A lower bound on the performance of dynamic curing policies for epidemics on graphs

Kimon Drakopoulos   Asuman Ozdaglar   John N. Tsitsiklis

Abstract— We consider an SIS-type epidemic process that evolves on a known graph. We assume that a fixed curing budget can be allocated at each instant to the nodes of the graph, towards the objective of minimizing the expected extinction time of the epidemic. We provide a lower bound on the optimal expected extinction time as a function of the available budget, the epidemic parameters, the maximum degree, and the CutWidth of the graph. For graphs with large CutWidth (close to the largest possible), and under a budget which is sublinear in the number of nodes, our lower bound scales exponentially with the size of the graph.

I. INTRODUCTION

We consider an SIS-type epidemic model with a common infection rate β and an endogenous, node specific curing rate $\rho_i(t)$. A network planner has a total curing budget, which is allocated to different nodes at each point in time according to a dynamic policy based on the history of the epidemic and the network structure.

In earlier work [7], [5] we propose a policy which achieves rapid containment for any set of initially infected nodes under some conditions on the total curing budget. More specifically, we prove that if the CutWidth of the underlying graph is much smaller than the available curing budget, then rapid containment of any epidemic is achieved by the proposed policy.

In this paper we study the case where the CutWidth of the graph is larger than the available budget. Specifically, we focus on a class of graphs for which the CutWidth of the graph is close to the largest possible value and obtain a lower bound on the optimal expected extinction time. When the curing budget scales sublinearly in the number of nodes, our bound implies that under any dynamic curing policy, the expected time to extinction scales exponentially in the number of nodes, in the worst case where all nodes are initially infected. This complements our results in [6], where we provide (weaker) lower bounds for the case where the CutWidth is larger than the available budget, linear in the number of nodes, but not necessarily close to the largest possible value.

A similar model, but in which the curing rate allocation is static (open-loop) has been studied in [4], [8], [3], [12], and the proposed methods were either heuristic or based on mean-field approximations of the evolution process. Closer

to our work, the authors of [2] let the curing rates be proportional to the degree of each node — but again independent of the current state of the network, which means that curing resources may be wasted on healthy nodes. On a graph with bounded degree, the policy in [2] achieves sublinear time to extinction, but requires a curing budget that is proportional to the number of nodes. Moreover, when the underlying graph is an expander and the curing budget is sublinear, it is established in [2] that the optimal expected extinction time scales exponentially in the size of the graph. We extend this result by obtaining an exponential lower bound for all graphs whose CutWidth is close to the largest possible. Furthermore, our result also applies to dynamic policies.

The rest of the paper is organized as follows. In Section II we present the details of our model. In Section III we study relevant graph-theoretic properties and present some key combinatorial results that are used in our analysis. In Section IV we present the lower bound on the optimal expected extinction time.

II. THE MODEL

We consider a network, represented by a connected undirected graph $G = (V, \mathcal{E})$, where V denotes the set of nodes and $\mathcal{E}$ denotes the set of edges. We use n to denote the number of nodes. Two nodes $u, v \in V$ are *neighbors* if $(u, v) \in \mathcal{E}$. We denote by Δ the maximum of the node degrees. To exclude trivial cases, and without loss of generality, we assume throughout that G is connected and, in particular, $\Delta > 0$.

We assume that the nodes in a set $I_0 \subseteq V$ are initially infected and that the infection spreads according to a controlled contact process where the rate at which infected nodes get cured is determined by a network controller. Specifically, each node can be in one of two states: *infected* or *healthy*. The controlled contact process — also known as the controlled SIS epidemic model — on G is a right-continuous, continuous-time Markov process $\{I_t\}_{t \geq 0}$ on the state space $\{0, 1\}^V$, where I_t stands for the set of infected nodes at time t. We refer to I_t as the *infection process*.

State transitions at each node occur independently according to the following dynamics.

a) The process is initialized at the given initial state I_0.
b) If a node v is healthy, i.e., if $v \notin I_t$, the transition rate associated with a change of the state of that node to being infected is equal to an infection rate β times the number of infected neighbors of v, that is,

$$\beta \cdot \left| \{(u, v) \in \mathcal{E} : u \in I_t\} \right|,$$

where we use $|\cdot|$ to denote the cardinality of a set. By rescaling time, we can and will assume throughout the paper that $\beta = 1$.

c) If a node v is infected, i.e., if $v \in I_t$, the transition rate associated with a change of the state of that node to being healthy is equal to a curing rate $\rho_v(t)$ that is determined by the network controller, as a function of the current and past states of the process. We are assuming here that the network controller has access to the entire past evolution of the process.

We assume a *budget constraint* of the form

$$\sum_{v \in V} \rho_v(t) \leq r, \tag{1}$$

at each time instant t, reflecting the fact that curing is costly. A *curing policy* is a mapping which at any time t maps the past history of the process to a curing vector $\rho(t) = \{\rho_v(t)\}_{v \in V}$ that satisfies (1). The continuous time nature of the problem allows us to restrict out attention to policies that focus on one node at each time, without loss of generality.

We define the *time to extinction* as the time until the process reaches the absorbing state where all nodes are healthy:

$$\tau = \min\{t \geq 0 : I_t = \emptyset\}. \tag{2}$$

The *expected time to extinction* (the expected value of τ) is the performance measure that we will be focusing on.

A. A perspective on our main result

If the graph G is complete, all policies that always allocate the entire curing budget to infected nodes are essentially equivalent, in the sense that the dynamics of $|I_t|$, the number of infected nodes, are identical under all such policies. Furthermore, I_t evolves as a birth-death Markov chain which is easy to analyze, and it is not hard to show that the expected time to extinction increases exponentially with n. On the other hand, for more general graphs with large CutWidth but bounded degree, an analysis using a one-dimensional birth-death chain or a simple Lyapunov function does not seem possible.

A related, and conceptually simple, way of deriving lower bounds for more general graphs is to try to show that the process must make consistent progress through configurations (subsets I of V) where the total curing rate is significantly lower than the total infection rate. Such progress must then be a low-probability event, implying an exponential lower bound on the time to extinction. Unfortunately, it is not clear whether this line of argument, based only on the instantaneous infection rates (namely, the "cuts" that are encountered — see Section sec:combinatorics for precise definitions) is possible for general graphs. Indeed, the lower bounds in [6], for more general regimes, rely on a much more sophisticated argument.

The main technical contribution in this paper is to show that in the regime examined (large CutWidth), the above outlined simple approach to deriving lower bounds is successful. Based on some nontrivial combinatorial properties

of the CutWidth, and the related concept of the *resilience* of a set of nodes, we will show that there is a sizeable part of the configuration space in which $|I_t|$ has a strong upward drift.

III. GRAPH THEORETIC PRELIMINARIES

In this section we consider a deterministic version of the problem in which we start with all nodes infected, and then cure them one at a time, deterministically, in a way that minimizes the maximum cut encountered in the process. This problem has been studied in various forms in the literature [9], [11] and its optimal value is called the *CutWidth* of a graph. Our analysis will be based on a related quantity, the *resilience* of a set of nodes, defined in a similar fashion as the CutWidth, except that we start with a given arbitrary set of nodes, and we are also allowed to deterministically infect nodes in the process. We study the properties of the resilience of a set of nodes and relate it to the size of the subset (Section III-B). Then, in Section III-C, we relate the resilience and the cut associated to a set of nodes, which is the key combinatorial result behind our analysis.

A. CutWidth and resilience

We first introduce some terminology: "*bags*" and "*crusades*."

Definition 1: A **bag** is a subset $A \subseteq V$ of the set of nodes V. We denote by $|A|$ the number of nodes in A.

We introduce two common operations on a bag A and we write

$$A + v = A \cup \{v\}, \qquad \text{for any } v \notin A,$$

and

$$A - v = A \setminus \{v\}, \qquad \text{for any } v \in A.$$

We also use the notations $A \setminus B = \{v \in A : v \notin B\}$ $A \triangle B = (A \setminus B) \cup (B \setminus A)$.

We next define the concept of a crusade. A crusade from A to B is a sequence of bags that starts from A and ends at B with the restriction that at every step of this sequence arbitrarily many nodes may be added to the previous bag, but at most one can be removed.

Definition 2: For any two bags A and B, an $(A - B)$-**crusade** ω is a sequence $(\omega_0, \omega_1, \ldots, \omega_k)$ of bags, of length $|\omega| = k + 1$, with the following properties:

(i) $\omega_0 = A$,

(ii) $\omega_k = B$, and

(iii) $|\omega_i \setminus \omega_{i+1}| \leq 1$, for $i \in \{0, \ldots, k-1\}$.

Property (iii) states that at every step of a crusade arbitrarily many nodes can be added to the current bag but *at most one* node may be removed from the bag. Note that the definition of a crusade allows *non-monotone* moves, since a bag at any step can be a subset, a superset or not comparable to the preceding bag. We denote by $\Omega(A - B)$ the set of all $(A - B)$-crusades.

We also consider a special case of crusades, the **monotone crusades** for which only removal of nodes is allowed at each step. Specifically, for any two bags A and B, $A, B \subseteq V$,

an **($A \downarrow B$)-monotone crusade** ω is an ($A - B$)-crusade $(\omega_0, \omega_1, \ldots, \omega_k)$ with the additional property: $\omega_i \supseteq \omega_{i+1}$, for $i \in \{0, \ldots, k-1\}$. We denote by $\Omega(A \downarrow B)$ the set of all ($A \downarrow B$)-crusades.

The number of edges connecting a bag A with its complement is called the cut of the bag. Its importance comes from the fact that it is equal to the total rate at which new infections occur, when the set of currently infected nodes is A.

Definition 3: For any bag A, its **cut**, $c(A)$, is defined as the cardinality of the set of edges

$$\big\{(u, v) : u \in A, \ v \in A^c\big\}.$$

In Proposition 1 below, we record, without proof, four elementary properties of cuts.

Proposition 1: For any two bags A and B, we have

(i) $c(A \cup B) \le c(A) + c(B) \le c(A) + \Delta \cdot |B|$.
(ii) If $A \subseteq B$, then for any $v \in A$,

$$c(A - v) - c(A) \le c(B - v) - c(B).$$

(iii) $c(A) \le \min\{|A| \cdot \Delta, (n - |A|) \cdot \Delta\}$
(iv) $|c(A) - c(B)| \le \Delta |A \triangle B|$.

Note that Proposition 1(ii) states the well-known submodularity property of the function $c(\cdot)$.

We define the width of a crusade ω as the maximum cut encountered during the crusade. Intuitively, this is the largest infection rate to be encountered if the nodes were to be cured deterministically according to the sequence prescribed by the crusade (assuming no new infections happen in between) .

Definition 4: The **width** $z(\omega)$ of an ($A - B$)-crusade $\omega = (\omega_0, \ldots, \omega_k)$ is defined by

$$z(\omega) = \max_{i \in \{1, \ldots, k\}} \{c(\omega_i)\}.$$

Note that in the definition above, the maximization starts after the first step of the crusade, i.e., we exclude the first bag ω_0 from the maximization.

We now formally define the CutWidth of a graph G as the minimum over all monotone crusades from V to the empty set of the corresponding crusade width.

Definition 5: For any given graph G, its **CutWidth** W is given by

$$W = \min_{\omega \in \Omega(V \downarrow \emptyset)} z(\omega). \tag{3}$$

Intuitively, this metric indicates the maximum cut that is encountered after the first step during an "optimal" monotone crusade that clears the graph.

The largest possible value of a cut, for graphs with maximum degree Δ, is $n\Delta/2$, and therefore the CutWidth is also upper bounded by $n\Delta/2$. For notational convenience, we define

$$E = \frac{2}{\Delta}\left(\frac{(n+2)\Delta}{2} - W\right), \tag{4}$$

and observe that $E \ge 2$. Note that "small" values of E indicate that the CutWidth is not too far from the largest possible value, $n\Delta/2$. In Section III we relate E to cuts and show that when E is small, then bags with large resilience, as defined below, also have a large cut.

Definition 6: The **resilience**, $\gamma(A)$, of a bag A is defined by

$$\gamma(A) = \min_{\omega \in \Omega(A - \emptyset)} z(\omega). \tag{5}$$

We denote by $\Omega^A \subseteq \Omega(A - \emptyset)$ the set of crusades ω that attain the minimum, i.e.,

$$z(\omega) = \gamma(A).$$

Crusades in Ω^A are referred to as A-optimal, or simply as optimal, when the set A is clear from the context.

Intuitively, $\gamma(A)$ indicates the maximum cut that is encountered, after the first step, during a crusade that clears a bag A. Note that in contrast to the definition of the CutWidth of a graph, the minimization, in the definition of resilience, is over all crusades, not just monotone crusades.

It can be seen that the resilience of a bag A satisfies the Bellman equation:

$$\gamma(A) = \min_{|A \setminus B| \le 1} \{\max\{c(B), \gamma(B)\}\}. \tag{6}$$

B. Properties of the resilience

This section explores some properties of the resilience. We first prove (Lemma 1(i)) that if A and B are two bags with $A \subseteq B$, then $\gamma(A) \le \gamma(B)$. Intuitively, one can construct a crusade from A to $\emptyset$ as follows: The crusade starts from A, then continues to the first bag encountered by a B-optimal crusade ω^B, and then follows ω^B. The constructed crusade and ω^B are the same except for the initial bag of the crusade. By the definition of the resilience, the initial bag does not affect the maximization and thus the width of the new crusade is equal to $\gamma(B)$.

We also prove (Lemma 1 (ii)) that if two bags differ by only one node v, then the corresponding resiliences are at most Δ apart. Intuitively, one can attach node v to the optimal crusade for the smaller of the two bags and achieve a maximum cut which is at most Δ different from the original.

Lemma 1: Let A and B be two bags.

(i) **[Monotonicity]** If $A \subseteq B$, then $\gamma(A) \le \gamma(B)$.
(ii) **[Smoothness]** If $B = A + v$, then $\gamma(B) \le \gamma(A) + \Delta$.
Proof:
(i) Let $\omega^B = (\omega_0^B, \ldots, \omega_k^B) \in \Omega^B$. Consider the sequence $\hat{\omega} = (\hat{\omega}_0, \ldots, \hat{\omega}_k)$ of bags for which $\hat{\omega}_0 = A$, and $\hat{\omega}_i = \omega_i^B$, for $i = 1, \ldots, k$. We claim that $\hat{\omega}$ is a crusade $\hat{\omega} \in \Omega(A - \emptyset)$. Indeed,

a) $\hat{\omega}_0 = A$;
b) $\hat{\omega}_k = \omega_k^B = \emptyset$
c) $|\hat{\omega}_0 \setminus \hat{\omega}_1| = |A \setminus \hat{\omega}_1| \le |B \setminus \omega_1^B| = |\omega_0^B \setminus \omega_1^B| \le 1$, where the first inequality follows from $A \subseteq B$ and $\hat{\omega}_1 = \omega_1^B$. Moreover, for $i = 1, \ldots, k$, we have $|\hat{\omega}_i \setminus \hat{\omega}_{i+1}| = |\omega_i^B \setminus \omega_{i+1}^B| \le 1$.

Clearly,

$$z(\hat{\omega}) = \max_{i \in \{1, \ldots, k\}} \{c(\hat{\omega}_i)\} = \max_{i \in \{1, \ldots, k\}} \{c(\omega_i^B)\} = \gamma(B).$$

Concluding, by the definition of $\gamma(A)$,

$$\gamma(A) = \min_{\omega \in \Omega(A - \emptyset)} z(\omega) \le z(\hat{\omega}) = \gamma(B).$$

(ii) Let $\omega^A = (\omega_0^A, \ldots, \omega_k^A) \in \Omega^A$. Consider the sequence $\hat{\omega} = (\hat{\omega}_0, \ldots, \hat{\omega}_{k+1})$ of bags with $\hat{\omega}_i = \omega_i^A \cup \{v\}$, for $i = 0, \ldots, k$, and $\hat{\omega}_{k+1} = \emptyset$. Clearly, $\hat{\omega}$ is a crusade $\hat{\omega} \in \Omega(B - \emptyset)$ for which

$$\gamma(B) \le z(\hat{\omega}) \le \max_{i \in \{1,\ldots,k\}} \{c(\omega_i^A)\} + \Delta = \gamma(A) + \Delta,$$

since the addition of one node can change the cut at each stage of the crusade by at most Δ (Proposition 1(i)). ∎

So far the notion of the resilience of a bag has not been related to the infection process under consideration. Recall that for the infection process defined in Section II, the number of infected nodes tends to increase at a rate equal to the cut of I_t. Therefore, in order to study the evolution of the set of infected nodes I_t through $\gamma(I_t)$, we relate $\gamma(I_t)$ to $c(I_t)$. To this end, we define a special class of bags, called *improvement* bags.

Definition 7: Let $\mathcal{C} = \{A \subseteq V : \text{there exists } v \in A \text{ for which } \gamma(A - v) < \gamma(A)\}$. Any bag in $\mathcal{C}$ is called an **improvement** bag.

Improvement bags have the important property that their cut can be approximately lower bounded by their resilience. Consequently, whenever an improvement bag with high resilience is encountered, the infection rate is also high. This observation will play a central role in subsequent sections.

Lemma 2: For any improvement bag $A \in \mathcal{C}$,

$$c(A) \ge \gamma(A) - \Delta.$$

Proof: Let v be such that $\gamma(A - v) < \gamma(A)$ and let $B = A - v$. Since $|A \setminus B| = 1$, Eq. (6) yields

$$\gamma(A) \le \max\{c(B), \gamma(B)\}. \tag{7}$$

Having assumed that $\gamma(B) < \gamma(A)$, Eq. (7) implies that $\gamma(B) < c(B)$. Hence, $\gamma(A) \le c(B) \le c(A) + \Delta$, where the last inequality follows from Proposition 1(i). ∎

The next step is to characterize the initial resilience $\gamma(I_0)$. The following, highly nontrivial theorem from [1] provides an answer when $I_0 = V$.

Theorem 1: For any graph, $\gamma(V) = W$

Theorem 1 implies that finding a minimum-width $(V - \emptyset)$-crusade is equivalent to finding a minimum-width $(V \downarrow \emptyset)$-crusade. This property is not true in general for a general initial set I_0; however, a related poperty will be established in Lemma 5.

Next, we explore the connection between the size of a bag and its resilience. We first obtain a bound on $\gamma(A)$ by considering a crusade which removes all nodes of A, one at a time, in an arbitrary order. We then obtain a related bound on W by constructing a crusade in $\Omega(V - \emptyset)$ that removes the nodes of the complement of A, one at a time, and then uses an A-optimal crusade ω^A. These two observations imply certain constraints (an "admissible region") for the pair $(\gamma(A), |A|)$ on the two dimensional plane, which are illustrated in Figure 1. Finally, using the properties of the function $\gamma(\cdot)$ that have been established so far, we obtain a
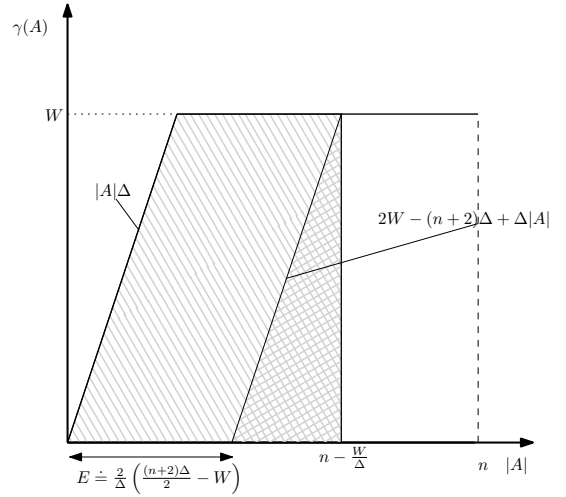

Fig. 1: Admissible region for the pair $(\gamma(A), |A|)$. If $\gamma(A) < W$, Lemma 3 implies that $(\gamma(A), |A|)$ belongs to the parallelogram shown in the figure. On the other hand, there is no restriction on the size $|A|$ of bags with $\gamma(A) = W$, and so the admissible region also includes the horizontal line segment at the top of the figure.

refinement of the admissible region, which is again illustrated in Figure 1.

Lemma 3: Consider a graph with $W \ge \Delta$ and a bag A. Let E be as defined in Eq. (4).

(i) $\gamma(A) \le |A|\Delta$;
(ii) If $\gamma(A) < W$, then $W \le (n - |A|)\Delta$;
(iii) If $\gamma(A) < W$, then $\gamma(A) \ge \Delta(|A| - E)$.

Proof:

(i) Consider some enumeration $(a_1, a_2, \ldots, a_{|A|})$ of the nodes of A. We construct a crusade $\hat{\omega} \in \Omega(A - \emptyset)$ by letting $\hat{\omega}_0 = A$, and $\omega_i = \omega_{i-1} \setminus \{a_i\}$ for $i = 1, \ldots, |A|$. By Proposition 1 (i), the maximum cut encountered by $\hat{\omega}$ is bounded by $|A|\Delta$. Therefore,

$$\gamma(A) \le z(\hat{\omega}) \le |A|\Delta.$$

(ii) Consider some enumeration $(a_1^c, a_2^c, \ldots, a_{n-|A|}^c)$ of the nodes of A^c, the complement of A. Let $\omega^A = (\omega_0^A, \ldots, \omega_k^A) \in \Omega^A$. We construct a crusade $\omega \in \Omega(V - \emptyset)$ by letting $\omega_0 = V$, $\omega_i = \omega_{i-1} \setminus \{a_i^c\}$ for $i = 1, \ldots, n - |A|$, and $\omega_i = \omega_{i-n+|A|}^A$ for $i = n - |A| + 1, \ldots, k + n - |A|$. Then,

$$W = \gamma(V) \le z(\omega) = \max\{\gamma(A), \max_{i \in \{1,\ldots,n-|A|\}} c(\omega_i)\} \tag{8}$$

$$\le \max\{\gamma(A), (n - |A|)\Delta\}.$$

The first equality above follows from Theorem 1; the second equality follows from the construction of ω; the last inequality follows from Proposition 1(i). Using the

assumption $\gamma(A) < W$, Eq. (8) implies that

$$W \le (n - |A|)\Delta.$$

(iii) Consider some bag A for which $\gamma(A) < W$. From (ii),

$$|A| \le n - W/\Delta. \qquad (9)$$

Let $C \subseteq V \setminus A$ be some nonempty bag with $|C| = n - \lfloor W/\Delta \rfloor - |A| + 1$. Note that Eq. (9) implies that

$$|C| = n - \lfloor W/\Delta \rfloor - |A| + 1 \ge n - W/\Delta - |A| + 1 \ge 1,$$

and that the assumption $W \ge \Delta$ implies that

$$|C| = n - \lfloor W/\Delta \rfloor - |A| + 1 \le n - |A|.$$

This shows the existence of a bag with the desired properties exists.

We define $F = A \cup C$. Note that

$$\begin{aligned} |F| = |A| + |C| &\ge n - \lfloor W/\Delta \rfloor + 1 \\ &> n - \lfloor W/\Delta \rfloor \ge n - W/\Delta, \end{aligned}$$

and thus

$$W > (n - |F|)\Delta.$$

Then, part (ii) of the Lemma implies that $\gamma(F) = W$. The resilience of F satisfies

$$W = \gamma(F) = \gamma(A \cup C) \le \gamma(A) + |C|\Delta,$$

where the inequality follows from applying Lemma 1(ii) $|C|$ times. Therefore,

$$\begin{aligned} \gamma(A) &\ge W - |C|\Delta \\ &= W - (n - \lfloor W/\Delta \rfloor - |A| + 1)\Delta \\ &\ge W - (n - W/\Delta + 1 - |A| + 1)\Delta \\ &\ge 2W - (n+2)\Delta + \Delta|A| \\ &= \Delta(|A| - E), \end{aligned}$$

which concludes the proof. ∎

C. Characterization of optimal crusades and some implications

In this section we prove that when E is small, i.e., when the CutWidth is close to the largest possible value, then bags with large resilience also have large cuts.

Lemma 4: Suppose that $W \ge \Delta$ and that the bag A satisfies $0 < \gamma(A) < W$. Then,

$$c(A) \ge \gamma(A) - 2(E+2)\Delta.$$

The rest of the section is devoted to proving this property. We start with a characterization of optimal crusades for a given bag A. Specifically, Lemma 5 states that for any bag A, there exists an optimal crusade which: (i) can add nodes, and potentially remove one node at the first step; (ii) cannot add nodes (i.e., is monotone) after the first step (parts (i)-(ii)). Moreover, we argue that except for trivial cases, an improvement bag must be encountered before the end of the crusade (part (vi)). These properties allow us to make a connection between resilience and cuts.

Lemma 5: For any nonempty bag A with $\gamma(A) > 0$, there exists a crusade $\hat{\omega} = (\hat{\omega}_0, \hat{\omega}_1, \ldots, \hat{\omega}_k) \in \Omega^A$ with the following properties:

(i) For $i \in \{1, \ldots, k\}$, $\hat{\omega}_i \ne \hat{\omega}_{i-1}$
(ii) For $i \in \{2, \ldots, k\}$, $\hat{\omega}_i \subset \hat{\omega}_{i-1}$.
(iii) For $i \in \{0, \ldots, k\}$, $\gamma(\hat{\omega}_i) \le \gamma(A)$.
(iv) $\gamma(\hat{\omega}_1) \ge \gamma(A) - \Delta$.
(v) $c(A) \ge c(\hat{\omega}_1) - \Delta(E+2)$.
(vi) Let $l = \min\{i \ge 0 : \hat{\omega}_i \in \mathcal{C}\}$. Then, $l < \infty$.

Proof: We assign to every $(A - \emptyset)$-crusade $\omega \in \Omega^A$ a value $P(\omega) = \left(\sum_{i=0}^{|\omega|-1}(c(\omega_i) + 1), \sum_{i=0}^{|\omega|-1}|\omega_i|\right)$. Let $\hat{\omega} \in \operatorname{argmin}_{\omega \in \Omega^A} P(\omega)$, where the minimum is taken with respect to the lexicographic ordering.

(i) We first prove that for all $i \in \{0, \ldots, k-1\}$,

$$\hat{\omega}_i \ne \hat{\omega}_{i+1}. \qquad (10)$$

For the purposes of contradiction, assume that for some $q \in \{1, \ldots, k-1\}$, $\hat{\omega}_q = \hat{\omega}_{q+1}$, and construct a crusade $\tilde{\omega} = (\tilde{\omega}_0, \ldots, \tilde{\omega}_{k-1})$ by setting $\tilde{\omega}_i = \hat{\omega}_i$ for all $i \le q$, and $\tilde{\omega}_i = \hat{\omega}_{i+1}$ for $i = q+1, \ldots, k-1$.
Clearly, $\tilde{\omega} = (\tilde{\omega}_0, \ldots, \tilde{\omega}_{k-1})$ is a crusade, i.e., $\tilde{\omega} \in \Omega(A - \emptyset)$. Moreover, $\tilde{\omega} \in \Omega^A$, because $\max_{1 \le i \le k-1} c(\tilde{\omega}_i) = z(\hat{\omega}) = \gamma(A)$. But $\sum_{i=0}^{k-1}(c(\tilde{\omega}_i) + 1) < \sum_{i=0}^{k}(c(\hat{\omega}_i) + 1)$, which implies that $P(\tilde{\omega}) < P(\hat{\omega})$, and contradicts the minimality of $\hat{\omega}$.

(ii) The idea of the proof of this property is borrowed from [1], and is based on the submodularity of $c(\cdot)$. We first argue that for all $i \in \{1, \ldots, k-1\}$,

$$c(\hat{\omega}_{i+1} \cup \hat{\omega}_i) \ge c(\hat{\omega}_i). \qquad (11)$$

For the purposes of contradiction, assume that there exists some $q \in \{1, \ldots, k-1\}$ such that

$$c(\hat{\omega}_{q+1} \cup \hat{\omega}_q) < c(\hat{\omega}_q), \qquad (12)$$

and construct the sequence of bags $\tilde{\omega} = (\tilde{\omega}_0, \ldots, \tilde{\omega}_k)$, by setting $\tilde{\omega}_i = \hat{\omega}_i$ for all $i \ne q$ and $\tilde{\omega}_q = \hat{\omega}_{q+1} \cup \hat{\omega}_q$. We first claim that $\tilde{\omega}$ is a crusade, i.e., $\tilde{\omega} \in \Omega(A - \emptyset)$. Indeed, since $\hat{\omega}$ is a crusade, we get $|\hat{\omega}_q \setminus \hat{\omega}_{q+1}| \le 1$ and $|\hat{\omega}_{q-1} \setminus \hat{\omega}_q| \le 1$. Therefore,

$$|\tilde{\omega}_{q-1} \setminus \tilde{\omega}_q| = |\hat{\omega}_{q-1} \setminus (\hat{\omega}_{q+1} \cup \hat{\omega}_q)| \le |\hat{\omega}_{q-1} - \hat{\omega}_q| \le 1,$$

where the first equality follows from the construction of $\tilde{\omega}$ and the second inequality from $\hat{\omega}_{q+1} \cup \hat{\omega}_q$. Furthermore,

$$|\tilde{\omega}_q \setminus \tilde{\omega}_{q+1}| = |(\hat{\omega}_{q+1} \cup \hat{\omega}_q) \setminus \hat{\omega}_{q+1}| \le |\hat{\omega}_q - \hat{\omega}_{q+1}| \le 1,$$

where the the first equality follows from the construction of $\tilde{\omega}$ and the second inequality from $\hat{\omega}_{q+1} \cup \hat{\omega}_q \supset \hat{\omega}_{q+1}$.
Moreover, we claim that $\tilde{\omega} \in \Omega^A$. Indeed

$$\begin{aligned} \max_{1 \le i \le k} c(\tilde{\omega}_i) &= \max\{c(\tilde{\omega}_q), \max_{1 \le i \le k, i \ne q} c(\hat{\omega}_i)\} \\ &\le \max_{1 \le i \le k} c(\hat{\omega}_i) = \gamma(A), \end{aligned}$$

where the inequality follows from (12).

On the other hand, it follows from (12) that $\sum_{i=0}^{k}(c(\tilde{\omega}_i)+1) < \sum_{i=0}^{k}(c(\hat{\omega}_i)+1)$ and thus $P(\tilde{\omega}) < P(\hat{\omega})$, which contradicts the minimality of $\hat{\omega}$. We have therefore established (11).

Using the submodularity of the cut as well as Eq. (11), we have that for all $i \in \{1, \ldots, k-1\}$,

$$c(\hat{\omega}_{i+1} \cap \hat{\omega}_i) \le c(\hat{\omega}_{i+1}). \tag{13}$$

We now prove that $|\hat{\omega}_{i+1} \cap \hat{\omega}_i| \ge |\hat{\omega}_{i+1}|$ for all $i \in \{1, \ldots, k-1\}$.

For the purposes of contradiction, assume that there exists some $q \in \{1, \ldots, k-1\}$ such that

$$|\hat{\omega}_{q+1} \cap \hat{\omega}_q| < |\hat{\omega}_{q+1}|. \tag{14}$$

Construct the sequence $\tilde{\omega}_i = \hat{\omega}_i$ for all $i \ne q+1$ and $\tilde{\omega}_{q+1} = \hat{\omega}_{q+1} \cap \hat{\omega}_q$.

We first claim that $\tilde{\omega}$ is a crusade, i.e. $\tilde{\omega} \in \Omega(A - \emptyset)$. Indeed, since $\hat{\omega}$ is a crusade we get $|\hat{\omega}_q \setminus \hat{\omega}_{q+1}| \le 1$ and $|\hat{\omega}_{q+1} \setminus \hat{\omega}_{q+2}| \le 1$. Therefore,

$$|\tilde{\omega}_q \setminus \tilde{\omega}_{q+1}| = |\hat{\omega}_q \setminus (\hat{\omega}_{q+1} \cap \hat{\omega}_q)| = |\hat{\omega}_q - \hat{\omega}_{q+1}| \le 1,$$

where the the first equality follows from the construction of $\tilde{\omega}$ and the second inequality from $\hat{\omega}_{q+1} \cap \hat{\omega}_q \subset \hat{\omega}_{q+1}$. Furthermore,

$$\begin{aligned}|\tilde{\omega}_{q+1} \setminus \tilde{\omega}_{q+2}| &= |(\hat{\omega}_{q+1} \cap \hat{\omega}_q) \setminus \hat{\omega}_{q+2}| \\ &\le |\hat{\omega}_{q+1} - \hat{\omega}_{q+2}| \le 1,\end{aligned}$$

where the the first equality follows from the construction of $\tilde{\omega}$ and the second inequality from $\hat{\omega}_{q+1} \cap \hat{\omega}_q \subset \hat{\omega}_{q+1}$. Moreover, we claim that $\tilde{\omega} \in \Omega^A$. Indeed

$$\begin{aligned}\max_{1 \le i \le k} c(\tilde{\omega}_i) &= \max\{c(\tilde{\omega}_q), \max_{1 \le i \le k, i \ne q+1} c(\hat{\omega}_{q+1})\} \\ &\le \max_{1 \le i \le k} c(\hat{\omega}_i) = \gamma(A),\end{aligned}$$

where the inequality follows from (13).

On the other hand, it follows from (13) that $\sum_{i=0}^{k}(c(\tilde{\omega}_i)+1) \le \sum_{i=0}^{k}(c(\hat{\omega}_i)+1)$ and from (14) that $\sum_{i=0}^{k}|\omega_i| < \sum_{i=0}^{k}|\tilde{\omega}_i|$. Therefore, $P(\tilde{\omega}) < P(\hat{\omega})$, which contradicts the minimality of $\hat{\omega}$.

Therefore we established that $|\hat{\omega}_{i+1} \cap \hat{\omega}_i| \ge |\hat{\omega}_{i+1}|$ for all $i \in \{i, \ldots, k-1\}$. The latter implies that for all $i \in \{1, \ldots, k-1\}$, $\hat{\omega}_{i+1} \subseteq \hat{\omega}_i$. Using part (i) of the lemma, it follows that that for $i \in \{1, \ldots, k-1\}$, $\hat{\omega}_{i+1} \subset \hat{\omega}_i$.

(iii) We prove the result by induction. First, observe that $\gamma(\hat{\omega}_0) = \gamma(A)$. Assume that $\gamma(A) \ge \gamma(\hat{\omega}_i)$. Moreover, by (6), for all $i \in \{1, \ldots, k-1\}$, $\gamma(\hat{\omega}_i) = \max\{\gamma(\hat{\omega}_{i+1}), c(\hat{\omega}_{i+1})\} \ge \gamma(\hat{\omega}_{i+1})$. Therefore, $\gamma(\hat{\omega}_{i+1}) \le \gamma(A)$.

(iv) We consider two cases. Assume that $\hat{\omega}_1 \supset A$. Then $\gamma(\hat{\omega}_1) \ge \gamma(A) \ge \gamma(A) - \Delta$. Otherwise, by the definition of a crusade we get $|A \setminus \hat{\omega}_1| \le 1$. Therefore, we can write $\hat{\omega}_1 = A \cup D - v$, for some set D (disjoint from A) and some $v \in A$. Using Lemma 1(ii), and then

Lemma 1(i), we obtain

$$\gamma(\hat{\omega}_1) \ge \gamma(A \cup D) - \Delta \ge \gamma(A) - \Delta.$$

(v) From (iii) we obtain $\gamma(\hat{\omega}_1) \le \gamma(A)$. Therefore, using Lemma 3(iii), we conclude that

$$|\hat{\omega}_1| \le \frac{\gamma(A)}{\Delta} + E. \tag{15}$$

Moreover, by Lemma 3(i), we get

$$|A| \ge \frac{\gamma(A)}{\Delta}. \tag{16}$$

We consider two cases. Assume that $\omega_1 \supset A$. Then,

$$|\hat{\omega}_1 \triangle A| = |\hat{\omega}_1| - |A|.$$

Otherwise, by the definition of a crusade we get $|A \setminus \hat{\omega}_1| \le 1$. Therefore, we can write $\hat{\omega}_1 = A \cup D - v$, where D is disjoint from A and $v \in A$. Thus,

$$|\hat{\omega}_1 \triangle A| = |D| + 1 = |\hat{\omega}_1| - |A| + 2.$$

Therefore, in both cases,

$$|\hat{\omega}_1 \triangle A| \le |\hat{\omega}_1| - |A| + 2.$$

We then use Eqs. (15) and (16) to obtain

$$|\hat{\omega}_1 \triangle A| \le E + 2.$$

The result follows by applying Proposition 1(iv).

(vi) Note that $\gamma(\hat{\omega}_k) = \gamma(\emptyset) = 0$. Note also that any single-element set B satisfies $\gamma(B) = 0$. Suppose that $\gamma(\hat{\omega}_1) > 0$. Then, there exists some $i \in \{1, \ldots, k-1\}$ such that $\gamma(\hat{\omega}_{i+1}) < \gamma(\hat{\omega}_i)$, and $\hat{\omega}_i$ is an improvement bag.

Suppose now that $\gamma(\hat{\omega}_1) = 0$. If $\hat{\omega}_1 = \emptyset$, then $\gamma(A) = 0$, which contradicts the assumption $\gamma(A) > 0$. If $\hat{\omega}_1$ is nonempty, then we must have $c(\hat{\omega}_2) = 0$, so that $\hat{\omega}_2$ is empty and $\hat{\omega}_1$ is a singleton. Since $\gamma(A) > 0$, the set A is not a singleton. Since at most one element can be removed in going from A to $\hat{\omega}_1$, it follows that A consists of two elements and that a single element was removed from A. In that case, $A = \hat{\omega}_0$ is an improvement bag.

In both cases, we see that there exists some i for which $\hat{\omega}_i$ is an improvement bag and therefore l is well-defined and finite. ∎

Proof: (of Lemma 4) Consider a crusade $\hat{\omega}_A$ with the properties in Lemma 5, and let $l \ge 0$ be such that $B = \hat{\omega}_l^A$ is the fist improvement bag encountered.

From Lemma 5(vi), l is well-defined and finite. We consider three cases:

(i) $l = 0$: If A is itself an improvement bag, then from Lemma 2, $c(A) \ge \gamma(A) - \Delta$.

(ii) $l = 1$: In this case, $\hat{\omega}_1$ is an improvement bag. From Lemma 2, $c(\hat{\omega}_1) \ge \gamma(\hat{\omega}_1) - \Delta$. Then, from Lemma 5(iv), we obtain

$$c(\hat{\omega}_1) \ge \gamma(A) - 2\Delta.$$

Moreover, from Lemma 5(v), we get

$$c(A) \geq c(\hat{\omega}_1) - (E+2)\Delta \geq \gamma(A) - (E+4)\Delta.$$

(iii) $l > 2$: In this case, by property (ii) in Lemma 5(ii), it follows that $B \subset \hat{\omega}_1$ and

$$|B \triangle \hat{\omega}_1| = |\hat{\omega}_1| - |B|.$$

Moreover, since B is the first improvement bag that is encountered, $\gamma(B) = \gamma(\hat{\omega}_1)$. We use Lemma 3(i) to obtain

$$|B| \geq \gamma(B)/\Delta = \gamma(\hat{\omega}_1)/\Delta,$$

and Lemma 3(iii) to obtain

$$|\hat{\omega}_1| \leq \gamma(\hat{\omega}_1)/\Delta + E.$$

Combining the above,

$$|B \triangle \hat{\omega}_1| = |\hat{\omega}_1| - |B| \leq E,$$

from which we conclude that

$$c(\hat{\omega}_1) \geq c(B) - E\Delta \geq \gamma(\hat{\omega}_1) - (E+1)\Delta.$$

where the first inequality follows from Proposition 1(iv) and the second from the fact that B is an improvement bag and $\gamma(B) = \gamma(\omega_1)$. Therefore, from Lemma 5(v), we obtain

$$c(A) \geq \gamma(\omega_1) - (2E+3)\Delta.$$

Finally, using Lemma 5(iv), we conclude that

$$c(A) \geq \gamma(A) - 2(E+2)\Delta,$$

which completes the proof of Lemma 4.

The combinatorial properties of the resilience derived in this section will be used next to obtain a lower bound on the expected extinction time, in the regime where $\gamma(I_0) \gg r$.

IV. EXPONENTIAL LOWER BOUND

In this section we state and prove our main result. Specifically, we use Lemma 4 to argue that the process must traverse a region in which the number $|I_t|$ of infected nodes has an upward drift, which in turn leads to the desired lower bound.

Theorem 2: Suppose that $\gamma(I_0) \geq \Delta(9E+12)+3r$. Then,

$$\mathbb{E}_{I_0}[\tau] \geq \frac{1}{r}\left(\left(\frac{\gamma(I_0) - (9E+12)\Delta}{3r}\right)^{\frac{\gamma(I_0)}{3\Delta}-1} - 1\right).$$

Proof: We define a process V_t which is coupled with the process I_t as follows.

$$V_t = \begin{cases} |I_t|, & \text{if } |I_t| \leq \left\lfloor \frac{2\gamma(I_0)}{3\Delta} \right\rfloor, \\ \left\lfloor \frac{2\gamma(I_0)}{3\Delta} \right\rfloor, & \text{if } |I_t| > \left\lfloor \frac{2\gamma(I_0)}{3\Delta} \right\rfloor. \end{cases} \quad (17)$$

The dynamics of V_t are as follows. If $i < \lfloor 2\gamma(I_0)/3\Delta \rfloor$, then

$$V_t : i \to i+1, \text{ with rate } c(I_t),$$
$$V_t : i \to i-1, \text{ with rate } r.$$

Furthermore, if $i = \lfloor 2\gamma(I_0)/(3\Delta)) \rfloor$, then

$$V_t : i \to i-1, \text{ with rate } r(I_t),$$

where $r(I_t) \leq r$.

Consider the stopping time

$$\tau^* = \inf\left\{t \geq 0 : |I_t| \leq \left\lfloor \frac{\gamma(I_0)}{3\Delta} \right\rfloor\right\}$$
$$= \inf\left\{t \geq 0 : V_t \leq \left\lfloor \frac{\gamma(I_0)}{3\Delta} \right\rfloor\right\}.$$

For every sample path, $\tau^* \leq \tau$. Therefore,

$$\mathbb{E}_{I_0}[\tau] \geq \mathbb{E}_{I_0}[\tau^*].$$

Suppose now that $|I_t|$ satisfies

$$\frac{\gamma(I_0)}{3\Delta} \leq |I_t| \leq \frac{2\gamma(I_0)}{3\Delta}.$$

Then, by parts (i) and (iii) of Lemma 3, we obtain that

$$\frac{\gamma(I_0)}{3} - E\Delta \leq \gamma(I_t) \leq \frac{2\gamma(I_0)}{3\Delta}.$$

Furthermore, Lemma 4 implies that

$$c(I_t) \geq \gamma(I_t) - 2(E+2)\Delta \geq \frac{\gamma(I_0)}{3} - (3E+4)\Delta.$$

It follows that the process V_t stochastically dominates a process Y_t, described in the Appendix, with parameters $\lambda = r$, $\mu = \gamma(I_0)/3 - (3E+4)\Delta$, and $L = \gamma(I_0)/3\Delta$. Therefore, using Eq. (20) in the Appendix,

$$\mathbb{E}_{I_0}[\tau] \geq \frac{1}{r}\left(\left(\frac{\gamma(I_0) - (9E+12)\Delta}{3r}\right)^{\frac{\gamma(I_0)}{3\Delta}-1} - 1\right). \quad \blacksquare$$

Note that when $3r < \gamma(I_0) - 9E\Delta - 12\Delta$, the optimal expected extinction time scales exponentially in the resilience of the set of the initially infected nodes. In [5] and [7], we focused on the case where $I_0 = V$ (the worst case) and proved that if the CutWidth of the graph is a sublinear function of the number of nodes, and if $r = o(n)$, then, the expected extinction time is $o(n)$. In contrast, the following result considers the case where W scales linearly in the number of nodes and provides an exponential lower bound on the expected extinction time. Specifically, using Theorem 1 to replace $\gamma(V)$ by W, and using also the definition of E, we can write our lower bound as

$$\mathbb{E}_V[\tau] \geq \frac{1}{2r}\left(\left(\frac{19W - 9n\Delta - 24\Delta}{3r}\right)^{W/3\Delta-1} - 1\right),$$

and we obtain the corollary that follows.

Corollary 1: Fix a constant $C > 1$ and consider those graphs for which $W > C9/19\Delta$. Moreover, assume that $r = o(n)$. Then,

$$\mathbb{E}_V[\tau] = \Omega(2^n).$$

REFERENCES

[1] D. Bienstock and P. Seymour. Monotonicity in graph searching. *Journal of Algorithms*, 12(2):239 – 245, 1991.

[2] C. Borgs, J. Chayes, A. Ganesh, and A. Saberi. How to distribute antidote to control epidemics, 2010.
[3] F. R. K. Chung, P. Horn, and A. Tsiatas. Distributing antidote using pagerank vectors. *Internet Mathematics*, 6(2):237–254, 2009.
[4] R. Cohen, S. Havlin, and D. ben Avraham. Efficient immunization strategies for computer networks and populations. *Phys. Rev. Lett.*, 91:247901, 2003. http://journals.aps.org/prl/abstract/10.1103/PhysRevLett.91.247901.
[5] K. Drakopoulos, A. Ozdaglar, and J. N. Tsitsiklis. An efficient curing policy for epidemics on graphs. In *2014 IEEE 53rd Annual Conference on Decision and Control (CDC)*, pages 4447–4454, Dec 2014.
[6] K. Drakopoulos, A. Ozdaglar, and J. N. Tsitsiklis. Network resilience against epidemic spread. *Working paper*, 2014.
[7] K. Drakopoulos, A. Ozdaglar, and J. N. Tsitsiklis. An efficient curing policy for epidemics on graphs. *IEEE Transactions on Network Science and Engineering*, 1(2):67–75, January 2015.
[8] E. Gourdin, J. Omic, and P. V. Mieghem. Optimization of network protection against virus spread. pages 86–93, 2011.
[9] A. S. LaPaugh. Recontamination does not help to search a graph. *J. ACM*, 40(2):224–245, Apr. 1993.
[10] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2006.
[11] T. D. Parsons. The search number of a connected graph. In *In Proceedings of the Ninth Southeastern Conference on Combinatorics, Graph Theory, and Computing (Florida Atlantic Univ., Boca Raton, Fla., 1978), Congress. Numer., XXI*, pages 549–554, 1978.
[12] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. J. Pappas. Optimal vaccine allocation to control epidemic outbreaks in arbitrary networks. *CoRR*, abs/1303.3984, 2013.
[13] J. Steele. *Stochastic calculus and financial applications*. Springer, 2000.

APPENDIX

We first recall a basic result on the standard continuous random walk on the integers. Specifically, let Z_t denote the state of a Markov process with the following dynamics. Fix some integers M and L with $0 < M < L$, and let $\mathbb{P}_M$ be the probabilty measure that describes the process when initialized at $Z_0 = M$. Let $\tau_L \inf\{t : Z_t = 0 \text{ or } Z_t = L\}$ denote the first time that the process Z_t visits state 0 or L, which is a *stopping time*. Moreover, let $\hat\tau \inf\{t : Z_t = 0\}$ denote the first time that Z_t hits 0. The following result is standard; see, e.g., Section 2.1 of [10] or Section 2.3 of [13].

Lemma 6: Consider the process Z_t and the stopping times τ_L and $\hat\tau$. Then,

$$\mathbb{P}_M(Z_{\tau_L} = L) = \frac{1 - (\lambda/\mu)^M}{1 - (\lambda/\mu)^L},\qquad(18)$$

Consider now a related Markov process Y_t, whose transition rates are as follows:

$$Y_t : i \to i+1, \text{ with rate } \mu,$$
$$Y_t : i \to i-1, \text{ with rate } \lambda,\qquad(19)$$

for $i \in \{1, \dots, L-1\}$ while

$$Y_t : i \to i-1, \text{ with rate } \lambda,$$

for $i = L$ and

$$Y_t : i \to i+1, \text{ with rate } \mu,$$

for $i = 0$.

We are looking for a lower bound on the expected time that it takes for the process Y_t to hit 0 for the first time, assuming that it starts at $L-1$. Let p be the probability that Y_t hits level L before hitting 0 starting from state $L-1$,

which is given by Lemma 6, with $M = L-1$. We consider the case where $\lambda < \mu$, so that $p > 1/2$. Each time that the process is at state $L-1$, the process regenerates, and we have a new trial, which succeeds in hitting state 0 before state L, with the same probability $1-p$. Let N be the number of trials and note that its expected value is $1/(1-p)$. In between trials, there needs to be a transition from state L to state $L-1$, whose expected time is $1/\lambda$. Thus, the total expected time elapsed until state 0 is hit for the first time is $\mathbb{E}[N-1]/\lambda = p/(1-p)\lambda$. Using Lemma 6 and some straightforward algebra, we obtain that this expected time is at least as large as

$$\frac{1}{2}\left(\left(\frac{\mu}{\lambda}\right)^{L-1} - 1\right)\frac{1}{\lambda}.\qquad(20)$$