

Curing Epidemics on Networks Using a Polya Contagion Model

Mikhail Hayhoe¹, Fady Alajaji², *Senior Member, IEEE*, and Bahman Ghahesifard³, *Member, IEEE*

Abstract—We study the curing of epidemics of a network contagion, which is modelled using a variation of the classical Polya urn process that takes into account spatial infection among neighbouring nodes. We introduce several quantities for measuring the overall infection in the network and use them to formulate an optimal control problem for minimizing the average infection rate using limited curing resources. We prove the feasibility of this problem under high curing budgets by deriving conservative lower bounds on the amount of curing per node that turn our measures of network infection into supermartingales. We also provide a provably convergent gradient descent algorithm to find the allocation of curing under limited budgets. Motivated by the fact that this strategy is computationally expensive, we design a suit of heuristic methods that are locally implementable and nearly as effective. Extensive simulations run on large-scale networks demonstrate the effectiveness of our proposed strategies.

Index Terms—Polya contagion urn scheme, epidemics on networks, non-stationary stochastic processes, supermartingales, curing strategies, gradient descent, node centrality.

I. INTRODUCTION

IN THIS paper we examine the problem of curing an epidemic using a network contagion model adapted from the Polya process [2], [3]. Here an epidemic can represent a disease [4], a computer virus [5], the spread of an innovation, rumour or idea [6], or the dynamics of competing opinions in a social network [7].

Epidemics on networks have been intensively studied in recent years, see [8]–[10] and references therein and thereafter. Our model is similar to the well-known susceptible-infected-susceptible (SIS) compartmental infection model [11], in the sense that initially, all nodes may be healthy or infected and as the epidemic spreads, nodes that are infected can be cured to become healthy, but any healthy node may become infected at any time, regardless of whether they have been cured previously. In contrast to the SIS model, our model

is motivated by the classical Polya contagion process, which evolves by sampling from an urn containing a finite number of red and black balls [12]–[14]. In the network Polya contagion model, introduced in [2], each node of the underlying network is equipped with an individual urn; however, instead of sampling from these urns when generating its contagion process, each node has a “super urn”, created by combining the contents of its own urn with those of its neighbours’ urns. This adaptation captures the concept of spatial infection, since having infected neighbours increases the chance that an individual is infected in the future. The stochastic properties of this model, along with a comparison to the traditional discrete-time SIS model, were examined in [2], [3]. Comparatively, the network Polya contagion process runs at a microscopic level to better model the behavior of the contagion process mathematically. Indeed, with finite memory our process becomes Markovian and exhibits similar behavior to the discrete-time SIS model, but in the case of infinite memory our model exhibits non-Markovian characteristics which may be truer to some real-life epidemics. In both cases, the microscopic viewpoint allows us to capture the behavior of the process and study its stochastic properties. With the explicit formulae provided for the network Polya contagion process, the joint and marginal probabilities of infection can be calculated exactly for any time without mean-field approximations or other moment closure techniques, as is typical for SIS models [15].

In this work, we study the problem of controlling the spread of contagion under the network Polya contagion model. This problem has been extensively studied in many different settings, including both deterministic and stochastic compartmental models such as SIS and SIR [8], expansions thereupon [16], models based on real epidemic data [17], as well as cascade models, among others [11]. Different problem formulations which have been explored include link removal via immunization [18], optimization of curing resources [19], message-passing approaches [20], optimal control paradigms which balance performance and resource usage [21], [22], as well as approaches involving competing or synergistic contagions on the same network [23], [24]. In this paper, we propose various natural ways to measure the total infection in this model, and examine conditions under which these measures have limits as time grows without bound. Using these measures, we pose an optimal control problem within the context of our model. As our first contribution, we characterize lower bounds on the allocation of curing to individuals which turn these infection measures into supermartingales.

Manuscript received November 2, 2017; revised November 18, 2018; accepted August 23, 2019; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor W. Wang. Date of publication September 25, 2019; date of current version October 15, 2019. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada. Parts of this work were presented at the 2018 American Control Conference [1]. (*Corresponding author: Mikhail Hayhoe.*)

M. Hayhoe is with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: mhayhoe@seas.upenn.edu).

F. Alajaji and B. Ghahesifard are with the Department of Mathematics and Statistics, Queen’s University, Kingston, ON K7L 3N6, Canada (e-mail: fa@queensu.ca; bahman.ghahesifard@queensu.ca).

This article has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors.

Digital Object Identifier 10.1109/TNET.2019.2940888

Our results provide conservative feasibility strategies for curing network epidemics. We next focus on realistic scenarios, where the curing budget is constrained. As our next contribution, we prove the convergence of a constrained gradient flow method for this problem and employ it to find near-optimal strategies under a fixed curing budget at each time step. In spite of its effectiveness, as we demonstrate, the gradient flow strategy is computationally expensive and only implementable in a centralized manner. These shortcomings motivate us to explore alternative strategies, considering notions of node centrality of the underlying network along with the composition of super urns at each time step. These strategies are less expensive computationally and may be carried out in a decentralized manner. Through extensive simulation results, we show that our proposed heuristic strategies perform well in curing epidemics.

The rest of the paper is organized as follows. Section II outlines some mathematical preliminaries that will be used throughout the paper. Section III contains the description of our network contagion process and the problem statement. Section IV discusses analytical results pertaining to the control of epidemics using this model, and Section V outlines strategies used to solve the problems posed. Section VI includes several simulation results. Finally, Section VII summarizes our contributions and proposes avenues for future work.

II. PRELIMINARIES

For a sequence $v_i = (v_{i,1}, \dots, v_{i,n})$, we use the notation $v_{i,s}^t$ with $1 \leq s < t \leq n$ to denote the vector $(v_{i,s}, v_{i,s+1}, \dots, v_{i,t})$, with $v_{i,0}^t = v_i^t$. Our technical results rely on notions from stochastic processes, some of which we recall here. Throughout, we assume that the reader is familiar with basic notions of probability theory.

Let (Ω, \mathcal{F}, P) be a probability space, and consider the stochastic process $\{Z_n\}_{n=1}^\infty$, where each Z_n is a random variable on Ω . We often refer to the indices of the process as “time” indices. We recall that the process $\{Z_n\}_{n=1}^\infty$ is *stationary* if for any $n \in \mathbb{Z}_{\geq 1}$, its n -fold joint probability distribution (i.e., the distribution of (Z_1, \dots, Z_n)) is invariant to time shifts. Further, $\{Z_n\}_{n=1}^\infty$ is *exchangeable* if for any $n \in \mathbb{Z}_{\geq 1}$, its n -fold joint distribution is invariant to permutations of the indices $1, \dots, n$. It directly follows from the definitions that an exchangeable process is stationary. Lastly, the process $\{Z_n\}_{n=1}^\infty$ is called a *martingale* (resp. *supermartingale*, *submartingale*) with respect to the filtration $\{\mathcal{F}_n\}_{n=1}^\infty$ if $E[|Z_n|] < \infty$ and $E[Z_{n+1} | \mathcal{F}_n] = Z_n$ almost surely (resp. less than or equal to, greater than or equal to), for all n . If the inequality is strict, we call the process a *strict supermartingale* or *strict submartingale*. Doob’s martingale convergence theorem [25] can then be used to show that $\{Z_n\}_{n=1}^\infty$ will have a limit as n grows without bound. Precise definitions of all notions, including that of *ergodicity*, can be found in standard texts (e.g., [25], [26]).

We now recall the classical version of the Polya contagion process [12], [14]. Consider an urn with $R \in \mathbb{Z}_{>0}$ red balls and $B \in \mathbb{Z}_{>0}$ black balls. We denote the total number of balls by T , i.e., $T = R + B$. At each time step, a ball is drawn from the urn. The ball is then returned along with $\Delta > 0$ balls of




Fig. 1. Illustration of the first draw for a classical Polya process. We drew a black ball and hence $Z_1 = 0$. Here $R = 2$, $B = 2$, and $\Delta = 2$.

the same color. To describe this process, we use a replacement matrix M_R :

$$M_R = \begin{bmatrix} \Delta & 0 \\ 0 & \Delta \end{bmatrix},$$

where $[M_R]_{1,1} = \Delta$, $[M_R]_{1,2} = 0$ means we add Δ red balls and 0 black balls when a red ball is drawn, and similarly $[M_R]_{2,1} = 0$, $[M_R]_{2,2} = \Delta$ means we add 0 red balls and Δ black balls when a black ball is drawn. We use an indicator Z_n to denote the color of ball in the n th draw (see Figure 1):

$$Z_n = \begin{cases} 1 & \text{if the } n\text{th draw is red,} \\ 0 & \text{if the } n\text{th draw is black.} \end{cases}$$

Let U_n denote the proportion of red balls in the urn after the n th draw. Then

$$\begin{aligned} U_n &:= \frac{R + \Delta \sum_{t=1}^n Z_t}{T + n\Delta} \\ &= \frac{\rho_c + \delta_c \sum_{t=1}^n Z_t}{1 + n\delta_c} \end{aligned}$$

where $\rho_c = \frac{R}{T}$ is the initial proportion of red balls in the urn and $\delta_c = \frac{\Delta}{T}$ is a correlation parameter. Since we draw balls from this urn at each time step, the conditional probability of drawing a red ball at time n , given $Z^{n-1} = (Z_1, \dots, Z_{n-1})$, is given by

$$\begin{aligned} P(Z_n = 1 \mid Z^{n-1}) &= \frac{R + \Delta \sum_{t=1}^{n-1} Z_t}{T + (n-1)\Delta} \\ &= U_{n-1}. \end{aligned}$$

It can be easily shown that $\{U_n\}_{n=1}^\infty$ is a martingale [27]. The process $\{Z_n\}_{n=1}^\infty$, whose n -fold joint distribution can be determined in closed form in terms of the parameters ρ_c and δ_c , is also exchangeable (hence stationary) and non-ergodic with both U_n and the process sample average $\frac{1}{n} \sum_{i=1}^n Z_i$ converging almost surely as $n \rightarrow \infty$ to a random variable governed by the Beta distribution with parameters $\frac{\rho_c}{\delta_c}$ and $\frac{1-\rho_c}{\delta_c}$ [27], [28]. The classical Polya process has been applied in many different contexts, including the modelling of communication channels with memory [28], image segmentation [29], as well as in biology, statistics and other areas (see [30]).

III. MODEL DESCRIPTION AND PROBLEM STATEMENT

A. Network Polya Contagion Process

In this section, we briefly recall the network Polya contagion process introduced in [2], [3]. Consider an undirected graph $\mathcal{G} = (V, \mathcal{E})$, where $V = \{1, \dots, N\}$ is the set of $N \in \mathbb{Z}_{\geq 1}$ nodes and $\mathcal{E} \subset V \times V$ is the set of edges. We assume that \mathcal{G}




Fig. 2. Illustration of a super urn in a network.

is connected, i.e. there is a path between any two nodes in \mathcal{G} . We use \mathcal{N}_i to denote the set of nodes that are neighbors to node i , that is $\mathcal{N}_i = \{v \in V : (i, v) \in \mathcal{E}\}$, and $\mathcal{N}'_i = \{i\} \cup \mathcal{N}_i$. Each node $i \in V$ is equipped with an urn, initially with $R_i \in \mathbb{Z}_{>0}$ red balls and $B_i \in \mathbb{Z}_{>0}$ black balls (we do not let $R_i = 0$ or $B_i = 0$ to avoid any degenerate cases). We let $T_i = R_i + B_i$ be the total number of balls in the i th urn, $i \in \{1, \dots, N\}$. We use $Z_{i,n}$ as an indicator for the ball drawn for node i at time n :

$$Z_{i,n} = \begin{cases} 1 & \text{if the } n\text{th draw for node } i \text{ is red,} \\ 0 & \text{if the } n\text{th draw for node } i \text{ is black.} \end{cases}$$

Thus we define the network contagion process as $\{Z_n\}_{n=1}^\infty$, where $Z_n = (Z_{1,n}, \dots, Z_{N,n})$. However, instead of drawing solely from its own urn, each node has a ‘‘super urn’’ created by combining all the balls in its own urn with the balls in its neighbours’ urns; see Figure 2. This allows the spatial relationships between nodes to influence their state. This means that $Z_{i,n}$ is the indicator for a ball drawn from node i ’s super urn, and not its individual urn. Hence, the super urn of node i initially has $\bar{R}_i = \sum_{j \in \mathcal{N}'_i} R_j$ red balls, $\bar{B}_i = \sum_{j \in \mathcal{N}'_i} B_j$ black balls, and $\bar{T}_i = \sum_{j \in \mathcal{N}'_i} T_j$ balls in total, and the network-wide initial proportion of red balls is $\rho = \frac{\sum_{i=1}^N R_i}{\sum_{i=1}^N T_i}$.

We allow the number of added balls to vary based on the colour drawn, the time index, and the node for which it was drawn; hence, the replacement matrix for node i at time t is

$$M_{R,i}(t) = \begin{bmatrix} \Delta_{r,i}(t) & 0 \\ 0 & \Delta_{b,i}(t) \end{bmatrix}.$$

We assume throughout that $\Delta_{r,i}(t) \geq 0, \Delta_{b,i}(t) \geq 0$, for all $t \in \mathbb{Z}_{\geq 1}$ and that there exists $i \in V$ and t such that $\Delta_{r,i}(t) + \Delta_{b,i}(t) \neq 0$; otherwise we are simply sampling with replacement. In the context of epidemics, the red and black balls in an urn, respectively, represent ‘‘infection’’ and ‘‘healthiness’’. For example, in the context of a rumour on a social network, adding a red ball may mean a user decided to share information that corroborates the rumour, while adding a black ball would mean a user instead shared information dismissing the rumour. This information is visible to their immediate neighbours, and hence the added balls enter the super urns of those neighbouring users. In both cases, $\Delta_{r,i}(t)$ and $\Delta_{b,i}(t)$

may represent the number of such posts or articles shared over a given time window. We refer the interested reader to [3] for a complete description of this relationship.

To express the proportion of red balls in the individual urns of the nodes, we define the random vector $U_n = (U_{1,n}, \dots, U_{N,n})$, where $U_{i,n}$ is the proportion of red balls in node i ’s urn after the n th draw, $i \in V$. For node i ,

$$U_{i,n} = \frac{R_i + \sum_{t=1}^n \Delta_{r,i}(t) Z_{i,t}}{X_{i,n}},$$

where

$$X_{i,n} = T_i + \sum_{t=1}^n \Delta_{r,i}(t) Z_{i,t} + \Delta_{b,i}(t) (1 - Z_{i,t}) \quad (1)$$

is the total number of balls in node i ’s urn after the n th draw, and the numerator of $U_{i,n}$ represents the total number of red balls in the same urn. We now define the random vector $S_n = (S_{1,n}, \dots, S_{N,n})$ as the proportion of red balls in each of the super urns of the nodes after the n th draw, so that $S_{i,n}$ is the proportion of red balls in node i ’s super urn after n draws. Hence, for node i ,

$$\begin{aligned} S_{i,n} &= \frac{\bar{R}_i + \sum_{t=1}^n \bar{Z}_{r,i,t}}{\bar{X}_{i,n}} \\ &= \frac{\sum_{j \in \mathcal{N}'_i} U_{j,n} X_{j,n}}{\bar{X}_{i,n}}, \end{aligned} \quad (2)$$

where

$$\begin{aligned} \bar{Z}_{r,i,t} &= \sum_{j \in \mathcal{N}'_i} \Delta_{r,j}(t) Z_{j,t}, \\ \bar{Z}_{b,i,t} &= \sum_{j \in \mathcal{N}'_i} \Delta_{b,j}(t) (1 - Z_{j,t}), \\ \bar{X}_{i,n} &= \bar{T}_i + \sum_{t=1}^n (\bar{Z}_{r,i,t} + \bar{Z}_{b,i,t}) = \sum_{j \in \mathcal{N}'_i} X_{j,n}. \end{aligned}$$

Note that $S_{i,0} = \bar{R}_i / \bar{T}_i$. In fact, $S_{i,n}$ is a function of the random draw variables of the network, and in particular of $\{Z_j^n\}_{j \in \mathcal{N}'_i}$, but for ease of notation, when the arguments are clear, we write $S_{i,n}(Z_1^n, \dots, Z_N^n) = S_{i,n}$. Then the conditional probability of drawing a red ball from the super urn of node i at time n given the complete network history, i.e. given all the past $n-1$ draw variables for each node in the network $\{Z_j^{n-1}\}_{j=1}^N = \{(Z_{1,1}, \dots, Z_{1,n-1}), \dots, (Z_{N,1}, \dots, Z_{N,n-1})\}$, satisfies

$$\begin{aligned} P(Z_{i,n} = 1 | \{Z_j^{n-1}\}_{j=1}^N) &= \frac{\bar{R}_i + \sum_{t=1}^{n-1} \bar{Z}_{r,i,t}}{\bar{X}_{i,n}} \\ &= S_{i,n-1}. \end{aligned} \quad (3)$$

That is, the conditional probability of drawing a red ball for node i given the entire past $\{Z_j^{n-1}\}_{j=1}^N$ is the proportion of red balls in its super urn, $S_{i,n-1}$. Since these random variables fully describe the evolution of the process, we say $\{\mathcal{F}_n\}_{n=1}^\infty$ is the natural filtration on $\{Z_i^{n-1}\}_{i=1}^N$ and by extension $\{U_{i,n}\}_{n=1}^\infty$ and $\{S_{i,n}\}_{n=1}^\infty$, for all $i \in V$. Thus note that in (3) we could have instead conditioned on \mathcal{F}_{n-1} .

Using the conditional probability established above, we next determine the n -fold joint probability of the entire network \mathcal{G} . Let $a_i^n \in \{0, 1\}^n$, where $i \in \{1, \dots, N\}$. We have

$$\begin{aligned} P_{\mathcal{G}}^{(n)}(a_1^n, \dots, a_N^n) &:= P(\{Z_i^n = a_i^n\}_{i=1}^N) \\ &= \prod_{t=1}^n P(\{Z_{i,t} = a_{i,t}\}_{i=1}^N \mid \{Z_i^{t-1} = a_i^{t-1}\}_{i=1}^N) \\ &= \prod_{t=1}^n \prod_{i=1}^N (S_{i,t-1})^{a_{i,t}} (1 - S_{i,t-1})^{1-a_{i,t}}, \end{aligned} \quad (4)$$

where $S_{i,t} = S_{i,t}(a_1^t, \dots, a_N^t)$ is defined in (3). The study of the asymptotic behaviour of each node's contagion process $\{Z_{i,n}\}_{n=1}^{\infty}$, $i \in V$ is established in [2], [3]. Our objective in this work is to demonstrate the implications of these results in designing curing strategies. With the above explicit joint distribution, it is possible to determine the distributions of each node's contagion process. More specifically, using (4), the n -fold distribution of node i 's process at time $t \geq n$ is

$$P_{i,t}^{(n)}(a_{i,t-n+1}, \dots, a_{i,t}) := \sum_{\substack{a_i^{t-n} \in \{0,1\}^{t-n} \\ a_j^t \in \{0,1\}^t, j \neq i}} P_{\mathcal{G}}^{(t)}(a_1^t, \dots, a_N^t).$$

It can be shown that the draw contagion process $\{Z_{i,n}\}_{n=1}^{\infty}$ of each individual node i is not stationary in general (and hence not exchangeable) [2], [3]. Thus the entire network contagion process $\{Z_n\}_{n=1}^{\infty}$ is not stationary.

In order to measure the spread of contagion in the network at any given time, we wish to see how likely it is, on average, for a node to be infected at that instant. We thus define the *average infection rate* in the network at time n as

$$\tilde{I}_n := \frac{1}{N} \sum_{i=1}^N P(Z_{i,n} = 1) = \frac{1}{N} \sum_{i=1}^N P_{i,n}^{(1)}(1).$$

Note that \tilde{I}_n is a function of the *network variables*, i.e., the network topology (V, \mathcal{E}) , the initial placement of balls R_i and B_i , the draw processes $\{Z_{i,t}\}_{t=1}^n$, and the number of balls added $\{\Delta_{r,i}(t)\}_{t=1}^n$ and $\{\Delta_{b,i}(t)\}_{t=1}^n$ for each node $i \in V$. Unfortunately for an arbitrary network, the above quantity does not yield an exact analytical formula (except in the simple case of complete networks). As such, in general it is hard to mathematically analyze the asymptotic behavior of \tilde{I}_n , which we wish to minimize when attempting to cure an epidemic. Instead we examine the asymptotic stochastic behavior of two closely related variables given by the network-wide average individual proportion of red balls at time n , namely

$$\tilde{U}_n := \frac{1}{N} \sum_{i=1}^N U_{i,n},$$

which we call the *network susceptibility*, and the network-wide average neighborhood proportion of red balls at time n ,

$$\tilde{S}_n := \frac{1}{N} \sum_{i=1}^N S_{i,n},$$

which we call the *network exposure*. Note that similarly to \tilde{I}_n , both \tilde{U}_n and \tilde{S}_n are functions of the network variables.

With the model in hand, we turn to the exploration of a curing problem. Our objective is to control the average infection rate \tilde{I}_n as n grows without bound; but when seeking analytic results, it might be more amenable to observe the asymptotic behavior of the network exposure \tilde{S}_n . These quantities are closely related; through (2) we see that if $U_{i,n}$ increases then this node-specific value causes $S_{j,n}$ to increase for every neighbour j of node i , and hence by (3) their conditional probabilities of drawing red balls increase. More specifically,

$$\begin{aligned} \uparrow U_{i,n} &\stackrel{(2)}{\Rightarrow} \uparrow S_{j,n} \text{ for all } j \in \mathcal{N}'_i \\ &\stackrel{(3)}{\Rightarrow} \uparrow P(Z_{i,n+1} = 1 \mid \{Z_j^n\}_{j=1}^N) \text{ for all } j \in \mathcal{N}'_i. \end{aligned} \quad (5)$$

Thus if \tilde{U}_n is high, then this average measure of individual nodes implies that the conditional probability of a node being infected is higher on average. Hence \tilde{U}_n can be understood as the *average node prevalence* of infection. The effect of the network exposure here is more direct, since (3) shows that \tilde{S}_n is in fact the network-wide average of the conditional probabilities of infection, which is a quantity that is intimately related to the state of infection in the neighbourhood of node i . Thus \tilde{S}_n represents the *average neighbourhood prevalence* of infection. We focus on these quantities over the average infection rate \tilde{I}_n for computational tractability. \tilde{U}_n and \tilde{S}_n provide us with proxy measurements of the ‘‘level of infection’’ in the network which may be used to provide mathematical justification for our exploration into optimal curing strategies.

B. Establishing a Control Problem

The quantities $\{\Delta_{b,i}(n)\}_{n=1}^{\infty}$, which denote the net number of ‘‘healthy’’ balls added to node i 's urn after each draw, can play the role of ‘‘healing or curing parameters’’. Our objective is to show that when these parameters are appropriately selected, one can steer the average infection rate towards a desirable level; the selection of curing parameters is, however, subject to an allowable budget on the maximal number of healthy balls that can be added in the network. Let us state this problem formally.

Problem 1 (Average Infection Rate Budget Constraint):
Minimize the limiting average infection rate \tilde{I}_t subject to a budget \mathcal{B} on the total healing at each time step:

$$\min_{\sum_{i=1}^N \Delta_{b,i}(t) \leq \mathcal{B} \quad \forall t} \limsup_{t \rightarrow \infty} \tilde{I}_t$$

Such optimal curing problems have been studied in many different contexts [16], [31]. For our model, the solution to Problem 1 would be an infinite horizon optimal control policy that would yield the best possible level of epidemic elimination, given the initial data. Finding such a policy in general appears to be difficult. Nevertheless, as we demonstrate in the upcoming sections, one can obtain interesting analytical results regarding the feasibility of this problem, and design algorithmic strategies to curtail the average infection rate.

IV. CONTROLLING EPIDEMICS: ANALYTICAL RESULTS

In order to determine when Problem 1 makes sense, we wish to examine when a limit exists. Such a result would describe

to us when our curing problem is feasible, and hence when the exploration of an optimal curing strategy for the network Polya contagion process is reasonable. As stated earlier, working with \tilde{I}_n can be difficult, and so we instead focus on the related measures of the network susceptibility \tilde{U}_n and network exposure \tilde{S}_n . Our next results will show how one can force these measures to form supermartingales by appropriately selecting the curing policies $\{\Delta_{b,i}(n)\}_{n=1}^\infty$, for all $i \in V$. In conjunction with Doob's martingale convergence theorem [32], these results show that $\{U_{i,n}\}_{n=1}^\infty$, $\{S_{i,n}\}_{n=1}^\infty$, and hence both $\{\tilde{U}_n\}_{n=1}^\infty$ and $\{\tilde{S}_n\}_{n=1}^\infty$, have limits. While the results presented herein do not obey the per-step budget constraint on the curing, these results in conjunction with the simulation results presented later show that strategies that fit within the framework of Problem 1 exist that reduce \tilde{I}_n on average. In particular, since it is possible to guarantee these quantities decrease on average, we argue that it is reasonable to assume there is an optimal method for forcing the average infection rate to be reduced on average.

An important assumption used herein is that the number of red balls to be added $\Delta_{r,i}(n)$ is known at least one step ahead of time, so that in particular $\Delta_{r,i}(n)$ is almost surely constant given \mathcal{F}_{n-1} . A sufficient, but not necessary, condition to satisfy this assumption is for $\{\Delta_{r,i}(n)\}_{n=1}^\infty$ to be set, for all $i \in V$, before the process begins.

Theorem 2 (Individual Urn Proportion Categories): In a general network $\mathcal{G} = (V, \mathcal{E})$, if we choose $\{\Delta_{b,i}(n)\}_{n=1}^\infty$ so that

$$\Delta_{b,i}(n) \geq \frac{\Delta_{r,i}(n)(1 - U_{i,n-1})S_{i,n-1}}{U_{i,n-1}(1 - S_{i,n-1})}$$

almost surely for all $n \in \mathbb{Z}_{\geq 1}$ and $i \in V$ (resp. equal to, less than or equal to) then $\{U_{i,n}\}_{n=1}^\infty$ is a supermartingale (resp. martingale, submartingale) with respect to the natural filtration $\{\mathcal{F}_n\}_{n=1}^\infty$, i.e.,

$$E[U_{i,n}|\mathcal{F}_{n-1}] \leq U_{i,n-1} \text{ almost surely } \forall n \in \mathbb{Z}_{\geq 1}.$$

Corollary 3 (Network Susceptibility Supermartingale): In a general network $\mathcal{G} = (V, \mathcal{E})$, if the curing policies $\{\Delta_{b,i}(n)\}_{n=1}^\infty$ obey the bound

$$\Delta_{b,i}(n) \geq \frac{\Delta_{r,i}(n)(1 - U_{i,n-1})S_{i,n-1}}{U_{i,n-1}(1 - S_{i,n-1})}$$

almost surely for all nodes $i \in V$, then the network susceptibility $\{\tilde{U}_n\}_{n=1}^\infty$, where $\tilde{U}_n = \frac{1}{N} \sum_{i=1}^N U_{i,n}$, is a supermartingale with respect to the natural filtration $\{\mathcal{F}_n\}_{n=1}^\infty$, i.e.,

$$E[\tilde{U}_n|\mathcal{F}_{n-1}] \leq \tilde{U}_{n-1} \text{ almost surely } \forall n \in \mathbb{Z}_{\geq 1}.$$

The proof for Theorem 2 is provided in Appendix A. While Corollary 3 is useful, the network exposure \tilde{S}_n is more closely related to the average infection rate \tilde{I}_n than the network susceptibility \tilde{U}_n , since our draws are taken from the super urn. It is with this in mind that we show the next results, which give us sufficient conditions for $\{S_{i,n}\}_{n=1}^\infty$ and $\{\tilde{S}_n\}_{n=1}^\infty$ to be supermartingales.

Theorem 4 (Super Urn Proportion Categories): In a general network $\mathcal{G} = (V, \mathcal{E})$, if the curing policy $\{\Delta_{b,i}(n)\}_{n=1}^\infty$

obeys the lower bound

$$\Delta_{b,i}(n) > \Delta_{r,i}(n) \frac{S_{i,n-1}}{1 - S_{i,n-1}} \max_{k \text{ s.t. } i \in \mathcal{N}'_k} \frac{1 - S_{k,n-1}}{S_{k,n-1}} \quad (\text{B1})$$

almost surely for all nodes $i \in V$, then the neighbourhood proportions of red balls $\{S_{i,n}\}_{n=1}^\infty$ are strict supermartingales with respect to the natural filtration $\{\mathcal{F}_n\}_{n=1}^\infty$, i.e.

$$E[S_{i,n}|\mathcal{F}_{n-1}] < S_{i,n-1} \text{ almost surely } \forall i \in V, n \in \mathbb{Z}_{\geq 1}.$$

Furthermore, if the curing policy $\{\Delta_{b,i}(n)\}_{n=1}^\infty$ obeys the upper bound

$$\Delta_{b,i}(n) < \Delta_{r,i}(n) \frac{S_{i,n-1}}{1 - S_{i,n-1}} \min_{k \text{ s.t. } i \in \mathcal{N}'_k} \frac{1 - S_{k,n-1}}{S_{k,n-1}} \quad (\text{B2})$$

almost surely for all nodes $i \in V$, then the neighbourhood proportions of red balls $\{S_{i,n}\}_{n=1}^\infty$ are strict submartingales with respect to the natural filtration $\{\mathcal{F}_n\}_{n=1}^\infty$.

Corollary 5 (Network Exposure Categories): In a general network $\mathcal{G} = (V, \mathcal{E})$, if the curing policies $\{\Delta_{b,i}(n)\}_{n=1}^\infty$ obey the lower bound (B1) almost surely for all nodes $i \in V$, then the network exposure $\{\tilde{S}_n\}_{n=1}^\infty$, where $\tilde{S}_n = \frac{1}{N} \sum_{i=1}^N S_{i,n}$, is a strict supermartingale with respect to the natural filtration $\{\mathcal{F}_n\}_{n=1}^\infty$, i.e.,

$$E[\tilde{S}_n|\mathcal{F}_{n-1}] < \tilde{S}_{n-1} \text{ almost surely } \forall n \in \mathbb{Z}_{\geq 1}.$$

Furthermore, if the curing policies $\{\Delta_{b,i}(t)\}_{t=1}^\infty$ obey the upper bound (B2) almost surely for all nodes $i \in V$, then the network exposure $\{\tilde{S}_n\}_{n=1}^\infty$ is a strict submartingale with respect to the natural filtration $\{\mathcal{F}_n\}_{n=1}^\infty$.

The proof of Theorem 4 is presented in Appendix B. While the duality of these bounds is interesting, in the context of curing we will focus on the lower bound (B1). It is important to note that the policy for $\{\Delta_{b,i}(n)\}_{n=1}^\infty$ used in Theorem 4 is not a tight lower bound on the curing resources which guarantee that the processes $\{S_{i,n}\}_{n=1}^\infty$ are supermartingales, and hence it is possible that less costly policies exist that still guarantee this property. In particular, strategies may exist which obey the fixed budget \mathcal{B} on the amount of curing resources that may be used. However, these results motivate the fact that the search for better policies makes sense, since we know that policies exist that will fight the infection and reduce it on average.

V. CONTROLLING EPIDEMICS: ALGORITHMIC STRATEGIES

The supermartingale results established in the previous section demonstrate the feasibility of a relaxed version of Problem 1, with no budget limitation. In this section, we establish numerical methods to find control policies that find efficient sub-optimal policies for Problem 1, under budget constraints and with having computational complexity in mind. We compare these strategies with the ones obtained from our supermartingale results. A summary of all strategies that will be discussed in this section is given in Table I.

Before we present these strategies in details, let us describe briefly how we have evaluated their performance. The simulation platform for these strategies is outlined in Algorithm 1. To achieve comparable results, independent

TABLE I
CURING STRATEGIES

(i)	Forcing all $U_{i,t}$ to be supermartingales (Theorem 4.1): $\Delta_{b,i}(t) = \frac{\Delta_{r,i}(t)(1-U_{i,t-1})S_{i,t-1}}{U_{i,t-1}(1-S_{i,t-1})}$
(ii)	Forcing all $S_{i,t}$ to be supermartingales (Theorem 4.3): $\Delta_{b,i}(t) = \Delta_{r,i}(t) \frac{S_{i,t-1}}{1-S_{i,t-1}} \max_{k \text{ s.t. } i \in \mathcal{N}'_k} \frac{1-S_{k,t-1}}{S_{k,t-1}}$
(iii)	Constrained gradient descent algorithm on a simplex: Find $\Delta_{b,i}(t)$ using Algorithm 2
(iv)	Ratio of degree, close./betw. cent. and super urn proportion: $\Delta_{b,i}(t) = \mathcal{B} \frac{ \mathcal{N}_i C_i S_{i,t-1}}{\sum_{j=1}^N \mathcal{N}_j C_j S_{j,t-1}}$
(v)	Uniformly allocate the budget to all nodes in the network: $\Delta_{b,i}(t) = \frac{\mathcal{B}}{N}$

trials of the process are ran with the same initial conditions $\vec{R} = (R_1, \dots, R_N)$, $\vec{B} = (B_1, \dots, B_N)$, and $\vec{\Delta}_r = (\Delta_{r,1}, \dots, \Delta_{r,N})$, for each curing strategy. The results for each strategy are then averaged to evaluate the expected performance. The full simulation results, along with discussions of their implications, are presented in Section VI.

Algorithm 1 Network Contagion Curing

$A \leftarrow$ adjacency matrix of the network
 $\vec{R}, \vec{B}, \vec{\Delta}_r \sim [\text{Uniform}((0, 10])]^N$
 $\mathcal{B} \leftarrow$ budget, $\sum_{i=1}^N \Delta_{r,i}$
 $C \leftarrow$ number of cases, each with a *strategy*
 $T \leftarrow$ number of trials to run for each case
 $K \leftarrow$ number of time steps for each trial
for $c = 1 : C$ **do**
 for $t = 1 : T$ **do**
 $\vec{Z}_{c,t} \leftarrow \text{RUNTRIAL}(A, \vec{R}, \vec{B}, \vec{\Delta}_r, \mathcal{B}, K, \text{strategy})$
 $\vec{Z}_c = \frac{1}{T} \sum_{t=1}^T \vec{Z}_{c,t}$
 procedure $\text{RUNTRIAL}(A, \vec{R}, \vec{B}, \vec{\Delta}_r, \mathcal{B}, K, \text{strategy})$
 Initialize $S_{i,0}, U_{i,0}$ using R_i and B_i for all $i \in V$
 for $s = 1 : K$ **do**
 Assign $\Delta_{b,i}(s)$ using *strategy*, obeying \mathcal{B} if required
 Generate $\vec{Y} \sim \text{Uniform}([0, 1])^N$
 if $Y_i \leq S_{i,s-1}$ **then**
 $Z_{i,s} = 1$
 else
 $Z_{i,s} = 0$
 Update $S_{i,s}, U_{i,s}$ using $\Delta_{r,i}, \Delta_{b,i}(s)$ for all $i \in V$
 based on A

A. Supermartingale Strategies

The supermartingales results that we have obtained in Section IV, specifically Theorems 2 and 4, naturally lead to a class of curing strategies. In particular, these strategies guarantee that \tilde{U}_n and \tilde{S}_n , respectively, are supermartingales. It is worth reminding that our theoretical results do not necessarily imply that average infection rate \tilde{I}_n forms a supermartingale. In spite of this, these strategies are still viable options for curing, as far as enough resources are available.

We next describe the differences between the strategy given by individual urn proportions, and the one given by super urn proportions.

By Corollary 3, we know that strategy (i) guarantees that the network susceptibility \tilde{U}_n will be a supermartingale. Hence we set the curing strategy for each node so that it will force its own individual urn proportion of red balls to be a supermartingale. Since draws are taken from the super urns and not the individual urns, the relationship between the reduction of \tilde{U}_n and \tilde{I}_n is not a strong one and our simulations suggests that this strategy does not appear to offer a large reduction in the average infection rate in general. In contrast, the curing strategy given by Corollary 5, where we choose our curing strategy to force the super urn proportions of red balls to be supermartingales for all nodes, performs reasonably well.

While these strategies guarantee a reduction in their respective measures, they use an arbitrary amount of curing resources to do so in general. In fact, as we will see later, these strategies always use a large amount of curing resources relative to the impact they have on reducing the average infection rate. However, the impact of these results are in proving achievability of the curing problem. It is with this in mind that we examine a numerical curing strategy that obeys a fixed budget on the per-step curing resources in order to stay within the framework of Problem 1.

B. Gradient Flow Methods

In this section, we employ the well-known gradient descent algorithm [33] for Problem 1. As discussed earlier, using \tilde{I}_n as a measure of infection is computationally expensive, and hence we instead focus on the network exposure \tilde{S}_n . While our suggested gradient descent algorithm will not provide the exact answer to Problem 1 for reducing \tilde{I}_n , we will show that it is guaranteed to provide the optimal policy to reduce the closely related measure \tilde{S}_n .


In Problem 1, our curing policy is constrained by a budget \mathcal{B} at each time step and so the feasible set, or set of valid curing policies, for our gradient descent is all policies which do not exceed \mathcal{B} . However, any optimal policy will make use of the whole budget, and so we consider our feasible set to be $\mathcal{X} = \left\{ \{\Delta_{b,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N \mid \sum_{i=1}^N \Delta_{b,i}(n) = \mathcal{B} \right\}$. Proposition 6 shows that for arbitrary initial conditions and network topologies, the problem under study for the expected network exposure $E[\tilde{S}_n | \mathcal{F}_{n-1}]$ is convex.

Proposition 6 (Gradient descent conditions are met): In a general network $\mathcal{G} = (V, \mathcal{E})$ with arbitrary initial conditions, the expected network exposure $E[\tilde{S}_n | \mathcal{F}_{n-1}]$ is convex with respect to the curing parameters $\{\Delta_{b,i}(n)\}_{i=1}^N$ for all n . Furthermore, the feasible set


$$\mathcal{X} = \left\{ \{\Delta_{b,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N \mid \sum_{i=1}^N \Delta_{b,i}(n) = \mathcal{B} \right\}$$

is convex and compact.


The proof for Proposition 6 can be found in Appendix C. The structure of the feasible set \mathcal{X} allows us to employ the simplex constrained gradient descent method, described




(a) Facebook group network with 1,363 nodes, 2,425 edges, and average degree 3.56.



(b) Plot of empirical average infection rate \tilde{I}_n compared to the initial level of infection $\rho = \frac{\sum_{i=1}^N R_i}{\sum_{i=1}^N T_i}$ (lower means less infection).



(c) Plot of empirical average usage of curing resources, $\sum_{i=1}^N \Delta_{b,i}(t)$ (lower means less curing resources used). Note that strategies (iii), (iv) and (v)'s usages are fixed at the per-step budget \mathcal{B} , and are overlaid.



(d) Plot of empirical average wasted curing, $\sum_{i=1}^N \sum_{t=1}^n \Delta_{b,i}(t) Z_{i,t}$ (lower means less curing resources assigned to nodes that did not use them).

Fig. 3. Comparison of all curing strategies presented in Table I on a Facebook group network. Here strategy (iv) used closeness centrality. Simulation results were averaged over 250 trials for each strategy, and altogether took approximately 49 hours on 10 cores of an Intel Xeon processor at 2.20GHz. Initial numbers of balls R_i and B_i , and numbers of red balls added $\Delta_{r,i}$ (which remained constant in time), were uniformly randomly assigned for each node but stayed consistent throughout all trials and strategies, while the assignments of $\{\Delta_{b,i}(t)\}_{t=1}^{\infty}$ were different for each strategy. Since the $\Delta_{r,i}$ are all constant, the budget was set as $\mathcal{B} = \sum_{i=1}^N \Delta_{r,i}$.

in Algorithm 2 [33, Chapter 2]. The complexity of this algorithm which runs at each time step is $O(Nsa)$, where N is the number of nodes, s is the number of iterations of the gradient descent, and $\frac{1}{a}$ is the granularity used to find the limit-minimized step size α_k . While Proposition 6 guarantees that the curing policy that this algorithm finds will be optimal for each individual step, it does not guarantee optimality over the entire time horizon. In spite of this, as the simulation results in Figure 3 show, this curing strategy still outperforms all other curing strategies studied in this paper. The downside of the gradient method is that it is computationally expensive to execute, as it requires intimate knowledge of the state of all nodes in the network. This motivates us to seek other methods which are computationally easier to execute, although they do not perform as well as the gradient descent strategy.

C. Heuristic Strategies

Both sets of strategies identified above come with challenges. The supermartingale strategies are accompanied by analytical results that guarantee that they will improve in general, but they do not obey a fixed budget, nor do they create a significant reduction in the average infection rate. The gradient flow method uses a fixed budget and is provably

Algorithm 2 Constrained Gradient Descent on a Simplex [33]

```

 $y_1 = (\mathcal{B}, 0, \dots, 0)$ 
 $T \leftarrow$  number of iterations to perform
for  $k = 1 : T$  do
     $i = \arg \min_{j \in V} \frac{\partial f}{\partial x_j}$ 
     $[\bar{y}_k]_i = \mathcal{B}$ , and  $[\bar{y}_k]_j = 0$  for all  $j \neq i$ 
     $\alpha_k = \arg \min_{\alpha \in [0,1]} f(y_k + \alpha(\bar{y}_k - y_k))$ 
     $y_{k+1} = y_k + \alpha_k(\bar{y}_k - y_k)$ 

```

optimal to reduce the expected network exposure $E[\tilde{S}_n | \mathcal{F}_{n-1}]$, but it is computationally costly and requires a large amount of information about the state of infection at every node, including the history of draws and values of the curing parameters. As a compromise between these strategies, we present the centrality-infection ratio, which is a heuristic centrality-based strategy designed to allocate the fixed per-step budget \mathcal{B} .

The idea is to create a ratio to split the budget between all nodes in the network, whose time complexity will be of the order $O(1)$. We consider three factors when determining how much curing a node should receive: local impact, topological position, and level of infection. Nodes with higher local impact

have more neighbours, and hence any healing they receive is immediately distributed to a larger number of nodes. Those with a better topological position are more centrally located within the network, in the sense that the distance from them to all other nodes is smaller. Lastly, nodes with a higher level of infection will need more curing resources to become healthy. Numerical results suggest that lack of consideration of any of these factors leads to a worse performance.

The resulting curing strategy, which we called the centrality-infection ratio, is

$$\Delta_{b,i}(n) = \mathcal{B} \frac{|\mathcal{N}_i| C_i S_{i,n-1}}{\sum_{j=1}^N |\mathcal{N}_j| C_j S_{j,n-1}}.$$

To measure local impact of node i , we use the degree, $|\mathcal{N}_i|$, which measures the number of neighbours for node i . Topological position is determined by calculating a centrality metric, either closeness centrality $C_{c,i}$ [34] or betweenness centrality $C_{b,i}$ [35], which, for node i , are defined as

$$C_{c,i} := \frac{1}{\sum_{j \in V} d(i,j)},$$

$$C_{b,i} := \sum_{s,t \in V, s \neq t, s,t \neq i} \frac{\sigma_{s,t}(i)}{\sigma_{s,t}},$$

where $d(i,j)$ is the length of the shortest path from node i to node j , $\sigma_{s,t}$ is the total number of shortest paths from s to t , and $\sigma_{s,t}(i)$ is the number of those shortest paths passing through i . Note that two nodes i,j for whom $d(i,j) = k$ take k time steps to feel the effects of one another. Hence, the closeness centrality encodes the ‘‘influence’’ nodes have on the network on average: those with higher closeness centrality will, on average, spread the effect of their curing resources more quickly. The betweenness centrality similarly measures how swiftly allocated curing resources will dissipate, but with a focus on shorter spreading time rather than average effect. Finally, to measure the level of infection, we use the super urn proportion of red balls $S_{i,n}$. From (3), we know that this quantity captures how likely it is for node i to be infected currently given the history of the process. Thus we give more importance to nodes who are more likely to be infected, in order to make them less likely to be infected in the future. This allocation ratio is a generalization of the best heuristic strategy presented in [2], which considered only node degrees.

The advantage of this heuristic strategy is twofold. Not only does it reduce computational time complexity at each time step from $O(Nsa)$ to $O(1)$, it is also somewhat distributed in the sense that it does not require constant information from the entire network. Unlike the gradient descent algorithm, strategy (iv) simply needs to know information about the network topology and the state of infection of each node. Since we assume that our network’s graph is constant in time, this topological information is only required initially and can be used thereafter. The only other information required from the network at large is the sum of the super urn ratios $\sum_{i=1}^N S_{i,n}$, and hence much less information needs to be communicated through the network for the implementation of this strategy.

Lastly, for comparison reasons we present the uniform curing strategy (v), which splits the budget \mathcal{B} equally to all

nodes in the network. This provides a benchmark to measure the improvement achieved by more intelligent strategies.

VI. SIMULATION RESULTS AND DISCUSSION

In order to confirm the results of Theorems 2 and 4, a number of simulations were performed; the pseudocode is outlined in Algorithm 1. While the simulations performed had the numbers of red balls added $\Delta_{r,i}$ vary between nodes, they were constant in time. This was done to simplify the choice of the per-step budget, and does not affect the execution of the simulations themselves. All initial conditions used in the simulations herein, as well as videos displaying the average performance of the curing strategies, are available online.¹

A. Simulation Setup

The network shown in Figure 3(a) was generated by using a tool [36] to crawl through 500 posts in a Facebook group. Individuals who created posts or interacted with others’ content are represented by nodes, while edges are created if individuals interacted with the post or comment of another (by commenting on the post, or liking the post or comment). The resulting graph has 1,363 nodes and 2,425 edges, and by design represents the topology of a real social network.

We now provide a detailed description of the simulation, as described in Algorithm 1. The values of R_i , B_i and $\Delta_{r,i}$ were uniformly randomly assigned for each node as integers between 1 and 10. These values remained consistent for all strategies and throughout all trials that were performed. Since the values for $\Delta_{r,i}$ were fixed over time, the per-step budget was set at $\mathcal{B} = \sum_{i=1}^N \Delta_{r,i}$. With the initial conditions set, a number of trials were performed for each strategy. Each trial was performed by successively drawing balls from super urns for a fixed number of time steps. At time t , we first assigned the curing $\Delta_{b,i}(t)$ based on the strategy selected. Then a uniform random variable on $[0, 1]$, Y_i , was generated for each node i and compared to the super urn proportion. If $Y_i < S_{i,t-1}$ then we say that a red ball was drawn and so $Z_{i,t} = 1$, otherwise we drew black and so $Z_{i,t} = 0$. Based on what was drawn, we added $\Delta_{r,i}$ red or $\Delta_{b,i}(t)$ black balls into node i ’s urn, and hence its super urn and those of its neighbours. At the end of each trial the draw variables were saved, and then averaged over all trials to produce the empirical performance of the curing strategy.

B. Discussion of Simulation Results

The comparisons of all strategies outlined in Section V can be seen in Figures 3–6. It is important to note that only strategies (iii), (iv) and (v) in Table I have a budget \mathcal{B} on their curing resources, while the other two strategies vary the total curing they use in time; the amount of resources each strategy consumes is shown in Figure 3(c). Figure 3(d) displays the average wasted curing resources for each strategy, which we will define later.

Figures 3(b) and 4 compare the performance of all strategies described in Section V on a Facebook network. Figure 3(b)

¹See: <http://bit.ly/2szl8PY>




Fig. 4. Plot of empirical average infection rate \tilde{I}_n on the network shown in 3(a) for a longer time frame. Strategies used are shown in Table I. The simulations presented here were performed identically to those described in Figure 3, with all initial conditions consistent between trials and strategies. Strategies were averaged over 1,000 trials, and altogether took approximately 30 hours on 12 cores of an Intel Xeon processor at 2.20GHz.

includes the gradient flow algorithm, while Figure 4 shows all other strategies over a longer time horizon. The benchmark uniform strategy (v) performs the worst, which is to be expected. Although (iii) is only proven to be optimal for the expected network exposure $E[\tilde{S}_n | \mathcal{F}_{n-1}]$, we observe that it is effective for the average infection rate \tilde{I}_n as well; as previously mentioned, this strategy outperforms all other curing strategies described in this paper. However, the heuristics strategy (iv) shows similar performance with dramatic improvements in computational complexity. The supermartingale strategies (i) and (ii) both reduce \tilde{I}_n below the initial average infection rate in the network ρ , but are less effective in doing so than the other two methods. This is likely due to their lack of consideration of the network topology, which plays a critical role in determining how curing resources spread across the network. Strategy (i) sees only an immediate small reduction in \tilde{I}_n , while strategy (ii) continuously decreases \tilde{I}_n . This empirical evidence follows (5), and further supports our assertion that we should focus on the network exposure \tilde{S}_n to reduce the average infection rate.


Figure 5(b) compares the performance of all strategies described in Section V on a denser Barabasi-Albert network, with average degree 19.55. Barabasi-Albert networks are randomly generated through preferential attachment and are widely used in the literature, as they have been shown to exhibit the properties of real social networks [37]. Interestingly, the gap between strategies (i) and (ii) is smaller relative to the sparser Facebook group network since strategy (ii) performs worse, highlighting the importance of considering topological information. Moreover, we observe that betweenness centrality was the correct metric to use in this setting, rather than closeness centrality (or other metrics including eigenvector, Pagerank, and degree centralities). Relative to the sparser network in Figure 3, there are many more paths between nodes in this context and so a simple ratio of their length, i.e., closeness centrality, fails to adequately measure the importance of nodes. Instead, the ratio of the number of such

paths in which a node appears, i.e., betweenness centrality, is a better measure of the influence of an individual. In other words, having shorter spreading times of curing resources to a larger number of nodes is more important than the average effect of the spreading of those curing resources.


In Figure 3(c) we examine the amount of curing resources used by each strategy. Since strategies (iii), (iv) and (v) all obey a per-step budget constraint their usages are fixed. Both supermartingale strategies, which may use arbitrary amounts of curing resources, initially use a larger amount of curing resources and then reduce their usage. This initial expenditure is the cost required to turn the measures into supermartingales, after which the strategies only need to maintain the property which requires less resources. Strategy (i)'s usage appears to decay exponentially to an almost constant amount, while strategy (ii) continues to decrease its usage in time. Further, strategy (i) uses almost 50% more curing resources than the budget \mathcal{B} initially, while strategy (ii)'s initial usage is only around 18% higher than \mathcal{B} . This is likely because strategy (i) is selfish; it only considers infection on an individual node level. In contrast, strategy (ii) considers the infection in local neighbourhoods of nodes, and hence is more judicious in applying resources to specific nodes.

The amount of curing resources wasted by each strategy is displayed in Figure 3(d). Waste is defined as curing resources which were assigned to nodes that did not use them since they displayed "infected" behaviour at that time, and is hence computed as $\sum_{i=1}^N \sum_{t=1}^n \Delta_{b,i}(t) Z_{i,t}$. We observe an intuitive correlation between the amount of resources wasted and curing performance: strategies which waste less resources tend to be more effective at reducing the average infection rate \tilde{I}_n . However, this does not tell the full story. The gradient flow algorithm has several spikes where it wastes more resources than the centrality-infection ratio (iv), but this does not appear to affect its curing performance. These likely occur because the gradient descent tends to focus on clusters of points when assigning curing, and hence the rest of the network becomes more infected. Then, when it switches to another cluster, the nodes are more infected and hence it wastes more resources until it sufficiently cures the infection in that cluster. Since the selection of these clusters depends heavily on the past evolution of the process there exists a higher variance in the wastage of curing resources, causing the line (iii) to be less smooth than those of other strategies. Strategy (i) initially wastes less than strategy (ii) even though it uses more curing resources, and it still performs worse with respect to reduction in \tilde{I}_n . This suggests that optimal curing strategies do not simply waste less, but also intelligently allocate their curing resources to make the best use of them.

Figure 6 shows the initial and final state of infection of all nodes in a randomly generated network for two different curing strategies. We define the state of infection for node i at time t by its individual proportion of red balls $U_{i,t}$. Here we use a small and sparse Barabasi-Albert network so the states at the node level are more visible. In Figure 6(a) we see that for such a network the centrality-infection ratio (iv) dramatically outperforms uniform curing (v), as was the case for the social network shown in Figure 3(a). After 1,000 time




(a) Barabasi-Albert network [37] with 200 nodes, 1,955 edges, and average degree 19.55.




(b) Plot of empirical average infection rate \tilde{I}_n compared to the initial level of infection $\rho = \frac{\sum_{i=1}^N R_i}{\sum_{i=1}^N T_i}$ (lower means less infection).


Fig. 5. Comparison of all curing strategies presented in Table I on a Barabasi-Albert network. Here strategy (iv) made use of betweenness centrality. Simulation results were averaged over 500 trials for each strategy, and altogether took approximately 20 hours on 4 cores of an Intel i5 processor at 3.40GHz. Simulations were initialized in the same way as those in Figure 3.




(a) Plot of empirical average infection rate \tilde{I}_n over 500 time steps.



(b) Plot of initial level of individual infection for each node $U_{i,0}$.



(c) Plot of final level of individual infection for each node $U_{i,n}$ using strategy (iv), the centrality-infection ratio.



(d) Plot of final level of individual infection for each node $U_{i,n}$ using strategy (v), uniform curing.

Fig. 6. Comparison of curing strategies (iv) and (v) on a Barabasi-Alberth network [37] with 100 nodes and 99 edges. Here strategy (iv) used closeness centrality. Blue represents total healthiness ($U_{i,n} = 0$) while red represents total infection ($U_{i,n} = 1$). Simulation results were averaged over 1,000 trials for each strategy, and altogether took approximately 5 minutes on 4 cores of an Intel Core i7 processor at 2.20 GHz. This simulation was performed identically to those described in Figure 3, with all initial conditions consistent between trials and strategies.

steps, strategy (iv) reduced the average infection rate \tilde{I}_n to about 15%, and no node is worse off than it was to begin with. In contrast, strategy (v) barely reduced \tilde{I}_n below the initial average infection rate ρ , and the individual infection of some nodes rose above 90%. This result illustrates the fact that intelligent allocation of curing resources is not only important to reduce the network-wide average infection rate, but the infection of individual nodes as well.

VII. CONCLUSION AND FUTURE WORK

In this paper we examined the problem of curing epidemics using a network contagion model adapted from the Polya urn process. We formulated an optimal control problem and provided analytical results that showed that finding solutions is a worthwhile endeavour. We used theoretical, numerical and heuristic curing strategies to attempt to cure the epidemic, and evaluated their performance using simulations.

Future work with this model could include the statement of different curing problems. A budget could be assigned over a finite time horizon instead of on a per-step basis, and strategies would need to judiciously use this limited supply to reduce infection as much as possible in the time window. The problem could even be reversed, so that some desirable healthiness threshold is provided and one could find the lowest possible budget that would guarantee that the average infection rate would be at or below the threshold. Such a problem could be examined for a per-step or fixed horizon budget. Applications of our model and curing results to control and contain contagion in real-world networks, including social networks, campaigning platforms and malware-prone computer networks, is another interesting future direction.

APPENDIX

A. Proof of Theorem 2

This result is a generalization of Theorem 4.6 in [3], since here we allow $\{\Delta_{r,i}(t)\}_{t=1}^{\infty}$ to vary in time. As such, some minor steps are omitted.

We will start with the case of a supermartingale. That is, we wish to show that almost surely for all $n \in \mathbb{Z}_{\geq 1}$,

$$E[U_{i,n} \mid \mathcal{F}_{n-1}] - U_{i,n-1} \leq 0 \Leftrightarrow E[U_{i,n} - U_{i,n-1} \mid \mathcal{F}_{n-1}] \leq 0,$$

since $U_{i,n-1}$ is almost surely constant given \mathcal{F}_{n-1} . Take $X_{i,n}$ as in (1). We then compute the difference $U_{i,n} - U_{i,n-1}$,

$$\begin{aligned} & U_{i,n} - U_{i,n-1} \\ &= \frac{R_i + \sum_{t=1}^n \Delta_{r,i}(t) Z_{i,t}}{X_{i,n}} - \frac{R_i + \sum_{t=1}^{n-1} \Delta_{r,i}(t) Z_{i,t}}{X_{i,n-1}} \\ &= \frac{\Delta_{r,i}(n) Z_{i,n}}{X_{i,n}} - \frac{(R_i + \sum_{t=1}^{n-1} \Delta_{r,i}(t) Z_{i,t})(X_{i,n} - X_{i,n-1})}{X_{i,n-1} X_{i,n}} \\ &= \frac{\Delta_{r,i}(n) Z_{i,n} - U_{i,n-1} (\Delta_{r,i}(n) Z_{i,n} + \Delta_{b,i}(n) (1 - Z_{i,n}))}{X_{i,n}}. \end{aligned}$$

Since $X_{i,n} > 0$ almost surely, for all $n \in \mathbb{Z}_{\geq 1}$, it will not change the sign of the inequality later on, and so we can ignore it to focus only on the numerator. Thus we wish to check if, almost surely,

$$E[\Delta_{r,i}(n) Z_{i,n} - U_{i,n-1} (\Delta_{r,i}(n) Z_{i,n} + \Delta_{b,i}(n) (1 - Z_{i,n})) \mid \mathcal{F}_{n-1}] \leq 0.$$

Now if the curing policy $\{\Delta_{b,i}(n)\}_{n=1}^{\infty}$ for node i satisfies the bound given:

$$\Delta_{b,i}(n) \geq \frac{\Delta_{r,i}(n) (1 - U_{i,n-1}) S_{i,n-1}}{U_{i,n-1} (1 - S_{i,n-1})},$$

then almost surely,

$$\begin{aligned} & E[\Delta_{r,i}(n) Z_{i,n} (1 - U_{i,n-1}) - U_{i,n-1} (1 - Z_{i,n}) \Delta_{b,i}(n) \mid \mathcal{F}_{n-1}] \\ & \leq E \left[\Delta_{r,i}(n) Z_{i,n} (1 - U_{i,n-1}) - U_{i,n-1} (1 - Z_{i,n}) \right. \\ & \quad \left. \times \frac{\Delta_{r,i}(n) (1 - U_{i,n-1}) S_{i,n-1}}{U_{i,n-1} (1 - S_{i,n-1})} \mid \mathcal{F}_{n-1} \right] \\ & = \Delta_{r,i}(n) (1 - U_{i,n-1}) \left[S_{i,n-1} - (1 - S_{i,n-1}) \frac{S_{i,n-1}}{1 - S_{i,n-1}} \right] \\ & = 0, \end{aligned}$$

where the second to last equality comes from the fact that $E[Z_{i,n} \mid \mathcal{F}_{n-1}] = P(Z_{i,n} = 1 \mid \mathcal{F}_{n-1}) = S_{i,n-1}$ almost surely by (3), and the last equality uses that $S_{i,n-1}$ is almost surely constant given \mathcal{F}_{n-1} . Thus as long as $\Delta_{b,i}(n)$ obeys this bound almost surely for all $n \in \mathbb{Z}_{\geq 1}$, $\{U_{i,n}\}_{n=1}^{\infty}$ is a supermartingale with respect to $\{Z_n\}_{n=1}^{\infty}$. Similarly, if $\Delta_{b,i}(n)$ is almost surely equal (resp. less than or equal) to this bound, $\{U_{i,n}\}_{n=1}^{\infty}$ is a martingale (resp. submartingale) with respect to $\{\mathcal{F}_n\}_{n=1}^{\infty}$. \square

B. Proof of Theorem 4

We will focus on the case of a supermartingale, since the proof for submartingales follows similarly. First, note that the question of $\{S_{i,n}\}_{n=1}^{\infty}$ being a strict supermartingale is equivalent to

$$E[S_{i,n} \mid \mathcal{F}_{n-1}] - S_{i,n-1} < 0$$

where $\{\mathcal{F}_n\}$ is the natural filtration (indeed, we can just condition on Z^{n-1}). Note, in particular, that $E[Z_{i,t} \mid \mathcal{F}_n] = Z_{i,t}$ almost surely, for all $i \in V$ and $t \in \{1, \dots, n\}$, and the same is true for $\{S_{i,t}\}_{t=1}^n$. Then almost surely, as in Theorem 2,

$$S_{i,n} - S_{i,n-1} = \frac{S_{i,n-1} (\bar{X}_{i,n-1} - \bar{X}_{i,n}) + \bar{Z}_{r,i,n}}{\bar{X}_{i,n}}.$$

Since $\bar{X}_{i,n} > 0$ almost surely for all $n \in \mathbb{Z}_{\geq 1}$ and all $i \in V$, we can ignore it. Further, since $S_{i,n-1}$ is almost surely constant, we need to only check if

$$\begin{aligned} & E[S_{i,n-1} (\bar{X}_{i,n-1} - \bar{X}_{i,n}) + \bar{Z}_{r,i,n} \mid \mathcal{F}_{n-1}] < 0 \\ & \Leftrightarrow E[(1 - S_{i,n-1}) \bar{Z}_{r,i,n} - S_{i,n-1} \bar{Z}_{b,i,n} \mid \mathcal{F}_{n-1}] < 0, \end{aligned}$$

since $\bar{X}_{i,n-1} - \bar{X}_{i,n} = -\bar{Z}_{r,i,n} - \bar{Z}_{b,i,n}$. Now let the lower bound (B1) be satisfied:

$$\Delta_{b,j}(n) > \Delta_{r,j}(n) \frac{S_{j,n-1}}{1 - S_{j,n-1}} \max_{k \text{ s.t. } j \in \mathcal{N}'_k} \frac{1 - S_{k,n-1}}{S_{k,n-1}}.$$

Notice $E[Z_{j,n} \mid \mathcal{F}_{n-1}] = S_{j,n-1}$ almost surely, so we have

$$\begin{aligned} & E[S_{i,n-1} (\bar{X}_{i,n-1} - \bar{X}_{i,n}) + \bar{Z}_{r,i,n} \mid \mathcal{F}_{n-1}] \\ & < E \left[(1 - S_{i,n-1}) \bar{Z}_{r,i,n} - S_{i,n-1} \sum_{j \in \mathcal{N}'_i} \Delta_{r,j}(n) \frac{S_{j,n-1}}{1 - S_{j,n-1}} \right. \\ & \quad \left. \times \max_{k \text{ s.t. } j \in \mathcal{N}'_k} \frac{1 - S_{k,n-1}}{S_{k,n-1}} (1 - Z_{j,n}) \mid \mathcal{F}_{n-1} \right] \\ & = \sum_{j \in \mathcal{N}'_i} \Delta_{r,j}(n) S_{j,n-1} (1 - S_{i,n-1}) - \Delta_{r,j}(n) \frac{S_{j,n-1}}{1 - S_{j,n-1}} \\ & \quad \times S_{i,n-1} \max_{k \text{ s.t. } j \in \mathcal{N}'_k} \frac{1 - S_{k,n-1}}{S_{k,n-1}} (1 - S_{j,n-1}) \\ & = \sum_{j \in \mathcal{N}'_i} \Delta_{r,j}(n) S_{j,n-1} \\ & \quad \times \left[1 - S_{i,n-1} \left(1 + \max_{k \text{ s.t. } j \in \mathcal{N}'_k} \frac{1 - S_{k,n-1}}{S_{k,n-1}} \right) \right] \\ & = \sum_{j \in \mathcal{N}'_i} \Delta_{r,j}(n) S_{j,n-1} \left[1 - \frac{S_{i,n-1}}{\min_{k \text{ s.t. } j \in \mathcal{N}'_k} S_{k,n-1}} \right]. \end{aligned}$$

Now note that $j \in \mathcal{N}'_i$ and hence, in particular, $\min_{k \text{ s.t. } j \in \mathcal{N}'_k} S_{k,n-1} \leq S_{i,n-1}$ almost surely, and all other quantities are non-negative. Thus, with our value of $\Delta_{b,j}(n)$ for all $j \in \mathcal{N}'_i$, we have almost surely

$$\begin{aligned} & E [S_{i,n-1}(\bar{X}_{i,n-1} - \bar{X}_{i,n}) + \bar{Z}_{i,n} \mid \mathcal{F}_{n-1}] \\ & < \sum_{j \in \mathcal{N}'_i} \Delta_{r,j}(n) S_{j,n-1} \left[1 - \frac{S_{i,n-1}}{\min_{k \text{ s.t. } j \in \mathcal{N}'_k} S_{k,n-1}} \right] \\ & \leq 0. \end{aligned}$$

Thus, for any $i \in V$, if $\{\Delta_{b,i}(n)\}_{n=1}^{\infty}$ obeys the lower bound (B1) almost surely, the neighbourhood proportion of red balls $\{S_{i,n}\}_{n=1}^{\infty}$ is a strict supermartingale.

For the case of a strict submartingale, notice if the upper bound (B2) is satisfied:

$$\Delta_{b,j}(n) < \Delta_{r,j}(n) \frac{S_{j,n-1}}{1 - S_{j,n-1}} \min_{k \text{ s.t. } j \in \mathcal{N}'_k} \frac{1 - S_{k,n-1}}{S_{k,n-1}},$$

then similarly to the case of a supermartingale,

$$\begin{aligned} & E [S_{i,n-1}(\bar{X}_{i,n-1} - \bar{X}_{i,n}) + \bar{Z}_{i,n} \mid \mathcal{F}_{n-1}] \\ & > E \left[(1 - S_{i,n-1}) \bar{Z}_{r,i,n} - S_{i,n-1} \sum_{j \in \mathcal{N}'_i} \Delta_{r,j}(n) \frac{S_{j,n-1}}{1 - S_{j,n-1}} \right. \\ & \quad \left. \times \min_{k \text{ s.t. } j \in \mathcal{N}'_k} \frac{1 - S_{k,n-1}}{S_{k,n-1}} (1 - Z_{j,n}) \mid \mathcal{F}_{n-1} \right] \\ & = \sum_{j \in \mathcal{N}'_i} \Delta_{r,j}(n) S_{j,n-1} \left[1 - \frac{S_{i,n-1}}{\max_{k \text{ s.t. } j \in \mathcal{N}'_k} S_{k,n-1}} \right] \\ & \geq 0, \end{aligned}$$

since $\max_{k \text{ s.t. } j \in \mathcal{N}'_k} S_{k,n-1} \geq S_{i,n-1}$ almost surely. Thus if $\{\Delta_{b,i}(n)\}_{n=1}^{\infty}$ obeys the upper bound (B2) almost surely, the neighbourhood proportion of red balls $\{S_{i,n}\}_{n=1}^{\infty}$ is a strict submartingale. \square

C. Proof of Proposition 6

First note that as a function of the parameters $x = (\Delta_{b,1}(n), \dots, \Delta_{b,N}(n))$, \tilde{S}_n is of the form $f_n(x) = \frac{1}{N} \sum_{i=1}^N \frac{c_i}{d_i + \sigma_i(x)}$, where from (2), we can see that

$$c_i = \bar{R}_i + \sum_{t=1}^{n-1} \bar{Z}_{r,i,t} + \sum_{j \in \mathcal{N}'_i} \Delta_{r,j}(n) Z_{j,n},$$

$$d_i = c_i + \bar{B}_i + \sum_{t=1}^{n-1} \bar{Z}_{b,i,t}, \text{ and}$$

$$\sigma_i(x) = \sum_{j \in \mathcal{N}'_i} x_j (1 - Z_{j,n}).$$

Note that the $Z_{i,t}$ are binary-valued random variables and we always assume $\Delta_{r,i}(t), \Delta_{b,i}(t) \geq 0$ for all i, t . Hence $c_i > 0$, $d_i > 0$, and $\sigma_i(x) \geq 0$ for all nodes i with any possible realization of the nodes' contagion processes $\{Z_{j,t}\}_{j=1}^N$. Now define the summand of $f_n(x)$ as $h_i(y) = \frac{c_i}{d_i + y}$, $y \in \mathbb{R}_{\geq 0}$. Note that $h_i(y)$ is convex on $\mathbb{R}_{\geq 0}$ since $d_i > 0$, and $\sigma_i(x)$ is linear in x . Hence, the composition $h_i(\sigma_i(x))$ is convex, and so $f_n(x) = \frac{1}{N} \sum_{i=1}^N h_i(\sigma_i(x))$ is convex. As a result \tilde{S}_n

is convex in the curing parameters $(\Delta_{b,1}(n), \dots, \Delta_{b,N}(n))$ for all n and any realization of $\{Z_{j,t}\}_{j=1}^N$; hence, its expectation $E[\tilde{S}_n \mid \mathcal{F}_{n-1}]$ is also convex. Lastly, the constraint set $\left\{ \{\Delta_{b,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N \mid \sum_{i=1}^N \Delta_{b,i}(n) = \mathcal{B} \right\}$ is clearly a finite-dimensional simplex and, hence, convex and compact. \square

ACKNOWLEDGEMENTS

The authors wish to acknowledge the Centre for Advanced Computing at Queen's University, whose computing cluster allowed the simulations presented herein to be performed.

REFERENCES

- [1] M. Hayhoe, F. Alajaji, and B. Ghahsifard, "Curing with the network poly contagion model," in *Proc. Amer. Control Conf.*, Jun. 2018, pp. 2644–2650.
- [2] M. Hayhoe, F. Alajaji, and B. Ghahsifard, "A Polya urn-based model for epidemics on networks," in *Proc. Amer. Control Conf.*, 2017, pp. 358–363.
- [3] M. Hayhoe, F. Alajaji, and B. Ghahsifard, "A Polya contagion model for networks," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 4, pp. 1998–2010, Dec. 2018.
- [4] L. Kim, M. Abramson, K. Drakopoulos, S. Kolitz, and A. Ozdaglar, "Estimating social network structure and propagation dynamics for an infectious disease," in *Proc. Int. Conf. Social Comput., Behav.-Cultural Modeling, Predict.* Cham, Switzerland: Springer, 2014, pp. 85–93.
- [5] M. Garetto, W. Gong, and D. Towsley, "Modeling malware spreading dynamics," in *Proc. 22nd Annu. Joint Conf. IEEE Comput. Commun. Societies*, vol. 3, Mar./Apr. 2003, pp. 1869–1879.
- [6] E. M. Rogers, *Diffusion of Innovations*, 5th ed. New York, NY, USA: Simon and Schuster, 2003.
- [7] E. Adar and L. A. Adamic, "Tracking information epidemics in blogspace," in *Proc. IEEE/WIC/ACM Int. Conf. Web Intell.*, Sep. 2005, pp. 207–214.
- [8] P. Van Mieghem, J. Omic, and R. Kooij, "Virus spread in networks," *IEEE/ACM Trans. Netw.*, vol. 17, no. 1, pp. 1–14, Feb. 2009.
- [9] C. Nowzari, V. M. Preciado, and G. J. Pappas, "Analysis and control of epidemics: A survey of spreading processes on complex networks," *IEEE Control Syst.*, vol. 36, no. 1, pp. 26–46, Feb. 2016.
- [10] W. Wang, M. Tang, H. E. Stanley, and L. A. Braunstein, "Unification of theoretical approaches for epidemic spreading on complex networks," *Rep. Prog. Phys.*, vol. 80, no. 3, 2017, Art. no. 036603.
- [11] D. Easley and J. Kleinberg, *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge, U.K.: Cambridge Univ. Press, 2010.
- [12] F. Eggenberger and G. Pólya, "Über die statistik verketteter vorgänge," *ZAMM-J. Appl. Math. Mech./Zeitschrift Angew. Math. Mech.*, vol. 3, no. 4, pp. 279–289, 1923.
- [13] G. Pólya and F. Eggenberger, "Sur l'interprétation de certaines courbes de fréquences," *Comptes Rendus De L'Académie Des Sci.*, vol. 187, pp. 870–872, Jul./Dec. 1928.
- [14] G. Pólya, "Sur quelques points de la théorie des probabilités," *Annales De L'institut Henri Poincaré*, vol. 1, no. 2, pp. 117–161, 1930.
- [15] E. Cator and P. Van Mieghem, "Second-order mean-field susceptible-infected-susceptible epidemic threshold," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 85, no. 5, 2012, Art. no. 056111.
- [16] C. Nowzari, V. M. Preciado, and G. J. Pappas, "Optimal resource allocation for control of networked epidemic models," *IEEE Trans. Control Netw. Syst.*, vol. 4, no. 2, pp. 159–169, Jun. 2017.
- [17] L. Wang and J. T. Wu, "Characterizing the dynamics underlying global spread of epidemics," *Nature Commun.*, vol. 9, no. 1, p. 218, 2018.
- [18] N. Madar *et al.*, "Immunitization and epidemic dynamics in complex networks," *Eur. Phys. J. B*, vol. 38, no. 2, pp. 269–276, 2004.
- [19] C. Borgs, J. Chayes, A. Ganesh, and A. Saberi, "How to distribute antidote to control epidemics," *Random Struct. Algorithms*, vol. 37, no. 2, pp. 204–222, 2010.
- [20] M. H. R. Khouzani, S. Sarkar, and E. Altman, "Optimal control of epidemic evolution," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 1683–1691.
- [21] V. M. Preciado, M. Zargham, and D. Sun, "A convex framework to control spreading processes in directed networks," in *Proc. 48th Annu. Conf. Inf. Sci. Syst.*, 2014, pp. 1–6.

- [22] K. Kandhway and J. Kuri, “How to run a campaign: Optimal control of SIS and SIR information epidemics,” *Appl. Math. Comput.*, vol. 231, pp. 79–92, Mar. 2014.
- [23] P. E. Paré, J. Liu, C. L. Beck, A. Nedić, and T. Başar, “Multi-competitive viruses over static and time-varying networks,” in *Proc. Amer. Control Conf.*, 2017, pp. 1685–1690.
- [24] Q.-H. Liu, L.-F. Zhong, W. Wang, T. Zhou, and H. E. Stanley, “Interactive social contagions and co-infections on complex networks,” *Chaos, Interdiscipl. J. Nonlinear Sci.*, vol. 28, no. 1, 2018, Art. no. 013120.
- [25] R. Ash and C. Doleans-Dade, *Probability and Measure Theory*. New York, NY, USA: Academic, 2000.
- [26] G. R. Grimmett and D. R. Stirzaker, *Probability and Random Processes*, 3rd ed. Oxford, U.K.: Oxford Univ. Press, 2001.
- [27] W. Feller, *An Introduction to Probability Theory and Its Applications*, vol. 2, 2nd ed. New York, NY, USA: Wiley, 1971.
- [28] F. Alajaji and T. Fuja, “A communication channel modeled on contagion,” *IEEE Trans. Inf. Theory*, vol. 40, no. 6, pp. 2035–2041, Nov. 1994.
- [29] A. Banerjee, P. Burlina, and F. Alajaji, “Image segmentation and labeling using the Polya urn model,” *IEEE Trans. Image Process.*, vol. 8, no. 9, pp. 1243–1253, Sep. 1999.
- [30] R. Pemantle, “A survey of random processes with reinforcement,” *Probab. Surv.*, vol. 4, no. 1, pp. 1–79, 2007.
- [31] E. Ramírez-Llanos and S. Martínez, “Distributed and robust fair optimization applied to virus diffusion control,” *IEEE Trans. Netw. Sci. Eng.*, vol. 4, no. 1, pp. 41–54, Jan./Mar. 2017.
- [32] R. Durrett, *Probability: Theory and Examples*, 2nd ed. Belmont, CA, USA: Duxbury Press, 1996.
- [33] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA, USA: Athena Scientific, 1995.
- [34] A. Bavelas, “Communication patterns in task-oriented groups,” *J. Acoust. Soc. Amer.*, vol. 22, pp. 725–730, Jun. 1950.
- [35] L. C. Freeman, “A set of measures of centrality based on betweenness,” *Sociometry*, vol. 40, pp. 35–41, Mar. 1977.
- [36] B. Rieder, “Studying Facebook via data extraction: The Netvizz application,” in *Proc. 5th ACM Web Sci. Conf.*, 2013, pp. 346–355.
- [37] R. Albert and A.-L. Barabási, “Statistical mechanics of complex networks,” *Rev. Mod. Phys.*, vol. 74, no. 1, pp. 47–97, 2002.



Mikhail Hayhoe received the B.Sc. degree in mathematics and engineering and the M.A.Sc. degree in applied mathematics from Queen’s University, Kingston, ON, Canada, in 2015 and 2017, respectively. He is currently pursuing the Ph.D. degree with the Department of Electrical and Systems Engineering, University of Pennsylvania.

His research interests include control and analysis of networked systems, spectral graph theory, machine learning, and optimization. He was a finalist for the Student Best Paper Award at the 2017 American Control Conference.



Fady Alajaji (S’90–M’94–SM’00) received the B.E. degree (Hons.) in electrical engineering from the American University of Beirut, Lebanon, in 1988, and the M.Sc. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, in 1990 and 1994, respectively.

He held a post-doctoral appointment at the Institute for Systems Research, University of Maryland, in 1994. In 1995, he joined the Department of Mathematics and Statistics, Queen’s University, Kingston, ON, Canada, where he is currently a Professor of mathematics and engineering. Since 1997, he has also been cross-appointed at the Department of Electrical and Computer Engineering, Queen’s University. From 2013 to 2014, he was an Acting Head of the Department of Mathematics and Statistics. From 2003 to 2008 and 2018 to 2019, he was the Chair of the Queen’s Mathematics and Engineering Program. His research interests include information theory, digital communications, error control coding, joint source–channel coding, and data compression.

Dr. Alajaji is an Associate Editor for Shannon Theory for the IEEE TRANSACTIONS ON INFORMATION THEORY. He served as an Area Editor for source-channel coding and signal processing from 2008 to 2015 and as an Editor for source and source-channel coding from 2003 to 2012 for the IEEE TRANSACTIONS ON COMMUNICATIONS. He has served as an organizer and a technical program committee member of several international conferences and workshops. He received the Premiers Research Excellence Award from the Province of Ontario.



Bahman Ghareisifard (M’13) received the B.Sc. degree in mechanical engineering and the M.Sc. degree in control and dynamics from Shiraz University, Iran, in 2002 and 2005, respectively, and the Ph.D. degree in mathematics from Queen’s University, Canada, in 2009.

He is an Associate Professor with the Department of Mathematics and Statistics, Queen’s University. He held post-doctoral positions with the Department of Mechanical and Aerospace Engineering, University of California, San Diego, CA, USA, from 2009 to 2012, and with the Coordinated Science Laboratory, University of Illinois at Urbana–Champaign, from 2012 to 2013. He held a visiting faculty position at the Institute for Systems Theory and Automatic Control, University of Stuttgart, in 2016. He has been an Alexander von Humboldt Research Fellow at the same institute since 2019. He received the 2019 CAIMS-PIMS Early Career Award, jointly awarded by the Canadian Applied & Industrial Math Society and the Pacific Institute for the Mathematical Sciences, a Humboldt Research Fellowship for experienced researchers from the Alexander von Humboldt Foundation in 2019, and an NSERC Discovery Accelerator Supplement in 2019. He was a Finalist (as an advisor) for the Best Student Paper Award at the American Control Conference in 2017. He received the Engineering and Applied Science First Year Instructor Teaching Award in 2015 and 2017 and was shortlisted for a Frank Knox Award for Excellence in Teaching in 2014. He also serves on the Conference Editorial Board of the IEEE Control Systems Society and as an Associate Editor for the IEEE CONTROL SYSTEMS LETTERS. His research interests include systems and controls, distributed control, distributed optimization and learning, geometric control theory, social and economic networks, game theory, geometric mechanics, and applied Riemannian geometry.