\rho + \rho')k \\ k &\rightarrow k - \text{Geom}(\beta/(\beta + \rho)) && \text{at rate } \rho + \rho'k. \end{aligned} \tag{5}$$

The first time this chain reaches a state $k \leq 0$ corresponds to the first time the chain (4) reaches the state $(1, 0)$ or $(0, 0)$, and the first time it reaches k_1 corresponds to the first time the chain (4) reaches the state $(1, k_1)$.

Consider the discrete time analog of the chain (5),

$$\begin{aligned} k &\rightarrow k + 1 && \text{with probability } 1 - p_k^{(1)} - p_k^{(2)} \\ k &\rightarrow k - 1 && \text{with probability } p_k^{(1)} \\ k &\rightarrow k - \text{Geom}(\beta/(\beta + \rho)) && \text{with probability } p_k^{(2)}, \end{aligned} \tag{6}$$

where

$$p_k^{(1)} = \frac{(\rho + \rho')k}{\beta n + (\rho + 2\rho' - \beta)k + \rho} \quad \text{and} \quad p_k^{(2)} = \frac{\rho + \rho'k}{\beta n + (\rho + 2\rho' - \beta)k + \rho}.$$

In order to analyze this chain, we will couple it to the chain

$$\begin{aligned} k &\rightarrow k + 1 && \text{with probability } 1 - p \\ k &\rightarrow k - 1 - \text{Geom}(\beta/(\beta + \rho)) && \text{with probability } p \end{aligned} \tag{7}$$

where

$$p = \frac{2(\rho + \rho')k_1}{\beta(n - k_1) + 2(\rho + \rho')k_1}.$$

Using the facts that the downstep in (7) bounds both downsteps in (6) and that $p_k^{(1)} + p_k^{(2)} \leq p$ as long as $k \leq k_1$, it is easy to see that the two chains can be coupled in such a way that so that (7) stays below (6), as long as the latter does not exceed the threshold k_1 .

Let X a random variable with the same distribution as the step distribution of (7), i.e., let X be equal to

$$\begin{aligned} &+1 && \text{with probability } 1-p \\ -1 - \text{Geom}(\beta/(\beta+\rho)) && \text{with probability } p, . \end{aligned}$$

We claim that the equation $\mathbb{E}[e^{-\theta X}] = 1$ has exactly two solutions: the solution $\theta = 0$, and the solution

$$\theta^* = -\log\left[1 - \frac{\beta}{\beta+\rho} + \frac{p}{1-p}\right] \quad (8)$$

In order to see this, we note that

$$\mathbb{E}[e^{-\theta X}] = (1-p)e^{-\theta} + p \sum_{j \geq 0} e^{(j+1)\theta} \left(1 - \frac{\beta}{\beta+\rho}\right)^j \frac{\beta}{\beta+\rho}.$$

If $(1 - \frac{\beta}{\beta+\rho})e^\theta \geq 1$, then the right hand side is infinity, implying that any solution to $\mathbb{E}[e^{-\theta X}] = 1$ is such that $(1 - \frac{\beta}{\beta+\rho})e^\theta < 1$. Under this condition, we sum the geometric series, giving

$$\mathbb{E}[e^{-\theta X}] = (1-p)e^{-\theta} + p \frac{\beta}{\beta+\rho} \left[e^{-\theta} - \left(1 - \frac{\beta}{\beta+\rho}\right) \right]^{-1}.$$

The equation $\mathbb{E}[e^{-\theta X}] = 1$ then turns into a quadratic equation for $e^{-\theta}$. This equation has two solutions: the solution $\theta = 0$ and the solution $\theta = \theta^*$, with θ^* given by (8). Since both solutions of the quadratic equation satisfies the condition $(1 - \frac{\beta}{\beta+\rho})e^\theta < 1$, these solutions are solutions of $\mathbb{E}[e^{-\theta X}] = 1$, which completes the proof of our claim.

Let $\phi(X) = e^{-\theta^* X}$. Then, $\phi(S_t)$ is a martingale. We want to estimate the probability P that S_t reaches a state $S_t \leq 0$ before reaching the state $S_t = k_1$. By the optional stopping theorem and the fact that $\phi(x) \geq \phi(0)$ for $x < 0$, we have

$$P\phi(0) + (1-P)\phi(k_1) \leq \phi(k)$$

and hence

$$P \leq \frac{\phi(k) - \phi(k_1)}{\phi(0) - \phi(k_1)} = \frac{1 - e^{-\theta^*(k_1-k)}}{1 - e^{-\theta^*k_1}} e^{-\theta^*k} \leq e^{-\theta^*k}.$$

Recall that we have coupled our original chain (6) and the chain (7) in such a way that the former stays above the latter as long as $k \leq k_1$ in the chain (6). Assume that the chain (6) reaches 0 before it reaches k_1 . Then the chain (7) must reach 0 before k_1 as well, implying that $P_{k,k_1}^{(1)} \leq P$.

To complete the proof of (i), we use the assumption of the lemma to bound

$$\frac{p}{1-p} = \frac{2(\rho + \rho')k_1}{\beta(n - k_1)} \leq \frac{1}{2} \frac{\beta}{\beta + \rho}.$$

This in turn implies that

$$e^{-\theta^*} \leq 1 - \frac{1}{2} \frac{\beta}{\beta + \rho} \leq e^{-\frac{1}{2} \frac{\beta}{\beta + \rho}}$$

which completes the proof of (i).

(ii) To prove (ii), we use (i) and the fact that the number of cured leaves during a center healthy is a geometric random variable with parameter $\beta/(\beta + \rho)$. As a consequence,

$$\begin{aligned} P_{k_1, k}^{(0)} &= \frac{\beta}{\beta + \rho} \sum_{j=0}^k P_{k_1, k-j} \left(1 - \frac{\beta}{\beta + \rho}\right)^j \leq \frac{\beta}{\beta + \rho} \sum_{j=0}^k e^{-\frac{1}{2} \frac{\beta}{\rho + \beta} (k-j)} e^{-\frac{\beta}{\rho + \beta} j} \\ &\leq \frac{\beta}{\beta + \rho} \frac{1}{1 - e^{-\frac{1}{2} \frac{\beta}{\rho + \beta}}} e^{-\frac{1}{2} \frac{\beta}{\rho + \beta} k} \leq 2 \left(1 + \frac{1}{2} \frac{\beta}{\rho + \beta}\right) e^{-\frac{1}{2} \frac{\beta}{\rho + \beta} k} \leq 3e^{-\frac{1}{2} \frac{\beta}{\rho + \beta} k}, \end{aligned}$$

where we used that $\frac{x}{1-e^{-x}} \leq 1+x$ for all $x \geq 0$ in the second to last step. \square

Proof of Theorem 1 (i) – (iii)

Let θ be as in Lemma 1, let

$$k_1 = \left\lfloor \frac{\beta^2 n}{4(\rho + \rho')(\rho + \beta) + \beta^2} \right\rfloor,$$

and let c and n_0 be such that $3e^{-\theta(k_1-1)} \leq e^{-4c\beta^3 n/\rho^3}$ if $n \geq n_0$.

Assume first that the chain starts with $k \geq k_1$ infected leaves. Since $(0, 0)$ is the only absorption point of the chain (4), it must at some point pass through a state with $k_1 - 1$ infected leaves. By Lemma 1, with probability at least $1 - e^{-2c\beta^3/n}$, the chain will return to k_1 at least $\lfloor e^{2c\beta^3/n} \rfloor$ times before going to the absorbing state $(0, 0)$. Since the transition rates of the chain (4) are bounded by $\beta n + 2(\rho + \rho')n \leq C\rho n$, the expected time it takes for these $\lfloor e^{2c\beta^3/n} \rfloor$ returns is at least $\lfloor e^{2c\beta^3/n} \rfloor / (C\rho n) \geq 2\tau_2$. Observing that the probability that the actual time is half the expected time decays exponential in $\lfloor e^{2c\beta^3/n} \rfloor$, we conclude that with probability at least $1 - e^{-2c\beta^3 n/\rho^3} - \exp(-c \lfloor e^{2c\beta^3/n} \rfloor) \geq 1 - 2e^{-2c\beta^3 n/\rho^3}$, the survival time is at least τ_2 .

Next we consider an initial condition with $k_2 = \lfloor (\rho/\beta) \log n \rfloor \leq k < k_1$ infected leaves. Applying Lemma 1 once more (and recalling that $\beta \leq \rho$), we conclude that with probability at least $1 - 3e^{-k\theta} \geq 1 - O(n^{-1/4})$, the chain reaches k_1 infected leaves, at which point with high probability, the survival time of the infections is at least τ_2 .

For the remainder of the proof, we distinguish the case where initially the center is healthy and the one where it is infected. We start with the latter. Assume thus that initially, the center is healthy, and $0 \leq k \leq k_2$ leaves are infected. A lower bound on the probability that the chain reaches the state $(1, k_2)$ is obtained by assuming that it goes through the sequence $(1, k), (1, k+1), \dots, (1, k_2)$ with no downsteps, and without curing of the center. The probability of this event is

$$\begin{aligned} \prod_{i=k}^{k_2-1} \frac{\beta(n-i)}{\beta(n-i) + (2\rho' + \rho)i + \rho} &\geq \prod_{i=0}^{k_2-1} \frac{1}{1 + \frac{C\rho}{\beta n} i} \geq \exp\left(-\sum_{i=0}^{k_2-1} \frac{C\rho}{\beta n} i\right) \\ &\geq \exp\left(-\frac{C\rho k_2^2}{2\beta n}\right) \geq 1 - \frac{C\rho k_2^2}{2\beta n} \end{aligned}$$

Inserting the value of k_2 and the lower bound on β , we see that this probability is $1 - O(n^{-3\epsilon}(\log n)^2)$. Together with the bounds proven so far and the fact that k_1 grows like $n^{1/3+2\epsilon}$ as $n \rightarrow \infty$ this proves statement (ii) of Theorem 1.

We are left with analyzing an initial condition with $1 \leq k < k_2$ infected leaves and healthy center (we will actually not use the assumption $k < k_2$ in most of the following argument). Since the downsteps in a center healthy phase are governed by a geometric distribution with parameter $\beta/(\beta + \rho)$, the probability that the infection dies out without ever infecting the center can be easily calculated; it is $(1 - \beta/(\beta + \rho))^k$. Conditioned on this event, the time it takes for extinction the random variable $T_k = \sum_{i=1}^k X_i$, where X_i is an exponential random variable with rate $i\rho$. Since the expectation of X_i is $1/i\rho$, Markov's bound gives that $T_k \geq \tau$ with probability at most $\frac{1}{\tau\rho} \sum_{i=1}^k \leq \frac{\log k + 1}{\tau\rho}$, which implies the first bound in (iii).

Before turning to the other bound in (iii), we would like to prove sharper bounds on the tail of T_k . Using standard large deviation techniques (and the fact that $e^{\alpha\rho X_i}$ has expectation $(1 - \frac{\alpha}{i})^{-1}$), this is not hard, leading an estimate of

$$P(T_k \geq \tau) \leq O(e^{-\alpha(\rho\tau - \log k - 1)})$$

where α can be chosen arbitrarily in $(0, 1)$, and the constant implicit in the O -symbol depends on α . This implies in particular that, conditioned on the event that the infection dies out without ever infecting the center, the survival time is at most $\frac{2}{\rho} \log n$ with probability $1 - O(n^{-\alpha})$. By contrast, conditioned on the event that the center gets infected at least once, the survival time is at least τ_2 with probability $1 - O(n^{-c})$ by the already proven statement (ii). This proves (i).

To prove the upper bound in (iii) we bound the event that the survival time is at most τ by

$$(1 - \frac{\beta}{\beta + \rho})^k P(T_k \leq \tau) + (1 - (1 - \frac{\beta}{\beta + \rho})^k) O(n^{-c}) \leq (1 - \frac{\beta}{\beta + \rho})^k + O(n^{-c})$$

where we used the fact that $\tau < \tau_2$, which implies that conditioned on an infected center, the probability that survival time of the infection $\leq \tau$ is bounded by $O(n^{-c})$.

4.3.3 Proof of short survival if $\beta = \rho n^{-1/3-\epsilon}$

The proof uses the following lemmas, which follow from well-known large deviation bounds for sums of i.i.d. Bernoulli and Geometric random variables.

Lemma 2. *Suppose X_1, X_2, \dots are i.i.d. Geometric(λ) random variables, so that $EX_1 = (1 - \lambda)/\lambda$. Then, for all $x < EX_1$, we have*

$$\frac{1}{t} \log P\left(\sum_{\ell=1}^t X_k \leq xt\right) \leq -(1+x)H\left(\frac{x}{1+x}; 1-\lambda\right),$$

$$\text{where } H(q; p) = q \log \frac{q}{p} + (1-q) \log \frac{1-q}{1-p}.$$

Proof. Let $p = 1 - \lambda$ and $q = \frac{x}{1+x}$. By Chernoff's inequality, we have

$$\frac{1}{t} \log P\left(\sum_{k=1}^t X_k \leq xt\right) \leq -\theta x + \log E[e^{\theta X_1}] = -\theta x + \log\left(\frac{\lambda}{1 - pe^{\theta}}\right)$$

for all $\theta \leq 0$. Optimizing over θ , we get $pe^\theta = q$ and

$$\frac{1}{t} \log P\left(\sum_{k=1}^t X_k \leq xt\right) \leq x \log \frac{p}{q} + \log\left(\frac{\lambda}{1-q}\right) = -(1+x)H(q;p).$$

Lemma 3. *Suppose X_1, X_2, \dots are i.i.d. Bern(p) random variables. Then*

$$P\left(\sum_{k=1}^t X_k < pt/2\right) \leq \exp\left(-pt \frac{1 - \log 2}{2}\right).$$

Proof. For Bernoulli random variables and $x \leq p$, Chernoff's inequality gives the large deviation estimate

$$\frac{1}{t} \log P\left(\sum_{k=1}^t X_k \leq xt\right) \leq x \log \frac{x}{p} + (1-x) \log \frac{1-p}{1-x} \leq x \log \frac{x}{p} - (p-x),$$

where in the last step we used the fact $\log y \leq y - 1$ for all $y > 0$. Setting $x = p/2$, this gives the claim of the lemma.

To prove Theorem 1 (iv), we first note that we may assume without loss of generality that $\varepsilon < 2/3$. Indeed, for $\varepsilon \geq 2/3$, we have that $\beta = O(n^{-1})$, implying that the probability of extinction during a center healthy phase is of order one, and the expected time for this event is bounded by $\log n$. On the other hand, for $\beta = O(n^{-1})$, the probability of curing the center is bounded from below by an n -independent, positive constant as well. This implies that the expected number of times the center gets cured before the epidemic dies out is of order one as well, so that the expected survival time is $O(\log n)$.

Assume thus without loss of generality that $0 < \varepsilon < 2/3$, and let $k_1 = \lceil (\beta n / \rho)^{3/4} \rceil$, so that $k_1 / (\beta n) \rightarrow 0$ and $k_1^2 / (\beta n) \rightarrow \infty$ as $n \rightarrow \infty$. The proof of the theorem proceeds in two steps. First, we prove that starting from a state with more than k_1 infected leaves, the time τ to reach k_1 leaves for the first time is at most $O(\log n)$ in expectation, and second we prove that starting with k_1 or fewer infected leaves, it takes at most time $O(\log n)$ in expectation to go extinct.

The next lemma will be used in both steps of the proof.

Lemma 4. *Given $\rho > 0$, $\rho' > 0$ and $0 < \varepsilon < 2/3$, there are constants $n_0 < \infty$ and $C < \infty$ (depending on ρ, ρ', ε) such that for $\beta = \rho n^{-1/3-\varepsilon}$, $n \geq n_0$ and $k \geq \sqrt{\beta n / \rho}$, the following holds: Starting from a state with more than k and less than $2k$ infected leaves, the expected number of transitions until the chain (4) first reaches a state with k infected leaves is bounded by Ck .*

Proof. We first note that it is enough to prove the lemma for an initial state with infected center. Indeed, if the center is not infected, the number of infected leaves will either go down to k in less than k steps without ever infecting the center, or the chain will end up

in a state with infected center and k_0 infected leaves, where k_0 is a random number in $I_k = \{k + 1, \dots, 2k\}$.

Assume therefore that we start with $k_0 \in I_k$ infected leaves and an infected hub, and let U be the number of steps it takes the chain (4) to reach k infected leaves. To bound U , we consider the discrete time chain S defined by (6). Assume that it takes K steps for the chain S to reach k or less infected leaves. Conditioned on this event, the total upward excursion of the original chain before it reaches k infected leaves is bounded by K since both S and (4) can only move up by one at each step. Consequently, the time U is bounded by $2K + k$. To prove the lemma, it is therefore enough to show that it takes the chain S at most $O(k)$ steps to reach k or fewer infected leaves.

To this end, we will couple the chain S to a chain \tilde{S} which starts at $\tilde{S}(0) = k_0$, and at each discrete time $t = 1, 2, \dots$ takes an upward step of $+1$ with probability $1 - p$, and a downward step of $\text{Geom}(\beta/\rho + \beta)$ with probability p , with

$$p = \frac{\rho'}{2(\rho + \rho')} \frac{1}{1 + \sqrt{\beta n/\rho}}. \quad (9)$$

We claim that the two chains can be coupled in such a way that $\tilde{S}(t) \geq S(t)$ as long as $S(t) \geq k$. To this end, we note that the probability of a downward move of the chain S is bounded from below by

$$p_{S(t)}^{(2)} \geq \frac{\rho' S(t)}{\beta n + 2(\rho + \rho') S(t)} \geq p,$$

provided $S(t) \geq k$. Since for $t = 0$, we have $S(0) = \tilde{S}(0) \geq k$, we can maintain the condition $\tilde{S}(t) \geq S(t) \geq k$ as long as $\tilde{S}(t) \geq k$.

Next we observe that in distribution

$$\tilde{S}(t) = \tilde{S}(t-1) + 1 - (Y(t) + 1)X(t),$$

where the $X(t)$ are i.i.d. $\text{Bern}(p)$ and independent of the $Y(t)$, which are i.i.d. $\text{Geom}(\beta/(\rho + \beta))$. We will show that for $t \geq k$ the probability that $\tilde{S}(t) - \tilde{S}(0) \geq -k$ is exponentially small in t , which in turn will allow us to give a bound on the expected number of steps the chain \tilde{S} needs to fall below k .

Indeed, by the independence of the $X(t)$ and $Y(t)$, and the non-negativity of the $Y(t)$, we have that for all $t \geq k$,

$$\begin{aligned} P(\tilde{S}(t) - \tilde{S}(0) \geq -k) &= P\left(\sum_{i=1}^t (1 + Y(i))X(i) \leq t + k\right) \\ &\leq P\left(\sum_{i=1}^t X(i) \leq pt/2\right) + P\left(\sum_{i=1}^{\lceil pt/2 \rceil} Y(i) \leq 2t\right). \end{aligned} \quad (10)$$

By Lemma 3, we have

$$P\left(\sum_{i=1}^t X(i) \leq pt/2\right) \leq \exp(-c_1 pt), \quad (11)$$

where $c_1 = (1 - \log 2)/2 > 0$. Moreover, by Lemma 2,

$$\frac{1}{\lceil pt/2 \rceil} \log P\left(\sum_{i=1}^{\lceil pt/2 \rceil} Y(i) \leq \frac{\lceil pt/2 \rceil \rho}{2\beta}\right) \leq -\left(1 + \frac{\rho}{2\beta}\right) H\left(\frac{\rho}{\rho + 2\beta}; \frac{\rho}{\rho + \beta}\right). \quad (12)$$

Now

$$\begin{aligned} H\left(\frac{\rho}{\rho + 2\beta}; \frac{\rho}{\rho + \beta}\right) &= \frac{\rho}{\rho + 2\beta} \log \frac{\rho + \beta}{\rho + 2\beta} + \frac{2\beta}{\rho + 2\beta} \log \frac{2(\rho + \beta)}{\rho + 2\beta} \\ &= \log \frac{\rho + \beta}{\rho + 2\beta} + \frac{2\beta}{\rho + 2\beta} \log 2 \\ &\geq \frac{\beta}{\rho + 2\beta} \log 4 - \frac{\beta}{\rho + \beta}, \end{aligned}$$

where we have used the inequality $\log(x) \leq x - 1$ to obtain the last inequality. Substituting this in (12), we get

$$\frac{1}{\lceil pt/2 \rceil} \log P\left(\sum_{i=1}^{\lceil pt/2 \rceil} Y(i) \leq \frac{\lceil pt/2 \rceil \rho}{2\beta}\right) \leq -\log 2 + \frac{\rho + 2\beta}{2(\rho + \beta)} \leq -c_2, \quad (13)$$

for some constant $c_2 > 0$ and all n sufficiently large. On the other hand,

$$\frac{\lceil pt/2 \rceil \rho}{2\beta} \geq \frac{pt\rho}{4\beta} = t \frac{\rho\rho'}{8(\rho + \rho')^2} \frac{\rho/\beta}{1 + \sqrt{\beta n/\rho}} \geq 2t$$

for all n sufficiently large. As consequence,

$$P\left(\sum_{i=1}^{\lceil pt/2 \rceil} Y(i) \leq 2t\right) \leq e^{-c_2 \lceil pt/2 \rceil} \leq e^{-c_2 pt/2} \quad (14)$$

provided n is sufficiently large. Combining (10), (11) and (14) we obtain that

$$P\left(\tilde{S}(t) - \tilde{S}(0) \geq -k\right) \leq e^{-c_1 pt} + e^{-c_2 pt/2}. \quad (15)$$

Let U denote the total number of steps for which the process $S(\cdot)$ remains above k . It is clear from the dominating arguments above that, if the random walk $\tilde{S}(\cdot)$ falls below $-k$ before time t , then the random walk S falls below k in this time interval as well. In other words, equation (15)) implies that

$$P(U_i > t) \leq e^{-c_1 pt} + e^{-c_2 pt/2} \quad \text{for all } t \geq k.$$

This in turn implies that

$$E[U_i] \leq k + O\left(\frac{1}{p}\right) = O(k) \quad (16)$$

where in the last step we used that by the assumption of the lemma $k \geq \sqrt{\beta n/\rho}$ to conclude that $1/p = O(\sqrt{\beta n/\rho}) = O(k)$, with the implicit constants in the O -symbols depending on ρ and ρ' . \square

The lemma immediately implies that starting from an arbitrary initial condition, it takes fewer than $O(\log n)$ steps to reach a phase with k_1 or fewer infected leaves. Indeed, let A be the set of states with k_1 or fewer infected leaves (and at least one infected node), let τ be the time it takes to first reach a state in A , and let B_i be the set of states with at least $k_1 2^{i-1} + 1$ and at most $k_1 2^i$ infected leaves. If the chain starts in phase A there is nothing to prove, so let us assume that the chain starts in some phase B_i for $i \geq 1$.

Let T_i the total time it takes the chain (4) to first enter the phase B_{i-1} , assuming it starts in B_i . It follows from (4) that the mean time for a step of this chain is bounded by $1/(\min\{\rho, \rho'\}k)$ when the number of infected leaves is k , and hence by $1/(\min(\rho, \rho')k_1 2^i)$ throughout phase i . Since the average number of steps in phase i is bounded by $Ck_1 2^i$, we conclude that for n large enough, $E[T_i] \leq C/(\min\{\rho, \rho'\})$. Since there are fewer than $\log_2 n$ phases, it follows $E[\tau] = O(\log n)$.

To prove fast extinction once we have reached the phase A , we distinguish the center healthy part of A , to be denoted by A^0 , and the center infected part of A , denoted by A^1 . We also note that the first time the chain leaves phase A , it will be either go extinct, or it will enter into the state $(1, k_1 + 1)$.

Lemma 5. *Given $\rho > 0$, $\rho' > 0$, and $0 < \varepsilon < 2/3$ there is a constant n_0 such that for $n \geq n_0$, $\beta = \rho n^{-1/3-\varepsilon}$ and $k_1 = \lceil (\beta n / \rho)^{3/4} \rceil$ the following holds: If the infection starts in the state $(1, k_1 + 1)$, the probability that at the first time the chain enters phase A , it enters it with a healthy center is $1 - o(1)$.*

Proof. Let C be the set of states with infected center, and a number of infected leaves which stays between $k_1 + 1$ and $2k_1 - 1$. If the chain (4) leaves phase C by curing the center, the number of infected leaves will be below $2k_1$, and since $\rho/\beta \gg k_1$, we conclude that with high probability, the chain will enter phase A through the state $(k_1, 0)$. We thus have to bound the probability that the chain (4) leaves the phase C through the states $(1, k_1)$ or $(1, 2k_1)$. To bound this probability, let us postpone the decisions on whether we cure the center, considering instead a chain with transition rates

$$\begin{aligned} k &\rightarrow k + 1 && \text{at rate } \beta(n - k) \\ k &\rightarrow k - 1 && \text{at rate } (\rho + \rho')k. \end{aligned} \tag{17}$$

In a second step, we will then look at the Poisson clocks for curing the center, deciding on whether we should actually have taken a step curing the center, instead of one of the steps of the chain (17). But before doing this, we bound the probability that the chain (17) falls to k_1 before it reaches the state $2k_1$. By a martingale argument similar to that used in the proof of Lemma 1, one easily shows that this probability is dominated by the probability of the event that the chain goes to k_1 in the very first step, leading to a bound of $O(\rho k_1 / \beta n)$ for this event.

We may therefore assume that the chain (17) leaves phase C through the state $2k_1$. Going back to the original chain (4), we now have to look at the Poisson clocks that the center gets cured instead of taking one of the steps needed to reach the state $2k_1$. Since the restricted chain (17) touches each state from $(1, k_1 + 1)$ to $(2, 2k_1)$ at least once, we conclude that conditioned on the event that the chain (17) leaves phase C through the state $2k_1$, the

probability of curing the center is at least

$$1 - \prod_{i=k_1+1}^{2k_1} \frac{\beta(n-i) + (\rho + \rho')i}{\beta(n-i) + \rho + (\rho + 2\rho')i} = 1 - \exp\left(-\frac{\rho'k_1^2}{2\beta n}(1 + o(1))\right)$$

Thus we have shown that with probability tending to one three events happen: the chain (17) walks to $2k_1$ before falling below $k_1 + 1$, the center gets cured when the chain (17) has between $k_1 + 1$ and $2k_1$ leaves, and at least k_1 leaves get cured directly after the center was cured. Together, these three events imply the event considered in the lemma. \square

Given the above two lemmas, the proof of Theorem 1 (iv) is now an easy exercise. Indeed, at time τ , the chain will either start in phase A^0 or in phase A^1 . If it starts in phase A^0 , the expected time it stays there is $O(\sum_{i=1}^{k_1} 1/i) = O(\log n)$, and the probability that it leaves this phase by curing all leaves before the hub is re-infected tends to 1 as n tends to infinity. If we start in phase A^1 , it will stay in this phase for an expected time which is $o(1)$, and it will either leave this phase by curing the center, entering phase A^0 , or it will leave this phase through the state $(1, k_1 + 1)$, at which point it will go to the phase A^0 with high probability, in expected time $O(1)$ (by Lemmas 5 and 4, respectively).

Thus with high probability, the infection will die out before the chain re-enters phase A^1 , and it will take expected time at most $O(\log n)$ before the infection dies out or re-enters phase A^1 . We therefore need a geometric number of visits to A^1 (with mean close to one) before the epidemic dies out. Since each of these moves takes expected time at most $O(\log n)$, the expected time before the epidemic dies out is of order $O(\log n)$ as well. This proves the last statement of Theorem 1.

Acknowledgement. The authors would like to thank David Wilson for insightful discussions and his comments on an earlier version of this paper.

References

- [1] W. Aiello, F. Chung, and L. Lu. A random graph model for power law graphs. *Experiment. Math.*, 10:53–66, 2001.
- [2] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [3] N. Berger, C. Borgs, J. Chayes, and A. Saberi. On the spread of viruses on the internet. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 301–310, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.
- [4] B. Bollobás, O. Riordan, J. Spencer, and G. Tusndy. The degree sequence of a scale-free random graph process. *Random Structures and Algorithms*, 18:279–290, 2001.
- [5] R. Christley, S. Robinson, R. Lysons, and N. French. Network analysis of cattle movement in great britain. 2005.

- [6] G. Forster and C. A. Gilligan. Optimizing the control of disease infestations at the landscape scale. 2007.
- [7] A. Ganesh, L. Massoulié, and D. Towsley. The effect of network topology on the spread of epidemics. In *Proceedings of IEEE Conference on Computer Communications (Infocom)*, 2005.
- [8] C. Gkantsidis, M. Mihail, and A. Saberi. Conductance and congestion in power law graphs. In *Proceedings of the 2003 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 148–159, 2003.
- [9] M. J. Keeling and K. T. D. Eames. Networks and epidemic models. *Journal of The Royal Society Interface*, 2005.
- [10] D. M. Kienzle and M. C. Elder. Recent worms: a survey and trends. In *Proceedings of the 2003 ACM workshop on Rapid Malcode*, 2003.
- [11] I. Z. Kiss, D. M. Green, , and R. R. Kao. Infectious disease control using contact tracing in random and scale-free networks. *Journal of The Royal Society Interface*, 2006.
- [12] T. M. Liggett. *Stochastic interacting systems: contact, voter and exclusion processes*, volume 324 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [13] F. Liljeros, C. R. Edling, L. A. N. Amaral, H. E. Stanley, and Y. Aberg. The web of human sexual contacts. *Nature*, 411:907, 2001.
- [14] L. A. Meyers, B. Pourbohloul, M. E. J. Newman, D. M. Skowronski, and R. C. Brunham. Network theory and sars: predicting outbreak diversity. *Journal of Theoretical Biology*, 2005.
- [15] M. Mihail, C. Papadimitriou, and A. Saberi. On certain connectivity properties of the Internet topology. pages 28–35, 2003. Extended version in the special issue of Journal of Computer and System Sciences.
- [16] D. Moore, C. Shannon, and J. Brown. Code-Red: a case study on the spread and victims of an Internet worm. In *Proc. of Internet Measurement Workshop 2002*, Nov 2002.
- [17] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Physical Review Letters*, 86:3200–3203, 2001.
- [18] R. Pastor-Satorras and A. Vespignani. Epidemic dynamics in finite scale-free networks. *Physical Review E*, 65, 2002.
- [19] D. J. Price. A general theory of bibliometric and other cumulative advantage processes. *J. Amer. Soc. Inform. Sci.*, 27:292–306, 1976.

- [20] A. Stacey. The contact process on finite homogeneous trees. *Probab. Theory Related Fields*, 121(4):551–576, 2001.
- [21] A. M. Stacey. The existence of an intermediate phase for the contact process on trees. *Ann. Probab.*, 24(4):1711–1726, 1996.
- [22] S. Staniford, V. Paxson, and N. Weaver. How to own the internet in your spare time. In *Proceedings of the 11th USENIX Security Symposium*, San Fransisco, California, USA, August 5-9 2002. USENIX Association.
- [23] L. S. Tsimring and R. Huerta. Modeling of contact tracing in social networks. *Physica A: Statistical Mechanics and its Applications*, 325, 2003.
- [24] A. Wagner and V. Anantharam. Designing a contact process: the piecewise-homogeneous process on a finite set with applications. *Stochastic Process. Appl.*, 115(1):117–153, 2005.