

Initialization and Curing Policies for Polya Contagion Networks*

Greg Harrington, Fady Alajaji and Bahman Gharesifard

Abstract

This paper investigates optimization policies for resource distribution in network epidemics using a model that derives from the classical Polya process. The basic mechanics of this model, called the Polya network contagion process, are based on a modified urn sampling scheme that accounts for both temporal and spatial contagion between neighbouring nodes in a network. We present various infection metrics and use them to develop two problems: one which takes place upon initialization and one which occurs continually as the Polya network process develops. We frame these problems as resource allocation problems with fixed budgets, and analyze a suite of potential policies. Due to the complexity of these problems, we introduce effective proxy measures for the average infection rate in each case. We also prove that the two-sided infection-curing game on the so-called expected network exposure admits a Nash equilibrium. In both the curing and initialization scenarios, we introduce heuristic policies that primarily function on the basis of limiting the number of targeted nodes within a particular network setup. Simulations are run for mid-to-large scale networks to compare performance of our heuristics to provably convergent gradient descent algorithms run on the simplified proxy measures.

Index Terms

Network epidemics; interacting Polya urn models; processes with reinforcement; temporal and spatial contagion; initialization and curing policies.

I. INTRODUCTION

The study of epidemics on networks is an active research topic, e.g., see [1]–[3] and the references therein. Real-life examples include the propagation of burst errors in a wireless communication channel [4], of a biological disease through a population [5], of malware in computer or smartphone systems [6], and the dissemination of rumors [7] and competing opinions [8] in social networks.

Controlling the spread of contagion in networks has been thoroughly investigated for various systems, including the well-known Susceptible-Infected-Susceptible (SIS) and Susceptible-Infected-Recovered (SIR) epidemics models [2] and extensions [9], models based on real epidemic data [10], as well as cascade models [1]. Other problem formulations address link removal via immunization [11], optimization of curing resources [12], message-passing methods [13], optimal control under performance and resource usage trade-offs [14], [15], and competing or co-evolving contagions within networks [16]–[20]. We herein study initialization and curing strategies and related game-theoretic problems for the (infinite-memory) Polya network contagion model introduced in [21], [22].

The classical Polya process [23]–[25] is a temporal contagion process that evolves through a method of sampling from a single urn containing a finite number of red and black balls, representing units of “infection” and “healthiness,” respectively. This model has been used in applications such as consensus dynamics [26] and generalized for various other purposes [27], [28]. The Polya network contagion process [21], [22], generalizes this sampling scheme to a general network, where each node is equipped with its own urn, by allowing for spatial interactions between neighbouring nodes. This is realized by introducing “super urns,” which are allocated for each network

*Department of Mathematics and Statistics, Queen’s University, Kingston, Ontario K7L 3N6, Canada (12gth@queensu.ca, fa@queensu.ca, bahman.gharesifard@queensu.ca).

The work of the authors was funded by the National Science Foundation of Canada. The last author wishes to thank the Alexander von Humboldt Foundation for their generous support.

node and contain red and black balls from both a node’s individual urn as well as those of neighbouring nodes [21], [22]. In addition to temporal contagion, this modification generates spatial contagion not present in the classical model.

The resulting Polya model for network contagion [21], [22] is similar to the SIS model [1], [29], [30]. The Polya model can be operated in two modes: an infinite-memory mode where the reinforcing balls added at each time step remain permanently in the underlying urn of each node, and a finite-memory mode where the added balls at a given time step remain in the system only for a fixed number of future time steps (this mode results in a finite-order Markov network contagion process). While it was empirically observed in [22] that the Polya network contagion model can mimic the behaviour of the SIS process, particularly in the finite-memory (Markovian) mode, there are some notable differences between the two models. For example, the Polya model benefits from exactly computable expressions for the joint and marginal probabilities of infection of its discrete-time contagion process, while the underlying (or exact) Markov process of the SIS model is typically represented and analyzed via a dynamical system using mean-field approximations; e.g., see [2], [30]–[32]. Furthermore, the Polya model captures the network contagion process at a microscopic level by associating an urn to each node, whose composition of red and black balls at each time step represents a granular profile for the degree of infection and healthiness of that node, which upon interaction with urns of nodes in proximity, results in infected/healthy states with each super urn draw. This allows us to minutely and analytically capture the stochastic evolution of the network-wide contagion process. Finally, the underlying discrete-time Markov chain of the SIS model has an absorbing (all healthy or disease-free) state [31], while the Polya model does not in general* as its contagion process is symbiotically generated from its initial state (of red/black ball mixtures in the individual urns of the network nodes) via the Polya ball sampling mechanism.

In this paper, we examine resource allocation problems for the (infinite-memory) Polya network contagion process. The first problem considers a one-time application of a control policy upon initialization of the network, while the second concerns the continual distribution of resources as the network contagion process develops. More specifically, the initialization problem concerns one party judiciously controlling the distribution of resources in an effort to minimize the average network-wide infection rate. This is a novel *preemptive* infection mitigation problem, not typically considered in SIS models, due to the Polya model’s inherent characteristics where contagion organically develops and propagates from initial network conditions. Resources in this case concern the distribution of red and black balls within a network, prior to any draws taking place. We focus on the one-sided finite horizon case, wherein a player controls the allocation of black balls according to a fixed budget, with the goal of minimizing the average infection rate at some fixed point in the future. We provide a series of results showing that optimal policies for this problem satisfy two conditions: nested nodes will receive no resources, and symmetric red ball initializations will yield symmetric black ball initializations (the specifics of what is meant by “symmetry” in this case is further detailed using fundamental ideas from graph theory). Since determining the average infection rate at time $n \geq 1$ becomes increasingly complex as n increases, we use the average infection rate at time $n = 1$ to simplify the problem, and we prove that under certain conditions this proxy measure admits an optimal policy.

*The Polya model is more general than the exact SIS Markov process in terms of allowing different asymptotic behaviours for the contagion process. For example, in the finite-memory M mode with homogeneous parameters, the Polya network contagion process is an M ’th order irreducible and aperiodic Markov chain with a unique stationary distribution that is not necessarily the all-healthy state with probability one.

An algorithm to obtain such an optimal policy using gradient descent is given and used as a base with which we compare our own heuristic policies. We detail three varieties of heuristic policies, each of which works on the basis of limiting our set of viable target nodes. The first set of policies target the set of “inner nodes”, the second works through a layered application of this inner node targeting technique, and the last works by iteratively targeting the most central nodes until full network coverage (either direct or indirect) is obtained. We provide two different algorithms for determining these target sets and variations based on previously tested heuristics (see [33]).

The second control problem, referred to as the “infection-curing problem” or the “Delta-curing problem,” concerns *reactive* intervention policies deployed during the progress of contagion in the network. In this setup, which is an extension of the contagion curing problem studied in [33], we define a two-player game on the average infection rate for a given network. Prior work on competitive dynamics over SIS and social networks include [16] and [17], respectively. In order to simplify the problem of finding optimal control policies, we consider a game on the proxy measure of the expected network exposure. We show that the expected network exposure is convex as a function of the curing parameters and concave as a function of the infection parameters. We prove that under budget constraints, there exists a deterministic Nash equilibrium [34] for the game on the expected network exposure, which can be determined using gradient descent algorithms. We also develop a set of heuristic policies for the Delta-curing problem using the same methods as in the initialization case. We next provide simulations to demonstrate the performance of the proposed heuristics for the initialization and curing problems on both sparse and dense large-scale networks. These policies are compared to optimal policies for the one-step proxy measures.

The rest of the paper is organized as follows. In Section II, we describe in sufficient detail the (infinite-memory) Polya network contagion process [22]. The initialization and infection-curing problems are theoretically investigated in Sections III and IV, respectively. Heuristic optimization strategies for the above problems are developed in Section V and simulation results are presented in Section VI. Finally, conclusions are stated in Section VII.

II. POLYA NETWORK CONTAGION

We devote this section to some background on the classical Polya urn process and the (infinite-memory) Polya network contagion process under study. Throughout, a number of results are derived using standard probability concepts, which can be found in various texts including [35], [36]. Herein, let (Ω, \mathcal{F}, P) be a probability space. Consider the stochastic process $\{Z_n\}_{n=1}^{\infty}$, where $Z_n = (Z_{1,n}, \dots, Z_{N,n})$ is a random vector on Ω . Here we think of the process indices as time indices, while the vector indices delineate spatial indices (to distinguish nodes within a network). For $i = 1, \dots, N$, we write the n -tuple $(Z_{i,1}, \dots, Z_{i,n})$ as Z_i^n , and more generally, we write $Z_{i,s}^t = (Z_{i,s}, Z_{i,s+1}, \dots, Z_{i,t})$ for $s < t$.

A. Classical Polya Process [23]–[25]

An urn initially contains $T = R + B$ balls of which $R \in \mathbb{Z}_{>0}$ are red and $B \in \mathbb{Z}_{>0}$ black. At each time, n , a ball is drawn from the urn and returned with $\Delta > 0$ balls of the same colour. The random variable Z_n is used to indicate the colour of the ball on the n th draw, where

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

We can then use U_n to denote the proportion of red balls in the urn after the n th draw, where

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

Here $\rho_c = \frac{R}{T}$ is the initial proportion of red balls in the urn and $\delta_c = \frac{\Delta}{T}$ is a correlation parameter. The conditional probability of drawing a red ball at time n , given the past history of draws $Z^{n-1} := (Z_1, \dots, Z_{n-1})$, is given by

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

Additionally, we note that the process $\{Z_n\}_{n=1}^{\infty}$ is exchangeable (thus stationary) and non-ergodic and that both U_n and the process sample average $\frac{1}{n} \sum_{t=1}^n Z_t$ converge almost surely (a.s.) to a random variable governed by a Beta distribution with parameters $\frac{\rho_c}{\delta_c}$ and $\frac{1-\rho_c}{\delta_c}$ [25], [37]. Variations on this model have been explored in a number of other works. We focus exclusively on an adaptation of this model first introduced in [21] and further explored in [22], [33], designed as a network process that introduces spatial contagion to the classical temporal contagion process presented above.

B. Polya Network Contagion Process

We now review the Polya network contagion process [21], [22]. Let $\mathcal{G} = (V, \mathcal{E})$ be a graph, where $V = \{1, \dots, N\}$ is the set of $N \in \mathbb{Z}_{>0}$ nodes and $\mathcal{E} \subset V \times V$ is the set of edges. Throughout, \mathcal{G} is taken to be undirected, i.e., $(i, j) \in \mathcal{E}$ if and only if $(j, i) \in \mathcal{E}$ for all $i, j \in V$, and is assumed to be connected, i.e., there exists a series of edges connecting any two nodes in \mathcal{G} . The set of neighbours to node i is given by $\mathcal{N}_i = \{v \in V : (i, v) \in \mathcal{E}\}$. We then define $\mathcal{N}'_i = \{i\} \cup \mathcal{N}_i$. Each node $i \in V$ is assigned an urn with $R_i \in \mathbb{Z}_{\geq 0}$ red balls and $B_i \in \mathbb{Z}_{\geq 0}$ black balls, with the total number of balls at node i denoted by $T_i = R_i + B_i$. We denote the black and red ball initializations by the vectors $\mathbf{B} := (B_1, \dots, B_N)$ and $\mathbf{R} := (R_1, \dots, R_N)$, respectively. We will refer to \mathbf{B} as the curing initialization and \mathbf{R} as the infection initialization.

We next introduce the concept of a super urn for a given node i , which contains the red and black balls at node i as well as the red and black balls of each of its neighbouring nodes. This concept is illustrated in Figure 1. For each node $i \in V$, we use $\bar{R}_i = \sum_{j \in \mathcal{N}'_i} R_j$ and $\bar{B}_i = \sum_{j \in \mathcal{N}'_i} B_j$ to denote the number of red and black balls, respectively, in node i 's super urn, and we assume throughout that $\bar{R}_i + \bar{B}_i > 0$ for all $i \in V$ (to ensure that the super urn associated to each node is non-empty). The total number of balls in the i th super urn is given by $\bar{T}_i = \bar{R}_i + \bar{B}_i$. Similar to the classical Polya process, a draw is conducted for each node at each time, and then a number of balls of the same colour is added to that node's individual urn. In this case, however, the draw is conducted on the super urn of each node. As well, we may allow for the number of added balls to vary based on which node the draw was for, the colour of the draw, and the time at which the draw occurred. Thus, for node i at time n , if a red ball is drawn, we add $\Delta_{r,i}(n)$ red balls to node i 's individual urn, and if a black ball is drawn, we add $\Delta_{b,i}(n)$ black balls to node i 's individual urn. We also assume that there exists $i \in V$ and n such that $\Delta_{b,i}(n) + \Delta_{r,i}(n) > 0$. We let $Z_{i,n}$ indicate the colour of the n th draw for node i :

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

Translating the above parameters to the framework of an epidemic in a population, black and red balls represent units of “healthiness” (such as white blood cells) and “infection” (such as bacteria or viruses), respectively [22]. In the super urn of a given node (such as a “person” or an “organism”), white blood cells (resp., bacteria) combine to improve (resp., impair) the overall health in the node’s neighborhood. Drawing red at time n from the super urn means that the bacteria in the neighborhood were successful in replicating, making the individual more infected or less immune; otherwise, they were healthier since the white blood cells reproduced (as balls of the the same color as the drawn balls are added). Thus, when $Z_{i,n} = 1$, we declare that node i is infected at time n , and if $Z_{i,n} = 0$, then it is healthy. We refer to $\{\Delta_{b,i}(n)\}_{n=1}^{\infty}$ as the curing parameters and $\{\Delta_{r,i}(n)\}_{n=1}^{\infty}$ as the infection parameters for node i . From an urn-sampling point of view, these Δ ’s are nonnegative integers; however, we allow them to be nonnegative real-valued numbers for mathematical convenience (e.g., when determining optimal control policies). We will likewise do the same for the initialization parameters \mathbf{B} and \mathbf{R} . Ultimately, the collection of draw values

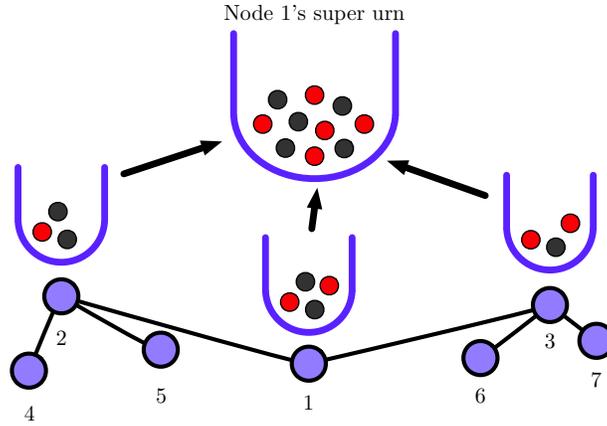


Fig. 1. Illustration of a super urn in a network [38].

at each time step form our Polya network contagion process, represented as $\{Z_n\}_{n=1}^{\infty}$, where Z_n represents the vector $(Z_{1,n}, \dots, Z_{N,n})$. If desired, we can also separate out the individual draw process $\{Z_{i,n}\}_{n=1}^{\infty}$ for a given node i . We introduce some metrics to measure the infection spread in the network.

Similar to the classical Polya process, we use $U_{i,n}$ to denote the proportion of red balls in node i ’s urn after the n th draw. Letting

$$X_{i,n} = T_i + \sum_{t=1}^n \Delta_{r,i}(t)Z_{i,t} + \Delta_{b,i}(t)(1 - Z_{i,t}) \quad (2)$$

with $T_i > 0$, represent the total number of balls in node i ’s urn after the n th draw, we can write

$$U_{i,n} = \frac{R_i + \sum_{t=1}^n \Delta_{r,i}(t)Z_{i,t}}{X_{i,n}}. \quad (3)$$

We note that $U_{i,0} = \frac{R_i}{T_i}$ represents the initial (at time $n = 0$) proportion of red balls in node i ’s individual urn. Though for the classical Polya process the individual urn proportion is commonly used in the analysis, in this case, the information we gain through this value is limited by the fact that our draws are performed on the super urns for each node, rather than the urns for the individual nodes themselves. We therefore also need a metric to measure the proportion of red balls in the super urn of node i at time n , which we denote by $S_{i,n}$. This proportion is given

by the equation

$$S_{i,n} = \frac{\bar{R}_i + \sum_{j \in \mathcal{N}'_i} \sum_{t=1}^n \Delta_{r,j}(t) Z_{j,t}}{\sum_{j \in \mathcal{N}'_i} X_{j,n}} = \frac{\sum_{j \in \mathcal{N}'_i} U_{j,n} X_{j,n}}{\sum_{j \in \mathcal{N}'_i} X_{j,n}}. \quad (4)$$

Like for the individual node proportions at time $n = 0$, we use the convention $S_{i,0} = \frac{\bar{R}_i}{T_i}$. This is simply representative of the proportion of red balls in node i 's super urn when the process is initialized. We next define

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

which we refer to as the *network susceptibility*, and

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

which we refer to as the *network exposure*. Both the network susceptibility and network exposure are functions of the underlying network contagion process $\{Z_n\}_{n=1}^\infty$, as are the respective measures for each node. In general, for notational ease, we will omit these arguments, unless otherwise noted. For example, for a given node i , we will write $S_{i,n}$ rather than $S_{i,n}(\{Z_j^n\}_{j=1}^N)$.

We next calculate infection probabilities within the network, i.e., the likelihood of drawing red balls. To start, the conditional probability of drawing a red ball for node i at time n , given the past history of draws for all the nodes

$$\{Z_j^{n-1}\}_{j=1}^N := \{(Z_{1,1}, \dots, Z_{1,n-1}), \dots, (Z_{N,1}, \dots, Z_{N,n-1})\},$$

is given by

$$P\left(Z_{i,n} = 1 \mid \{Z_j^{n-1}\}_{j=1}^N\right) = \frac{\bar{R}_i + \sum_{j \in \mathcal{N}'_i} \sum_{t=1}^{n-1} \Delta_{r,j}(t) Z_{j,t}}{\sum_{j \in \mathcal{N}'_i} X_{j,n-1}} = S_{i,n-1}. \quad (5)$$

This is analogous to (1) of the classical Polya process, except that for the network contagion process this conditional probability is represented by the super urn proportion rather than the individual urn proportion. We also need an unconditioned version of this equation. The n -fold joint probability of the network \mathcal{G} , which is a function defined for each specific draw history $a_i^n \in \{0, 1\}^n$, $i \in V$, can be written as

$$\begin{aligned} P_{\mathcal{G}}^{(n)}(a_1^n, \dots, a_N^n) &:= P(\{Z_i^n = a_i^n\}_{i=1}^N) \\ &= \prod_{t=1}^n P(\{Z_{i,t} = a_{i,t}\}_{i=1}^N \mid \{Z_i^{t-1} = a_i^{t-1}\}_{i=1}^N) \\ &= \prod_{t=1}^n \prod_{i=1}^N (S_{i,t-1})^{a_{i,t}} (1 - S_{i,t-1})^{1-a_{i,t}}. \end{aligned} \quad (6)$$

It is important to note here that the draw value of different nodes at a given time t are independent when conditioned on the history of the process up to that point, thus allowing us to write $P_{\mathcal{G}}^{(n)}$ in the form above. Finally, we introduce

the metric

$$\tilde{I}_n := \frac{1}{N} \sum_{i=1}^N P(Z_{i,n} = 1), \quad (7)$$

which we call the *average infection rate at time n*. Since $P(Z_{i,n} = 1)$ represents the marginal probability that a node i is infected at time n , the average infection rate gives us the network-wide average of seeing infection within a node. Thus, this value represents the marginal probability that a randomly selected node within the network is infected at time n . In most cases where we look to study the asymptotic behaviour of the network contagion process, we do so in regard to this metric.

We can further break down (7) by noting that

$$\begin{aligned} P(Z_{i,n} = 1) &= \sum_{\{a_j^{n-1}\}_{j=1}^N} P(Z_{i,n} = 1 | \{Z_j^{n-1} = a_j^{n-1}\}_{j=1}^N) P(\{Z_j^{n-1} = a_j^{n-1}\}_{j=1}^N) \\ &= \sum_{\{a_j^{n-1}\}_{j=1}^N} S_{i,n-1}(\{a_j^{n-1}\}_{j=1}^N) P_G^{(n-1)}(\{a_j^{n-1}\}_{j=1}^N), \end{aligned} \quad (8)$$

where the summation is over all possible draw histories $a_j^{n-1} \in \{0, 1\}^{n-1}$, $j \in V$.

Each of these measures, \tilde{U}_n , \tilde{S}_n , and \tilde{I}_n , are closely related to one another and can be useful in various situations. However, all of these values are network-specific as in general all three depend in some way on the underlying network topology and the initial ball distributions. The quantity \tilde{I}_n is particularly difficult to analyze, and hence we often resort to using alternative measures that allow us to simplify our analysis or gain insight into how \tilde{I}_n behaves.

By inspecting (3) and (4), it is clear that an increase in the node-specific proportion of red balls increases the super urn proportion of red balls in all neighbouring nodes. Likewise, if we consider how (5) relates to (8), it would make sense that an increase in the super urn proportion would also lead to an increase in the marginal network-wide infection rate; we further discuss these observations in later sections. For now, we have that

$$\uparrow U_{i,n} \implies \uparrow S_{j,n} \text{ for all } j \in \mathcal{N}'_i,$$

and since \tilde{I}_n is essentially the marginal version of the conditional \tilde{S}_n , we often start our analysis by using \tilde{S}_n in place of \tilde{I}_n , without losing significant information.

III. INITIALIZATION PROBLEMS

When we consider the problem of limiting the spread of infection for the Polya network contagion model, there are different ways to implement a curing policy. One method is to control the allocation of our curing parameters $\{\Delta_{b,i}(n)\}_{i=1}^N$ at each time n , which we explore in Section IV. We can, alternatively, control how the process is initialized. This is where we currently draw our focus for discussion.

The initialization parameters $\mathbf{B} \in \mathbb{R}_{\geq 0}^N$ and $\mathbf{R} \in \mathbb{R}_{\geq 0}^N$ can be tailored to alter the evolution of the Polya network contagion process for a given network. The goal is to determine how to allocate such resources so as to achieve a more desirable result, and determine whether or not a theoretically optimal initial allocation policy exists. Such optimizations may be performed for either or both sets of initialization parameters subject to an initialization budget

$\mathcal{B} \in \mathbb{R}_{\geq 0}$. We start with a formal setup for the one-sided initialization problem, wherein we aim to minimize the average infection rate for a given time n over the black ball initialization parameters \mathbf{B} subject to a budget \mathcal{B}_b . Throughout this section we assume the initial distribution of red balls \mathbf{R} is fixed such that $\bar{R}_i > 0$ for all $i \in V$.

Problem 1: (One-Sided Finite Horizon Constrained Budget Initialization): For a fixed time n , minimize the average infection rate \tilde{I}_n subject to a budget \mathcal{B}_b on the curing initialization $\mathbf{B} = (B_1, \dots, B_N)$, i.e., find:

$$\min_{\{B_i\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N : \sum_{i=1}^N B_i \leq \mathcal{B}_b} \tilde{I}_n.$$

We remark that the problem above can be extended to the infinite horizon case, but for the purposes of our work, we limit ourselves to a finite horizon. Moreover, one can establish a *two-sided* initialization problem, wherein the objective is to minimize the average infection rate for a given n , but we do so over both sets of initialization parameters \mathbf{B} and \mathbf{R} . We, however, only focus on Problem 1 in this work.

We first consider Problem 1 for $n = 1$. We have

$$\tilde{I}_1 = \frac{1}{N} \sum_{i=1}^N P(Z_{i,1} = 1) = \frac{1}{N} \sum_{i=1}^N \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_i}. \quad (9)$$

As a function of the initialization parameters, \tilde{I}_1 is well-defined over $\{(\mathbf{B}, \mathbf{R}) \in \mathbb{R}_{\geq 0}^{2N} | \bar{B}_i + \bar{R}_i \neq 0, \forall i \in V\}$. Furthermore, \tilde{I}_1 is convex in \mathbf{B} and concave in \mathbf{R} . We use the form of (9) to provide some simplifications regarding solutions to Problem 1 for $n = 1$. We will then extend these results to a general time n .

A. Outer Node Allocation

We first show that for the one-sided initialization problem (Problem 1) with $n = 1$, nodes with a nested neighbourhood (i.e., nodes with neighborhood \mathcal{N}' contained in the neighborhood of another node) can be ignored. We refer to such nodes as “outer nodes” due to their topological location.

Lemma 1: For a general network $\mathcal{G} = (V, \mathcal{E})$ equipped with the Polya network contagion model, if for any nodes $i, j \in V$ we have $\mathcal{N}'_i \subset \mathcal{N}'_j$, then it can be assumed that $B_i = 0$ when minimizing \tilde{I}_1 over the initial distribution of black balls. Likewise, we can assume $R_i = 0$ when maximizing \tilde{I}_1 over the initial distribution of red balls.

Proof: Consider a given initial distribution of black balls $\mathbf{B} = (B_1, \dots, B_N)$, and red balls $\mathbf{R} = (R_1, \dots, R_N)$. As written in (9), we have

$$\tilde{I}_1 = \frac{1}{N} \sum_{k=1}^N \frac{\bar{R}_k}{\bar{R}_k + \bar{B}_k}.$$

Without loss of generality, assume that $\mathcal{N}'_1 \subset \mathcal{N}'_2$, and let the distribution $\mathbf{B}^* = (0, B_1 + B_2, B_3, \dots, B_N)$ be an alternative curing initialization. For a given node $k \in V$, we have 3 cases to consider:

- 1) $k \in \mathcal{N}'_1 \implies k \in \mathcal{N}'_2$;
- 2) $k \notin \mathcal{N}'_1$ and $k \in \mathcal{N}'_2$;
- 3) $k \notin \mathcal{N}'_2 \implies k \notin \mathcal{N}'_1$.

For cases 1 and 3, note that $\sum_{l \in \mathcal{N}'_k} B_l = \sum_{l \in \mathcal{N}'_k} B_l^*$. For case 2, we have $\sum_{l \in \mathcal{N}'_k} B_l \leq \sum_{l \in \mathcal{N}'_k} B_l^*$ with equality if

and only if $B_1 = 0$. Denoting $\bar{B}_k^* = \sum_{l \in \mathcal{N}'_k} B_l^*$, we have

$$\frac{\bar{R}_k}{\bar{R}_k + \bar{B}_k} \geq \frac{\bar{R}_k}{\bar{R}_k + \bar{B}_k^*}, \forall k \in V.$$

Therefore, for \tilde{I}_1 as a function of the initial distribution of black balls we have that $\tilde{I}_1(\mathbf{B}^*) \leq \tilde{I}_1(\mathbf{B})$. The proof for the maximization case follows similarly. ■

To extend Lemma 1 for general n , we derive the following intermediate results.

Lemma 2: Consider two sequences of draw values $\{a_j^n\}_{j=1}^N$ and $\{b_j^n\}_{j=1}^N$, where both sequences are equal, except for some $(k, s) \in V \times \{1, \dots, n\}$, where $a_{k,s} = 1$ and $b_{k,s} = 0$. Then $P(Z_{i,n+1} = 1 | \{Z_j^n = a_j^n\}_{j=1}^N) \geq P(Z_{i,n+1} = 1 | \{Z_j^n = b_j^n\}_{j=1}^N)$, for all $i \in V$.

Proof: Assume that $k \in \mathcal{N}'_i$ (note that if this is not the case, the two values we wish to compare are equal).

Let

$$y = \sum_{j \in \mathcal{N}'_i} \sum_{t=1}^n a_{j,t} \Delta_{r,j}(t) \quad x = \sum_{j \in \mathcal{N}'_i} \sum_{t=1}^n (1 - a_{j,t}) \Delta_{b,j}(t).$$

Then letting $y^* = y - \Delta_{r,k}(s)$, and $x^* = x + \Delta_{b,k}(s)$, we obtain

$$\begin{aligned} P(Z_{i,n+1} = 1 | \{Z_j^n = a_j^n\}_{j=1}^N) &\geq \frac{\bar{R}_i + y}{\bar{R}_i + \bar{B}_i + y + x^*} \\ &\geq \frac{\bar{R}_i + y^*}{\bar{R}_i + \bar{B}_i + y^* + x^*} \\ &= P(Z_{i,n+1} = 1 | \{Z_j^n = b_j^n\}_{j=1}^N). \end{aligned}$$

■

Intuitively, this result states that the process is self-reinforcing; i.e., higher rates of infection increase the likelihood of the infection recurring. We will use this lemma in the proof of Lemma 3.

Given the general form of \tilde{I}_n given in (7), we next compare the performance of two curing initializations, \mathbf{B} and \mathbf{B}^* , in limiting the spread of infection. In order to differentiate between the resultant probabilities, we will use $*$ to denote any probability obtained using the curing initialization \mathbf{B}^* , e.g., $\tilde{I}_n^* = \tilde{I}_n(\mathbf{B}^*)$.

Lemma 3: For a general network $\mathcal{G} = (V, \mathcal{E})$ equipped with the Polya network contagion model with two different curing initializations $\mathbf{B} = (B_1, \dots, B_N)$ and $\mathbf{B}^* = (B_1^*, \dots, B_N^*)$. If $S_{i,t}^*(a_1^t, \dots, a_N^t) \leq S_{i,t}(a_1^t, \dots, a_N^t)$ for all $(i, t) \in V \times \{1, \dots, n\}$ and any sequence of draws $\mathbf{a}^n = (a_{1,s}, \dots, a_{N,s})_{s=1}^n \in \{0, 1\}^{N \times n}$, then $\tilde{I}_{n+1}^* \leq \tilde{I}_{n+1}$.

Proof: Let $\{Y_{i,t} : t = 1, \dots, n, i = 1, \dots, N\}$ be independent random variables each uniformly distributed on the unit interval $[0, 1]$. For all $i = 1, \dots, N$, let $Z_{i,t}$ and $Z_{i,t}^*$ be inductively defined for $t = 1, \dots, n$ by

$$Z_{i,t} = \begin{cases} 1 & \text{if } Y_{i,t} \leq S_{i,t-1}(\{Z_{i,i}^{t-1}\}_{i=1}^N) \\ 0 & \text{otherwise,} \end{cases} \quad (10)$$

and

$$Z_{i,t}^* = \begin{cases} 1 & \text{if } Y_{i,t} \leq S_{i,t-1}^*(\{Z_{i,i}^{*,t-1}\}_{i=1}^N) \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

Here we assume that $S_{i,0} \geq S_{i,0}^*$ for all $i = 1 \dots, N$. Note that since

$$P(Z_{i,t} = 1 | \{Z_i^{t-1}\}_{i=1}^N) = S_{i,t-1}(\{Z_i^{t-1}\}_{i=1}^N),$$

and $P(Z_{i,t}^* = 1 | \{Z_i^{*,t-1}\}_{i=1}^N) = S_{i,t-1}^*(\{Z_i^{*,t-1}\}_{i=1}^N),$

for all $i = 1, \dots, N$ and $t = 1, \dots, n$, the processes have the required Polya network process distribution with initializations \mathbf{B} and \mathbf{B}^* , respectively.

We claim that if for some $t \in \{1, \dots, n-1\}$ the sequence $\{Z_i^{t-1}\}_{i=1}^N$ dominates $\{Z_i^{*,t-1}\}_{i=1}^N$, in the sense that $Z_{i,s} \geq Z_{i,s}^*$ for all $i = 1, \dots, N$ and $s = 1, \dots, t-1$, then $\{Z_i^t\}_{i=1}^N$ dominates $\{Z_i^{*,t}\}_{i=1}^N$. To prove this claim, note that by repeated application of Lemma 2, we have that

$$S_{i,t-1}(\{Z_i^{t-1}\}_{i=1}^N) \geq S_{i,t-1}^*(\{Z_i^{*,t-1}\}_{i=1}^N),$$

for all $i = 1, \dots, N$; specifically, since $Z_{i,s} \geq Z_{i,s}^*$ means that either $Z_{i,s} = Z_{i,s}^*$ or $Z_{i,s} = 1$ and $Z_{i,s}^* = 0$, Lemma 2 is applied $K = |\{(i, s) : Z_{i,s} \neq Z_{i,s}^*\}|$ times. This, (10), and (11) then give $Z_{i,t} \geq Z_{i,t}^*$, which in turn implies that $\{Z_i^t\}_{i=1}^N$ dominates $\{Z_i^{*,t}\}_{i=1}^N$, finishing the proof of the claim.

To proceed with the proof of the result, by using the conclusion of the claim for $t = 0$, which vacuously holds by the assumption that $S_{i,0} \geq S_{i,0}^*$ for all $i = 1, \dots, N$, we obtain $Z_{i,1} \geq Z_{i,1}^*$, $i = 1, \dots, N$. Applying now the result of the claim inductively, for $t = 2, \dots, n$, we obtain that

$$Z_{i,t} \geq Z_{i,t}^*, \quad i = 1, \dots, N, \quad t = 1, \dots, n$$

with probability one. Thus the event $\{Z_{i,t}^* = 1\}$ implies $\{Z_{i,t} = 1\}$ and so $P(Z_{i,t}^* = 1) \leq P(Z_{i,t} = 1)$, for all $i = 1, \dots, N$ and $t = 1, \dots, n$. Therefore $\tilde{I}_t \leq \tilde{I}_t^*$, for all $t = 1, \dots, n$. ■

We can now use this result to extend Lemma 1 to the case of a general time n .

Theorem 1: For a general network $\mathcal{G} = (V, \mathcal{E})$ equipped with the Polya network contagion model, if for any nodes $i, j \in V$ we have $\mathcal{N}'_i \subset \mathcal{N}'_j$, then it can be assumed that $B_i = 0$ when minimizing \tilde{I}_n over the curing initialization. Likewise, it can be assumed that $R_i = 0$ when maximizing \tilde{I}_n over the infection initialization.

Proof: Without loss of generality, assume $\mathcal{N}'_1 \subset \mathcal{N}'_2$. Consider the general form of $S_{i,n}(a_1^n, \dots, a_N^n)$:

$$\frac{\bar{R}_i + \sum_{j \in \mathcal{N}'_i} \sum_{t=1}^n a_{j,t} \Delta_{r,j}(t)}{\bar{R}_i + \bar{B}_i + \sum_{j \in \mathcal{N}'_i} \sum_{t=1}^n (a_{j,t} \Delta_{r,j}(t) + (1 - a_{j,t}) \Delta_{b,j}(t))}. \quad (12)$$

Letting $B_1^* = 0$, and $B_2^* = B_1 + B_2$, as in the proof of Lemma 1, we have for any $i \in V$ that $\bar{B}_i^* \geq \bar{B}_i$. Plugging into (12) yields that $S_{i,n}^* \leq S_{i,n}$ for any realization of the draw process. This will hold for any $i \in V$ and for any time t . Applying Lemma 3 yields the desired result. ■

B. Symmetry

The next simplification we make regarding solutions to Problem 1 has to do with networks that possess ‘‘symmetry’’ properties, in a sense that we make precise shortly. We start with some preliminaries, referring the reader to [39]–[41] for more details.

Definition 1: An automorphism of a graph $\mathcal{G} = (V, \mathcal{E})$ is a permutation σ of V , such that $(u, v) \in \mathcal{E}$ if and only if $(\sigma(u), \sigma(v)) \in \mathcal{E}$. The set of automorphisms of a given graph forms a group under the composition operation, called the automorphism group of \mathcal{G} and denoted $Aut(\mathcal{G})$.

Definition 2: For a graph $\mathcal{G} = (V, \mathcal{E})$, if $H \leq Aut(\mathcal{G})$ is a subgroup of automorphisms of \mathcal{G} , then $u, v \in V$ are similar under H if there exists an automorphism in H which maps u to v . Equivalence classes defined by similarity under H are called orbits of the graph \mathcal{G} by H . The sets of orbits by H form a partition of the vertices of \mathcal{G} , called an orbit partition.

If there exists an automorphism which maps between two different nodes, it means that those nodes are distinguishable only by their labels. For our purposes, ‘‘symmetry’’ between nodes is characterized by graph automorphisms, and an orbit partition allows us to divide a network based on these symmetries. Our goal is to prove that under certain conditions, sets of nodes that are similar under some $H \leq Aut(\mathcal{G})$ will be treated equivalently; i.e., they will all be given the same number of resources by an optimal policy. We start with a preliminary result.

Lemma 4: Consider a general network $\mathcal{G} = (V, \mathcal{E})$ with a given initialization $(\mathbf{B}, \mathbf{R}) = (B_1, \dots, B_N, R_1, \dots, R_N)$. Suppose there exists a subset of nodes $V' \subset V$ such that $\bar{R}_i = \bar{R}_j$ for all $i, j \in V'$. Now suppose we can find some alternate curing initialization $\mathbf{B}^* = (B_1^*, \dots, B_N^*)$ such that $\bar{B}_i^* = \frac{1}{|V'|} \sum_{j \in V'} \bar{B}_j$ for all $i \in V'$. Then $\sum_{i \in V'} P^*(Z_{i,1} = 1) \leq \sum_{i \in V'} P(Z_{i,1} = 1)$.

Proof: We prove this result by induction on $|V'|$. Start with the base case, where $|V'| = 2$. Consider two nodes $i, j \in V$ such that $\bar{R}_i = \bar{R}_j$. Let $\bar{B}_i^* = \bar{B}_j^* = \frac{1}{2}(\bar{B}_i + \bar{B}_j)$. We have

$$P(Z_{i,1} = 1) + P(Z_{j,1} = 1) = \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_i} + \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_j}, \quad (13)$$

$$P^*(Z_{i,1} = 1) + P^*(Z_{j,1} = 1) = \frac{2\bar{R}_i}{\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j)}. \quad (14)$$

The difference between (13) and (14) yields

$$\begin{aligned} & \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_i} + \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_j} - \frac{2\bar{R}_i}{\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j)} \\ &= \bar{R}_i \left[\frac{2\bar{R}_i + \bar{B}_i + \bar{B}_j}{(\bar{R}_i + \bar{B}_i)(\bar{R}_i + \bar{B}_j)} - \frac{2}{\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j)} \right] \\ &= \bar{R}_i \left[\frac{(2\bar{R}_i + \bar{B}_i + \bar{B}_j)(\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j)) - 2(\bar{R}_i + \bar{B}_i)(\bar{R}_i + \bar{B}_j)}{(\bar{R}_i + \bar{B}_i)(\bar{R}_i + \bar{B}_j)(\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j))} \right] \\ &= \bar{R}_i \left[\frac{2\bar{R}_i^2 + 2\bar{R}_i(\bar{B}_i + \bar{B}_j) + \frac{1}{2}(\bar{B}_i + \bar{B}_j)^2 - 2\bar{R}_i^2 - 2\bar{R}_i(\bar{B}_i + \bar{B}_j) - 2\bar{B}_i\bar{B}_j}{(\bar{R}_i + \bar{B}_i)(\bar{R}_i + \bar{B}_j)(\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j))} \right] \\ &= \bar{R}_i \left[\frac{\frac{1}{2}(\bar{B}_i^2 + \bar{B}_j^2 - 2\bar{B}_i\bar{B}_j)}{(\bar{R}_i + \bar{B}_i)(\bar{R}_i + \bar{B}_j)(\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j))} \right] \\ &= \bar{R}_i \left[\frac{\frac{1}{2}(\bar{B}_i - \bar{B}_j)^2}{(\bar{R}_i + \bar{B}_i)(\bar{R}_i + \bar{B}_j)(\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j))} \right] \geq 0. \end{aligned}$$

Now, consider a subset of nodes $V' \subset V$ such that $\bar{R}_i = \bar{R}_j$ for all $i, j \in V'$, and suppose $|V'| = n$. Without loss of generality, let $V' = \{1, \dots, n\}$. We assume the result holds for any subset of $n - 1$ of these nodes. Thus,

we can write that

$$\sum_{i=1}^{n-1} P(Z_{i,1} = 1) = \sum_{i=1}^{n-1} \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_i} \geq (n-1) \frac{\bar{R}_1}{\bar{R}_1 + \frac{1}{n-1} \sum_{i=1}^{n-1} \bar{B}_i}.$$

Let $\bar{B} = \frac{1}{n-1} \sum_{i=1}^{n-1} \bar{B}_i$, and let $\bar{B}_i^* = \frac{1}{n}((n-1)\bar{B} + \bar{B}_n)$, for all $i \in V'$. We note that $\bar{B}_i^* = \frac{1}{n} \sum_{j \in V'} \bar{B}_j$ for all $i \in V'$. We have

$$\sum_{i \in V'} P(Z_{i,1} = 1) \geq (n-1) \frac{\bar{R}_1}{\bar{R}_1 + \bar{B}} + \frac{\bar{R}_1}{\bar{R}_1 + \bar{B}_n}, \quad (15)$$

$$\sum_{i \in V'} P^*(Z_{i,1} = 1) = n \frac{\bar{R}_1}{\bar{R}_1 + \frac{1}{n}((n-1)\bar{B} + \bar{B}_n)}. \quad (16)$$

To prove the result, we take the difference between the right-hand sides of (15) and (16), obtaining

$$\begin{aligned} & \bar{R}_1 \left[\frac{(n-1)(\bar{R}_1 + \bar{B}_n) + (\bar{R}_1 + \bar{B})}{(\bar{R}_1 + \bar{B})(\bar{R}_1 + \bar{B}_n)} - \frac{n}{\bar{R}_1 + \frac{1}{n}((n-1)\bar{B} + \bar{B}_n)} \right] \\ &= \bar{R}_1 \left[\frac{n\bar{R}_1 + (n-1)\bar{B}_n + \bar{B}}{(\bar{R}_1 + \bar{B})(\bar{R}_1 + \bar{B}_n)} - \frac{n}{\bar{R}_1 + \frac{1}{n}((n-1)\bar{B} + \bar{B}_n)} \right] \\ &= \bar{R}_1 \left[\frac{(n\bar{R}_1 + (n-1)\bar{B}_n + \bar{B})(\bar{R}_1 + \frac{1}{n}((n-1)\bar{B} + \bar{B}_n)) - n(\bar{R}_1 + \bar{B})(\bar{R}_1 + \bar{B}_n)}{(\bar{R}_1 + \bar{B})(\bar{R}_1 + \bar{B}_n)(\bar{R}_1 + \frac{1}{n}((n-1)\bar{B} + \bar{B}_n))} \right] \\ &= \bar{R}_1 \left[\frac{\frac{n-1}{n}(\bar{B}^2 + \bar{B}_n^2 - 2\bar{B}\bar{B}_n)}{(\bar{R}_1 + \bar{B})(\bar{R}_1 + \bar{B}_n)(\bar{R}_1 + \frac{1}{n}((n-1)\bar{B} + \bar{B}_n))} \right] \\ &= \bar{R}_1 \left[\frac{\frac{n-1}{n}(\bar{B} - \bar{B}_n)^2}{(\bar{R}_1 + \bar{B})(\bar{R}_1 + \bar{B}_n)(\bar{R}_1 + \frac{1}{n}((n-1)\bar{B} + \bar{B}_n))} \right] \geq 0. \end{aligned}$$

■

We now use this result to prove that for networks with symmetry, solutions to Problem 1 will be symmetric in nature as well. Suppose for a graph $\mathcal{G} = (V, \mathcal{E})$ we have an automorphism σ of \mathcal{G} . Assuming \mathcal{G} is finite, we have that $Aut(\mathcal{G})$ is a finite group, and thus we can use σ to generate a cyclic subgroup $\langle \sigma \rangle$ of $Aut(\mathcal{G})$. The subgroup $\langle \sigma \rangle$ is necessarily finite and of finite order. We also have that for any node i in V there exists some integer $k \in \mathbb{Z}_{>0}$ such that $i = \sigma^k(i)$. We will denote by k_i the smallest such integer for which this holds for a given $i \in V$.

Theorem 2: For a network $\mathcal{G} = (V, \mathcal{E})$ with a given initialization (\mathbf{B}, \mathbf{R}) , consider an automorphism σ of \mathcal{G} . Let m be the order of the cyclic subgroup $\langle \sigma \rangle$ of $Aut(\mathcal{G})$. Suppose $R_i = R_{\sigma(i)}$, for all $i \in V$. Let \mathbf{B}^* be an alternate curing initialization such that $B_i^* = \frac{1}{m} \sum_{j=1}^m B_{\sigma^j(i)}$ for all $i \in V$. Then, $\tilde{I}_1^* \leq \tilde{I}_1$.

Proof: Note that for a given $i \in V$, $\bar{R}_i = \sum_{j \in \mathcal{N}'_i} R_j$. We have that if $j \in \mathcal{N}'_i$ then $\sigma^k(j) \in \mathcal{N}'_{\sigma^k(i)}$ for any $k \in \mathbb{Z}_{>0}$. Since by assumption $R_j = R_{\sigma(j)} = R_{\sigma^k(j)}$ for all $j \in \mathcal{N}'_i$, we have that $\bar{R}_i = \bar{R}_{\sigma^k(i)}$ for all $i \in V$ and $k \in \mathbb{Z}_{>0}$. From Lemma 4, we need only to prove that $\bar{B}_i^* = \frac{1}{k_i} \sum_{j=1}^{k_i} \bar{B}_{\sigma^j(i)}$, for all $i \in V$. We start with

$$\bar{B}_i^* = \sum_{j \in \mathcal{N}'_i} B_j^* = \sum_{j \in \mathcal{N}'_i} \frac{1}{m} \sum_{l=1}^m B_{\sigma^l(j)}.$$

We note that for a given positive integer p ,

$$p \sum_{l=1}^{k_j} B_{\sigma^l(j)} = \sum_{l=1}^{k_j} (B_{\sigma^l(j)} + B_{\sigma^{l+k_j}(j)} + \cdots + B_{\sigma^{l+(p-1)k_j}(j)}) = \sum_{l=1}^{pk_j} B_{\sigma^l(j)},$$

which follows from the definition of k_j . Since m is the order of $\langle \sigma \rangle$, we have $\sigma^m(j) = j$, and since k_j is the smallest such integer for which this holds, we must have $k_j | m$ (i.e., k_j divides m). Choosing $p = \frac{m}{k_j}$ and using the identity between the left-most and right-most terms in the above equation, we have that

$$\frac{1}{k_j} \sum_{l=1}^{k_j} B_{\sigma^l(j)} = \frac{p}{pk_j} \sum_{l=1}^{k_j} B_{\sigma^l(j)} = \frac{1}{pk_j} \sum_{l=1}^{pk_j} B_{\sigma^l(j)} = \frac{1}{m} \sum_{l=1}^m B_{\sigma^l(j)},$$

and thus

$$\begin{aligned} \bar{B}_i^* &= \sum_{j \in \mathcal{N}_i'} \frac{1}{m} \sum_{l=1}^m B_{\sigma^l(j)} = \frac{1}{m} \sum_{l=1}^m \sum_{j \in \mathcal{N}_i'} B_{\sigma^l(j)} = \frac{1}{m} \sum_{l=1}^m \sum_{j \in \mathcal{N}_{\sigma^l(i)}} B_j \\ &= \frac{1}{m} \sum_{l=1}^m \bar{B}_{\sigma^l(i)} = \frac{1}{k_i} \sum_{l=1}^{k_i} \bar{B}_{\sigma^l(i)}. \end{aligned}$$

The result then follows from Lemma 4. ■

Another way to think about Theorem 2 is that the optimal allocation distributes resources equally within sets of the orbit partition by $\langle \sigma \rangle$. We assume that infection resources are distributed evenly within this orbit partition, a consequence of assuming that $R_i = R_{\sigma(i)}$ for all $i \in V$. We then prove redistributing curing resources within each set of the partition, i.e., letting $B_i^* = \frac{1}{m} \sum_{j=1}^m B_{\sigma^j(i)}$ for all $i \in V$, will result in a decrease in the overall average infection rate.

We finish this part with few remarks. First, due to the complicated nature of \tilde{I}_n , it is difficult to analytically prove this result for time $n > 1$. Also, while Theorem 2 can be applied to a given automorphism, one can push this idea further. Consider a network $\mathcal{G} = (V, \mathcal{E})$, and the orbit partition of \mathcal{G} by $\text{Aut}(\mathcal{G})$, denoted $P = \{V_1, \dots, V_q\}$, where $q \leq N$. For any given initialization, (\mathbf{B}, \mathbf{R}) , for which $\bar{R}_i = \bar{R}_j$ if i and j are similar under $\text{Aut}(\mathcal{G})$, we postulate that we can reduce the average infection rate \tilde{I}_n by redistributing curing resources evenly within each set of the partition. To elaborate, for each $i \in V$, if $i \in V_k$, we conjecture that taking $B_i^* = \frac{1}{|V_k|} \sum_{j \in V_k} B_j$ yields that $\tilde{I}_n(\mathbf{B}^*) \leq \tilde{I}_n(\mathbf{B})$, for any n . This conjecture is supported empirically, but for reasons of space, we omit including a sample simulation.

C. Gradient Descent

As in [22], and further explored in Section IV, one can implement a gradient descent algorithm [42] to find solutions to Problem 1. We first verify that this strategy can be employed for our initialization problem. For the curing problem presented in [22], the objective was to minimize the limiting average infection rate. As this measure is rather complicated to analyze, a proxy measure given by the expected average network exposure, $E[\tilde{S}_n | \mathcal{F}_{n-1}]$, is used instead. For the initialization problem the equivalent method would be to study this value for time $n = 1$, at which point this simply reduces to minimizing the average infection rate at time one, given by \tilde{I}_1 , as a function of the curing initialization \mathbf{B} . In Problem 1, we assign curing resources subject to a budget \mathcal{B}_b . If for a general

Algorithm 1 Constrained gradient descent on a simplex [42]

 $f(x_1, \dots, x_N) \leftarrow$ function to be minimized

Start at an arbitrary node:

 $y_1 = (y_{1,1}, \dots, y_{1,N}) = (\mathcal{B}_b, 0, \dots, 0)$ **for** $k = 1$: *stoptime* **do**

Find the direction of steepest descent:

 $i = \arg \min_{j \in V} \frac{\partial f}{\partial x_j} |_{y_k}$

Move only in that direction:

 $\bar{y}_{k,i} = \mathcal{B}_b$, and $\bar{y}_{k,j} = 0$ for all $j \neq i$

Select the step size using the limit minimization rule:

 $\alpha_k = \arg \min_{\alpha \in [0,1]} f(y_k + \alpha(\bar{y}_k - y_k))$

Perform the gradient descent:

 $y_{k+1} = y_k + \alpha_k(\bar{y}_k - y_k)$ **end for**

time n we instead use \tilde{I}_1 as a simple proxy, it is easy to use gradient flow techniques to find an optimal policy for this one-step problem. We note that any optimal policy will make full use of the budget \mathcal{B}_b .

Proposition 1: (Gradient descent conditions): For a general network $\mathcal{G} = (V, \mathcal{E})$ equipped with the Polya network contagion model, if for a given infection initialization \mathbf{R} we have that $\bar{R}_i > 0$ for each $i \in V$, then the average infection rate at time one, \tilde{I}_1 , is convex with respect to the curing initialization parameters \mathbf{B} . Moreover, the feasible set

$$\mathcal{X} = \left\{ \{B_i\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N \mid \sum_{i=1}^N B_i = \mathcal{B}_b \right\}$$

is convex and compact.

The proof of this proposition is straightforward, with the matter of convexity of \tilde{I}_1 having been discussed earlier in this section. Of note, however, is the fact that we require $\bar{R}_i > 0$ for each $i \in V$, in order to ensure our feasible set \mathcal{X} remains valid. If this is satisfied, we can then use a constrained gradient descent method [42, Chapter 2], as described in Algorithm 1. This algorithm will converge to an optimal solution for the one-step optimization problem, but not necessarily for the n -step case.

IV. CURING-INFECTION PROBLEMS

Another technique to control network infection for the Polya network model is to control the allocation of our curing parameters, $\{\Delta_{b,i}(n)\}_{i=1}^N$, subject to a fixed budget, \mathcal{B} . As in the initialization setup, this can be done for a finite or infinite horizon. Previously, [33] examined the problem of minimizing the limiting average infection rate and introduced a number of heuristic optimization policies. We begin by refocusing the problem in a game-theoretic setup.

We consider a two player game, where player one minimizes the average infection rate, \tilde{I}_n , while player two maximizes this same value. Player one controls the distribution of curing parameters, $\{\Delta_{b,i}(n)\}_{i=1}^N$, in accordance with a budget, \mathcal{B}_b , while player two controls the distribution of the infection parameters, $\{\Delta_{r,i}(n)\}_{i=1}^N$ in accordance with a budget, \mathcal{B}_r . The budget in each case is fixed for all time steps, and the allocation of resources for a given time n is determined prior to any draws being made at that time. Thus, if either player allocates resources to a node for which the opposite colour ball is drawn, those resources will go to waste.

Formally, player one's objective is to find:

$$\min_{\{\{\Delta_{b,i}(k)\}_{i=1}^N: \sum_{i=1}^N \Delta_{b,i}(k) = \mathcal{B}_b\}, k=1, \dots, n} \tilde{I}_n, \quad (17)$$

assuming the minimum exists, while player two's objective is to find:

$$\max_{\{\{\Delta_{r,i}(k)\}_{i=1}^N: \sum_{i=1}^N \Delta_{r,i}(k) = \mathcal{B}_r\}, k=1, \dots, n} \tilde{I}_n, \quad (18)$$

assuming the maximum exists. We refer to (17)-(18) as the Delta-curing problem.

Since finding an optimal control policy for either player is not tractable for a general network, we simplify the problem by adopting the expected network exposure, $E[\tilde{S}_n | \mathcal{F}_{n-1}] = E[\tilde{S}_n | \sigma(\{Z_j^{n-1}\}_{j=1}^N)]$, as the metric to optimize instead of the average infection rate. Using this (strongly correlated) proxy measure drastically facilitates the analysis. We thus consider a two player zero-sum game, where player one minimizes the expected network exposure over the curing parameters $\{\Delta_{b,i}(n)\}_{i=1}^N$ and player two maximizes the same value over the infection parameters $\{\Delta_{r,i}(n)\}_{i=1}^N$. We begin with an important result about the expected network exposure.

Proposition 2: (Convexity-Concavity of Network Exposure): Let $\mathcal{G} = (V, \mathcal{E})$ be a general network, and consider the Polya network contagion model on \mathcal{G} , with arbitrary initial conditions. Then, the expected network exposure $E[\tilde{S}_n | \mathcal{F}_{n-1}]$ is convex with respect to the curing parameters $\{\Delta_{b,i}(n)\}_{i=1}^N$ and concave with respect to the infection parameters $\{\Delta_{r,i}(n)\}_{i=1}^N$, for all n .

Proof:

Using (2) and (4), we consider $E[\tilde{S}_n | \mathcal{F}_{n-1}]$ as a function of the vectors

$$\begin{aligned} \mathbf{x} &= (x_1, \dots, x_N)^T = (\Delta_{b,1}(n), \dots, \Delta_{b,N}(n))^T \\ \mathbf{y} &= (y_1, \dots, y_N)^T = (\Delta_{r,1}(n), \dots, \Delta_{r,N}(n))^T, \end{aligned}$$

by reformulating (4) as follows:

$$S_{i,n} = f_{i,n}(\mathbf{x}, \mathbf{y}, Z_n) = \frac{c_i + \delta_i(\mathbf{y}, Z_n)}{c_i + d_i + \sigma_i(\mathbf{x}, Z_n) + \delta_i(\mathbf{y}, Z_n)},$$

where

$$\begin{aligned} c_i &= \bar{R}_i + \sum_{t=1}^{n-1} \sum_{j \in \mathcal{N}'_i} \Delta_{r,j}(t) Z_{j,t}, & d_i &= \bar{B}_i + \sum_{t=1}^{n-1} \sum_{j \in \mathcal{N}'_i} \Delta_{b,j}(t) (1 - Z_{j,t}) \\ \delta_i(\mathbf{y}, Z_n) &= \sum_{j \in \mathcal{N}'_i} y_j Z_{j,n}, & \sigma_i(\mathbf{x}, Z_n) &= \sum_{j \in \mathcal{N}'_i} x_j (1 - Z_{j,n}). \end{aligned}$$

Alternatively, we let A represent the adjacency matrix, including self-loops, for our given network \mathcal{G} , where $A_{ij} = 1$ if and only if $i \in \mathcal{N}'_j$, and $A_{ij} = 0$ otherwise. We then construct an $N \times N$ square matrix, D , where $D_{ij} = A_{ij}(1 - Z_{j,n})$. Letting D_i represent the i^{th} row of the matrix D , we have that $\sigma_i(\mathbf{x}, Z_n) = D_i \mathbf{x}$. Likewise, we can construct an $N \times N$ square matrix, C , where $C_{ij} = A_{ij} Z_{j,n}$. Letting C_i represent the i^{th} row of the matrix C , we have that $\delta_i(\mathbf{y}, Z_n) = C_i \mathbf{y}$, and hence

$$f_{i,n}(\mathbf{x}, \mathbf{y}, Z_n) = \frac{c_i + C_i \mathbf{y}}{c_i + d_i + C_i \mathbf{y} + D_i \mathbf{x}},$$

where we have dropped the dependencies of the matrices C and D on Z_n for simplicity.

Taking the expectation of $S_{i,n}$ given the history of the contagion process up to time $n - 1$, we have

$$\begin{aligned} E[S_{i,n}|\mathcal{F}_{n-1}] &= E[f_{i,n}(\mathbf{x}, \mathbf{y}, Z_n)|\mathcal{F}_{n-1}] \\ &= \sum_{z_n \in \{0,1\}^N} f_{i,n}(\mathbf{x}, \mathbf{y}, z_n) P(Z_n = z_n | Z^{n-1} = z^{n-1}). \end{aligned} \quad (19)$$

We note that $P(Z_n = z_n | Z^{n-1} = z^{n-1})$ is independent of our choice of \mathbf{x} , and for any fixed realization $z_n = (z_{1,n}, \dots, z_{N,n}) \in \{0,1\}^N$, $f_{i,n}(\mathbf{x}, \mathbf{y}, z_n)$ is convex in \mathbf{x} over $\mathbb{R}_{\geq 0}^N$ (the proof is given in [33]). Hence, it follows that $E[S_{i,n}|\mathcal{F}_{n-1}]$ is convex in \mathbf{x} over $\mathbb{R}_{\geq 0}^N$, and thus so too is $E[\tilde{S}_n|\mathcal{F}_{n-1}]$. Concavity in \mathbf{y} follows from a symmetry argument since

$$1 - f_{i,n}(\mathbf{x}, \mathbf{y}, Z_n) = \frac{d_i + D_i \mathbf{x}}{c_i + d_i + C_i \mathbf{y} + D_i \mathbf{x}}$$

is convex in \mathbf{y} , thus showing $f_{i,n}(\mathbf{x}, \mathbf{y}, Z_n)$ is concave in \mathbf{y} . The rest of the proof follows as in the convexity argument for \mathbf{x} . \blacksquare

The natural symmetry of this model makes this result rather intuitive, and allows us to establish a nice setup for the game theoretic problem. We have proven that $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ is convex in \mathbf{x} and concave in \mathbf{y} over $\mathbb{R}_{\geq 0}^N \times \mathbb{R}_{\geq 0}^N$. In order to get a better understanding of this function, we consider the partial derivatives with respect to \mathbf{x} and \mathbf{y} . We note that

$$\begin{aligned} \nabla_{\mathbf{x}} f_{i,n}(\mathbf{x}, \mathbf{y}, z_n) &= \frac{-\nabla_{\mathbf{x}} D_i \mathbf{x} (c_i + C_i \mathbf{y})}{(c_i + d_i + C_i \mathbf{y} + D_i \mathbf{x})^2} \\ \nabla_{\mathbf{y}} f_{i,n}(\mathbf{x}, \mathbf{y}, z_n) &= \frac{\nabla_{\mathbf{y}} C_i \mathbf{y} (d_i + D_i \mathbf{x})}{(c_i + d_i + C_i \mathbf{y} + D_i \mathbf{x})^2}, \end{aligned}$$

and furthermore,

$$\frac{\partial}{\partial x_j} D_i \mathbf{x} = A_{ij}(1 - z_{j,n}), \quad \text{and} \quad \frac{\partial}{\partial y_j} C_i \mathbf{y} = A_{ij} z_{j,n}.$$

Since A_{ij} and $z_{j,n}$ can only take values in $\{0,1\}$, we have that $\frac{\partial}{\partial x_j} f_{i,n}(\mathbf{x}, \mathbf{y}, z_n) \leq 0$ and $\frac{\partial}{\partial y_j} f_{i,n}(\mathbf{x}, \mathbf{y}, z_n) \geq 0$. Using this fact in conjunction with (19), we determine that over the space $\mathcal{X} \times \mathcal{Y} = \mathbb{R}_{\geq 0}^N \times \mathbb{R}_{\geq 0}^N$, $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ has no saddle point. Given our fixed allocation budgets \mathcal{B}_b and \mathcal{B}_r , we restrict ourselves to considering sets of the form $\mathcal{X} \times \mathcal{Y} = \{ \{ \Delta_{b,i}(n) \}_{i=1}^N \in \mathbb{R}_{\geq 0}^N \mid \sum_{i=1}^N \Delta_{b,i}(n) \leq \mathcal{B}_b \} \times \{ \{ \Delta_{r,i}(n) \}_{i=1}^N \in \mathbb{R}_{\geq 0}^N \mid \sum_{i=1}^N \Delta_{r,i}(n) \leq \mathcal{B}_r \}$. Returning to our game over the expected network exposure, we remark that for any given n , the sets $\mathcal{X} = \{ \{ \Delta_{b,i}(n) \}_{i=1}^N \in \mathbb{R}_{\geq 0}^N \mid \sum_{i=1}^N \Delta_{b,i}(n) \leq \mathcal{B}_b \}$ and $\mathcal{Y} = \{ \{ \Delta_{r,i}(n) \}_{i=1}^N \in \mathbb{R}_{\geq 0}^N \mid \sum_{i=1}^N \Delta_{r,i}(n) \leq \mathcal{B}_r \}$ are convex and compact. This gives rise to the following result.

Theorem 3: (Nash Equilibrium for Network Exposure): Let $\mathcal{G} = (V, \mathcal{E})$ be a general network equipped with the Polya network contagion model, with arbitrary initial conditions. For a given time n , consider a two-player zero-sum game where player one tries to minimize $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ over the parameters $\{ \Delta_{b,i}(n) \}_{i=1}^N$ and player two tries to maximize $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ over the parameters $\{ \Delta_{r,i}(n) \}_{i=1}^N$. Then, if we take our set of allowable policies to be of the form $\mathcal{X} \times \mathcal{Y} = \{ \{ \Delta_{b,i}(n) \}_{i=1}^N \in \mathbb{R}_{\geq 0}^N \mid \sum_{i=1}^N \Delta_{b,i}(n) \leq \mathcal{B}_b \} \times \{ \{ \Delta_{r,i}(n) \}_{i=1}^N \in \mathbb{R}_{\geq 0}^N \mid \sum_{i=1}^N \Delta_{r,i}(n) \leq \mathcal{B}_r \}$, the resulting game admits a Nash equilibrium. Moreover, the equilibrium policy will satisfy $\sum_{i=1}^N \Delta_{b,i}(n) = \mathcal{B}_b$ and $\sum_{i=1}^N \Delta_{r,i}(n) = \mathcal{B}_r$.

Proof: Since the function is convex-concave and over a compact set, the existence follows from the classical minimax theorem, see [43]. By the definition of $E[\tilde{S}_n|\mathcal{F}_{n-1}]$, and since the function has no saddle point in the interior of its domain, the optimal policy will utilize the full budget. ■

The equilibrium policy from Theorem 3 can be determined numerically using gradient descent algorithms [42], as we are optimizing a convex/concave function over a simplex, but for large networks such algorithms can be computationally expensive, especially when considering the complexity of $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ as seen in (19). We next develop efficient heuristic strategies for both the initialization and Delta-curing problems (studied in Section III and the current section, respectively). Naturally, their effectiveness will be measured via the original metric intended for the problems, which is average network infection rate.

V. HEURISTIC STRATEGIES

We provide here heuristic strategies based on the analysis performed thus far, while also integrating some of the previously tested techniques in [38], [33].

When developing heuristics in the initialization case, we make heavy use of Theorem 1, which proved that any optimal solution to Problem 1 would allocate no resources to outer nodes (i.e., nodes with a nested neighbourhood). For the general Delta-curing problem, we have no such result, as it is possible to manufacture situations where the equilibrium policy as in Theorem 3 will assign resources to outer nodes. Such cases arise under specific circumstances, and it is uncommon for these nodes to receive significant resources. Thus, we introduce a similar set of resource allocation policies for both the initialization and Delta-curing problems of Sections III and IV, respectively. While any of these resource allocation policies can be similarly applied to a player controlling the distribution of infection resources, for the purposes of our analysis here, we will consider these strategies from the perspective of the player controlling the curing resources.

A. Interior Node Targeting

We now discuss our first strategy. Given a network, we determine which nodes are inner nodes, i.e., which nodes have a neighbourhood that is not a strict subset of another node's neighbourhood. We then distribute initialization resources uniformly among these nodes. If we let \mathcal{B}_b represent our curing initialization budget, and let $V' = \{i \in V | \mathcal{N}'_i \not\subset \mathcal{N}'_j, \forall j \in V\}$, we allocate our resources according to the following scheme:

$$B_i = \begin{cases} \frac{1}{|V'|} \mathcal{B}_b & \text{if } i \in V' \\ 0 & \text{else.} \end{cases} \quad (20)$$

For the Delta-curing problem, a similar policy is employed as follows:

$$\Delta_{b,i}(n) = \begin{cases} \frac{1}{|V'|} \mathcal{B}_b & \text{if } i \in V' \\ 0 & \text{else,} \end{cases} \quad (21)$$

We can also combine these schemes with the heuristic strategies in [38], [33]. Among these strategies, the best performing heuristic strategy allocates resources based on a combination of centrality measures, node degree and closeness centrality, as well as the infection level of each node's super urn. In the case of initialization policies, as

they are implemented before the onset of contagion, we implement an allocation strategy based on only the node degree and closeness centrality.

The degree of a node is simply measured as the size of a node's neighbourhood, $|\mathcal{N}_i|$. The closeness centrality of a node, C_i , is a measure used to determine a node's topological position within a network, defined as

$$C_i = \frac{1}{\sum_{j \in V} d(i, j)}, \quad (22)$$

where $d(i, j)$ is the shortest path length from node i to node j . Thus the node with the highest closeness centrality has the shortest average path length to every node within the network. We can distribute resources as follows:

$$B_i = \begin{cases} \frac{|\mathcal{N}_i| C_i}{\sum_{j \in V'} |\mathcal{N}_j| C_j} \mathcal{B}_b & \text{if } i \in V' \\ 0 & \text{else.} \end{cases} \quad (23)$$

In the Delta-curing setup, we also account for the evolution of the super urn proportion, $S_{i,n-1}$, of each node. As before, we let C_i represent the closeness centrality of a given node i , calculated by (22). We then allocate resources as follows:

$$\Delta_{b,i}(n) = \begin{cases} \frac{|\mathcal{N}_i| C_i S_{i,n-1}}{\sum_{j \in V'} |\mathcal{N}_j| C_j S_{j,n-1}} \mathcal{B}_b & \text{if } i \in V' \\ 0 & \text{else.} \end{cases} \quad (24)$$

Alternatively, similar strategies based on ‘‘betweenness centrality’’ can be used [33].

B. Minimized Node Targeting

We can improve upon our initialization and curing policies by using the idea of minimizing the number of nodes we target at once. We introduce a method that attempts to target nodes in such a manner that all nodes have resources within their immediate neighbourhood \mathcal{N}'_i (i.e., ensuring $\bar{B}_i > 0, \forall i \in V$), while targeting the fewest possible nodes. We use Algorithm 2 to determine which nodes will be allocated resources.

This algorithm starts by locating outer nodes in the network and then targets nodes directly adjacent these outer nodes. All nodes that are within the neighbourhood of the targeted nodes are then removed from the set of nodes being evaluated. Within this new, smaller set of nodes, we once again locate outer nodes and repeat the process. In the case that there are no outer nodes within the set of nodes being evaluated, all remaining nodes are added to the set of targeted nodes. The purpose of this algorithm is to effectively employ Theorem 1 and ensure that our curing policy does not target outer nodes, while still requiring that $\bar{B}_i > 0$ for these nodes. Once the set of outer nodes is excluded from directly receiving resources, the algorithm then works inwards, attempting to minimize the number of targeted nodes.

Once the set of nodes to be targeted V' has been determined, we can once again distribute resources in one of two manners as in the previous strategy set. We can distribute resources to these nodes uniformly according to (20). Alternatively, we can distribute resources to more central nodes among those targeted using (23). We note that these heuristics are entirely based on network structure, and by design will symmetrically allocate resources in accordance with Theorem 2. Furthermore, by limiting the number of nodes that are targeted, we limit the computational complexity of determining resource allocation.

Algorithm 2 Node Targeting Algorithm

```

 $V \leftarrow$  set of all nodes
 $V' \leftarrow$  target set
 $V_{test} \leftarrow$  nodes to test for this loop
Start with  $V' = \emptyset$ ,  $V_{test} = V$ 
while  $V_{test} \neq \emptyset$  do
  Identify outer nodes:
  Set  $V_{outer} = \{i \in V_{test} \mid \mathcal{N}'_i \subset \mathcal{N}'_j, \text{ for some } j \in V_{test}\}$ 
  if  $V_{outer} = \emptyset$  then
     $V' = V' \cup V_{test}$ 
    break
  end if
  Target nodes adjacent to outer nodes:
  Set  $V_{inner} = V_{test} \setminus V_{outer}$ 
  Set  $V_{added\_targets} = \{i \in V_{inner} \mid i \in \mathcal{N}'_j \text{ for some } j \in V_{outer}\}$ 
  Set  $V' = V' \cup V_{added\_targets}$ 
  Update parameters for next loop:
  Set  $V_{coverage} = \{i \in V_{test} \mid i \in \mathcal{N}'_j \text{ for some } j \in V_{added\_targets}\}$ 
  Set  $V_{test} = V_{test} \setminus V_{coverage}$ 
end while

```

We can apply the same techniques for the Delta-curing setup, where we consider a uniform allocation within the target set via (21), or use a strategy that accounts for the centrality and super urn proportion of the targeted nodes according to (24).

C. Dense Networks

We now consider a limitation of the heuristics presented thus far. Our previous sets of solutions depend on the assumption that a network will possess a set of nodes whose respective neighbourhoods are strict subsets of some other node's neighbourhood. While this may be a useful tactic for sparse networks, it is a stringent requirement for dense networks. To account for this, we introduce a final node targeting strategy that does not have this limitation, but works based on similar principles.

Our previous methods utilized the idea of leveraging Theorem 1, in order to prevent suboptimal node targeting while still ensuring $\bar{B}_i > 0$, for every node $i \in V$. We introduce an alternative method for defining inner and outer nodes when the conditions of Theorem 1 are not met.

This technique, presented in Algorithm 3, begins by ranking the nodes by a centrality measure (closeness centrality for our purposes), from highest to lowest. Nodes are added one-by-one to the set of targeted nodes, V' , in order of rank. Once the set of targeted nodes provides total coverage of the network (i.e., the neighbourhoods of these nodes contain the entire network), the algorithm stops. For an additional reduction, we go through the list of targeted nodes in reverse order and remove any nodes not required to maintain full coverage of the network.

Once the target set has been determined, we can distribute resources either uniformly among these nodes via (20), or preferentially according to centrality as per (23). Likewise for the Delta-curing case, we can allocate resources uniformly among these nodes using (21) or based on centrality and super urn proportion of the targeted nodes via (24).

Algorithm 3 Node Targeting Algorithm for Dense Networks

```

 $V \leftarrow$  set of all nodes
 $V' \leftarrow$  nodes to be targeted
 $C_i \leftarrow$  centrality score of given node
Without loss of generality, assume  $C_1 > C_2 > \dots > C_N$ , where  $N$  is the network size
Start with  $V' = \emptyset$ ,  $V_{coverage} = \emptyset$ ,  $n=0$ 
while  $V_{coverage} \neq V$  do
  Add most central node to target set:
   $V' = V' \cup \{n\}$ 
  Update loop parameters:
  Set  $V_{coverage} = \{i \in V \mid i \in \mathcal{N}'_j, \text{ for some } j \in V'\}$ 
   $n++$ 
end while
Optional:
for  $n = |V'| : 1$  do
  Let  $V_{coverage} = \{i \in V \mid i \in \mathcal{N}'_j, \text{ for some } j \in V' \setminus \{n\}\}$ 
  if  $V_{coverage} = V$  then
    Set  $V' = V' \setminus \{n\}$ 
  end if
end for

```

VI. SIMULATION RESULTS

We evaluate the performance of the heuristic policies detailed in Section V, and compare their performance to supposed ‘optimal’ strategies obtained through gradient descent (note that we say ‘optimal’ because the gradient descent only converges to an optimal policy for the one-step proxy measures we introduced in both the initialization and Delta-curing cases). Algorithm 4 provides the general format of the simulations, while specifics of both the initialization and Delta-curing setups are provided in Sections VI-B and VI-C, respectively.

In order to ensure effectiveness of different strategies, we test these policies for a few different networks, each of which is depicted in Figure 2. These networks include algorithmically generated 100 node Barabasi-Albert networks, depicted in Figures 2(a) and 2(b), with average densities of 1.98 and 18.9 connections per node, respectively. The third network, depicted in Figure 2(c), was generated by a tool [44] that crawled through posts in a Facebook group in order to establish connections between users based on interactions on user posts (either via commenting, or ‘likes’ on the post or comments). The resulting 1,363 node graph contains an average density of 3.56 connections per node, and provides a topology for a real-world social network.

A. Targeting Algorithms

Since the majority of the following analysis applies curing policies based on targeting subsets of nodes, we provide a brief comparison of the different node targeting algorithms.

Figure 3 depicts a comparison of the different targeting algorithms for the sparse network seen in Figure 2(a). For this specific network, the target set in Figure 3(b) is a subset of the set of inner nodes in Figure 3(a), which itself is a subset of the target set shown in Figure 3(c). While, this is not a general rule, it is a common trend for very low-density networks. Algorithm 3 is not designed for use over sparse networks, and does not successfully leverage Theorem 1. Thus, we know that distributing resources among this set of target nodes will result in inefficiencies. If

Algorithm 4 Simulation Setup

```

 $A \leftarrow$  adjacency matrix of the network
 $numTrials \leftarrow$  trials to run for given policy
 $steps \leftarrow$  number of time steps for each trial
 $initPolicy \leftarrow$  initialization policy (uniform in no control case)
 $curePolicy \leftarrow$  Delta-curing policy (uniform in no control case)
Assign  $(\mathbf{B}, \mathbf{R})$  using  $initPolicy$  (under initialization budget)
for  $s = 1 : numTrials$  do
   $\vec{Z}_s \leftarrow$  RUNTRIAL( $A, \mathbf{R}, \mathbf{B}, \mathcal{B}_r, \mathcal{B}_b, steps, curePolicy$ )
end for
RunTrial $A, \mathbf{R}, \mathbf{B}, \mathcal{B}_r, \mathcal{B}_b, steps, curePolicy$ 
Initialize  $S_{i,0}$  using  $R_i$  and  $B_i$  for all  $i \in V$ 
for  $t = 1 : steps$  do
  Assign  $\Delta_{b,i}(t), \Delta_{r,i}(t)$  using  $curePolicy$  (under curing budget)
  Generate  $\vec{Y} \sim \text{Uniform}([0, 1]^N)$ 
  if  $Y_i \leq S_{i,t-1}$  then
     $Z_{i,t} = 1$ 
  else
     $Z_{i,t} = 0$ 
  end if
  Update  $S_{i,t}$  using  $\Delta_{r,i}$  and  $\Delta_{b,i}(t)$  for all  $i \in V$ 
end for

```

we compare this to how Algorithm 3 applies to denser networks, as shown in Figure 3(d), we observe a significant reduction in size of the set of target nodes. Algorithm 2, on the other hand, will be entirely ineffective at determining a subset of nodes to target for this same network, due to the lack of nested nodes.

B. Initialization Trials

We start by comparing solutions to Problem 1, and evaluating their effectiveness for various network compositions. Table I provides an overview of the considered initialization strategies. More detailed explanations of these strategies can be found in Sections III and V. It is important to note that we do not include any trials that use the optional component of Algorithm 3, simply because we have the same target set as Algorithm 2 when applied to the networks in 2(a) and 2(c), and it does not provide any further reduction to the target set when applied to the network in Figure 2(b). The same reasoning applies to our curing strategies in Table II.

We follow the same process for each of the three networks in Figure 2. We begin with equal initialization budgets for both red and black resources (i.e., take $\mathcal{B}_b = \mathcal{B}_r$), and let the infection initialization be uniform over the network (i.e., take $R_i = \frac{\mathcal{B}_r}{N}$, for all $i \in V$). We choose our curing initialization \mathbf{B} in accordance with the desired strategy from Table I. Then, we run the Polya network contagion process with $\Delta_{r,i}(n) = \Delta_{b,i}(n) = \Delta$ fixed for each node $i \in V$ and each time n . We record the draw value, $Z_{i,n}$, for each node i at each time n . The draw values are averaged over the number of trials that are run and then averaged over the set of nodes in the network to obtain an empirical measure for the average infection rate \tilde{I}_n at a given time n . The resulting output of these trials can be seen in Figure 4.

We start our comparison with a few observations. First, as expected, the solution determined using the gradient descent algorithm results in the lowest empirical average infection rate at time $n = 1$. For time $n > 1$, however,

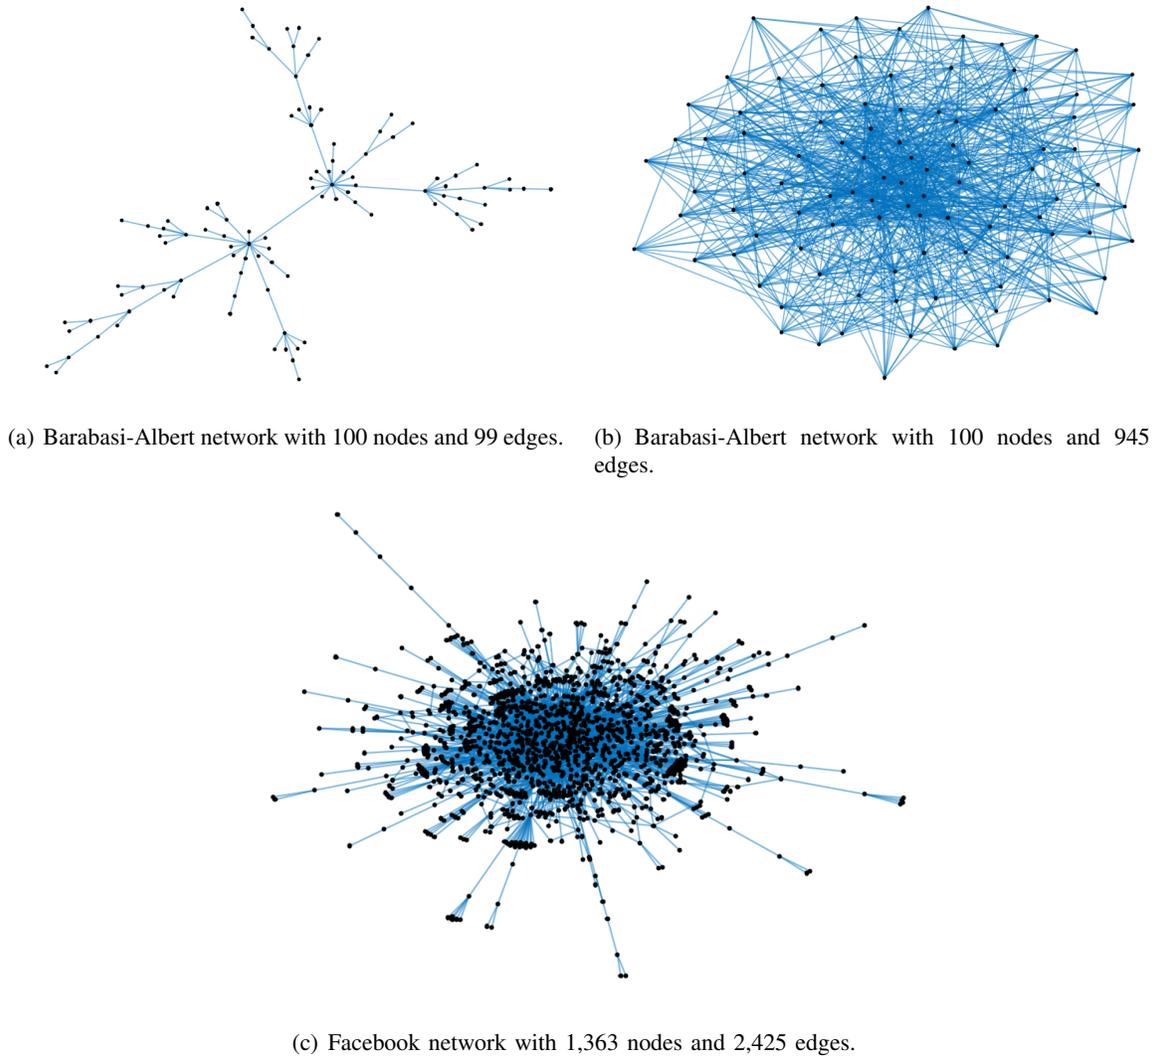


Fig. 2. Mid-to-large scale networks used for simulation purposes. Adjacency matrices for each of these networks is available online at <https://bit.ly/2ygLEqg>.

this is not always the case. Figure 4(c) is of particular interest, as we clearly see that the gradient descent strategy begins with a much lower infection rate, but over time is surpassed by strategy (vi) , which targets a reduced subset of nodes in accordance with Algorithm 2, with resources being focused more around central nodes.

The fact that the best strategy at time $n = 1$ is not optimal as n increases supports our hypothesis that the optimal solution will change for increasing values of n and as $n \rightarrow \infty$. It is important to note, however, that our gradient algorithm merely converges to the optimal solution at time $n = 1$, and thus the solution we have from using this algorithm is not necessarily optimal, as the algorithm has not been run indefinitely. Indeed, we note that had we chosen to either use a finer filter when implementing the limit minimization rule, as described in Algorithm 1, or increased the stopping time, or even chosen a step size that allowed for faster convergence, we would have seen an improved performance using this strategy (at least at time $n = 1$).

Regardless, we consistently see that targeting a reduced subset of nodes, whether via Algorithms 2 or 3, results in performance that is close to or better than that obtained using gradient descent. These strategies, (vi) and $(viii)$, always provide the best results of all the heuristic strategies being compared. This illustrates that we obtain near

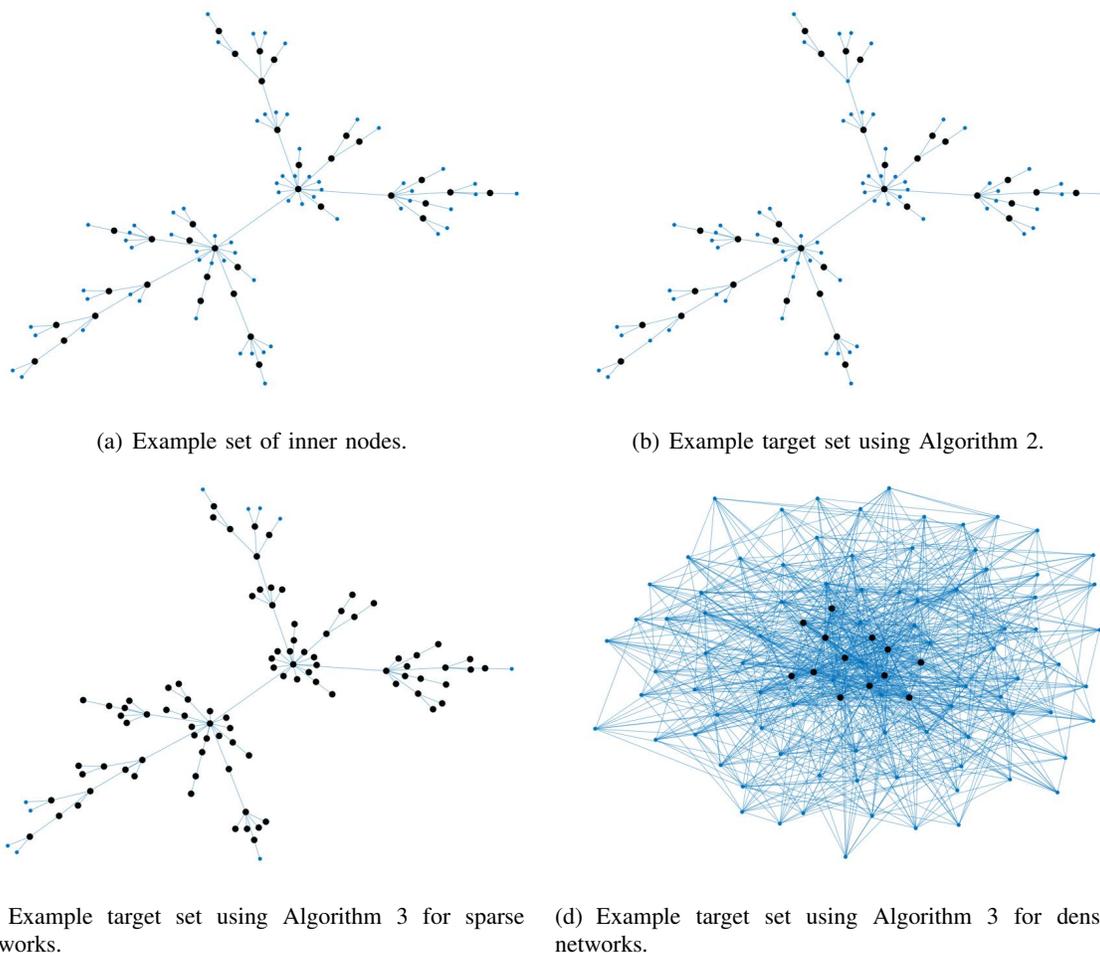


Fig. 3. Comparison of target sets for various node targeting algorithms. In each case, the set of targeted nodes is enlarged and coloured black. It is important to note that Figure 3(c), as depicted here, does not apply the optional component of Algorithm 3. In fact, if we instead choose to apply that option, we end up getting the same target set as shown in Figure 3(b).

optimal performance while only distributing resources amongst a greatly reduced subset of nodes. Performance between the other heuristic strategies varies greatly depending on the network, though we note that even just targeting the inner nodes using strategy *(iii)* sees a significant improvement over the no control case (i.e., the results obtained using a uniform distribution of initialization resources).

We also remark that the general curve for each of these trials remains relatively flat. This is mostly due to the fact that Δ is assigned a fixed value. At a given time for a given node the expected number of red balls added to that particular node's super urn will be proportional to the average super urn proportion of the neighbouring nodes. This causes the super urn proportion of a given node to trend towards the average super urn proportion of its neighbours, which often takes a similar value because these nodes have overlapping sets of neighbours.

We next examine whether these same results carry over to the general curing case, or whether there are any notable differences. We should note that the curing and initialization policies are by nature not identical. They do, however, mirror one another in such a way that we can at least provide a surface-level comparison between the two setups.

TABLE I
INITIALIZATION STRATEGIES

(i)	Constrained gradient descent on a simplex: Find B_i using Algorithm 1 on the function \tilde{I}_1
(ii)	Uniformly allocate the budget to all nodes: $B_i = \frac{\mathcal{B}_b}{N}$
(iii)	Uniformly allocate the budget to inner nodes: $B_i = \begin{cases} \frac{1}{ V' } \mathcal{B}_b & \text{if } i \in V' \\ 0 & \text{else} \end{cases}$, where $V' = \{k \in V \mathcal{N}'_k \not\subseteq \mathcal{N}'_j, \forall j \in V\}$
(iv)	Ratio of degree and closeness centrality (inner): $B_i = \begin{cases} \frac{ \mathcal{N}_i C_i}{\sum_{j \in V'} \mathcal{N}_j C_j} \mathcal{B}_b & \text{if } i \in V' \\ 0 & \text{else} \end{cases}$
(v)	Allocate as in (iii) using V' from Algorithm 2
(vi)	Allocate as in (iv), using V' from Algorithm 2
(vii)	Allocate as in (iii) using V' from Algorithm 3 (without Optional)
(viii)	Allocate as in (iv), using V' from Algorithm 3 (without Optional)
(ix)	Ratio of degree and closeness centrality: $B_i = \frac{ \mathcal{N}_i C_i}{\sum_{j \in V} \mathcal{N}_j C_j} \mathcal{B}_b$

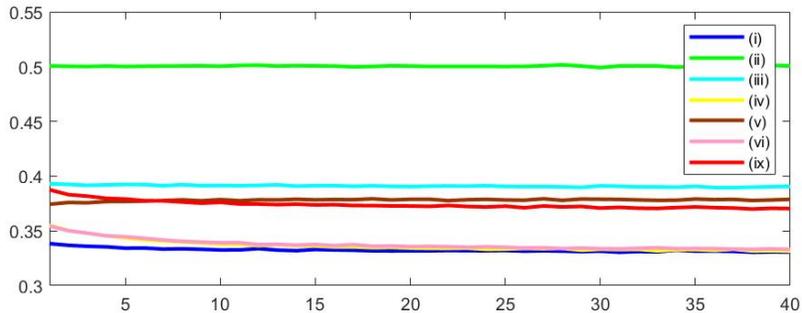
TABLE II
DELTA-CURING STRATEGIES

(i)	Constrained gradient descent on a simplex: Find $\Delta_{b,i}(n)$ using Algorithm 1 on the function $E[\tilde{S}_n \mathcal{F}_{n-1}]$
(ii)	Uniformly allocate the budget to all nodes: $\Delta_{b,i}(n) = \frac{\mathcal{B}_b}{N}$
(iii)	Uniformly allocate the budget to inner nodes: $\Delta_{b,i}(n) = \begin{cases} \frac{1}{ V' } \mathcal{B}_b & \text{if } i \in V' \\ 0 & \text{else} \end{cases}$, where $V' = \{k \in V \mathcal{N}'_k \not\subseteq \mathcal{N}'_j, \forall j \in V\}$
(iv)	Ratio of degree, closeness centrality, and super urn proportion (inner): $\Delta_{b,i}(n) = \begin{cases} \frac{ \mathcal{N}_i C_i S_{i,n-1}}{\sum_{j \in V'} \mathcal{N}_j C_j S_{j,n-1}} \mathcal{B}_b & \text{if } i \in V' \\ 0 & \text{else} \end{cases}$
(v)	Allocate as in (iii) using V' from Algorithm 2
(vi)	Allocate as in (iv), using V' from Algorithm 2
(vii)	Allocate as in (iii) using V' from Algorithm 3 (without Optional)
(viii)	Allocate as in (iv), using V' from Algorithm 3 (without Optional)
(ix)	Ratio of degree, closeness centrality and super urn proportion: $B_i = \frac{ \mathcal{N}_i C_i S_{i,n-1}}{\sum_{j \in V} \mathcal{N}_j C_j S_{j,n-1}} \mathcal{B}_b$

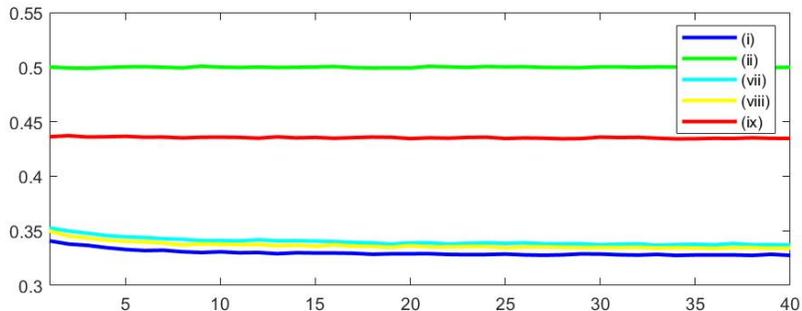
C. Delta-Curing Trials

We consider the one-sided curing problem of trying to find a solution for (17), for some time n . As in the initialization case, we compare the effectiveness of the various curing policies presented in Section V. A brief overview of these policies is found in Table II.

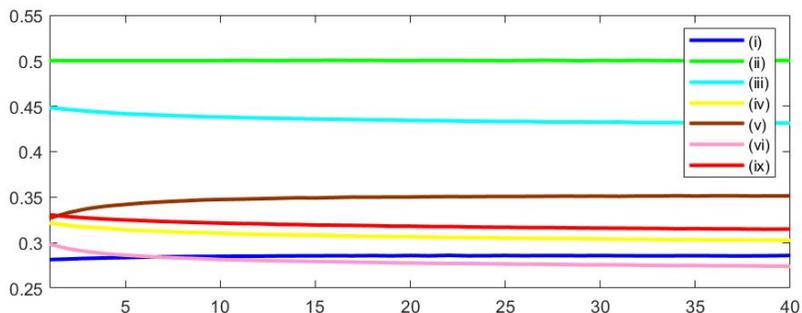
We begin with a uniform initialization, where each node is allocated a fixed number of red balls, R , and black



(a) Results of initialization trials for low density Barabasi-Albert network depicted in 2(a)



(b) Results of initialization trials for high density Barabasi-Albert network depicted in 2(b)

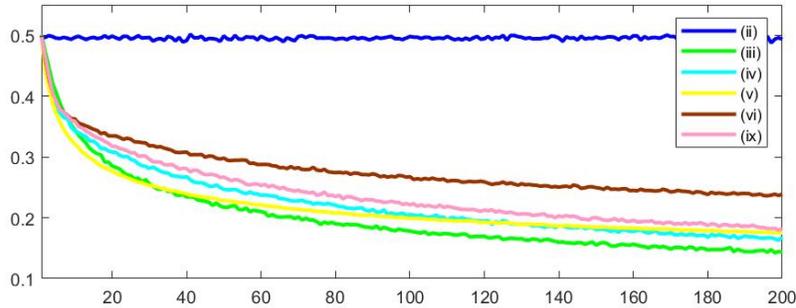


(c) Results of initialization trials for Facebook network depicted in 2(c)

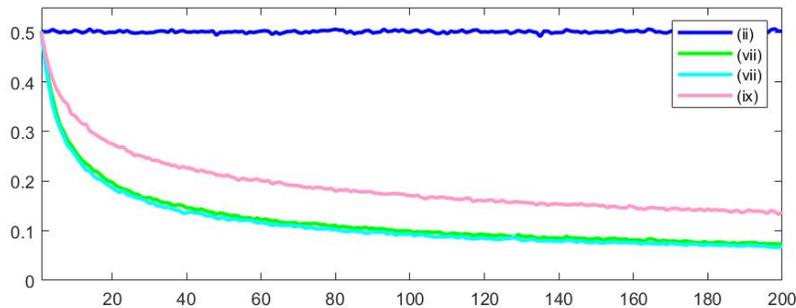
Fig. 4. Plots of empirical average infection rate, \tilde{I}_n , for various initialization strategies across multiple networks. Lower values indicate lower levels of infection. In each case the network was given a fixed initialization budget of $\mathcal{B}_r = \mathcal{B}_b = 10N$, with $R_i = 10 \forall i \in V$, and the Polya network contagion process was run over 1000 trials with a fixed $\Delta = 5$.

balls, B , with $R = B$. We then run the Polya network contagion process. Before each draw is performed, we assign a number of curing resources to each node $\Delta_{b,i}(n)$ in accordance with the desired policy from Table II, while the infection resources are assigned uniformly, $\Delta_{r,i}(n) = \frac{\mathcal{B}_r}{N}$ (with a fixed budget $\mathcal{B}_b = \mathcal{B}_r$ for each time n). As in the initialization case, we record the draw value, $Z_{i,n}$ for each node $i \in V$ at each time n . The draw values are averaged over the number of trials and the set of nodes to obtain a measure for the empirical average infection rate, \tilde{I}_n . The results of these simulations are depicted in Figure 5.

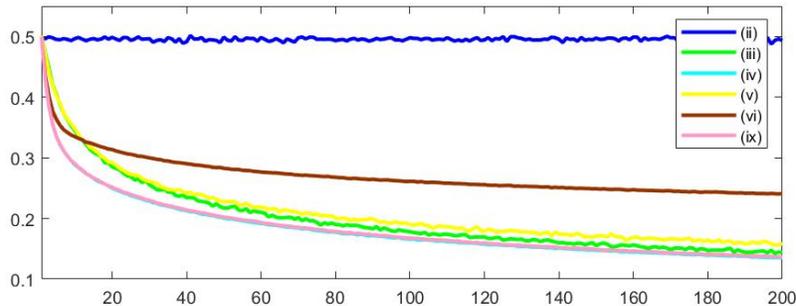
We first remark that unlike the initialization case, targeting a reduced subset of nodes based on centrality does not always lead to the best performance. This is evidenced by both Figures 5(a) and 5(c), where strategy (vi) appears to perform worst of all the heuristics apart from the no control case. In fact, between the three network setups, the heuristics vary wildly in performance, with the best performing heuristic from [38] performing extremely well



(a) Results of curing trials for low density Barabasi-Albert network depicted in 2(a)



(b) Results of curing trials for high density Barabasi-Albert network depicted in 2(b)



(c) Results of curing trials for Facebook network depicted in 2(c)

Fig. 5. Plots of empirical average infection rate, \bar{I}_n , for various curing strategies across multiple networks. Lower values indicate lower levels of infection. In each case the network was initialized with $R_i = B_i = 10, \forall i \in V$, and the Polya network contagion process was run over 500 trials with a fixed budget $\mathcal{B}_r = \mathcal{B}_b = 10N$, with $\Delta_r = 10$ fixed.

in Figure 5(c), but performing significantly worse than the other techniques in 5(b). The various strategies also perform to different degrees of success for different values of n (see strategy (vi) in Figure 5(c)).

Between the three networks, the strategies that seem to perform the most consistently well are strategies that target only the inner set of nodes (strategies (iii) , (iv) and the analogous strategies for denser networks, (vii) and $(viii)$). It is interesting that these strategies would perform most consistently for the Delta-curing problem, but would perform far from optimally for the one-time initialization policy as seen in Figure 4. The main difference between these setups is that the solutions to the Delta-curing problem are applied continually as the infection evolve, while solutions to the initialization problem are applied once. Obviously, some of the benefits of targeting a greatly reduced set of nodes, as in seen in Figure 3(b), are lost when they do not properly account for changes in the urn compositions. Interestingly, for denser networks, targeting a greatly reduced set of nodes as in Figure 3(d) was still

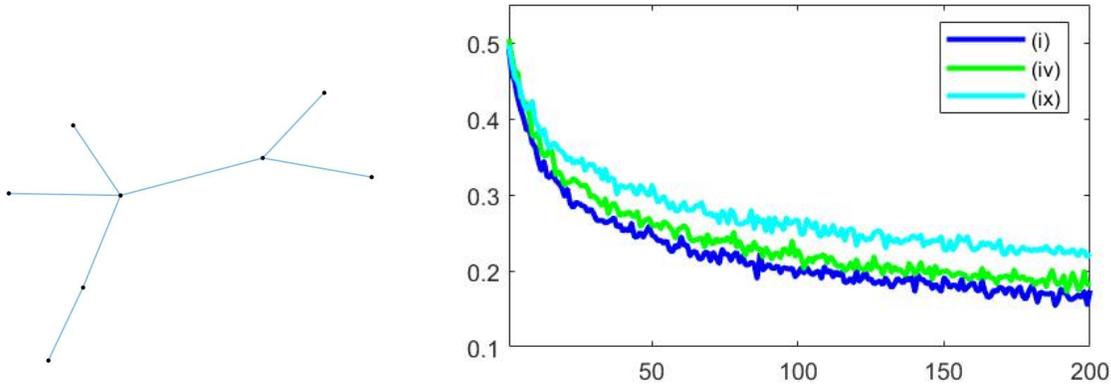


Fig. 6. Empirical average infection rate, \tilde{I}_n , for small 8 node test network. The performance of gradient descent over the expected network exposure, $E[\tilde{S}_n|\mathcal{F}_{n-1}]$, was compared to the two best performing heuristics from Figure 5(a). The network was initialized with $R_i = B_i = 10$, $\forall i \in V$, and the Polya network contagion process was run over 500 trials with a fixed budget $\mathcal{B}_r = \mathcal{B}_b = 80$, with $\Delta_r = 10$ fixed.

very effective, most likely as a direct result of the high network density.

One notable exclusion from Figure 5 is the fact that none of the three subplots show the performance obtained using gradient descent, in part due to the highly complex nature of the expected network exposure $E[\tilde{S}_n|\mathcal{F}_{n-1}]$, which is the function we are minimizing over the curing parameters. While any large scale gradient-descent simulations are somewhat infeasible as a result, it is still possible to provide a comparison on a smaller scale, in order to see how far from ‘optimal’ our various heuristics are.[†] We provide such a comparison in Figure 6. While it is clear that the gradient descent does marginally improve the empirical average infection rate, we see that strategy (iv) still performs comparably, at a fraction of the computational cost.

VII. CONCLUSION

We considered the (infinite-memory) Polya network contagion model and formulated (preemptive and reactive) control policies to mitigate the spread of infection. We framed these problems in two lights: both as a one-time initialization problem, and as a continual control problem that is an extension of contagion curing methods in prior work. We provided simplified metrics for the one-sided finite-horizon version of these problems and proved the existence of optimal policies. We developed effective heuristics that functioned primarily by modifying the scope of solutions to only a limited number of nodes. We then compared these heuristics to the one-step optimal policies that were obtained through use of gradient descent. We established underlying results for the initialization problem proving that optimal initialization strategies would only allocate resources to inner nodes. We showed how network symmetry impacts the allocation of resources for optimal policies. We also demonstrated that the proposed heuristics perform comparably to the one-step optimal policy in certain situations. Finally, we framed the Delta-curing problem as a two-person zero sum game on the expected network exposure, and we proved the existence of a Nash equilibrium. We borrowed ideas from the initialization problems in order to implement heuristics for the Delta-curing problem and demonstrated that these strategies can perform near-optimally at a fraction of the computational expense. Future directions include analogous optimization problems when underlying resource distribution is hidden

[†]As noted earlier, we say ‘optimal’ since the gradient-descent algorithm only converges to an optimal solution for the one-step problem of minimizing $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ over the curing resources.

from the observer, studying the initialization problem in a game-theoretic framework, and investigating the spread of the Polya contagion process in networks with unknown topologies.

ACKNOWLEDGEMENT

The authors sincerely thank Tamás Linder for considerably shortening the proof of Lemma 3.

REFERENCES

- [1] D. Easley and J. Kleinberg, *Networks, Crowds and Markets: Reasoning about a Highly Connected World*. Cambridge Univ. Press, 2010.
- [2] P. V. Mieghem, J. Omic, and R. Kooij, “Virus spread in networks,” *IEEE/ACM Trans. Netw.*, vol. 17, no. 1, pp. 1–14, 2009.
- [3] C. Nowzari, V. M. Preciado, and G. J. Pappas, “Analysis and control of epidemics: A survey of spreading processes on complex networks,” *IEEE Control Systems Magazine*, vol. 36, pp. 26–46, 2016.
- [4] F. Alajaji and T. Fujia, “A communication channel modeled on contagion,” *IEEE Trans. Inf. Theory*, vol. 40, no. 6, pp. 2035–2041, 1994.
- [5] 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 Computing, Behavioral-Cultural Modeling, and Prediction*. Springer, 2014, pp. 85–93.
- [6] M. Garetto, W. Gong, and D. Towsley, “Modeling malware spreading dynamics,” in *Proc. Conf. Comp. Comm.*, vol. 3, 2003, pp. 1869–1879.
- [7] E. M. Rogers, *Diffusion of Innovations*, 5th ed., 2003.
- [8] E. Adar and L. A. Adamic, “Tracking information epidemics in blogspace,” in *Proc. IEEE/WIC/ACM Int. Conf. Web Intelligence*, 2005, pp. 207–214.
- [9] C. Nowzari, V. M. Preciado, and G. J. Pappas, “Optimal resource allocation for control of networked epidemic models,” *IEEE Trans. Cont. Netw. Sys.*, vol. 4, no. 2, pp. 159–169, 2017.
- [10] L. Wang and J. T. Wu, “Characterizing the dynamics underlying global spread of epidemics,” *Nature Comm.*, vol. 9, no. 1, pp. 1–11, 2018.
- [11] N. Madar, T. Kalisky, R. Cohen, D. Ben-Avraham, and S. Havlin, “Immunization and epidemic dynamics in complex networks,” *The European Physical Journal B*, vol. 38, no. 2, pp. 269–276, 2004.
- [12] C. Borgs, J. Chayes, A. Ganesh, and A. Saberi, “How to distribute antidote to control epidemics,” *Rand. Struct. & Alg.*, vol. 37, no. 2, pp. 204–222, 2010.
- [13] M. Khouzani, S. Sarkar, and E. Altman, “Optimal control of epidemic evolution,” 2011, pp. 1683–1691.
- [14] V. M. Preciado, M. Zargham, and D. Sun, “A convex framework to control spreading processes in directed networks,” 2014, pp. 1–6.
- [15] K. Kandhway and J. Kuri, “How to run a campaign: Optimal control of SIS and SIR information epidemics,” *Applied Mathematics and Computation*, vol. 231, pp. 79–92, 2014.
- [16] P. E. Par, J. Liu, C. L. Beck, A. Nedi, and T. Baar, “Multi-competitive viruses over static and time-varying networks,” in *2017 American Control Conference (ACC)*, May 2017, pp. 1685–1690.
- [17] W. Mei and F. Bullo, “Competitive propagation: Models, asymptotic behavior and quality-seeding games,” *IEEE Transactions on Network Science and Engineering*, vol. 4, no. 2, pp. 83–99, 2017.
- [18] Q.-H. Liu, L.-F. Zhong, W. Wang, T. Zhou, and H. E. Stanley, “Interactive social contagions and co-infections on complex networks,” *Chaos*, vol. 28, no. 2, 2018.
- [19] W. Wang, Q.-H. Liu, J. Liang, Y. Hu, and T. Zhou, “Coevolution spreading in complex networks,” *Physics Reports*, vol. 820, pp. 1–51, 2019.
- [20] A. R. Hota and S. Sundaram, “Game-theoretic protection against networked SIS epidemics by human decision-makers,” *IFAC-PapersOnLine*, vol. 51, pp. 145–150, 2019.
- [21] M. Hayhoe, F. Alajaji, and B. Ghahsifard, “A Polya urn-based model for epidemics on networks,” *Proc. 2017 American Cont. Conf.*, 2017.
- [22] ———, “A Polya contagion model for networks,” *IEEE Trans. Cont. Netw. Sys.*, vol. 5, pp. 1998–2010, 2018.
- [23] F. Eggenberger and G. Polya, “Über die Statistik verketteter Vorgänge,” *Z. Angew. Math. Mech.*, vol. 3, no. 4, pp. 279–289, 1923.
- [24] G. Polya and F. Eggenberger, “Sur l’interprétation de certaines courbes de fréquences,” *Comptes Rendus C. R.*, vol. 187, pp. 870–872, 1928.
- [25] G. Polya, “Sur quelques points de la théorie des probabilités,” *Annales de l’institut Henri Poincaré*, vol. 1, no. 2, pp. 117–161, 1930.
- [26] A. Fazeli and A. Jadbabaie, “On consensus in a correlated model of network formation based on a Polya urn process,” 2011.
- [27] M. Chen and M. Kuba, “On generalized Polya urn models,” *Journal of Appl. Probab.*, vol. 50, no. 4, pp. 1169–1186, 2013.
- [28] F. Chung, S. Handjani, and D. Jungreis, “Generalizations of Polya’s urn problem,” *Ann. Combin.*, vol. 7, pp. 141–153, 2003.
- [29] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos, “Epidemic spreading in real networks: An eigenvalue viewpoint,” in *Proceedings 22nd IEEE International Symposium on Reliable Distributed Systems*, 2003, pp. 25–34.
- [30] H. J. Ahn and B. Hassibi, “Global dynamics of epidemic spread over complex networks,” in *52nd IEEE Conference on Decision and Control*, 2013, pp. 4579–4585.
- [31] N. A. Ruhi, H. J. Ahn, and B. Hassibi, “Analysis of exact and approximated epidemic models over complex networks,” *arXiv preprint arXiv:1609.09565*, 2016.

- [32] W. Mei, S. Mohagheghi, S. Zampieri, and F. Bullo, "On the dynamics of deterministic epidemic propagation over networks," *Annual Reviews in Control*, vol. 44, pp. 116–128, 2017.
- [33] M. Hayhoe, F. Alajaji, and B. Ghahsifard, "Curing epidemics on networks using a Polya contagion model," *IEEE/ACM Transactions on Networking*, vol. 27, no. 5, pp. 2085–2097, 2019.
- [34] J. Nash, "Equilibrium points in n -person games," in *Proceedings of the National Academy of Sciences USA*, vol. 36, 1950, pp. 48–49.
- [35] R. Ash and C. Doléans-Dade, *Probability and Measure Theory*, 2000.
- [36] G. R. Grimmett and D. R. Stirzaker, *Probability and Random Processes*, 3rd ed. Oxford Univ. Press, 2001.
- [37] W. Feller, *An Introduction to Probability Theory and its Applications*, 2nd ed. New York: Wiley, 1971, vol. 2.
- [38] M. Hayhoe, "A Polya urn stochastic model for the analysis and control of epidemics on networks," Master's thesis, Queen's University, 2017.
- [39] J. A. Bondy and U. S. R. Murty, *Graph theory with applications*. Macmillan London, 1976, vol. 290.
- [40] N. Biggs, *Algebraic Graph Theory*, 2nd ed. Cambridge University Press, 1994.
- [41] C. D. Godsil and G. F. Royle, *Algebraic Graph Theory*, ser. Graduate Texts in Mathematics. Springer, 2001, vol. 207.
- [42] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA: Athena Scientific, 1999.
- [43] D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar, *Convex Analysis and Optimization*, 1st ed. Belmont, MA: Athena Scientific, 2003.
- [44] B. Rieder, "Studying Facebook via data extraction: the Netvizz application," *Proc. 5th ACM Web Science Conf.*, pp. 346–355, 2013.