

Infection-Curing Games over Polya Contagion Networks

Greg Harrington Fady Alajaji Bahman Gharesifard

Abstract—We investigate infection-curing games on a network epidemics model based on the classical Polya urn scheme that accounts for spatial contagion among neighbouring nodes. We first consider the zero-sum game between competing agents using the cost measure for the average infection in the network. Due to the complexity of this problem we define a game on a proxy measure given by the so-called expected network exposure, and prove the existence of a Nash equilibrium that can be determined numerically using gradient descent algorithms. Finally, a number of simulations are performed on small test networks to provide empirical evidence that a Nash equilibrium exists for games defined on the average network infection.

Index Terms—Network epidemics, Polya contagion, Nash equilibrium.

I. INTRODUCTION

The study of epidemics on networks is an active research topic (e.g., see [1], [2], [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.

In this paper, we examine two player zero-sum games over Polya contagion networks as introduced in [9], [10]. The setup for this model is an adaptation of the classical Polya contagion process [11], [12], [13] to allow for spatial interactions between neighbouring nodes. This model is similar to the well known class of Susceptible-Infected-Susceptible (SIS) epidemics models [1]. Indeed in [10], it was empirically shown that the network Polya contagion model can mimic the behaviour of the SIS process. Furthermore, SIS models have been studied in a game-theoretic context in [14], [15] and references therein. We herein establish a similar problem for the network Polya contagion model.

First, we define a two-player game on the average infection rate for a connected network. This setup is an extension of the contagion curing problem that was studied extensively in [16]. 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 establish 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 Nash equilibrium for the game on the expected network exposure, which can be determined using gradient

descent algorithms. Finally, we run a series of simulations on small test networks, to provide empirical evidence that a Nash equilibrium exists for the infection-curing game on the average infection rate, and to support the claim that such an equilibrium policy can be closely approximated using the equilibrium policy for the expected network exposure.

The rest of the paper is organized as follows. Section II familiarizes the reader with important preliminary material, including a brief description of the classical Polya contagion process as well as the network Polya contagion process. Section III defines the two-player games that are studied in this paper, and proves important convexity/concavity results as well as the existence of a Nash equilibrium for the expected network exposure game. Section IV presents illustrative simulation results. Finally, Section V concludes the paper and highlights a few areas for potential future research.

II. PRELIMINARIES

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 . We represent by $\{\mathcal{F}_n\}_{n=1}^{\infty}$, the natural filtration on the process $\{Z_n\}_{n=1}^{\infty}$. We write the sequence $(Z_{i,1}, \dots, Z_{i,n})$ as Z_i^n , and more generally we use the notation $Z_{i,s}^t = (Z_{i,s}, Z_{i,s+1}, \dots, Z_{i,t})$, where $1 \leq s < t \leq n$. Further clarification of standard probability concepts can be found in various texts such as [17], [18].

A Nash equilibrium is a solution for a non-cooperative multiplayer game from which no player may improve its payoff by unilaterally changing its strategy [19], [20]. In this paper, we consider games with deterministic (pure) strategies. In the context of a two player zero-sum game, denote by \mathcal{X} and \mathcal{Y} the set of allowable strategies for player one and two, respectively. More precisely, consider the case where player one receives a payoff $f(x, y)$, with $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ strategies of player one and two, respectively. Then, $(x^*, y^*) \in \mathcal{X} \times \mathcal{Y}$ is a Nash equilibrium if $\forall x \in \mathcal{X}, \forall y \in \mathcal{Y}$:

$$f(x, y^*) \leq f(x^*, y^*) \leq f(x^*, y).$$

A. Classical Polya Contagion Process

We recall the classical Polya contagion process [11], [12], [13]. An urn contains $R \in \mathbb{Z}_{>0}$ red balls and $B \in \mathbb{Z}_{>0}$ black balls. We use $T = R + B$ to denote the total number of balls in the urn. At each time, n , a ball is drawn from the urn and returned with $\Delta > 0$ balls of the same colour. From an urn sampling point of view the parameter Δ is integer-valued; however, we allow them to be real-valued for mathematical convenience (e.g., when determining optimal policies). The random variable Z_n is used to indicate the colour of the ball

The authors are with the Department of Mathematics and Statistics at Queen's University, Kingston, ON, Canada 12gth@queensu.ca, fady@mast.queensu.ca, bahman@mast.queensu.ca. This work was partially supported by the Natural Sciences and Engineering Research Council of Canada.

on the n th draw, where $Z_n = 1$ if the drawn ball was red and $Z_n = 0$ if the drawn ball was black. We can then use U_n to denote the proportion of red balls in the urn after the n th draw, where

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

where $\rho_c = \frac{R}{T}$ is the initial proportion of red balls in the urn and $\delta_c = \frac{\Delta}{T}$ is a correlation parameter. 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

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

B. Network Polya Contagion Process

We now review the network Polya contagion process introduced in [9], [10]. Let $\mathcal{G} = (V, \mathcal{E})$ be a network, 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, and is assumed to be connected. 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 $R_i \in \mathbb{Z}_{>0}$ red balls and $B_i \in \mathbb{Z}_{>0}$ black balls, with the total number of balls at node i being denoted by $T_i = R_i + B_i$.

We next introduce the concept of super urns for 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. 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 t , if a red ball is drawn we add $\Delta_{r,i}(t)$ red balls to node i 's individual urn and if a black ball is drawn we add $\Delta_{b,i}(t)$ black balls to node i 's individual urn. We use an indicator, $Z_{i,n}$ to denote the colour of the n th draw for node i , where $Z_{i,n} = 1$ if the draw was red and $Z_{i,n} = 0$ if the draw was black. 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 .

For this process, we can represent the proportion of red balls in node i after the n th draw by $U_{i,n}$, where

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

Here

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

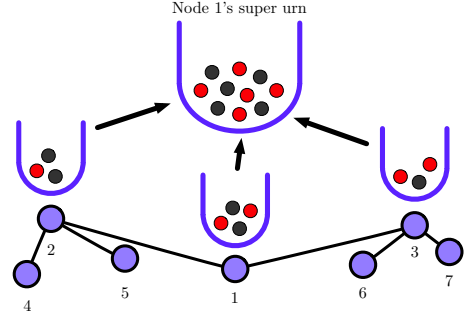


Fig. 1. Illustration of a super urn in a network.

represents the total number of balls in node i 's urn after the n th draw. We represent the proportion of red balls in node i 's super urn after the n th draw as $S_{i,n}$, where

$$S_{i,n} = \frac{\bar{R}_i + \sum_{j \in \mathcal{N}'_i} \sum_{t=1}^n \Delta_{r,j}(t) Z_{j,t}}{\sum_{j \in \mathcal{N}'_i} X_{j,n}}. \quad (2)$$

We take $S_{i,0} = \frac{\bar{R}_i}{\bar{T}_i}$, which is just the initial super urn proportion of red balls for node i . We define

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

which we refer to as the *network exposure*.

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

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

Finally, we refer to

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

as the *average infection rate*.

III. GAME-THEORETIC SETUP

Previously, [16] has looked at the problem of minimizing the limiting average infection rate subject to a budget \mathcal{B} on the curing parameters $\{\Delta_{b,i}(n)\}_{i=1}^N$ at each time step. In this setup, \mathcal{B} is fixed for each time step and all other values are assumed to be given. We now consider a two player game, where player one looks to minimize the average infection rate, while player two looks to maximize this same value. Player one controls the distribution of curing parameters in accordance with a budget, \mathcal{B}_b , while player two controls the distribution of the infection parameters in accordance with a budget, \mathcal{B}_r . In both cases the budget is fixed for all time steps, and for any given time step the allocation of the budget is performed prior to the draw. 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 as follows:

$$\underset{\{\{\Delta_{b,i}(k)\}_{i=1}^N: \sum_{i=1}^N \Delta_{b,i}(k) = \mathcal{B}_b\}, k=1, \dots, t\}}{\text{minimize}} \quad \tilde{I}_t, \quad (3)$$

assuming the minimum exists, while player two's objective takes the form:

$$\underset{\{\{\Delta_{r,i}(k)\}_{i=1}^N: \sum_{i=1}^N \Delta_{r,i}(k) = \mathcal{B}_r\}, k=1, \dots, t\}}{\text{maximize}} \quad \tilde{I}_t, \quad (4)$$

assuming the maximum exists.

For a general network, finding an optimal control policy for either player can be complicated, and the behaviour of \tilde{I}_n is difficult to study. As discussed in [16], there is a strong correlation between the behaviour of \tilde{S}_n and \tilde{I}_n , and thus we can simplify our problem by first considering a game on the expected network exposure for a given time step, rather than the limiting average infection rate.

We consider the expected network exposure $E[\tilde{S}_n | \mathcal{F}_{n-1}]$ as a two player zero-sum game, where player one looks to minimize the expected network exposure over the curing parameters $\{\Delta_{b,i}(n)\}_{i=1}^N$ and player two looks to maximize 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 3.1: (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 (1) and (2), 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 (2) 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, let A represent the adjacency matrix (including self connections) for our given network \mathcal{G} , where

$A_{ij} = 1$ iff $i \in \mathcal{N}'_j$, and $A_{ij} = 0$, otherwise. 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 get $\sigma_i(\mathbf{x}, Z_n) = D_i \mathbf{x}$, where

$$D_i \mathbf{x} = [A_{i1}(1 - Z_{1,n}) \quad \dots \quad A_{iN}(1 - Z_{N,n})] \begin{bmatrix} x_1 \\ \vdots \\ x_N \end{bmatrix}.$$

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 get the following:

$$\delta_i(\mathbf{y}, Z_n) = C_i \mathbf{y} = [A_{i1} Z_{1,n} \quad \dots \quad A_{iN} Z_{N,n}] \begin{bmatrix} y_1 \\ \vdots \\ y_N \end{bmatrix}.$$

Allowing us to alternatively write

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

noting of course that the matrices C and D are functions of Z_n .

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

$$\begin{aligned} E[S_{i,n} | \mathcal{F}_{n-1}] &= E[f_{i,n}(\mathbf{x}, \mathbf{y}, Z_n) | \mathcal{F}_{n-1}] \\ &= \sum_{z_n \in \{0,1\}^N} f_{i,n}(\mathbf{x}, \mathbf{y}, z_n) P(Z_n = z_n | Z^{n-1}). \end{aligned} \quad (5)$$

We can note that $P(Z_n = z_n | Z^{n-1})$ 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 $\mathbb{R}_{\geq 0}^N$. 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 wherein we note

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

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 for 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

$$\frac{\partial}{\partial x_j} D_i \mathbf{x} = A_{ij} (1 - z_{j,n})$$

$$\frac{\partial}{\partial y_j} C_i \mathbf{y} = A_{ij} z_{j,n}.$$

Since A_{ij} and $z_{j,n}$ can only take values in $\{0, 1\}$, we get that $\frac{\partial}{\partial x_i} f_{i,n}(\mathbf{x}, \mathbf{y}, z_n) \leq 0$ and $\frac{\partial}{\partial y_j} f_{i,n}(\mathbf{x}, \mathbf{y}, z_n) \geq 0$. Using this fact in conjunction with equation (5) 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. This lack of saddle point is demonstrated in Figure 2, which depicts the shape of $E[\tilde{S}_n | \mathcal{F}_{n-1}]$ for a single node network. However, given our fixed allocation budgets \mathcal{B}_b and \mathcal{B}_r , we restrict ourselves to considering sets of the form $\mathcal{X} \times \mathcal{Y} = \{\{\Delta_{b,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N | \sum_{i=1}^N \Delta_{b,i}(n) \leq \mathcal{B}_b\} \times \{\{\Delta_{r,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N | \sum_{i=1}^N \Delta_{r,i}(n) \leq \mathcal{B}_r\}$.




Fig. 2. Sample plot of $E[\tilde{S}_n | \mathcal{F}_{n-1}]$ for a single node network. A saddle point will not exist over $\mathbb{R}_{\geq 0}^N \times \mathbb{R}_{\geq 0}^N$ regardless of the network structure.

Returning to our game over the expected network exposure, we can then note that for any given n , the sets $\mathcal{X} = \{\{\Delta_{b,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N | \sum_{i=1}^N \Delta_{b,i}(n) \leq \mathcal{B}_b\}$ and $\mathcal{Y} = \{\{\Delta_{r,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N | \sum_{i=1}^N \Delta_{r,i}(n) \leq \mathcal{B}_r\}$ are convex and compact. This gives rise to the following result:

Theorem 3.2: (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 $\mathcal{X} \times \mathcal{Y} = \{\{\Delta_{b,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N | \sum_{i=1}^N \Delta_{b,i}(n) \leq \mathcal{B}_b\} \times \{\{\Delta_{r,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N | \sum_{i=1}^N \Delta_{r,i}(n) \leq \mathcal{B}_r\}$, 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 [21]. 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 3.2 can be determined numerically using gradient descent algorithms [22]

(as we are optimizing a convex/concave function over a simplex), but for large networks such algorithms can be computationally expensive.

IV. SIMULATIONS

We have proven that there exists a Nash equilibrium when we consider the two player infection-curing game for the expected network exposure $E[\tilde{S}_n | \mathcal{F}_{n-1}]$. This however, does not guarantee that the same holds true when considering a similar setup for the average infection rate as in (3) and (4). We need to check first of all, that we can find policies for each player that mimic the behaviour of a hypothetical equilibrium policy. As well, we want to see if the optimal policies for the network exposure provide a good estimation for our hypothetical equilibrium policy.


We consider three basic networks, as depicted in Figures 3(a), 3(b), and 3(c). Each network is assigned a uniform initial distribution for both red and black balls. A fixed budget of $\mathcal{B}_b = \mathcal{B}_r$ is given for both players at each time step. At any given time, the optimal policy for the infection-curing game over the expected network exposure is calculated for both the curing and infection parameters. The first set of simulations uses this optimal policy for both players. In the second set of simulations only the curing parameters are assigned using this optimal policy, while the infection parameters are distributed uniformly. In the final set of simulations, the infection parameters are assigned using the optimal policy, while the curing parameters are uniformly distributed. At each time n , the draw values $\{Z_{i,n}\}_{i=1}^N$ are averaged over 1000 trials, and then averaged over all nodes to obtain the empirical average infection rate \tilde{I}_n . The plots in Figure 3 display the resulting output for each network plotted against time.

For all three networks, we observe that case 3 generally leads to the highest average infection rate and case 2 results in the lowest average infection rate. While we do not know what the optimal policy for either player would be in the case of \tilde{I}_n , we can see that the behaviour observed suggests that a Nash equilibrium exists. This is less obvious when we consider the network in Figure 3(c). For this network, it makes sense that a hypothetically optimal policy would be close to uniform, and so we observe less distinct separation between each of the cases in Figure 3(f) as compared to Figure 3(d) and Figure 3(e). The results in Figure 3(f) do not, however, preclude the possibility that there exists an equilibrium policy for both players.


In general, the results support the notion that a Nash equilibrium exists for the game over the average network infection, \tilde{I}_n , and that using the solutions obtained from the gradient descent over the expected network exposure may provide a good estimation as to what such an equilibrium policy may look like. However, the desire to guarantee such claims necessitates further follow up, as these simulations provide only a narrow view.

V. CONCLUSIONS AND FUTURE DIRECTIONS


In this paper, we explored two player zero-sum games over Polya contagion networks, verifying analytically the




(a) Seven node line network.




(b) Six node star network.




(c) Six node circular network.



(d) Plot of empirical average infection rate \tilde{I}_n for line network depicted in Figure 3(a).



(e) Plot of empirical average infection rate \tilde{I}_n for star network depicted in Figure 3(b).



(f) Plot of empirical average infection rate \tilde{I}_n for circular network depicted in Figure 3(c).

Fig. 3. Comparison of simulation results for three test networks. Three cases were compared over each network. Case one used the equilibrium policies for the infection-curing game on the expected network exposure. These policies were obtained through gradient descent algorithms [22]. Case two used a uniform distribution for only the infection parameters. Case three used a uniform distribution for only the curing parameters. All simulations were performed with a uniform initial distribution of 10 red and black balls per node. A budget of $\mathcal{B}_r = \mathcal{B}_b = 10N$ was used in all cases, where N is the number of nodes in the network.

existence of a Nash equilibrium for an infection-curing game on the expected network exposure, and providing empirical evidence to support the existence of an equilibrium point for the infection-curing game on the average network infection. Future work may further investigate the existence of such an equilibrium point. Considering variations of this problem may also be of interest; for example, it may be useful to explore the infection-curing game under the restriction that only a single node be provided resources, or to consider the case where the underlying ball distribution is unknown to the participating players. Such problems could be considered for both the expected network exposure and average infection rates. The infection-curing game on the average infection rate can also be further broken down into exploring both finite and infinite horizon cases.

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. Fuja, "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. Koltz, 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*, pp. 85–93, Springer, 2014.
- [6] M. Garetto, W. Gong, and D. Towsley, "Modeling malware spreading dynamics," in *Proc. Conf. Comp. Comm.*, vol. 3, pp. 1869–1879, 2003.
- [7] E. M. Rogers, *Diffusion of Innovations*. 5 ed., 2003.
- [8] E. Adar and L. A. Adamic, "Tracking information epidemics in blogspace," in *Proc. IEEE/WIC/ACM Int. Conf. Web Intelligence*, pp. 207–214, 2005.
- [9] M. Hayhoe, F. Alajaji, and B. Ghahesifard, "A Polya urn-based model for epidemics on networks," *Proc. 2017 American Cont. Conf.*, 2017.
- [10] M. Hayhoe, F. Alajaji, and B. Ghahesifard, "A Polya contagion model for networks," *IEEE Trans. Cont. Netw. Sys.*, vol. 5, pp. 1998–2010, 2018.
- [11] F. Eggenberger and G. Polya, "Über die statistik verketteter vorgänge," *Z. Angew. Math. Mech.*, vol. 3, no. 4, pp. 279–289, 1923.
- [12] G. Polya and F. Eggenberger, "Sur l'interprétation de certaines courbes de fréquences," *Comptes Rendus C. R.*, vol. 187, pp. 870–872, 1928.
- [13] G. Polya, "Sur quelques points de la théorie des probabilités," *Annales de l'institut Henri Poincaré*, vol. 1, no. 2, pp. 117–161, 1930.
- [14] A. R. Hota and S. Sundaram, "Game-theoretic protection against networked sis epidemics by human decision-makers," *IFAC-PapersOnLine*, vol. 51, 03 2017.

- [15] P. E. Paré, J. Liu, C. L. Beck, A. Nedić, and T. Başar, "Multi-competitive viruses over static and time-varying networks," in *2017 American Control Conference (ACC)*, pp. 1685–1690, May 2017.
- [16] M. Hayhoe, F. Alajaji, and B. Ghahsifard, "Curing with the network Polya contagion model," *Proc. 2017 American Cont. Conf.*, 2018.
- [17] R. Ash and C. Doléans-Dade, *Probability and Measure Theory*. 2000.
- [18] G. R. Grimmett and D. R. Stirzaker, *Probability and Random Processes*. Oxford Univ. Press, 3 ed., 2001.
- [19] J. Nash, "Equilibrium points in n -person games," in *Proceedings of the National Academy of Sciences USA*, vol. 36, pp. 48–49, 1950.
- [20] J. Nash, "Non-cooperative games," *The Annals of Mathematical Statistics*, vol. 54, pp. 286–295, 1951.
- [21] D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar, *Convex Analysis and Optimization*. Belmont, MA: Athena Scientific, 1st ed., 2003.
- [22] D. P. Bertsekas, *Nonlinear Programming*. Belmont, MA: Athena Scientific, 2nd ed., 1999.